

Computing Community Consortium Reverse Site Visit

Dec 5-6, 2019

Review Panel Report

I. Executive Summary

This report documents the findings and recommendations of a Reverse Site Visit held for the Computing Community Consortium (CCC) at the National Science Foundation on Dec 5, 2019. The review panel had previously read the 2017 CCC proposal, the annual report, the recommendations of the previous reverse site visit, and associated materials. At the Reverse Site Visit, the CCC leadership presented its progress including addressing topics that the panel requested. Further they answered questions in a detailed Q&A session.

The panel was impressed by the CCC's presentation of the progress over the past 2 years. It has been a time of significant activity and success. In particular, the "20 Year Community Roadmap for Artificial Intelligence Research in the US" was an unexpected activity that the CCC engaged in a 13-month intense sprint to bring to fruition. This effort was one of the CCC's biggest commitments in time and effort, and also correspondingly has had great visibility and likely will have significant impact. The committee applauds the effort the CCC put into this, and regards it as a success. The activity had buy-in from key stakeholders in the AI community and funding agencies, and follow up with funding agencies and beyond is a model of the kind of impact the panel feels the CCC is uniquely positioned to catalyze.

Beyond the AI Road map, the panel commends the CCC on making significant progress on multiple fronts and adding value to the community:

- Increasing engagement of early career faculty through participation in their workshops and white paper reports,
- Broadening engagement in the Council and among workshop participants. The gender balanced CCC Council membership and the balanced profile of attendees (gender, geography) in workshops are commendable,
- Reaching out to the American Association for the Advancement of Science for greater engagement which we strongly encourage to continue,
- The report on Industry-Academia relations which was timely, and
- Creating "Gems" as indicators of impact that can make the case for the CCC's value going forward.

We thank the CCC Chair, Vice Chair, and staff for their excellent work, leadership, and service to the community, and we are confident the incoming Chair will be able to take CCC's impact to the next level.

We agree with the articulated desired outcomes of the CCC: increasing federal agency awareness, engaging the computing-research community, creating high-impact tangible resources, growing awareness of CCC, and growing leadership and community capacity.

Our recommendations going forward are outlined below. The detailed recommendations are given in subsequent sections.

1. **Strategic planning of portfolio.** The CCC has had successes, and now is in a position to define a more strategic approach to create successful impact. They should develop an intentional strategic plan about how to balance activities in size and scope ranging from big activities (like the AI roadmap) to visioning workshops, to white papers. The RSV Panel recommends that the CCC attempt to engage in a large-scale activity approximately once every two or three years.
2. **Coverage of CS.** Some areas of computer science seem to have been missing from visioning workshops and the CCC council. This is natural due to the limited slots. Nevertheless, the CCC should make a conscious effort to rotate areas of coverage to address gaps over time.
3. **Follow-through on activities.** Each activity the CCC engages in should include a clear plan for follow-through activities to have an impact (e.g., lining up agency stakeholders ahead of time with the intent of disseminating workshop outcomes to them, having attendees/workshop chairs run town halls in conferences, etc.). Always committing to that kind of follow through is critical.
4. **Communication.** The CCC should learn from best practices of similar institutions for how to maximize impact; clearly specify post-workshop expectations for workshop chairs and participants; and follow our suggestions on communication below.
5. **Outcomes.** CCC must find ways to assess its success in achieving the desired outcomes, be intentional in aiming to achieve them, and develop written best practices. Indicators of impact should be developed and maintained for **all** activities (workshops, workshop reports, white papers, etc.). A retrospective analysis might be valuable to develop future strategies.

II. CCC Goals, Strategies and Priorities

The high-level goals and strategies of the CCC have remained largely unchanged since its inception, with a focus on *catalyzing research in the computing community*, primarily through the articulation of research agendas in white papers and at workshops with associated reports. In addition to articulating new research agendas, the CCC has also developed reports on the computing research community itself, such as its report on the evolving academia/industry relations in computing research. In the last several years, the CCC has begun expanding its capacity through the creation of task forces and working groups in selected areas of interest, to amplify the efforts of its Council members in those areas.

Numerous audiences have appetites for the CCC's outputs, including: computing researchers; government agencies that fund computer-science research, the computing industry, and the general public. Producing outputs that impact *all* of these communities is a challenging task, and so it is natural that CCC prioritizes some audiences over others in attempting to maximize its impact. When asked about this prioritization,

the CCC leadership indicated that researchers and government agencies who fund them are the audiences of primary interest, which seemed entirely appropriate to the RSV Panel.

The CCC leadership reported on their initial efforts to engage members of the computing industry in CCC activities, specifically on the CCC Council and in CCC workshops. The leadership expressed concerns about the possibility of industrial engagement influencing the workshop outcomes to benefit the relevant companies.

