



Finding and Advancing Visions in Computer Science & Engineering

Computing Community Consortium

Version 1: June 16, 2010¹

New ideas arising from computer science and engineering research have dramatically changed the world. One metric that documents this impact is the number of ideas that have been the basis for billion dollar markets, e.g. time-sharing, parallel computing and virtual reality. Today, the opportunities are numerous as the intellectual content of the field is re-shaped frequently. Research investment, public support and venture investment to bring new ideas to market fruition are all available. And formulating new challenges will attract the best and brightest minds of a new generation to our field.

The Computing Community Consortium (CCC) has the objective of catalyzing the computing research community to:

- envision long-range, more audacious & inspiring research challenges
- state them in compelling ways
- build momentum around such visions
- move them towards funded initiatives

Other organizations catalyze visioning activities; we seek to complement them.

Specifically, CCC supports small groups with an audacious visioning plan. It does not pick winners and losers. Rather, CCC supports the pursuit of new visions in many sub-areas in parallel because, frankly, we believe that the field is idea-limited at present.

The purpose of this document is to discuss CCC's view of the basic process of visioning end-to-end. It provides anecdotal observations of what we have found to work and what has not. It is intended to be a resource for a new-comer to CCC's visioning process who wants to be as effective as possible. It is also intended to be a resource for CCC Council members as they take on liaison responsibility with a group doing visioning.

Defining the Visioning Process

First you create the vision, and then you execute on it.

1. Create the Vision. New ideas come from the researchers in computer science and engineering, as well as from researchers in other areas.

CCC facilitates the discussions of such ideas by

- bringing together people with a shared interest,
- including representative people with diverse ideas and perspectives, and
- typically, doing so in one (or a few) small workshop(s).

¹ Contact: Erwin Gianchandani, Director, Computing Community Consortium (202-266-2936; erwin@cra.org). For the most recent version of this document, visit <http://www.cra.org/ccc/vision.php>.

It is critical to have the “right people in the room”. You are seeking a collective vision, not individual visions. It is desirable to have the most knowledgeable and experienced people and the most inventive. Discussions should be structured to tease out new ideas, to constructively evaluate them and to build consensus.

CCC appoints a Council member to act as a liaison between the Council and each visioning group. This liaison is a person

- who has observed other visioning activities,
- who is sympathetic to the particular proposed vision,
- who can act as an independent “sounding board”,
- who facilitates relations between visioning group leaders and CCC,
- who, once a vision is articulated, is willing to work with the visioning group to develop a strategy for advancing that vision, and to help facilitate execution of that strategy (including its dissemination to relevant policy makers and funding agents).

You don’t have a vision of future research unless you can write it down. So, a (successful) visioning process should produce some form of document that communicates your vision.

The form of the written product will depend on who participated in visioning, what happened at the workshop(s), and who the intended audience for your work is. Often visioning groups will document a clear notion of the problems/challenges that need to be addressed; the visioning document will then discuss each one in a compelling way. In other cases, the group might sketch and document research approaches that, if they succeed, will meet the challenge or solve (even partially) the problem. Lastly, some groups create an explicit roadmap, with clear goals to be achieved by a specific date.

There is not just one style of written product from a CCC-supported visioning exercise. The contents and form of report will depend on what consensus was built during the visioning process. But CCC insists on having a very short report that it can post on its web site, preferably with various URL pointers to related community sites with more in depth material. This helps ensure that the results of visioning exercises are available to the community, initially in a very compact, digestible form but with additional detail available.

2. Follow-through on a stated vision: communicate the vision; validate it; and develop a broader community to support the vision.

2.1. **Communicate** with the wider community. Visioning is an activity to get broader buy-in by the research community, perhaps even building a community of people who will identify with that vision. Stakeholders need to **validate** the vision as being meaningful. CCC itself is not a body that validates vision statements. Validation comes from the interested community.

2.2. Determine a strategy that will put in place what is needed for the research to be done.

- at an appropriate level
- at an appropriate time – should this vision be pursued “now”? Is it really actionable?
- by experts who have the ability to do the research
- are there existing research programs under whose aegis this vision could be accomplished? Does this vision provide new direction for an existing set of people and programs? For instance, how could we take existing funding programs and extend them to provide the necessary resources and support needed? Does the vision require altogether new funding programs? How does it “compete” with other visions in a closely allied area?

Typically, the strategy involves finding one or more sources of funding that can be used to translate a vision into an effective research effort.

