

Computing Community Consortium Reverse Site Visit 2/8/2012

Review Panel Report

Review Panelists

- Dr. Michelle Effros, California Institute of Technology
- Dr. Jose Fortes, University of Florida
- Dr. John Guttag, Massachusetts Institute of Technology (Chair)
- Dr. Mary Jane Irwin, Penn State University
- Dr. Roy Levin, Microsoft Research
- Dr. Maja Mataric, University of Southern California
- Dr. Muthu Muthukrishnan, Rutgers University

1. Introduction

The Computing Community Consortium (CCC) can be thought of as a highly successful startup. Like many startups, it ended up becoming successful in part because it was flexible enough to change its initial business plan. The CCC did not deliver everything that was promised in the initial proposal, but it delivered enormous value to the computer science community, the NSF, and the Nation. NSF's return on its investment was large.

The current proposal reflects the evolution of the CCC. The investigators have done a careful job of analyzing both the successes and the disappointments of the past few years, and have plausible plans for amplifying the successful activities and rethinking the less successful ones.

Much of the text below is about changes that we think the CCC should consider. Most of the suggestions reflect the panel's admiration of the CCC's past success, our belief in the continuing importance of the CCC's mission, and our hope that the CCC can expand its reach even further.

2. Intellectual Merit

Computation and information science and engineering (CISE) continues to be a rapidly evolving field. Unlike traditional natural sciences, where research is driven by the search for improved understanding of the physical world, research in CISE is mostly driven by an ever-changing landscape of societal needs. As the field gets deeper and wider, and its researchers more specialized, the value of a consortium that can think about the needs of the entire field becomes even greater.

In some fields, e.g., high-energy physics, the community is brought together by the need to share physical infrastructure. Fortunately for society at large, research in CISE has traditionally not been dependent on large infrastructure investments. Even the Internet was relatively inexpensive (compared to, for

example, the Hadron Collider) to get going. The CISE community depends upon an intellectual infrastructure that can help disparate researchers to coalesce around important problems. In the early days of the CS boom, ARPA and a small number of well-funded industrial research laboratories provided this intellectual infrastructure. Today, the CCC is attempting to fulfill that role.

The CCC has organized itself so as to remain agile and capable of responding to opportunities quickly, as they arise. The leadership of the CCC has been able to coordinate information gathering, document generation, and meeting activities as needed in order to meet the specific requests by its various constituencies. This agility is a valued quality in an organization that represents such a large community.

3.0 Specific Achievements to Date

The CCC has been highly productive during its first 3.5 years of existence. The following is a listing of the major accomplishments of the organization, some of which were in its original list of goals, some of which were added as the CCC evolved.

- *Facilitating the mission of the NSF:* The CCC has been highly effective as a bridge between the NSF and the computing research community. It has served as a means of identifying, defining, and bringing the NSF's attention to new areas of interest to the computing research community. It has also been effective in helping with communication between the NSF and other agencies.
- *Providing timely information to the legislative and executive branches:* The CCC leadership has leveraged its connections with key leaders in Congress and in the President's administration in order to serve as a voice for the computing community and to convey the importance of the field in an accessible and compelling manner to that audience. Specifically, the CCC leadership (including the members of the CCC Council) has been able to provide fast feedback to the administration and the Congress in response to requests.
- *Establishing lasting relationships with policy makers:* The CCC has established and maintained strong relationships with staffers in Congress and the Executive Branch. While some researchers may take the time to visit D.C. occasionally, until the CCC there was a lack of a sustained voice that kept a steady stream of visits, contacts, and information flow on behalf of the community at large. The CCC fills this important and unique role effectively.
- *Successful CIFellows program:* To address the economic climate that resulted in a surplus of PhDs in computing fields relative to academic positions, the CCC, working with the CRA and the NSF, established the Computing Innovation Fellows program, a mentoring postdoctoral experience that matched top selected PhDs in computational areas with top faculty across the nation. This program has been highly sought out and continues to address a national need.
- *Bringing researchers together and jump-starting initiatives:* The CCC has provided resources to the research community to develop visioning and road mapping activities. The robotics visioning project was the most measurably

successful of these activities. Henrik Christensen was sponsored by the CCC to hold workshops and to develop a roadmap for robotics. With his leadership, the field was effective at doing so, scaling up its effort and resulting in a cross-agency National Robotics Initiative funded at \$75M in the first year and expected to continue. The CCC also organized a successful workshop in Health IT that has produced new collaborations and discussions toward continued teaming in this area of national priority. An important contributor to the successes of these efforts has been the recognition of the importance of interdisciplinary teams consisting of discipline scientists and computational scientists working together to solve important problems.

