

FIND Observer Panel Report

Vint Cerf Bruce Davie Albert Greenberg Susan Landau David Sincoskie

April 9, 2009

Executive Summary

A panel of five external senior researchers (Cerf, Davie, Greenberg, Landau, Sincoskie) was tasked to (1) gauge the quality, the potential for game-changing impact, and the intellectual merit of the NSF Future Internet Design (FIND) program, and to (2) recommend whether and how NSF should continue with FIND.

Our findings are as follows:

First, the panel is pleased and encouraged by the results presented on over 30 projects, over the course of the April 6–7, 2009 workshop sponsored by NSF. FIND has had a refreshing and liberating impact on network architecture research — refreshing in the sense that architectural gaps in the Internet architecture have been identified, and liberating in the sense that researchers are working on long-term target of a better Internet, not short term myopic improvements. New ground is being broken over a wide range of core networking areas, e.g., naming, addressing, routing, monitoring, mobility, network management, access and transport technologies, sensing, content and media delivery, and networked applications.

FIND has now run for three years, and has ramped up to its current level of funding 49 projects, each roughly at level of 500 thousand to one million dollars over three to four years. The outcomes are important research contributions to a wide set of point problems. This work illuminates components of what the future Internet might look like. Common understanding of requirements and competing designs are starting to emerge for these components (e.g., naming), though we are far from the level of understanding needed to pick coherent architectural alternatives among the ideas and prototype designs.

The panel has three recommendations to NSF:

- Continue.
- Additional focus on security and network management. Security and network management represent architectural gaps in today's Internet, and the industry's band-aids are inadequate. Foster research to tackle the interdisciplinary and hard problems of getting the network to reach a level of security and robustness expected of critical infrastructure.
- Integrate. It is time now to ask the community to come up with self-forming teams that integrate a large number of components of the future Internet. This is *in addition*

to the basic research work on components or point solutions. The time is right for members of the networking research community to develop a set of coherent architectural alternatives and to *implement* them. The panel envisions as many as four to five teams, funded roughly at a level of 10 to 20 million dollars per team over the three-to-four year period¹².

1 FIND status and evaluation

The panel's overall evaluation of the FIND program was strongly positive. In the following paragraphs, we highlight first those aspects of the program that we found most encouraging; subsequently, we discuss some areas where more work or a different focus is required.

1.1 Positive aspects

It was clear from the research reports that we heard and read, and from comments made by many of the researchers, that FIND has effectively put architectural research "back on the map". The majority of funded projects tackle significant architectural issues. Furthermore, most of the projects seem to have taken the "clean slate" approach, freeing the researchers from thinking only about incremental improvements to today's technologies.

Not only has FIND enabled a significant amount of architectural research, it has also led to some advances in developing theories of network architecture. We believe a stronger theoretical underpinning for architectural research will bring considerable benefits to the field of networking in the future.

By encouraging a greater number of researchers to undertake architectural work, and providing a venue in which they can meet with each other, the program has also increased the opportunities for collaboration. While this is providing some benefits already, active steps should be taken to ensure greater collaboration in the later phases of FIND.

We have heard in the past of a concern that architectural research that is not incremental might be considered "too risky" to be undertaken by junior faculty members. This was not perceived to be an issue by the workshop attendees. By providing a solid funding stream and a community engaged in architectural research, FIND has enabled a generation of younger faculty to pursue research focused on network architecture.

Several researchers have noted that the architectural thinking that is encouraged by FIND has directly influenced their teaching in a positive way. Rather than teaching "from the RFCs" (that is, teaching only how networks work today), researchers are teaching their students to think about how networking technologies and protocols could be designed and integrated. Students are learning architectural principles and how to solve problems in networking rather than just studying today's artifacts. We urge

¹This number is very rough, but recognizes the likelihood that significant implementation projects will require teams with staff as well as P.I.s and student.

