

Computing Community Consortium

Reverse Site Visit, April 3, 2017

Review Panel Report

Executive Summary

This report documents the outcome of a Reverse Site Visit held for the Computing Community Consortium (CCC) at the National Science Foundation on April 3, 2017. The review panel had previously read the CCC's proposal and associated material. At the Reverse Site Visit, CCC leadership presented its proposal, then engaged in discussion with the panel and responded to the panel's questions.

The CCC's central goal is to act as a catalyst for the computing research community, bringing the community together to envision audacious future research challenges, and to see these ideas developed as funded research programs. Its successes to date have ranged from the National Robotics Initiative to Aging In Place and the BRAIN initiative. The CCC's current leadership has been very strong, and the CCC has done a good job of transitioning from its "start-up" mode to an ongoing organization.

Some aspects of the CCC's external environment are clearly changing, and the leadership of CCC is doing a commendable job of beginning to adapt to the uncertainties associated with these changes. The panel is glad to see the leadership's plans for increased engagement with industry, and initial thoughts for engagement with foundations and non-profits.

While the CCC does face challenges, the panel believes it is well prepared to deal with them. Our recommendations, therefore, provide suggestions for further strengthening the CCC as it goes forward. Detailed recommendations are given in subsequent sections, but may be summarized at a high level as follows.

Summary recommendations:

- **Select topics more strategically.** We suggest that CCC think strategically to seek a more balanced selection of task force topics. Recent task forces have focused on application-driven (top-down) topics; this should be balanced with more topics dealing with emerging technologies (that is, bottom-up thinking), which will require deeper and broader engagement with the research community and industry.
- **Increase engagement with professional societies.** As the CCC proposal noted, the CCC still has less visibility within the research community than could be desired, and the same is true for industry and foundations as well. We encourage the CCC to go beyond blogs and web pages, currently its primary means of engagement, to increase its visibilities in these communities. For example, strong partnerships with professional societies (ACM, IEEE, etc.) and a diverse set of special interest groups could significantly

increase the visibility of the CCC itself, its calls for visioning workshops, and the reports and white papers.

- **Broaden participation.** The CCC must continue its efforts to broaden active participation from the research community, as well as ramp up participation from industry and foundations. The CCC's current organizational structure appears sound and does not need modification, but we suggest that augmenting it with an Advisory Board might be a suitable means for broadening participation in CCC thinking. Champions from industry will be particularly important going forward.

CCC Goals, Strategies and Priorities

The CCC's central goal is to act as a catalyst for the computing research community, bringing the community together to envision audacious future research challenges, and to see these ideas developed as funded research programs. It does so in a highly agile and flexible way that has proven very good at fast responses to requests and ideas as they arise, with successes to date ranging from the National Robotics Initiative to Aging In Place and the BRAIN initiative. The panel commends the CCC leadership, staff, and participants on these successes.

That said, some aspects of the CCC's external environment are clearly changing. The interests and objectives of the US Government continually evolve, and the CCC's interactions with industry, and perhaps non-profit foundations, may well grow considerably. In our review, the panel has discussed with the CCC leadership its plans for adjusting its strategies accordingly. There are many challenges here, but the panel was happy to see that the CCC has recognized them and is considering how it will adapt in response.

Another important issue going forward involves CCC strategies for picking topic areas for new visioning activities and workshops. Recent task forces have focused on application-driven (top-down) topics. While there have been some workshops with a technology focus (e.g., nanotechnology), the resulting topics have generally been application-centric and covered a limited number of areas. Identifying forward-looking ideas across a diverse set of areas, including topics dealing with emerging technologies, will require a deeper engagement with both the broader research community and industry.

Recommendations:

- The panel is glad to see the leadership's plans for increased engagement with industry. This is an important direction and is a good strategy to enable the CCC to continue to increase its impact over the next four years. The panel recommends that CCC might usefully engage with industry by targeting a set of specific interest domains. This might be via industrial consortia for pre-competitive activities, researcher access to advanced industry platforms and data, or joint sponsorship or visioning efforts.

- Similarly, the panel recommends that the CCC thinks strategically to seek a more balanced selection of task force topics. The panel recommends that the CCC complements application-driven (top-down) topics with technology-driven topics (bottom-up), leading to a more diversified portfolio. Reaching out to research and industry leadership in a wide range of areas may help to seed such visioning activities. These approaches may enhance the CCC's ability to be more proactive and audacious in selecting visioning topics.

CCC Leadership and Succession Plan

The CCC's current leadership has been very strong, and the CCC has done a good job of transitioning from its "start-up" mode to an ongoing organization. The panel appreciated the Chair's openness and transparency in discussions during the Reverse Site Visit. The succession plans for the Chair and Vice Chair which have been developed since the last RSV are clear and reasonable. Going forward, it will be important for the future leadership to continue to adapt to conditions and to be flexible.