- *Leadership training workshop:* With an understanding that the CCC alone cannot carry the field in its communication with the administration, and in recognition of the need to train younger researchers toward filling these roles, the CCC initiated a Leadership in Science Policy Institute, which held a workshop for mid-career researchers interested in connecting with policymakers.
- *Media outreach:* Members of the CCC have used their connections with the media to promote computation research. For example, the New York Times had a series of articles on computing in 2011 that drew heavily on information supplied by the CCC.

4. Sustainability

For the reasons outlined above, we anticipate that the need for a CCC-like organization will persist for many years. Consequently, attention should be devoted to ensuring its long-term sustainability.

4.1. Financial and Administrative Sustainability

From a financial perspective, the CCC is almost entirely dependent on NSF funding. While recognizing the difficulty in finding alternative or complementary sources of funding, it is prudent to address potential long-term financial challenges. While this is not urgent and should not distract the CCC leadership from pursuing the goals of the current proposal, the CCC leadership should start investigating the viability of other sources and models of funding, such as

- other funding agencies: this would reflect enlarged recognition of the CCC as the voice, catalyst, and enabler of the computer research community by those agencies; and
- industry: the CCC's mission can benefit from engaging industry more broadly, including industrial membership/sponsorship of activities.

In addition to providing financial stability, raising funds from these sources has the potential of increasing the "buy-in" of other funding agencies and industry, thus increasing the impact of the CCC.

From an administrative perspective, the CCC's management structure and

operational processes are in place and adequate for both the current stage of the CCC and the make-up of its leadership and management teams. The CCC is administratively housed within the Computing Research Association (CRA), which seems appropriate.

4.2. Leadership Sustainability

Like most startups, the success of the CCC has been built on the backs of its founders. It is too soon for those founders to disengage. Success over the next few years will continue to depend upon their extensive involvement, but it is not too soon to start planning for a succession.

The CCC recognizes the need for succession planning for the present leadership, especially Ed Lazowska. The leadership team expects that a successor for the role of Chair will have been identified within two years' time and that in three years' time the transition to a new Chair will occur. At that time, Lazowska will move into a "past chair" role for an unspecified, but probably significant, period of time. We believe this is a reasonable plan, and that the early identification and "grooming" of a successor is an essential and challenging endeavor. We agree that the "future chair, present chair, past chair" model can be an effective way to effect orderly succession, although it is not the only way. We expect that the CCC will approach the succession process thoughtfully and choose an appropriate mechanism. We also recognize that the appropriate process for this first transition might not be appropriate for subsequent transitions.

We do have some concerns about whether the current leadership structure is sustainable over the long run. The current leadership has been "heroic," but an organization that depends upon heroes is rarely sustainable over the long haul. Consequently, we recommend that:

- The CCC should always have a crisp and clear statement of its vision and mission that should inform its organizational structure and outcomes assessment.
- The CCC should define specific portfolios and expectations for Council members and make them transparent to the community. The web site and communications of the CCC should specify what areas/communities each council member purports to represent, thereby also clarifying what types of service that individual may provide to members of the community.
- The CCC should consider, where possible, a codification of procedures and networks used for outreach to communities and stakeholders, and for engagement with policymakers, agencies and communities. This would reduce the impact of inevitable leadership transitions on the ability of CCC to fulfill its mission.
- The CCC does not have a formal executive committee, but there is clearly a "kitchen cabinet" that has played an important role. We recommend that the kitchen cabinet be extended to include individuals being considered as possible successors to the current Chair and Vice Chair.

No matter how talented and deep the leadership, the CCC will not be sustainable without a broad and deep engagement with the computing research community and with government agencies. The CCC thus needs to cast a broader net when recruiting individuals for CCC activities. The CCC cannot rely entirely on Council members, since there will never be enough of them and they cannot possibly represent all relevant CCC communities. We recommend that the CCC explore several avenues to identify Council members and find a way to free up enough of their time so that they can devote a sustained effort to advancing the mission of the CCC. Perhaps funds could be found to “buy” people out of teaching or to support sabbaticals devoted to serving the nation or the CCC community. Such activities might be supported with funds acquired by the CCC, possibly through a supplemental proposal to NSF. It is also possible that supportive industrial research organizations that already fund faculty for sabbatical visits would be willing to “donate” part of the sponsored sabbatical time for CCC activities.

5. Broadening the Reach of the CCC

While we believe that the CCC would continue to contribute considerable value were it to proceed along its current trajectory, we believe that it can and should broaden its reach, as follows.