²Dave Clark has written a paper on the required elements. The panel received these requirements favorably, and has added some of its requirements and desired targets.

the FIND community to convey this approach to teaching networking to the larger networking community.

1.2 Areas of Growth

1.2.1 Security

One of the lessons learned from the Internet is that security cannot be an add-on. While in the FIND program, it is understood that any Internet redesign **must** be secure, not all the research efforts have placed security front and center. In addition, it seems that many proposed models have not been subject to a thorough security analysis. Security needs to be significantly more central in the FIND research effort.

1.2.2 Network Management

Currently Internet management is characterized by a lack of information of network status and health, a deluge of data (at once voluminous, ambiguous, incomplete, and inconsistent), and blunt or imprecise actuators or control actions whose impact is hard to predict. Silent failures and hidden dependencies are common. To transform today's status quo, a future Internet requires deeply ambitious research in network management.

Tools that can be used include simple declarative policy specifications to assure correct behavior at massive scale, statistical machine learning (which has been successfully applied to program and computer architecture fault analysis), and massive automation based on recovery oriented computing (recovery through restart, reboot, re-image, return approaches). Whether through the use of these methods, through the creation of self-diagnosing protocols, through more fine-grained and informative data collection, through cross-layer event tracing or monitoring, or through unanticipated new research, we must achieve the same gains in network management.

1.2.3 Incentives

In today's Internet, end users and their applications can clash badly with the network infrastructure, leading to unfortunate outcomes. Comcast's actions in blocking BitTorrent, and the subsequent reaction, is one notable example of such conflict. Yet with proper understanding of the provider's scarce resources, the incentives of the providers, and the goals of the users and their applications, better win/win solutions might be found, which improve things along all dimensions simultaneously (delay, bandwidth, user experience, etc). Mechanisms based on the understanding of the incentives of all players are now emerging.

More research is needed, however. The network architecture provides little motivation for network users and providers to openly reveal their costs and goals. The lack of clarity raises the barrier of entry for new services, and thus limits the Internet's evolution. Existing theories (based on game theory and economics) do not always map well onto Internet realities. The panel recommends more research on understanding the incentives and the economics of networks for developing mechanisms to guide protocol design for the future Internet.

1.2.4 Users

The network user behavior needs to be considered as a holistic part of the entire network system, almost as an OSI layer 8. Applications can and do include users behavior and such behavior can be understood as closing certain feedback loops that are important to network function. Classic examples from telephony include the role of sidetone in adjusting user speaking volume, and delayed dialtone as a way of influencing demand for circuits. Diurnal patterns in network traffic loads are often taken into account when engineering network capacity. Mass events drive mobility and hence network load. Privacy considerations are affecting location based services. However, most research focuses on the behavior of the machines (communications and computational) that constitute the network, while leaving the user behavior largely out of the picture. Users behavior can have significant impact upon the performance and function of the network as a whole. For example, many of the security attacks propagated across the internet today involve some form of social engineering, identity spoofing, or other forms of abuse of user trust to enable the attack. Research that includes or models the user as a part of the entire network system should be emphasized in future work.

2 FIND Next Steps and Challenges

In order to create the future Internet from a clean slate, FIND has taken a “let a thousand flowers bloom” approach. The signal that came through most clearly to the panel is the support from the community for that approach. In the panel’s judgment, the ideas and prototypes that have emerged from FIND are still preliminary. The panel strongly recommends that FIND continue to support this fundamental research.

The panel proposes that FIND expand to fund a small number of new integrative team efforts that take a coherent approach toward attacking the list of research questions in the following section. Specifically, we envision funding as many as 4 to 5 teams at a level of 10 to 20 million dollars each. This would roughly triple the FIND budget.

The type of research in Internet architecture being sponsored in FIND is complex and difficult to transfer into practice. Also, it often requires consensus (at least rough consensus) and experimental verification (running code) as validation of the research and a necessary prerequisite for further consideration. In short, this is a highly experimental discipline, often requiring complex experiments. However, the academic research community in FIND faces many institutional and cultural obstacles to organizing into large, interdependent teams that are often necessary to mature the research. Incentives, or programmatic mechanisms that encourage self-organization may thus be in order.

