

# THE COMPUTING COMMUNITY CONSORTIUM

## September 2006

### A. WHY DO WE NEED A COMPUTING COMMUNITY CONSORTIUM?

What questions shape our intellectual future? What attracts the best and brightest minds of a new generation? For biologists, it is a deep understanding of life and its processes. For physicists, it is deducing the structure of the universe and a grand unified theory of the fundamental forces. These are profound questions, older than history, with some potential solutions now within our grasp. Moreover, this knowledge has and is enabling diverse advances, from cancer treatment to engineering design.

As computing researchers, we work in that most malleable of media, information and its processes. Our ideas continue to enable and transform all of science, reshape the world's economy and change our culture. What are the next big computing ideas, the ones that will define the future of computing, galvanize the very best students, and catalyze research investment and public support?

The purpose of the Computing Community Consortium (hereafter CCC) is to create compelling research visions and the mechanisms to realize these visions. Other research communities effectively establish consensus and national research agendas. For instance, astronomers and physicists use decadal committees under the auspices of the National Academies; an example is the physics decadal study.<sup>1</sup> A typical result of such studies is the 2003 report *Connecting Quarks with the Cosmos*,<sup>2</sup> which enunciates eleven physics and astronomy challenges for the new century, including such fundamental questions as “What is dark matter?” and “Was Einstein right about gravity?” These fields organize themselves to facilitate community interaction and produce documents that clearly state consensus choices and define research priorities and initiatives.

Compelling visions take many forms. History has amply demonstrated the importance of entrepreneurial, grassroots efforts as creative engines in computing research. History has also demonstrated the value of large teams, large facilities, and large amounts of funding. Many see an increasing need for shared research facilities and teams in our field to allow us to tackle certain “grand challenge” problems. Planning for large-scale research should not, and need not, harm smaller-scale efforts or place them at a disadvantage.

With this background in mind:

*The challenge for the Computing Community Consortium is to catalyze the computing research community to debate longer range, more audacious research challenges; to build consensus around research visions; to articulate those research visions; to evolve the most promising visions toward clearly defined initiatives; and to work with funding organizations to move the challenges and visions toward funding initiatives. The CCC will do this without harming the research environment that has created the computing world of today.*

### B. WHY CRA?

Who should form the Computing Community Consortium? Fortunately, there is an existing organization whose role is to strengthen research and advanced education in computing and allied fields: the Computing Research Association (CRA).

---

<sup>1</sup> [http://www7.nationalacademies.org/bpa/projects\\_physics\\_2010.html](http://www7.nationalacademies.org/bpa/projects_physics_2010.html)

<sup>2</sup> <http://www.nap.edu/catalog/10079.html>

CRA was formed in 1972 as the Computer Science Board (CSB) to provide a forum for the chairs of Ph.D. granting computer science departments to discuss issues and share information. In 1986 CSB was incorporated as the Computing Research Board (CRB). In 1990, CRB was given its present name and a permanently staffed office was opened in Washington DC.

Today, CRA's membership includes nearly all major stakeholders in the community, including more than 200 academic departments, 30 industry and government research labs and six affiliated professional societies that together represent more than 150,000 computing professionals. For decades, CRA has sponsored the biennial *Snowbird Conference*, a leadership summit of computing research societies and leaders to discuss computing policy and community needs.

CRA organizes events and workshops for federal agencies. Recent examples include a cyber-learning workshop series for NSF,<sup>3</sup> the NITRD-funded *Workshop on the Roadmap for the Revitalization of High End Computing*,<sup>4</sup> NSF Town Hall Meetings on the Global Environment for Network Innovation (GENI), and a series of workshops highlighting the NSF Broadening Participation in Computing Program. CRA has also sponsored a series of Computer Science Grand Challenge conferences,<sup>5</sup> supported by the National Science Foundation, whose goal is to develop community research agendas in computing.

However, individual conferences are insufficient to build a truly community-wide consensus and to work with research agencies toward funded initiatives. What is needed is a way to coalesce research visions and initiatives, and move them toward funded programs. It is precisely for this purpose that the CRA, with the support of NSF, is creating the CCC.

The CCC will be funded initially via a \$6M, three year award from the National Science Foundation to the CRA announced in September 2006. For details, see <http://www.cra.org/ccc/>.