5.1. Outreach to Federal Government Agencies and Industry

The CCC has done an excellent job of developing strong connections with NSF/CISE. The CCC leadership has been highly responsive to requests from the CISE AD at the top and from individual Program Directors in the trenches. The CCC has also developed some connections with DARPA and is working on developing connections with DoE(nergy), DHS, and NIH. We recognize that developing such two-way relationships is a time-consuming effort, and as leadership in the agencies changes these connections often have to be rebuilt. Thus, it is understandable that the list of federal government agencies with which the CCC has developed strong ties is not longer. However,

- It is important that the CCC build ties to potential IT funding agencies beyond just NSF/CISE and DARPA.

The CCC should be strategic in building connections within the “NITRD tent” agencies that align with major CCC initiatives, e.g., DoE(ducation) in Learning Technologies, NSA in Security/Privacy, ARPA-E in Energy/Sustainability. There are also important agencies outside the “NITRD tent” that need (whether they know it or not) to become more engaged with computing research. For example, the FDA, Commerce, and Treasury are all grappling with “big data” issues.

- It is in the national interest for the CCC to raise awareness of the contributions that advancements in computing can bring to bear on the problems faced by these agencies, even if no funding flows to the community.

The CCC already has good connections with Microsoft Research, Intel, IBM/Watson, and Google. The CCC has done less well in connecting with less R&D-oriented IT companies, and even less well building connections to companies that, while not typically thought of as “tech” companies, depend heavily upon advances in the technical areas covered by the CCC. Such industries include finance, pharmaceutical, healthcare, energy, and automotive. We recommend that the CCC

- Invest more effort in establishing strong connections between both tech and non-tech industries and the CCC community.

The CCC workshops, which have done a good job getting representation from government agencies and industry, provide a mechanism for opening dialogs, but more follow-up would be useful. A step in this direction would be constructing a well-organized database to facilitate further connections.

5.2. Outreach to the CISE Community

For all the good that it has done, the CCC is not well known within the community of CISE researchers. It is time to do more to brand the CCC and showcase successes so as to further motivate the community to participate and engage with the CCC.

The field of computation/information is large and amorphous. With its growing success, it is permeating increasingly vast areas of study, and therefore its borders are difficult to draw. For the purposes of the CCC, which aims to be a voice of a community, it is important to be clear about what that community is. Currently, it is referred to as “the computing community” but in practice this community is relatively narrow. We are concerned that there are areas served by CISE with which the CCC has little involvement, and therefore

- We suggest that the CCC make more of an effort to engage what one might call the “information science” and “computer engineering” parts of CISE. These areas include information theory, signal processing, communications and other areas that are typically housed in electrical engineering rather than in computer science departments.

There is a natural tension between wanting to serve all relevant constituencies and being effective at the service provided. The CCC needs to both better define its scope and better communicate with its selected constituency. This means clearer messaging on the web site, more inclusive communications, and broader representation in the leadership (i.e., the council, subcommittees, etc.).

Even researchers whose community has been well represented by the CCC are often unaware of the impact of the CCC. The CCC should make an effort to do a better job of spreading its message. It might, for example,

- Identify a CCC representative at each university engaged in relevant research. This individual could carry the CCC message to their colleagues, and provide the CCC with additional insight into what is going

on at the local level.

- Acquire a database of past and present CISE PIs and use that list to communicate directly with the CISE community.

The CCC has enjoyed success in reaching out to the community in a series of specialized workshops. We encourage this. The CCC has also reached out through “visioning” sessions at conferences. These have met with mixed success and should perhaps be rethought and better targeted. In contrast, the Leadership in Science Policy Institute was enormously successful and should be continued in its current form and widely publicized in advance.

5.3. Outreach to University Leadership

Many universities are currently coming to grips with the role of computational and information-related research on their campuses. There is wide-spread appreciation of the need to hire faculty who are comfortable using computational techniques in their research. However, there is ample confusion about the difference between users of computation and true computer and information scientists.

- We suggest that the CCC work with the CISE leadership to develop mechanisms for bringing the appropriate message to university deans, provosts, and presidents, to ensure the growth and recognition of the field, rather than its dilution.

6. Conclusion

After thoroughly reviewing the proposal, meeting privately with the CISE Director and other NSF staffers, and engaging in a spirited discussion with the proposed principal investigators, the review panel recommends, in the strongest possible terms, that the Computing Community Consortium (CCC) proposal be funded in full.

To manage the expanded goals listed above, it will be critical to set up specific milestones to address the tension between the expansion of breadth (broader community, more federal agencies) and the need for measurable impact, and to prioritize so as to maximize impact given the available resources.

Our major concern is that the requested funds may be insufficient for the CCC to realize its full potential. Buying substantial time for the “volunteer” leadership and having sufficient permanent staff are of critical importance. We encourage the CCC to seek supplemental funding (from the NSF or elsewhere) as needed, to fund specific initiatives.