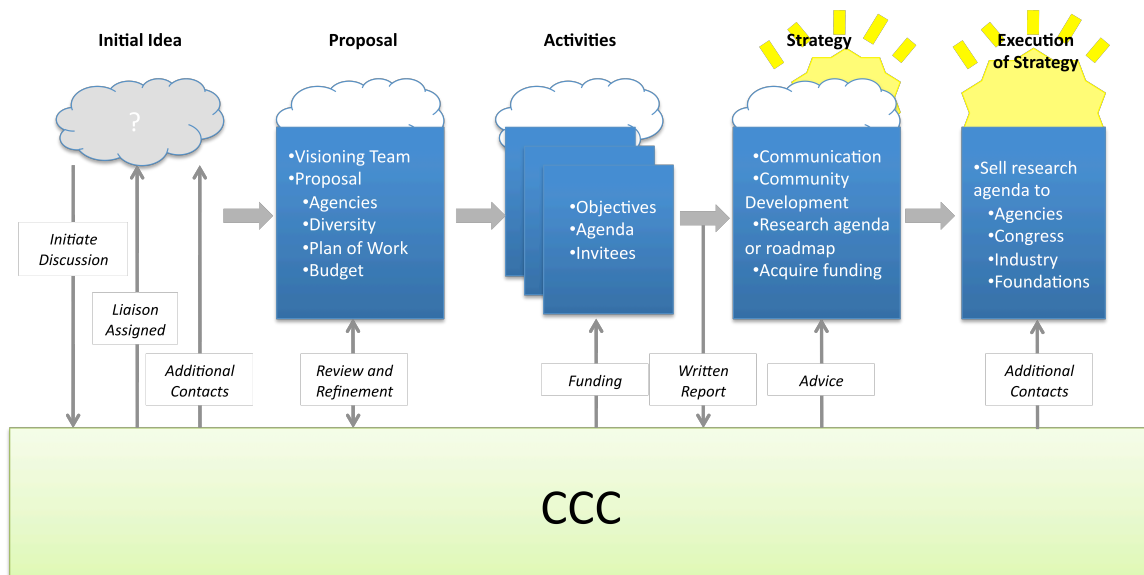
2.3. EXECUTE the strategy. The visioning team, plus the CCC liaison, will attempt to **execute** that strategy. This can involve

- discussions with program officers in various funding agencies
- interaction with industry
- interaction with Congress

Most computer scientists are familiar with the various funding agencies and industrial players. But there are also Congressional caucuses with information technology interests. CCC has access to them; they include:

- Congressional Internet Caucus
- Congressional Robotics Caucus
- House STEM Ed Caucus
- UAV Caucus
- Congressional Research and Development Caucus

Visioning, as practiced by CCC, is a process that moves from creation of ideas to the commencement of the desired research program, as illustrated by the diagram below.



Lessons Learned

The first portion of this document has defined the visioning process, including its expected outputs. This second portion gives a list of “lessons learned” by people who have engaged in visioning under the auspices of CCC sponsorship. While every visioning exercise is different, there are lessons to learn by knowing what went “wrong” or “right” in another sub-field; this information might help you to be more effective.

For the sake of discussion, we assume that a few leaders in some research area wish to convene the right people and to produce a vision document. They approach CCC to financially support either part or the entire workshop.

Formulating the proposal to CCC to convene a workshop:

The proposal you write to CCC is not a research proposal. Nor is it a proposal to assemble a group to write a research proposal. It's a proposal to define a direction or to nucleate or grow a community.

Be sure that there is a plan to engage a diverse community, preferably through an open process. Ensure that an adequate number of junior people are invited alongside the more established researchers. And be sure that the content will be diverse as to intellectual approaches within the field. Seek diversity in specialties and points of view. No one should look at the prospective attendees and complain that this is a meeting of a “thought clique”.

Make the case that there's an opportunity for progress now. It's no good to just say "we'll meet and chew it over." It is better to have some ideas, even preliminary or controversial ones, to catalyze the discussion.

Think about what it would take for NSF or some other agency to develop a "program" around the direction you're proposing.

Proposal Budget:

Typically, a budget for a CCC visioning exercise will support participant travel, lodging, and food, as well as the cost of the meeting venue. Rough budget estimates should be sufficient, e.g., there could be a need for 1/3 of the attendees to travel across the country, 1/3 to travel a short distance, and 1/3 with no travel costs of note. All stay at the workshop hotel. Estimate food costs. A corporation or a university may underwrite venue expenses. Budgets for 30-person workshops have been in the range of \$25,000-\$40,000.