### **C. WHO WILL BE INVOLVED?**

The CCC will be broadly inclusive of the computing research community. We anticipate diverse participation from the community in a multiplicity of visioning activities. Any computing researcher who wishes to be involved will be encouraged to be involved.

Facilitating the CCC's activities will be a CCC Council (hereafter Council), intended to be a group with the stature, diversity, and longevity needed for the CCC to be effective. The Council's role is to stimulate and facilitate visioning. The Council is responsible not for *doing* the visioning, but for *putting processes into place* that stimulate and facilitate visioning by the computing research community. These processes will include conferences, workshops, task forces, white papers, and a wide variety of other vehicles, all of which will be widely advertised and open to the broadest possible spectrum of the computing research community.

Guided by the Council, the CCC will foster evolution of the most promising visions toward major funding initiatives. Some funding initiatives will require significant instrumentation; others will not. The Council will work closely with appropriate members of NSF and other funding agencies to advance the interests of the community.

The Council will be a standing committee of the CRA and the CRA Board will select a Chair of the Council. The Council will consist of 12–15 people serving three-year terms, with the appointments

---

<sup>3</sup> <http://www.cra.org/Activities/workshops/cyberlearning/>

<sup>4</sup> <http://www.cra.org/Activities/workshops/nitrd/>

<sup>5</sup> <http://www.cra.org/grand.challenges/>

staggered so that one third of the positions are open every year. Members will be eligible to serve no more than two consecutive terms. The selection committee for potential Council members will consist of the Council Chair, the CRA Board Chair and representatives from the current Council and the CRA Board. All appointments will come after broad participation by the community and NSF to ensure that the CCC Council will exhibit the multidimensional diversity of our field.

#### **D. HOW WILL THE CCC OPERATE?**

The visioning process begins with a community activity to identify fundamental questions in computing. These questions are not program or facility specific, but may ultimately encompass multiple programs or facilities. A community-wide effort to develop the current list of fundamental questions will occur once every five years. However, a less intensive activity to sustain the list by producing updates and progress reports will occur every year. The CCC will communicate the output of this activity to the national community. This list of fundamental questions will provide the framework and rationale for large initiatives.

In support of the development of the fundamental questions list, the Council will charter visioning activities, which will identify potential major opportunities, set priorities or establish grand scientific or engineering challenges for the field. These visioning activities will be based on a topical interest area either proposed (formally or informally) by members of the computing research community or formulated by Council members. Such proposals may be community generated or may result from workshops and study groups organized by the Council.

As visioning activities gather momentum, the Council will establish *Visioning Task Forces*, whose members will be recruited from the proposers, based on interest and expertise, with oversight by the Council. These task forces will conduct workshops and meetings, ideally in conjunction with related conferences. Some task force activities may be conducted in foreign venues to ensure international participation.

Task force members will generate a public report that describes the prospects for this research area and that estimates the resources required to stimulate sustained activity. We expect these reports, either individually or in collected editions published by the CRA, to constitute authoritative statements of the scope and benefits for major computing research initiatives. The reports will also form a basis for consensus building, helping establish an agenda for future initiatives and community thinking around audacious research goals. Because research is an international activity, the broader community will participate in these activities.

A possible outcome of Visioning Task Forces will be the identification of ideas for major instrumentation or research initiatives that enjoy widespread community support and that address deep challenges and problems in computing. The Council, in such cases, will place the initiative in the context of the computing community's key research questions, and will seek agreement from appropriate funding agencies that the idea is worthy of further exploration. Based on such agreement, the Council will work with the Task Force to form an *Initial Planning Group*.

The charter of the Initial Planning Group will be to formulate a plan that outlines major strategic thrusts, identifies possible sources and types of funding and identifies the portion of the scientific community that should participate. The CCC will assist the Initial Planning Group in presenting their findings to appropriate funding agencies and help them establish committed prospects for funding.

