

Computing Community Consortium Response to the Feb. 8 Reverse Site Visit Review Panel Report

Version 7, March 10, 2012

Introduction:

We thank the members of the review panel for the time and effort they invested in assessing the Computing Community Consortium (CCC) in the context of its renewal proposal.

We are pleased with the panel's overall judgment of the CCC (page 1 of the panel's report):

“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.”

We further concur that the intellectual merit of the CCC resides in the “intellectual infrastructure” that it is providing to the computing research community:

“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... 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)... 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.”

We embrace the report and underlying recommendations, and we provide a point-by-point response below, as appropriate.

Sustainability:

In Section 4 of the report, the panel finds that “the need for a CCC-like organization will persist for many years,” and, “consequently, attention should be devoted to ensuring its long-term sustainability.” On this point, we could not agree more. The CCC serves as a catalyst and enabler for the computing research community, and increasingly is bringing the community together to envision future research needs and thrusts in a way that will help the field continue to be visionary and impactful. We therefore regard institutionalizing the CCC as an extremely high priority as we continue to transition from “start-up” phase to a more mature organization. The review panel has provided several suggestions for aiding in this process, and we will address these below.

Financial and Administrative Sustainability:

In terms of financial sustainability, the panel suggests, “The CCC leadership should start investigating the viability of other sources and models of funding, such as other funding agencies ... and industry.” We recognize the importance of financial sustainability for the CCC, but we also strongly agree with the panel’s finding that this “should not distract the CCC leadership from pursuing the goals of the current proposal.” As we approach Federal agencies and our colleagues in industry, there will be opportunities for co-funding of workshops and other activities, as there have been in the past. [For example, the CCC’s workshop on Discovery and Innovation in Health Information Technology, held in October 2009, was supported by multiple agencies beyond NSF, and the recent symposium on the Federal Networking and Information Technology Research and Development (NITRD) Program was funded by a crosscut of the agencies participating in NITRD.] In other words, we believe we should be opportunistic in our fundraising.” We realize this means we must continue to be cognizant about being funded principally by NSF but being an advocate for the computing research community and a link between that community and a wide range of other Federal agencies.

With respect to the CCC’s relationship with the Computing Research Association (CRA), we agree that it remains appropriate for the CCC to be administratively housed within the CRA. Members of the CCC Council and CRA Board are increasingly comfortable with this arrangement.

Leadership Sustainability:

We concur with the panel’s assessment that succession planning is critical for the CCC’s longevity. Any stable organization must have a well-defined leadership transition plan, as well as a succession plan for its current leadership. Moreover, any organization

benefits from constant intellectual renewal and diversity, including among its leadership. Most importantly, as the panel notes, “The current leadership has been ‘heroic’, but an organization that depends upon heroes is rarely sustainable over the long haul.” We agree, and as part of the planning that we will undertake, we will consider alternative operational structures to ensure that the success and vitality of the CCC is not dependent upon a small set of individuals.

To begin addressing this issue, we have, with the most recent Council rotation, become more deliberate about recruiting potential future leadership candidates.

We further recognize the importance of the panel’s recommendations:

“The CCC should always have a crisp and clear statement of its vision and mission that should inform its organizational structure and outcomes assessment.”

We presented a revised version of our mission statement to the reverse site visit panel:

- CCC is a catalyst and enabler for the computing research community by:
- Bringing the community together to contribute to shaping the future of the field;
 - Providing leadership for the community, encouraging revolutionary, high-impact research;
 - Encouraging the alignment of computing research with pressing national priorities and national challenges (many of which cross disciplines);
 - Working with policy makers to facilitate the translation of these important research directions into funded programs;
 - Giving voice to the community, communicating to a broad audience the many ways in which advances in computing will create a brighter future; and
 - Growing new leaders for the computing research community.

We realize this statement needs to be made more succinct, and it will evolve over time.

Importantly, as noted previously, while NSF principally funds the CCC, and while this is structured as a Cooperative Agreement, we carry out our mission on behalf of the computing research community and therefore must exhibit some level of independence from NSF.