A major activity for CCC this past year was the authoring of a roadmap for artificial intelligence research. To the CCC's credit, it managed to pull this off in just one year. The RSV Panel felt that this effort was a model for CCC to try to replicate in the future (albeit not necessarily on such an aggressive timeline!). In particular, this activity is referenced in the recent call-for-proposals for NSF AI Research Institutes, which promise to be major drivers for AI research in the US for the foreseeable future.

A significant discussion concerned how workshop topics are identified and developed to advance CCC goals. While some topics are clearly timely and strategic, there was some concern about uneven coverage. For example, some topics seem to have in-depth coverage through a series of workshops, while other prominent topics such as computer vision are not covered.

CCC leadership has identified the process of topic selection as falling into three bins: "top down", "sideways" and "bottom-up". The AI Roadmap is an example of the top-down approach. The Wide-Data Analytics topic was developed through the sideways approach via interactions spanning the Systems and Architecture task forces, whereas the Thermodynamic Computing topic welled up from the community, and thus through the bottom-up process. Going forward, we see this reflection on the process of how topics are identified as a good starting point for further development, especially when combined with the key parameters that are predictors of success, such as timeliness, agency/sponsor interest, etc.

An element of the CCC's strategy to maximize its impact is a communications plan that has grown to include AAAS engagement (e.g., panels) and podcasts, alongside its traditional focus on workshops, reports, and white papers. As a stated goal of the communications plan is increasing engagement from and influence in the computing community broadly defined, the RSV Panel encourages the CCC leadership to explore new means to reach and involve subcommunities of computing researchers that have traditionally been underserved by the CCC, such as computational biology.

Recommendations:

1. The CCC leadership should deliberate as to the proper balance of visioning activities in terms of size and scope (i.e., ranging from very ambitious, such as the AI roadmap, to more specialized) and in terms of the "direction" from which these activities are derived (top-down, bottom-up, sideways). Part of this evaluation should include the cost vs. benefit of the series-of-3 workshops as a method for conducting more ambitious visioning activities (e.g., Assured Autonomy), coupled with significant follow-through for such a commitment. The RSV Panel understands that

large-scale initiatives are often driven by external demand and so cannot always be planned, but the RSV Panel recommends that the CCC attempt to recruit such a large-scale activity approximately once every two or three years. The impacts of such efforts are compelling.

2. Some areas of computer science seem to have been missing from visioning workshops, white papers, etc. This is natural, as a limited number of visioning workshops per year cannot cover all of computer science every year. The CCC should make a conscious effort (as they explained at the RSV Panel they actually do) to rotate areas of coverage through membership in the CCC Council so that over a multiyear period these gaps will be addressed.
3. The CCC leadership should reinvigorate their efforts to engage the computing industry in their workshops and report production. The CCC team raised concerns about the perception of corporate influence on the report outcomes. The RSV panel understands that in certain cases (e.g., where it specifically presents perception problems due to the topic of the report), industry participation might not be solicited.

III. CCC Leadership and Management

The CCC's current leadership has been very strong, and the CCC has done a great job of seizing the opportunity of the AI Roadmap Project and navigating it to a very high-impact conclusion. The leadership-transition plan, whereby the Vice Chair serves for two years, then serves as Chair for two years, and then serves as Chair Emeritus for one year allows for ramp-up, fresh ideas, and continuity.

The CCC leadership believes that the CCC has the right organization, but needs some staffing enhancement at the junior level to free time for the senior staff. The Panel is sympathetic to this argument, but would like to see the CCC specify explicitly how the staff could not merely "do more workshops" but actively help the leadership, the Council, Task Forces and the like be more effective in achieving the Desired Outcomes spelled out in the CCC Strategic Plan. Developing such an explicit plan would help the CCC leadership assess whether their organization and staff is adequate for the task.

The CCC has a smooth working relationship with CRA. The CCC has a tight working relationship with the NSF management team under the Cooperative Agreement. The Panel recommends a mutual evaluation of the form of this relationship in order to enhance its effectiveness. For example, it may be useful to bring into regular conversations program officers or management from different CISE divisions.

Recommendations:

1. The CCC must formulate written guidelines regarding how the leadership, the Council, the Task Forces, and the staff could work together to be more effective in achieving the Desired Outcomes spelled out in the CCC Strategic Plan. Specifically, the desired outcomes are: increase federal agency awareness, engage the computing research community, create high-impact tangible resources, grow the awareness of the CCC, and grow leadership and community capacity. The CCC must find ways to assess its success in achieving these outcomes, be intentional in aiming to achieve them, and develop written best practices.

2. The CCC leadership and NSF should review together the form of their working relationship under the Cooperative Agreement.

