

Meeting report
FIND Investigators Meeting
April, 2008

Summary:

This report is a summary of the FIND PI meeting held in Los Angeles on April 8/9, 2008.

The original planning for the FIND program suggested that new sort of research activities (what have been called “Phase 2 activities”) be added to the research portfolio over time. Under the assumption that it was now time to begin these activities, the FIND Program Committee organized this meeting to discuss what new sorts of activities should be undertaken to encourage the growth and success of the FIND program. It is hoped that this report will stimulate the community to contemplate and propose research of the sorts discussed here.

The attendees at the meeting concluded that the success of the phase two activities will depend on what signals NSF sends to the research community about the range of activities they will encourage and foster in the next solicitation. This report contains a catalog of possible sorts of research activities; NSF must indicate which of these it wants to encourage. As the report makes clear, the posture of the community is a mix of excitement and well-justified uncertainties, including fears about taking career risks in daunting times by doing something that is more unfamiliar, more uncertain, and thus less predictable. There must be clarity as to what sorts of research proposals will be encouraged in order to mitigate the risk of participation in FIND.

The discussions at the meeting took place in the context of some larger questions, most centrally how the community is prepared to proceed with the larger challenge of FIND, as outlined below. The vision of FIND includes both intellectual and structural transformation of the research agenda surrounding future networks. The research supported in the first phase of FIND is intellectually distinctive, distinguished by the requirement for long-range concepts unfettered by the requirement for immediate application to the current Internet. Phase one work was intended to inform and educate the community, and build a foundation for the future research, but it is phase two that will mark the point where the structural distinctions of FIND emerge. Phase two thus challenges the research community to put forward research ideas of a different sort and character.

Background to the meeting:

The FIND program, as framed in its early days, was a challenge to the research community to envision what a global network 15 years in the future should look like. FIND was formulated to meet a perceived national and international need—to propose new architectures that will shape the future of networking, and allow those networks to better meet future demands. NSF has stated that through FIND its research community can make a difference in this arena. It was a challenge to set a goal—a vision of a future network—and work toward it.

In order to achieve this larger agenda, FIND was framed as having three phases, each with increasing breadth of activities and ambitions¹. In phase one, the community would be encouraged to argue about requirements for a new architecture, to propose, develop and demonstrate new ideas that might be a part of a new architecture, and to pursue individual research. When the second phase began, some number of architecture teams would begin to assemble integrated and coherent proposals for candidate architectures, drawing on the insights from the work put forward in phase one. In the third phase, implementation teams would reduce selected architectures to practice—fleshing out the details of protocols where this remains undone, and producing code that could run on an evaluation infrastructure such as GENI.

At the time FIND was proposed, these next steps, the articulation and demonstration of candidate future networks, were seen as an integral and critical part of doing long-range disruptive research. Ideas that are disruptive are often initially rejected by industry (and others with short-term interests). One cannot take long-range disruptive ideas and “throw them over the fence” by publishing them in a conference and expect to have a big impact. As FIND was framed, taking these next step was part of a commitment to the research community: FIND would include integration of new ideas into future networks in order to improve the chances that the various ideas put forward would find fertile ground. Without this commitment to sustain the new ideas that might emerge from long-range thinking, there might be too much career risk for individuals to think long-range.

There have now been three solicitations for the first phase of FIND. The FIND planning committee concluded that it was time to discuss what the second phase should look like: what new sorts of activities should be added to the research portfolio in order to move FIND forward. We wanted to challenge and elaborate the initial simple conception of phase two. The agenda of this meeting was constructed in order to allow this discussion. There was an initial plenary session to launch the meeting, and two sequential breakout sessions for discussion, each followed by an extended plenary discussion.

This summary does not recapitulate the different breakout sessions, but provides an integrated summary of what was discussed.

Defining phase two: While the early concept of phase two had researchers starting to think about integrated proposals for future networks, this meeting invited a fresh, unconstrained discussion about phase two. The challenge question was posed as follows: what sorts of additional activities should be added to the sorts of research now being funded in order to move FIND forward, and lead to the eventual accomplishment of its long-term goals.

At a high level, an answer is that phase two must encourage and support **integrative design**, moving toward an eventual outcome where one or more proposals for a future network emerge. But the simple answer of integrative design does not capture the range of activities that will constitute (and support) that outcome.