The panel recommended:

“The CCC should define specific portfolios and expectations for Council members and make them transparent to the community...”

In the past year, the CCC has established subcommittees and task forces to (a) improve the structure and operation of the CCC organizationally (subcommittees) and (b) advance specific areas of interest for the field (task forces). A key goal of these subcommittees

and task forces has been to define specific portfolios and expectations for Council members. Subcommittees include Visioning; Government and Public Outreach; Research Community and Student Outreach; and Council Membership. Task forces include Health Information Technology; Computational Sustainability; and Big Data.

Whereas a subcommittee is expected to last for extended periods of time or even be permanent as it covers a key function or processes of or relating to the CCC, we expect a task force will be successfully disbanded once it achieves a majority of the goals articulated in its charter. A task force may exist for as little as a few months. Similarly, new task forces will be created, particularly as new topics or interests arise for the CCC.

We will continue to build upon (and expand) this model for defining portfolios and expectations for Council members, and we will publicize the subcommittees and task forces – something that we have simply neglected to do until now.

In addition, we are considering implementing a short “orientation session” for new Council members, similar to a session that precedes each CRA Board meeting and serves to bring incoming Board members up to speed about responsibilities, expectations, and so on.

“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.”

We agree with this recommendation, and will begin codifying our procedures and networks. This is critical to institutionalizing the CCC. Of course, certain procedures and networks will naturally evolve over time.

“We recommend that the ‘kitchen cabinet’ be extended to include individuals being considered as possible successors to the current Chair and Vice Chair.”

Our consciousness has been raised to this issue, and we intend to explore alternate organizational and operational structures that achieve the right balance between the loose engagement of many and the tight engagement of a few. We agree that this matter is of significant importance to the continued success of the activity.

“The CCC ... 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 have always strived to cast a broad net via our visioning activities, white papers, and other activities (e.g., the Leadership in Science Policy Institute, the Library of Congress Symposium, and the NITRD Symposium). Participation of non-Council members in these activities – both as members of organizing/steering committees and activity

participants – has been critical to success, and has helped grow members of our community into leaders (e.g., Henrik Christensen with the robotics visioning activity).

“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.”

In addition to the collection of activities noted above, improved communications and outreach to the community (e.g., direct communication with CISE PIs, as described below) will help us identify future Council members.

We have also requested additional staff support to further advance the mission of the CCC.

We are, however, concerned about the prospect of “buying” the time of Council members or others (beyond the current CCC leadership), either to “buy out” their teaching assignments or “support” sabbaticals devoted to CCC activities. The CCC Council Chair and Vice Chair are currently supported at 50 percent time and 25 percent time, respectively, and we believe extending financial support beyond these positions creates a dangerous precedent that will detract from the CCC’s long-term fiscal sustainability. (Moreover, as noted previously, other structural models that may reduce the Chair and Vice Chair time may be necessary as we work to institutionalize the organization in the coming years.) Instead, we will work to identify other ways for incentivizing engagement of Council members and others.

Broadening the Reach of the CCC:

In Section 5 of the report, the panel notes that the CCC “can and should broaden its reach” to Federal agencies, industry, the computing research community, and university leadership. We believe communicating with each of these stakeholders falls within the purview of our mission, and we address each of the panel’s recommendations below. As a reminder, we have proposed to bring on a communications specialist who will work to streamline and optimize our communications so that we can maximize our effectiveness – delivering the right content to the right stakeholders.

Outreach to Federal Government Agencies and Industry:

We completely agree with the panel’s assessment that “developing two-way relationships [with agency officials] is a time-consuming effort, and as leadership in the agencies changes these connections often have to be rebuilt.” As the panel noted, we have made significant strides in the past two years in our efforts to go beyond NSF and DARPA. For example, our Health IT Task Force has forged relationships with key officials at the National Library of Medicine, Office of Behavioral and Social Sciences Research, and National Cancer Institute within the NIH. The CCC director has delivered a presentation

to the NITRD Health IT Senior Steering Group. And we are cementing relationships with the Department of Energy.