IV. CCC Activities and Budget

The CCC advances its core mission by organizing workshops and developing white-papers on key topics. Governance is provided by the Council and the key activities are developed with the guidance and engagement of “task forces,” where each task force focuses on a specific thematic area; currently, there are 7 task forces covering topics such as AI, Cybercrime and cybersecurity, Health and computing, Industry relations, etc.

The CCC is to be commended for maturing their workshop strategy and putting in place best practices for seeking out topics and supporting the organizers to be effective by developing strong reports and leveraging those for advocacy. A key strength is the support provided to early career workshop organizers. Nevertheless, post-workshop activities need strengthening. These include the uptake of reports or report summaries and materials derived from the reports for research advocacy and community building. It is notable that several workshops have resulted in reports that have already demonstrated potential to influence future research directions and programs. These outcomes should be expanded through focused effort to develop a program of post-workshop advocacy activities. A major success is the recent development of a “20-Year Community Roadmap for Artificial Intelligence Research in the US” which is already being referenced in the newly announced AI Centers solicitation. Further, the report is comprehensive and thorough, mapping out in detail key thematic areas of importance within AI and recommending approaches to develop them. As these areas advance in the next two years, the CCC should seek to leverage opportunities for effective advocacy and stakeholder engagement.

The distribution of CCC activities, both programmatically and budgetarily are largely appropriate. Additionally, with over 12 years of successful development, there is an opportunity to leverage lessons learned to inform the development of a portfolio of activities that can drive outcomes and guide efficient operations and effective staff development.

The RSV Panel observes that an important component is missing from the activity and the budget: organized promotion of report results, through engagement with appropriate agencies, communities, and press. These activities are not planned for and not budgeted for; see “Learn the best practices” (Section V, Recommendation 1) below for suggestions of activities.

In summary, the RSV Panel is pleased with recent successes and recommends that the CCC leadership undertake an assessment and portfolio planning process to guide the effective management of resources in support of overall CCC strategic priorities. Such a process can enable proper coverage of topical areas, amplify outcomes, and lay the foundation for future growth.

V. CCC Communication Strategy

Impact is achieved by doing outstanding work at an opportune time and by communicating the results of this work effectively to various stakeholders. Effective communication is thus central to impact.

The CCC has made significant progress in the past few years, increasing and structuring communication with the goals of increasing impact, broadening participation, and growing engagement.

The impact achieved by the “20-Year Community Roadmap for Artificial Intelligence Research in the US” is a sterling example. This very substantial project required intensive effort, but was accomplished in only 13 months from conception to delivery. The Association for the Advancement of Artificial Intelligence (AAAI) was brought in as a partner - one of many examples of the CCC’s outreach and inclusiveness. The project chairs and workshop participants included the leadership of the AI field as well as a broad cross-section of researchers from academia and industry. The report was widely briefed, including a town hall at AAAI meetings and briefings at OSTP, NSF, DARPA, NITRD AI-IWG, DIB, NSCAI, ITIF, Lincoln Labs, NSF Big Data Innovation Hubs, and Hill staff¹. CCC leadership and study leadership devoted extraordinary effort not only to the workshop and report, but also to the dissemination of the findings in ways that have strongly influenced research direction and investment.

The involvement of AAAI in the AI Roadmap is one example of the engagement of other organizations in CCC activities. Such engagement builds bridges, increases the visibility of the CCC, and increases impact. There are multiple forms of engagement that establish important communication channels and deliver short-term and long-term value:

- “Partner engagement” with other organizations in the computer science field, such as AAAI’s engagement in the AI Roadmap project.
- “Horizontal engagement” that brings computer scientists together with leaders of other fields to illuminate new collaborative opportunities at the boundaries, such as the engagement of economists in the “Algorithmic and Economic Perspectives on Fairness” project.
- “Vertical engagement” that brings together researchers at various levels of the “stack” in situations where collaboration is essential; “Next Steps in Quantum Computing: Computer Science’s Role” is an example where the “stack” from physics to architecture to software to algorithms.

The consistency and quality of workshop reports is another area of notable progress. Not every workshop report has the breadth and depth of the AI Roadmap report, of course, but it is clear that considerable

¹ Association for the Advancement of Artificial Intelligence, Office of Science and Technology Policy, National Science Foundation, Defense Advanced Research Projects Agency, Networking and Information Technology Research and Development AI Interagency Working Group, Defense Innovation Board, National Security Commission on Artificial Intelligence, Information Technology and Innovation Foundation,

effort has been devoted to ensuring that workshop chairs, workshop participants, and CCC staff produce high quality artifacts that document the results of each workshop.