¹ The term “phase” caused some confusion, and probably should be replaced with a different word. “Scope”, or “stage” were proposed. The issue is that “phase” implies to some a strict timing: phase one would end when a phase two starts. This strict timing was never intended, as the early documentation makes clear, but the word out of context might mislead. The idea is that new sorts of activities would be added to the mix of research over time, while the original activities would continue to be funded.

The group, in trying to explore this process of integrative design, used the word “modalities”, as distinct from “topics” to characterize phase two. The view was that phase one of FIND addressed a variety of network topics in a very traditional modality—single PI grants or small collaborations that had specific research objectives that would seem familiar to a normal merit review process. In this respect, FIND (in its first phase) is not that structurally different from other NSF programs—while we have meetings and encourage community-building, the individual research grants are framed in a familiar way. It will be as we identify and add new sorts of activities—what we called new modalities—that FIND will take on a different structural character.

While phase two is about integrative design, we should not see phase two as the simple weaving together of ideas from the first phase of research, but a process that is perhaps “informed by” phase one research (and other research) and which then puts forward the key ideas that will define a new network. We should recognize both the importance and the challenge of this next phase of thinking. However, it may not resemble “traditional” research, and so there is a question of how NSF would write a solicitation that would invite this sort of work.

Here are the examples of research modalities that were identified in this workshop.

- **Design—the core task:** Some members of the community seem prepared to start a project with the direct goal of integrative design, either looking at a full range of requirements and design elements, or a subset that advances the process of integrative thinking. While some attendees seemed prepared to step up to this challenge, it was not clear that these folks were clear as to what this activity would look like, in terms of scope, participation, and structure.

Many of the activities/modalities listed below can be seen, in some sense, as supporting the core design challenge. However, they are distinctive and need to be recognized and encouraged in their own right.

- **Finding common principles:** Common principles are basic—they come before consideration of objectives such as improved security, and provide constraints for the discussion of those objectives. Using examples from the past, the Internet is defined by the use of packets as a common transport mechanism. Perhaps after the fact, the end-to-end argument was identified as a core principle of the Internet. It might be that we will find new such principles as well as validating old ones, as we contemplate a new network.
- **Cross-cutting objectives:** Some of the objectives for a future network, such as security, management, or availability, were seen as cross-cutting: they cut across layers, different types of applications, and different research communities. So they require special attention. FIND needs to make sure that these are given special study, to see if we can find high-level approaches and frameworks to address them. Then we need to see how they interact (or perhaps conflict with) each other. These activities may well require a considerable amount of intellectual effort.
- **Multi-disciplinary studies on non-technical issues:** The group sees that many of the problems and constraints faced by the Internet are non-technical, and will require solutions that at best are only partially technical. Understanding and resolving these issues will require a range of skills. The FIND solicitation has always welcomed multi-disciplinary work to address these issues, but aside from collaboration with economists,

the portfolio is thin. This sort of work must be encouraged as we expand the FIND portfolio.

- **Tools to validate architecture proposals:** As we begin to propose designs for a future network, we need at the same time to think about how to validate these designs. As part of phase two, we should begin to think about this. The group identified a number of approaches, from the informal (structured discussion and analysis), formal methods, simulation, and prototyping. The latter, which involves building and trying real systems, will demand significant resources, but in the end, promising ideas will have to be built to be evaluated.
- **Larger development projects:** Some projects may include prototyping and significant design efforts. These will require the use of staff and post-docs, as well as graduate students and traditional levels of faculty time. Some faculty may desire to buy out of some teaching responsibilities in order to concentrate on this research. NSF needs to recognize that certain sorts of phase two activities will require higher levels of funding, and support for soft-money personnel.
- **Design of systems to support restricted application classes:** to the extent that the community is not yet ready to propose a “general purpose” future network (or to the extent that virtualization reduces the power of a general network) it may be useful to go through the exercise of designing systems that address the full range of requirements, but in the context of a restricted set of applications with a restricted set of requirements.
- **“Pollination”:** this perhaps cute term was coined to describe a process of community building in which certain folks take on the task of visiting research sites in turn, picking up good ideas and passing them on as they travel. It was seen as a good tool to help big ideas form out of smaller ones, and for potential collaborations to form.

This group was not prepared to say that every one of these activities were critical at this point. Rather, this list should be seen as scoping the potential range of new modalities that should be contemplated as NSF constructs a new FIND solicitation.