Many such mechanisms exist and there are precedents for use of these mechanisms in the history of research in this field. The panel recommends that NSF require the integrated teams demonstrate system operation — on a regular technical basis — as a condition for continued funding. To be frank, integration should be real.

As FIND moves into its integrative phase, there will be a clear need for infrastructure. At this point, it is probably too early to say exactly what that infrastructure should

be. We believe that integrated FIND experiments need not be limited to GENI for infrastructure support. The panel heard that integrative teams may want to pursue other alternatives.

We are aware that unlike the original work that resulted in the Internet, future Internet use will occur in an international context. Cultural differences will affect the acceptance of such seemingly technical decisions as naming, privacy protections, etc. Thus international input early in the design effort is very important. Although NSF funds U.S. based research and thus the agency is limited in how it can enable international collaboration, we suggest that NSF find ways to encourage and enable international discussion and collaboration in the FIND effort.

Some FIND researchers have said that because their work is architectural design, they have had difficulty in having their work appear at security conferences. Yet such evaluation and serious vetting by the security community is necessary early in the design process. One possibility would be the creation of a joint annual workshop providing a venue in which both architecture design and security researchers meet.

3 Research Questions

FIND was created to explore the design of a Future Internet, motivated by a number of real shortcomings of the current Internet. The following list³ is representative of the motivating factors:

- Security
- Availability and resilience
- Better management
- Economic viability
- Longevity
- Meet societys needs
- Support for tomorrows computing
- Exploit tomorrows networking
- Support tomorrows applications
- Fit for purpose (it works)

For further explanation of these items we refer the reader to the presentation and to the in-progress document “Toward the design of a Future Internet” by David Clark.

The following paragraphs list a set of topics the FIND community should be tackling. Note that many of these are already receiving considerable attention in the existing FIND program. Ideally, the integrative projects should address as many of these topics

³This list is drawn from the presentation “Architecture from the Top Down” made by Dr. Clark at the April 6–7 meeting.

as possible. We consider these topics important but recognize that equally important topics may have been omitted.

- **Security.** Security is one of the fundamental reasons for a clean-slate Internet design. The goals of the FIND work should be an available, robust architecture enabling confidentiality, integrity, privacy, and strong authentication. A Red team approach to security may be appropriate in some cases.
- **Identity.** In order to handle issues ranging from spam to DoS attacks, and trust, a future internet must include appropriate mechanisms for identity. The key word here is “appropriate”; in some instances identity should be available at the packet level, in other instances, it should only be available to the communicating endpoints, and the architecture should be flexible enough to enable this. This is an area in which FIND researchers must collaborate with both security researchers and with researchers well versed in the social aspects of the research (e.g., lawyers, public policy researchers).
- **Privacy.** With the combination of mobile devices, sensors, and ubiquitous communication in modern life, privacy concerns permeate the FIND work and thus mechanisms for preservation of privacy should be directly addressed in the FIND research.
- **Scaling.** At the risk of stating the obvious, research on a Future Internet must take account of the likely scaling issues, with billions of users and probably orders of magnitudes more devices than users. It is important that research efforts under the FIND initiative identify limits to designs as various parametric values (e.g. latency) reach extremes.
- **Mobility.** A future Internet should provide first-class support for mobile devices, including persistence of processes, connections, references, etc.
- **Theoretical Foundation.** There is already some impressive theoretical work being undertaken as part of FIND. As an example, a deeper understanding of Internet traffic analysis to the “Erlang” formulas for telephony would be a significant advance.
- **Realistic economic models.** Many economic models are simplified to make them more mathematically tractable. To the extent that research on economic modeling is part of the FIND effort, it is important to find ways to validate these models through prediction and measurement, where this is possible.
- **Cloud Computing.** Cloud computing has opened up new challenges and opportunities to rethink computing, storage, and networking infrastructure. We need better mechanisms to map users and applications to cloud infrastructure (data centers, proxies, caches, etc.) to hide latency and provide a user experience on par with on-site infrastructure. We need information and network resilience at every level; operating under the assumptions that every element is on its way to failure, and that information is partitioned, replicated and in motion. The value

proposition of the cloud rests on achieving positive economics, which together with the heightened degree of control over much of the infrastructure, should drive innovative design.