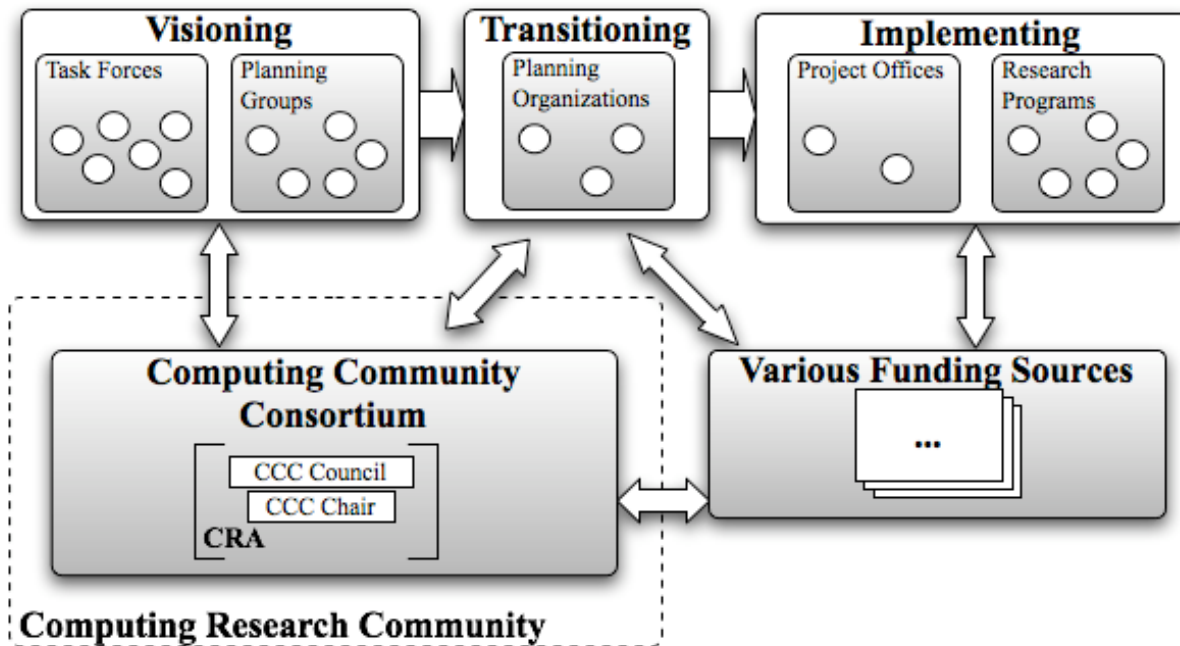
While many of these thrusts will primarily require interest and initiation through new coordinated funding programs at agencies, others will require the development of large-scale instrumentation and are more suitable for programs such as the NSF Major Research Equipment and Facilities Construction (MREFC)

program.<sup>6</sup> For future research initiatives requiring the construction of large-scale shared resources through the MREFC process, CCC, through the Initial Planning Group, will play a key role in the preliminary stages. Upon successful completion of this MREFC conceptual design phase, the CCC will work with NSF to establish a *Planning Organization* to continue the preliminary design phase. As this process proceeds, the NSF will become more directly involved in oversight, though the CCC will continue to represent the community in evaluating whether the instrument being planned and built by the Project Office is meeting the research needs it was created to serve.

Tailored, small groups constituted by CCC will seek to reduce the time needed not just to formulate a consensus around a research vision in the community, but to aid the funding agencies in exploring alternative initiative formulations that reduce the time between vision setting and program initiation. By identifying and sustaining a consensus suite of the computing research community’s fundamental research questions, the CCC will provide a ready source of motivation for innovative research programs and a broadly based rationale for their funding.

As the community becomes more adept at formulating visionary research and instrumentation initiatives, the CCC will oversee the prioritization of these initiatives. Because the computing community is significantly larger and more diverse than, for example, the astronomy community, a single mechanism such as a decadal study is unlikely to suffice. Prioritization will occur in response to the output of multiple visioning efforts. As in the astronomy community, the reports describing these visions will be widely circulated and discussed so as to build consensus.

The figure below shows both the intended organization and the operation of the CCC:



<sup>6</sup> This pool of funding is distinct from that used by CISE to support research grants and fellowships. At present, the computing community does not draw funds from the MREFC pool. Such large-scale MREFC projects require approval by the National Science Board and are line items in the NSF Congressional budget.