Given the small size of the CCC staff, both planned and unplanned changes in senior personnel can have a significant impact. In order ensure a sustainable organization, we encourage the CCC to document procedures and lessons learned to train new key personnel, minimizing the impact of changes in personnel.

The leadership of CCC has been doing a commendable job of planning appropriate adaptations in response to shifting government priorities, for example through increased engagement with industry and consortia. The panel also commends the willingness of the CCC leadership to experiment with changes in the CCC organization (such as recruiting non-council members for task forces). Lastly, the description of the coordination between CRA and CCC gives evidence of agility that will stand CCC in good stead as it evolves.

Recommendation:

- The panel agrees with the CCC leadership that creation of an Advisory Board to advise the CCC leadership is a good idea. The proposed Advisory Board will broaden the community input that the leadership receives, serve as a sounding board for envisioned changes, provide breadth in area coverage, provide additional ideas for effective evolution, and help promote the CCC externally. An optimal Advisory Board will include industry members from key companies or domains as well as people knowledgeable on science policy.

CCC Communication and Engagement Strategy

The CCC produces white papers, workshop reports, videos and other supporting documents and makes them publicly available on its web page. These reports aim to be the voice of the

community and are based on detailed notes from the workshops that the CCC runs. They have had an impact on many government agencies and on informing new initiatives. The CCC also reaches out to the community through partnering with various professional organizations to run Blue Sky sessions in some conferences, reaching out through council members and the core CCC leadership.

The committee noted that the quality of the workshop reports is highly variable. While some are excellent and comprehensive (like the Brain) others are narrow and lack a clear message. This inconsistency in quality hurts the CCC brand and can significantly reduce the impact of the reports. A related concern is that reports seem to have different goals (e.g., workshop summary versus defining a research agenda), so they may target a different audience and require a different format, but the Workshop Reports web page does not make this clear. We encourage the CCC to improve quality control so that reports are consistently of high quality, both in terms of presentation and content, in a way that is consistent with their purpose.

The panel felt that it might be possible to reach a broader audience by making public videos of talks from the visioning workshops. We understand that since the workshops are quite open-ended, full recordings may not be ideal or may significantly hurt an open exchange among attendees. We therefore encourage the CCC to try to structure workshops so that at least have some part that can be shared with the research community, and then recording these parts and making them public and visible. We believe this will also help in making CCC's efforts more visible to the general research community.

The committee liked many aspects of CCC's current communication strategy. The white papers seem to be valuable in many cases; examples like the Robotics and BRAIN white paper were cited as being influential. The CCC is currently responding to requests for white papers in targeted areas from agencies to help the agencies make more informed decisions and policies. These efforts are valuable and should continue in the future.

However, the panel agreed with survey respondents quoted in the proposal that CCC is not particularly visible to the members of the computing community, both in research and industry. This concern should be carefully assessed. While the CCC blog and web pages seem to be effective in engaging the agencies, there are clear indications that they are much less effective in engaging other communities. The committee has several recommendations to go beyond the current communications to broaden engagement with the research community and also industry. Achieving those objectives may require additional staff resources, focusing specifically on actively engaging the research community and its professional societies, in addition to pure "communications" functions such as blogging and posting reports on the web.

Recommendations:

- The panel felt that the visibility and recognition of the CCC within the academic and research community is limited. We recommend that greater efforts be made to reach out to the research community by increasing (a) outreach to a wider range of universities, (b) expanding the reach to more research areas than is covered by the

current task forces and council members, and (c) proactively reaching out to conferences and SIGs, and (d) industry.

- CCC should aim for broader dissemination of its reports and visioning workshop calls. The CCC path of sending a report to appear as a CACM viewpoint is a good step in that direction. The panel recommends that the CCC aim to produce that level of output and impact for all workshops, in partnership with ACM, IEEE, and other professional societies. A partnership with the ACM to put reports on the ACM DL could broaden its outreach. Further, it would be valuable to disseminate the visioning workshop calls to more conferences, journals, SIGs and other venues.

CCC Organization, Management and Budget

The CCC has a strong management history and plans to continue this structure into the next phase of the project. This was well displayed in the reverse site visit where the leadership was clear, authoritative, and responsive.

The CCC management structure is very effective for the CCC's current size. When planning for the future, the CCC should carefully consider how its size impacts its mission. While scaling up would allow the CCC to increase its coverage of a larger and more diverse set of topics, but, at a certain scale, it would require a new management structure and more administrative processes that could threaten the agility that has been a major strength of the CCC.