However, we also agree we must continue to be strategic in “[building] ties to [other] potential IT funding agencies” within the Federal NITRD Program. To this end, we have begun working with prospective allies in the Department of Education on the basis of our Learning Technologies visioning activity, and we are including ARPA-E as part of our outreach in energy and sustainability.

We welcome the panel’s suggestion to reach out to NSA in security/privacy – at our most recent Council meeting, we have begun discussions about possible visioning in the privacy realm – as well as important Federal agencies that may not be within the NITRD Program but could nevertheless result in mutually beneficial connections with the computing research community. The CCC director has recently held meetings with colleagues at the FTC and FCC (on privacy), and there are opportunities to expand further to include the FDA (on medical devices) as well as the Commerce and Treasury Departments (on big data/computation for financial systems). Among other things, we agree we can attempt to influence government agency hiring of information technology experts that complement existing expertise and improve decision making, and also make our community fully understand these agencies’ computing-related challenges.

As for industry, we have Council members from traditional “tech” companies but agree we could reach out to “companies that, while not typically thought of as ‘tech’ companies, depend upon advances in the technical areas covered by the CCC.” As with other outreach efforts, we must be opportunistic here.

Finally, when SRI International conducted its assessment of the CCC, we provided SRI with a list of participants of all CCC activities. We will resurrect that list to “[construct] a well-organized database” that allows us to track and facilitate further connections with prospective industry partners – and to aid with our outreach to the CISE community more broadly.

Outreach to the CISE Community:

The panel offers several recommendations here:

“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.).”

Our conversation with the reverse site visit panel and subsequent discussions among CCC Council members has prompted us to take a look at our current scope. We have always thought of ourselves as focused on computing research broadly construed. We do not consciously eliminate any specific area(s), but we are opportunistic – responsive to agency requests, visioning proposals from the community, etc.

“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.”

This was a terrific suggestion by the panel – something that simply did not occur to us – and we will move forward implementing it. We have an aggregate network of nearly 200 people comprising current and former Computing Innovation Fellows (CIFellows), former Council members, and graduates of our recent Leadership in Science Policy Institute. One possibility would be to sign up these individuals and encourage them to represent the CCC to their home institutions. We will work to ensure these positions are not burdensome in the time or effort required.

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

We have already taken steps to acquire a list of NSF/CISE PIs, scraping 3,000 names of active PIs from the NSF website.

“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.”

At our most recent Council meeting, we discussed specific features that have contributed to successful challenges and visions tracks in hopes of replicating these in the future. In addition, we are working to identify “whitespace areas” that have not yet been explored with sessions to date, and we will pursue these at the appropriate research conferences.

“...[The] Leadership in Science Policy Institute was enormously successful and should be continued in its current form and widely publicized in advance.”

Indeed, we plan to continue and expand the Leadership in Science Policy Institute in the years ahead.

Beyond the panel’s recommendations, there are other communications channels that we will explore, e.g., a regular column in *Communications of the CACM* and/or *IEEE’s Spectrum*, a greater presence at non-CS conferences, etc.

Outreach to University Leadership:

In Section 5.3 of the report, the panel suggests:

“Many universities are currently coming to grips with the role of computational and information-related research on their campuses... 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...”

We have not previously targeted university leadership as part of our outreach, but we will take this recommendation under serious consideration.

Conclusion:

In sum, we were delighted with the panel’s conclusion (page 7 of the report):

“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.”

We take seriously the concern about “scaling.” The reverse site visit panel concludes (as did SRI International in its evaluation a year ago) that a CCC-like activity is necessary and that it should be both broad and deep. We must – and will – think carefully about alternative operational structures that will allow us to continue to scale while continuing to remain nimble and cost-effective.

In general, the proposal writing and review process over the past year was extremely beneficial to the CCC – both in honing the CCC’s own thinking, and in allowing it to benefit from the thinking of other leaders in the community. We appreciate the efforts of the review panel and the counsel of our colleagues at NSF and CRA during this time. We

are reenergized by the panel's findings, and we look forward to continuing the mission of the CCC on behalf of the computing research community.