- **Virtualization.** While it has been a popular notion to virtualize physical resources there are side-effects to treating all physical resources as somehow equivalent. Specific physical parameters (bandwidth, memory, processor speeds, etc.) may have a dramatic impact on the performance of an ensemble of diverse virtual components. Research is needed to assess how to characterize virtual resources to achieve optimal utilization while preserving the apparent equivalence of the virtual components.
- **Internationalization** The Internet is used by about 23% of the world's population and incorporating the scripts needed to support a significant part of the world's languages in various parts of the Internet is important. While there is existing IETF work in this space, elaboration of the use of non-ASCII scripts in various parts of a re-designed Internet is an important goal and poses significant design challenges.
- **Open Platforms.** It is liberating and game changing to have largely unfettered access to switching and routing platforms, whether through APIs, SDKs, or source code. Nascent efforts have emerged that provide some of this from Click, Xorp, Quagga, OpenFlow, NetFPGA, VINI, etc. These efforts put important new tools in the hands of researchers, allowing real experiments that would otherwise be unthinkable. Moreover, researchers are starting to see which primitives uncovered through these experiments should be committed to high speed hardware. We see this sort of flexibility as a large enabler of the FIND program, in particular to support larger integrative efforts.
- **Deployable QoS.** While there has probably been an overabundance of QoS research in the past, QoS deployment in the Internet has been almost non-existent. Future QoS research must take account of issues of economics, regulation, incentives etc.
- **Role of Layering.** It is important to reconsider the role of layers in network architecture. Recent work has looked at explicit cross layer communication on wireless networks, for example. Taking better advantage of links that naturally support broadcast is another example of research in this area.
- **Applications-focused Network Research.** Applications are often left out of architectural considerations and their effective operation may depend strongly on underlying network features. Application architecture research has the potential to improve application inter-working and ability to use underlying network resources and functionality.
- **Operational Aspects.** It is difficult as a researcher to get large scale, commercial network operations experience, and yet such experience is invaluable in maturing a researchers understanding of the hard research questions. Its recommended

that the community develop mechanisms that allow access by researchers to hard industry problems faced by commercial entities.

- **Access to real data.** Access to networking data is becoming increasingly difficult to get within the research community. How will the community ensure some level of access to large datasets?
- **Triggers for change.** What are the trigger mechanisms for architectural change? Large scale networks require enormous capital investments. Causing change in these investments is complex and difficult, and will not happen purely based upon technical merit. Research into predicting the trigger mechanisms (and not purely economic triggers) might be encouraged.
- **Societal context.** As noted by David Clark, network researchers need to be aware that networks are now part of the fabric of society. Research should take account of the societal context. An example: a tussle exists between those who wish to have unfettered access to all data, those who wish to censor communications, and those who wish to limit access to “objectionable” material. As noted above, it is important not to “bake in” a fixed set of social norms to the network architecture.

4 Conclusions

We live in a world in which the Internet has become the communications medium not only for business and private citizens, but also for governments and critical infrastructure. The risks for society are great if the fragility and security problems of the current Internet are not understood and technologies for addressing them developed.

The goals of the FIND work — designing a viable⁴ architecture for a secure, available, resilient network enabling confidentiality, integrity, and privacy — are highly ambitious. The FIND program has the potential to mitigate many of the risks. Our conclusions are that current FIND work is good, but still preliminary. More research is needed, but simultaneously work should begin on testing some of the ideas being proposed.

⁴Viable includes not only technically viable, but also economically so.