Workshop leadership:

It takes significant leadership to develop and lead the development of a research vision, especially in a community that has widely different and valid schools of thought about its future and its research methodology. That doesn't mean it shouldn't be done.

It takes leadership, but it also takes a team; that team must include actively engaged thought leaders.

Understand the motivation behind a visioning effort. Is the workshop intended primarily to grow the field in numbers or in scope? Is it to change the direction of the field or re-define its challenges/problems? Do you have to build a new community, or does one already exist?

Get the vision and scoping right, only then pursue the appropriate funding channels. Keep the horse is before the cart. It is very difficult to make the opposite ordering “come together”.

If the vision is intellectually broad, develop a leadership team with cross-cutting experts. Get the very best folks in the adjacent areas.

Tap wisdom (e.g., from CCC and NSF) but make sure the intellectual drive comes from the research community. In one case, our CCC liaison and some key NSF folks played a critical role in helping structure the effort and drive it towards conclusion.

Think BIG and have clear national ambitions, i.e., connect your BIG thoughts to *national priorities*. In BIG documents, no single researcher should expect to see his/her current research agenda directly reflected.

Workshop logistics:

Engage the research community with care, so as to maximize the value of their participation.

Build in consensus mechanisms from the outset through mechanisms such as public forums and web-based feedback to achieve broad agreement on issues.

Use telephone conferences as a way to engage a broader community.

Don't take too long!

There is a tension between assembling the established players in a field to discuss its evolution, and growing the field by inviting interested new people and bringing them up to speed. New people, with their new values and perspectives, are often critical to the success of a new vision. So, introduce the new players early in the schedule.

The Hadoop Summit that preceded the Data-Intensive Scalable Computing (DISC) Symposium had an open invitation policy and allowed many people to become acquainted with DISC work. But, it quickly became over-subscribed for the available venue and the organizers had to scramble to accommodate the demand. Be warned. Don't have too many speakers or too tight an agenda. You'll wear out your audience and not leave adequate discussion time. Have plenty of breaks.

In some CRA-sponsored Grand Challenge conferences that did visioning, but that predated CCC, the program committee only created a meta-schedule before the conferences began. Each night the program committee designed the agenda for the next day, based on what the committee had heard. It is pretty scary to adopt this scheduling style, but it actually worked quite well.

Presentations at the workshop should have as their objective **looking forward, looking broadly, and mobilizing new action** --- not exhibiting deep details or incremental results.

Organizers may wish to forbid speakers from "playing their tapes," or performing "station identification", i.e., giving a variant of their "standard research talks". Talks should not be "research results" presentations.

A wrap-up should try to extract some ideas for moving forward.

Written products:

Publish as completely as possible the presentations or discussion notes from the workshop. This is not only part of our responsibility to NSF, but it's also a good way to nurture and expand the research community.

Similarly, a summary should be published within a short period. This summary needs to be accessible and inspirational to any computer scientist. It should be posted on the CCC site, <http://www.cra.org/ccc/>.

If the group concludes that a new vision is at hand and that vision should be actively pursued now, then a vision statement (the written product) needs to be written in a compelling form to educate a larger community and gain validation. Importantly, it should be written in a way that speaks not just to the community but also to Federal policy makers or funding agents.

Before the workshop(s) is (are) over, a few identified individuals need to be given the responsibility for communicating that vision more broadly and for creating a strawman strategy for “next steps”.

In some cases, the conclusion of the workshop(s) will be that no vision arose, or that it is not timely to pursue it, or even that other visions are just more important. This is fine, too.

Communicate vision product; achieve “buy-in”:

English prose is not the only way to communicate or enable a vision.

Open source software assets are a good way to induce wide participation.

The cyber-physical systems (CPS) community held a workshop that was part industry and part academic in order to build bridges to the industry that will build CPS, because technology transfer is a key component of their vision. This is an example of sculpting the community involved.

Strategy for moving a vision forward:

A vision requires a research work-program, the resources to execute the program, and a willing set of experts to accomplish the work program.

There are few funding sources for research: the feds, the states, philanthropy and industry. A set of potential sources – well matched to the vision – needs to be identified.