The CCC has a concrete plan for leadership succession. The organizational structure in which there is a Chair and Vice-Chair who are aided by a full-time Director and two full-time staff seems generally appropriate and has been working reasonably well to develop the goals and execute the projects of the organization. The Director and staff currently report to the Executive Director of CRA.

The panel felt that the budget was appropriate for the scope of the proposed effort.

Recommendations:

- To streamline decision making and improve lines of communication, CCC staff should report directly to the CCC Chair rather than to the CRA Executive Director.

Outcomes, Metrics, and Assessment

The CCC has a good track record of achieving its goals of supporting the research community in developing new research directions and disseminating them, particularly to relevant government funding agencies. This is supported by a number of compelling qualitative "success stories," such as the Robotics Roadmap that contributed to the National Robotics Initiative, and the BRAIN Initiative.

As is widely acknowledged, it is difficult to assess the impact of innovative research ideas and

their presentation in white papers, workshops, blog entries, and the like. By the very nature of “high-risk, high-payoff” ideas, impact takes time, and a significant number of ideas will fail. (Indeed, if not enough are failing, perhaps they are not taking enough risk!) However, as years go by, there is enough history to provide a track record. Early workshop reports, for example the 2009 Robotics Roadmap, have had major impact, especially by influencing the creation of the NRI funding program.

That said, much of the presented evaluation was quantitative but not enlightening. Major parts of the evaluation consisted of numbers of papers written, workshops conducted, and attendees hosted, which are essentially “effort measures.” What is really needed is “outcome measures” that line up with the stated goals of the CCC. An example is the results of surveys conducted with attendees at the end of each workshop. The surveys capture the opinions of attendees immediately after the workshop about possible future outcomes, but the actual impact of the visioning workshops can only be measured after a considerably longer period of time. More work is needed to devise adequate outcome measures for CCC’s goal of articulating “far reaching and strategic research visions.”

Citation counts (e.g. from Google Scholar) are often used as outcome measures, reflecting the influence of a paper on other researchers. However, based on a quick search, citation counts do not seem to adequately reflect the influence of CCC task force reports or white papers. In section 2.3.3 of the proposal, the discussions of Outcome 1 provide clear qualitative evidence of the impact of particular CCC workshops and white papers, on science policy and on investment strategies of major funding agencies. In Outcome 2, the number of visioning activities is an effort measure, but many of the quotes provided are qualitative outcome measures. In Outcomes 3, 4, and 5, the descriptions mix effort measures (“we did a lot”) with outcome measures (“what we did had impact”).

Recommendation:

- The CCC should devise clearer and more compelling outcome measures, reflecting such things as influence on agency policies, new programs initiated, influence on individual research directions, funding acquired by individuals, or increased leadership roles in years after LiSPI participation compared with matched non-participants, etc. (Perhaps the problem of assessing innovation ideas could be the topic for a workshop.) Even when impact cannot be assessed for a number of years, relevant qualitative and quantitative data should still be collected for later analysis and quantitative evaluation of the impact of CCC.

Plans for Moving Forward

Computer Science has increasing impact in other disciplines - the value of viewing the sciences, social sciences, and even the humanities through a computing lens is well known and supported through a range of NSF and other programs.

In our campuses CS+X is driving research and education. This speaks to increased coordination

with other organizations in these domains. It is hard to go at it alone and it will be good to partner with other organizations and agencies to accomplish this. CCC is already doing some of this (e.g. MForesight), and it is recommended that CCC approach this systematically and strategically.

The panel agrees with increasing outreach and interactions with industry and not-for-profit stakeholders. These communities need to be tapped for technical inputs and potentially also funding. It is good to see that industry investment is broadly defined, and includes getting datasets for research use, access to research platforms for the community. Further, making the industry the driver in some of the visioning workshops is a good idea. However, the panel heard limited details for engagement with industry. Clearly this is a work in progress, and this area needs to be thought out much more.

Recommendation:

- As CCC increases its engagement with industry, it will be important to have a few CCC champions from industry who are thought leaders and decision makers. These could be in the advisory council, or council members, or informal ambassadors/advisors. These will be critical to provide the level of support that CCC has successfully built up over the years on the government side.

In conclusion, the review panel strongly recommends that NSF fully fund the CCC proposal. The CCC fulfills a unique and important role in bringing together the computing community and in energizing it to advance the nation's priorities. The panel's recommendations in this report aim to assist the CCC leadership in fulfilling this role even more effectively. These recommendations include that the CCC select topics so as to respond rapidly to emerging technologies, increase engagement with professional societies and key industries, and broaden active participation in CCC, including through an Advisory Board.