Other improvements in communication in recent years include creating a visible and reliable presence on arXiv, reworking the CCC website, institution of the “Catalyzing Computing” podcast and its translation into podcast engagements with the American Association for the Advancement of Science and the Heidelberg Laureate Forum, an annual session at the AAAS Annual Meeting, and the “Blue Sky” conference tracks (this latter item going back a number of years).

All of this represents significant progress. There are, however, many opportunities for further improvement.

In general, we feel that there has been a great deal of progress in specifying the processes, procedures and expectations up to and including the production of the workshop report, but far less progress in specifying the processes, procedures and expectations for the strategic communication of results in order to maximize impact. CCC needs a communication, engagement and outreach plan that is strategic and executable within available resources. Prioritization is necessary - prioritization of “targets” (at the RSV the CCC suggested the community, the federal government and federal funding agencies, and in some cases public media) and prioritization of mechanisms.

Recommendations:

1. Learn the best practices from neighboring agencies for how to maximize impact. Liz Bradley, as incoming CCC Chair, should consult with the Chair (Farnam Jahanian) and Senior Director (Jon Eisenberg) of the National Academies Computer Science & Telecommunication Board to explore how CSTB and similar NASEM Boards and Committees structure their communication activities. What are the vectors of communication? What is the role of staff? What is the role of study chairs, and of study committee members? We recognize that NASEM and CSTB are at different scales than the CCC and CRA, but they have been at this for a long time with considerable success, and there will be things to be learned.
2. Whereas the 2018 communication plan is impractical to execute and poorly prioritized, and whereas in any future proposal there must be a better plan and evidence that the CCC can execute it, the RSV Panel recommends that CCC rewrite a prioritized strategic communication plan *now* and practice executing it.
3. The post-workshop expectations for workshop chairs and participants should be carefully documented and communicated, as should the contributions that chairs and participants can expect from CCC staff. Prospective participants in each workshop should commit to fulfilling these expectations at the time they agree to attend the workshop, and each workshop should allocate an hour at the end for a focused discussion of followup activities.

Tactical suggestions for communication with the research community:

- Targeted: There should be explicit outreach to relevant ACM SIGs as part of the followup for each CCC activity. SIGBoard meetings provide a good forum for meeting many SIGs.
- Targeted: There should be explicit outreach to all leaders of the relevant field(s) (for example, program committee members and authors at relevant conferences) as part of the followup for each CCC activity.
- Broad: Explore with CACM editorship possible arrangements under which technical/viewpoint material from CCC might appear regularly.
- Broad: Designate a CCC liaison at each CRA member organization. Create a mailing list for these people. Send them a monthly or quarterly update with links to outputs and with a reminder to forward to faculty / research staff members. Alternatively, perhaps CRA has a mailing list that reaches all individuals at member organizations?
- Broad: Explore ways in which CRA can amplify CCC visibility (e.g., at the Snowbird Conference).

VI. Plans for Moving Forward

The proposed plans (for activities in the remaining duration of the grant) appear reasonable. Because the AI Roadmap was a bulge in expenditures (fully justified by its successful outcome), within the current budget the number of visioning workshops will need to be curtailed. The CCC has demonstrated that it can adapt to shifting priorities (viz., AI Roadmap), but within the original remaining budget there is limited flexibility to adapt if another such opportunity arises. There is a current supplementary proposal that addresses these issues.

Recommendations:

1. The Annual Report and the RSV presentation itself list “Desired Outcomes” that are appropriate, but the PIs and staff should more systematically evaluate and document whether the activities (workshops, workshop reports, white papers, etc.) have achieved these outcomes. This evaluation need not take the form of quantitative “metrics”: there can be other *indicators* of these outcomes. Indicators can include:
 - Consideration of CCC reports in research-priority or funding decisions by Federal agencies (of course, the CCC is correctly modest in hesitating to claim credit for decisions that are “in the zeitgeist” anyway, but evidence that the reports *contributed* to wise Federal decision-making can be reported as an *indicator* of useful outcome).
 - Publication of (an article about) a CCC report in a refereed venue such as CACM.
 - Other useful kinds of indicators, which CCC has already documented in the form of “Gems.”
2. As part of the CCC’s strategic planning (especially if there are more workshop proposals than can be funded), CCC management might consider a retrospective evaluation of the last few years’ workshops: which ones were *more successful* (by any appropriate variety of criteria) and which

were less successful; were there any characteristics of the workshops that (in hindsight) could have predicted success; and can these criteria be used for the (admittedly harder) task of predicting the future, in selecting which visioning workshops will have the most impact.

3. We recommend, for *each* workshop funded by the current grant (since 2017) the CCC list these indicators of success, and assess, “was this a less successful / more successful / home run” activity.

In summary, the RSV Panel values the incredible asset the CCC is to the computing community. We applaud the success of the past two years, and look forward to seeing the CCC build on this success going forward.