Contemplation of this range of activities in turn triggered several further discussion topics.

Fear of proposing phase two projects: Because phase two is seen as embracing this broad set of new research modalities, there were serious concerns that unless the selection/evaluation criteria for phase two research are very clear, there would be too much uncertainty that proposals in these new research modalities would be rejected by the merit review process. These modalities might seem unfamiliar to a typical merit review panel, and even if these activities require serious intellectual engagement and generate important new ideas, the activity might not resemble “traditional” research. So the necessary risk of writing a proposal that earns a low rating is compounded by the concern that the ideas will be rejected as unfamiliar. NSF must mitigate these fears or phase two may go unheeded. In particular, if phase two FIND activities are included in the “three per year” rule, unless NSF can really reduce the uncertainty of writing phase two proposals, they will not be submitted. Resolving this problem is an example of where NSF must show leadership to ensure the success of FIND.

NSF leadership: The attendees encouraged NSF to take on a leadership role as FIND develops. This issue generated considerable discussion. At one level, there is agreement that NSF is not a mission agency, and should not be picking a winner among options for a future network. NSF has always encouraged the emergence of multiple ideas: it is this multiplicity that is part of the motivation for the generality of GENI. But the group felt that some leadership and guidance would be required for the goals of FIND to be achieved. The central issue we discussed is how the FIND program, or specific proposals within FIND, can build themselves a balanced research portfolio to achieve their goals. As we look at the phase two activities, and as we look into the future at how individual ideas come into focus, the group felt that if successful proposals are selected based only on individual assessment of merit, without regard for how they fit into an overall portfolio, then the activities will of necessity be unbalanced.

There was discussion of ways this has been addressed previously and might be addressed now: cooperative agreements that allow for subcontracts, well-constructed centers, special evaluation criteria that allow the merit review committee to take into account the overall objectives of FIND as well as the absolute merit of each proposal, and so on. The group had no opinion about preferred options, but encouraged NSF to consider what might be the best ways to allow a vigorous research portfolio and strong leadership to emerge within FIND.

Gauging the success of FIND: While the FIND program is really only two years old, probably too soon to evaluation success in any meaningful way, we spent some time discussing our initial reactions about the program, and how success might be measured in the future.

Several folks said that at a “gut level”, FIND was being successful, but we need to agree on metrics to answer that question, and after two years, it was too soon to be asking for this sort of assessment. One topic of discussion was how our work would get attention: how it would have impact and gain recognition. The original idea of FIND is that NSF would fund not just the basic research and writing of papers, but the realization of the ideas in running code and the deployment of suitable experimental infrastructure. We discussed the question of whether this path continues to be the preferred and practical one.

One view is FIND will be measured by impact and impact is measured with respect to industry. So we should invite industry to our meetings, and seek advice from them about what problems need to be solved. We should arrange for our graduate students to interact with industry so that they learn first hand about real issues, as opposed to learning indirectly through their advisors.

Another view is that our goal is fundamental knowledge and understanding, and the benefit to industry will emerge indirectly over time based on this knowledge. It was also argued that we were on a different time frame from industry.

The diversity of opinion included the position that direct and immediate impact on industry should be seen as a negative signal for a project with this time horizon and objectives. This argues that we must continue to support the concept that good ideas in FIND will be implemented as running systems on distributed experimental infrastructure.

Related to the question of impact is the question of whether we are well-enough informed to be solving the right problems. Do we know what problems are “real”. Again, some argue that industry involvement is important to confirm that we are solving real problems, others feel that we have a strong base of experience in the research community. It was said that “all the problems are right there in front of us”.

We discussed the makeup of our community and the diversity of skills and inclinations. FIND should be constructed in such a way that we exploit the range of skills of the community, which implies first that it should support a range of research modalities and second that as FIND progresses, new players with new skills may enter the program, and should be explicitly sought out.

Continued community building: NSF contemplates that the sort of grants now being made as part of FIND (so-called phase one grants) will continue, which implies that new grant recipients will be entering the community. Additionally, we need to encourage people who may not have chosen to participate in the early part of FIND to consider participation as phase two emerges. So there is a continued need to build a community. As the community grows, “all hands PI” meetings may become less effective. Perhaps the collection of phase two events can serve as a basis for community building, but this issue, while discussed, was not resolved, and remains as a challenge for the Planning Committee.