Federal funding for science and engineering research comes mainly from 5 agencies: NIH, DoD/Intelligence community (including DARPA, AFOSR, ARO, IARPA, NSA and ONR), NSF, NASA and DoE, plus some much smaller, but important players such as the Departments of Agriculture and Homeland Security. Most of these agencies fund research as seen through the filter of the mission of the agency. Parts of a specific vision may be better matched or motivated by the mission of one agency than another. This should be analyzed. (See comments on *Research Communities* below.)

The funding agencies execute different processes for developing new programs. The vision strategists need to involve someone who knows well how a target agency is best approached, and understand what their reaction would be and why. Often, the agency will be willing to have one of their program managers participate in your visioning workshop. But with or without agency involvement in your workshop, **it is important to involve someone who knows the funding organization** – in particular, somebody who knows

- current ongoing programs,
- the specific office within the agency that is best matched to the vision,
- how funding allocations for the specific office are made – and when in the annual cycle,
- the relevant program officers in the specific office that is targeted, and
- any other relevant insight into what constraints those program officers work under.

Develop a plan for approaching the “right” agency program officers. It is advantageous – but not mandatory – that some of the people involved in developing the approach to the agency actually know, and are known by, the agency program officers.

Working across two agencies is difficult, and it is best if you do not attempt to coordinate a program across two agencies. The government people may try to coordinate with each other, and it is great if they succeed. They may also solicit your involvement in coordinating amongst themselves (e.g., in drafting language indicating the community’s desire to work on research projects funded by the multiple agencies), but again, your involvement in this activity should come at their direction.

The roles and responsibilities of the CCC Council liaison:

The liaison should be a sage advisor on any issue involved in visioning, including helping to define the specific visioning process, selecting venues, crafting workshop formats, and structuring the written product.

Most importantly, the liaison should advise on how best:

- to communicate the vision to both the broader community and to the Council,
- to craft the strategy to realize the needed research,
- to identify ways to approach funding organizations, and
- to invite other well-positioned leaders to aid in approaching the funding organizations.

The liaison is not in charge of the visioning activity; not all of his/her feedback may be adopted by the leaders/organizers. However, it is important that he/she be actively involved in providing guidance, in steering the visioning activity in right directions at critical junctures as needed, in keeping the CCC Council informed of progress, and, ultimately, in ensuring that the visioning activity is completed end-to-end in a timely manner.

Some thoughts on research communities

Oversimplifying only a little, there are two different kinds of research communities. First are those dedicated to *fundamental research* problems (e.g., networking, operating systems, VLSI design, data mining, etc.). These communities are traditionally “reinforced by” computer science and engineering education, where there are one or more courses for each such long-standing community. At times, a whole new fundamental research community can be stood up. For example, social computing is one that has emerged recently, and there is already a body of algorithms and heuristics that demonstrate the considerable intellectual depth of this new area.

Second are communities dedicated to *integrative problem solution research* (e.g., robotics, unmanned air and sea vehicles, and synthetic biology). Communities have grown up around each. There is typically much more industry involvement, in at least the “D” of R&D related to integrative problems. There are people involved in these integrative communities who are also part of a fundamental research community, because

integrative problem solutions rely on expertise from multiple such sub-fields. Note that the integrative problem community may be more attuned to directly addressing a specific mission of an agency.

There is an often repeated, evolving debate in the computing field over what the field does and/or should "include." For example, in the 1990s some distinguished people in the computing research field felt that "digital libraries" was not really part of the field, and concerns were raised about substantial projects funded by NSF (CISE), NASA, and DARPA on digital libraries. In retrospect we can see that at least three billion-dollar information technology industries came from or were influenced by that work: large-scale search, GPS-enabled consumer products, and the digitally-enabled publishing business. So, today, that specific debate is concluded. But, debates about other domains arise; it is part of the vitality of the field.

The CCC visioning projects have embraced topics such as IT in economic development (especially in developing countries) and application of IT to K-12 education. The CCC Council has deliberately taken a "big-tent" view of the field. It is best to explore a new domain, rather than rule it out based on definitions of the past.

Growing a new community takes time. CCC aims to nurture both new and old communities. Yet a community needs to be defined with a scope that makes it possible for that community to find robust funding for the research that it seeks to perform.

Metrics for Visioning

CCC is just starting to explore metrics for evaluating the quality of both individual visioning efforts, as well as the overall CCC visioning activity. Metrics include:

- number of and quality of new research challenges, problems, or directions proposed
- new teams of researchers established
- students/teaching programs affected
- existing research programs affected
- number of agencies engaged
- new research programs proposed and/or implemented
- Federal dollars affected