## **E. HOW WILL RESEARCH THRUSTS BE ADVANCED?**

Computing has a very diverse range of research thrusts, which the CCC will reflect in the consensus suite of fundamental research questions. At any time, many research objectives are being actively pursued. One CCC objective is to catalyze the formulation of new research thrusts – more rapidly than they have formed in the past. For example, we believe that current digital libraries and the learning research efforts could have been formulated more rapidly if there had been a change agent outside the funding agencies, such as CCC, actively assembling a community of interest to define clearly and rapidly the research promise of a new thrust. This change agent would then work actively with funding agencies to find a source of funds for the new ideas. Some of these thrusts will grow and prosper; others will not gain a community consensus or will not appeal to funding sources and will disappear.

Because we see value in multiplicity, the CCC routinely will pursue multiple thrusts simultaneously. Strength of community interest will communicate the priority of a candidate thrust to the funding agencies. In addition to serving in this matchmaking role, the CCC will serve in a high-level oversight role ensuring that the scientific mission of the program is serving the broad computing community.

However, the community will need to choose among many nascent ideas to decide how to structure its activities. The Council will permanently maintain a list of candidate ideas for new avenues of research to pursue. Each idea will be described in as compelling a way as possible. At least quarterly, the Council will review this candidate list and determine whether to sponsor a new effort to explore an idea. The Council will maintain this list on a public web site and encourage commentary on the candidate topics. For each candidate topic that the Council selects for active support, a small proponent committee of 2-3 people will plan the activities and the community interaction to explore the idea. The committee will post both its plans and the status of activities on the web site so that transparency is maintained and all who are interested can volunteer to participate in topic development.

## **F. WHAT IS CCC'S ROLE IN THE GENI PROJECT?**

In the Program Solicitation for establishing the Computing Community Consortium, NSF focused on the role of the CCC to “ensure broad community engagement in the identification of compelling research agendas and in the subsequent identification and refinement of related shared use infrastructure requirements.” The Program Solicitation also stated:

*One of the first responsibilities of the CCC will be guiding the design of the Global Environment for Networking Innovations (GENI). GENI is a facility concept already being explored by the research community, including investigators from the disciplines supported by CISE. GENI will complement ongoing CISE research investments in networking, distributed systems and other areas.*

In accord with the solicitation, the CCC Council will establish a GENI Science Council to provide broad community involvement for GENI as it continues to move through the NSF MREFC process. MREFC initiatives are dependent on identification of a committed group of researchers who will champion the planning effort and the actual implementation and operation of the project. In the case of GENI, such a group is already in place. For details on GENI see <http://www.geni.net/>.

## **G. WHAT ARE THE BROADER IMPACTS OF THE CCC?**

The majority of the gain in economic productivity in the U.S. over the past decade has been due to innovations in information technology. Underlying those innovations are prior decades of computing research in algorithms, architecture, networking, software systems, telecommunications and other subfields, much of it funded by NSF and other government agencies. The CCC will promote continued innovation by enhancing the ability of the scientific community to envision and pursue long-term, audacious computing research goals rather than incremental ones. By so doing, CCC will further the

value of research investments. The CCC also will interface with the broader scientific community to develop visions for the role of computing and information science and technology in tomorrow's world, and also will suggest research strategies for how to realize those visions.

Education of future scientists is critical. By advancing new programs, infrastructure, and instrumentation, CCC will provide new training opportunities for today's students, who will become tomorrow's innovators. CCC will involve a diverse set of students from a broad spectrum of institutions in workshops and other venues where they will be able to participate with leading scientists and interact on critical research issues. Although K-12 and undergraduate education are not directly within the CCC's mission, we expect that the ideas and visions that emanate will help entice more students to enter computing-related fields.

Part of the CCC's efforts as will include communication with the public. In partnership with CRA, the CCC will develop materials that show the dramatic societal impact of information technology research, as well as its even greater potential. Such efforts will help foster understanding and support for how computing research will continue to improve the quality of life and standard of living for people in the U.S. and worldwide.

#### **H. WHAT WILL THE COMPUTING COMMUNITY CONSORTIUM ACCOMPLISH?**

The opportunity for the CCC is dramatic – we are at a time when the computing research community is ready to take the next step and assume more responsibility for its own success through the creation of funding programs and instrumentation to drive the next generation of researchers. Moreover, it is clear that the impact of computing on the nation's economy and our citizens' lives will continue to grow dramatically. Together we can:

1. Bring the computing research community together to discuss, prioritize and envision our 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.