Readings on alternative futures for the Internet

Feb 15, 2006

We have discussed in class the fact that the Internet seems so well-established, that it is hard to conceive of alternative structures for the Internet—different technical designs with different implications for the larger social and economic context.

Here are two readings that might help stimulate the imagination.

The first is a list of discussion topics that will define the working agenda of an upcoming multi-disciplinary workshop on the future of the Internet. The topics were picked because they emphasize the need for cross-discipline assessment and discussion.

The second was prepared to illustrate to the National Science Foundation what sort of research might be undertaken as part of the recently-announced Future Internet Design (FIND) project. This document was prepared for a more technical audience, and the list of issues will vary as to what their impact would be in the larger context. It might be an interesting exercise to go through this list and look for policy issues that would arise as a part of these projects.

Discussion topics for a workshop on Future Internet Design

The goal of this workshop is to explore a new space of multi-disciplinary design, and the related question of process—how can a multi-disciplinary team function effectively to contemplate design alternatives for a future Internet?

The workshop will consider several case studies, which will be the basis for a series of mini-design sessions. Each case study will describe a set of design alternatives, which we will consider in terms of both the technical options and the social or economic implications. Below is a list of possible questions that this workshop might consider.

Identity "in the network"

In the future, to what extent should an observer "in the network" be able to identify the sender of a packet? What are the social and economic implications of different degrees of identity and anonymity in the network? How would the presence of identity information change current uses of and interactions on the Internet? Should this ability be used to develop a policing system or "surveillance cameras" for cyberspace?

Today, the source IP address in the packet gives some hint as to the identity of the sender. Precisely, it indicates only the network location (not the physical location) of the sender; however, this information can often be translated into some better signal of identity. This link to identity may get much weaker. With increasing mobility of end-nodes, the IP address will become more of a transient indicator of location. The location may change rapidly over time, even for an ongoing connection. Given this trend, there is essentially nothing visible in a packet that can link it back to the identity of the sender.

There will, of course, be identifying information handed back and forth between the end-points of a connection. At a minimum, there will need to be some sort of identifier that the end points use to keep track of each other as their location changes. Some end-points may demand very robust identification from the other end-points. But while this information is carried among the end-points inside packets, it may not be visible—it may not be in a standard recognized format and it may be encrypted. Further, it may not be in *all* the packets of the connection. So while the end-nodes may know each other's identity, observers in the network may not. Given this fact, a "surveillance camera" in cyberspace may be able to see nothing except encrypted packets going between transient locations. We could choose to include explicit identity information in packets, but how might we minimize the chilling effect that is associated with loss of anonymity and analogs such as identity cards and travel papers?

Hence the question: are there significant reasons for packets to carry some sort of identity that is visible "in the network".

Technical form of the question: Should there be a field in the packet header that identifies the sender of a packet? If ISPs want to give different users different services, what information in the packet should be used to signal that discrimination—session setup or packet state? If the packet had an identity field in the header, do we need to fully specify its semantics, or would it make sense to use the field in different ways at different times and places?

Identity in e-mail

In the Future Internet, when person A sends an email to person B, what level of information should B expect to have about A? With what level of certainty? How would acting to strengthen confidence/knowledge of identity in the Future Internet affect users? Would the benefits of stronger identity outweigh the possible negative consequences?

Today, users of email identify themselves to each other using their email addresses, strings such as user@example.net. Between cooperating users, these work well, but they can be forged. So a certain amount of spam is sent using a forged sender email address, and the mail classified as "phishing" has false sender email addresses. There have been calls to "improve" this situation, and give person B a higher level of confidence that person A is actually person A, and a higher level of knowledge as to who person A actually is.

If it is generally agreed that identity needs to be strengthened, further questions remain as to how this would be accomplished. One option is that all users should be issued a "top-down" identity—that is, an identifier issued by some trusted party, who is in turn validated by some higher-level trusted party, perhaps leading to a chain of trustworthy identities that can be traced back to a nation-state. This identity could be used to identify the sender to the receiver. In a system like this, the question is: What attributes of identity are included (and verified) by this issued identity? Would it be just the given name of the person? Name, address and years at that address? Age? Gender?

Another view is that we use "bottom-up" identities. In a system like this, users issue themselves identifiers. If properly designed, these cannot be forged, but of course since they are self-issued, they do not tell the receiver anything about the sender. All they allow is a sender to prove that he is the same sender as in the last message. They allow a user to connect a series of interactions, and "get to know" someone by that means.

Technical form of the questions: Should we move toward a regime of mail signed with public/private key pairs? If so, should we presume a single PKI, several unrelated PKIs, or self-signed certificates. How can we present identity information in the user interface of the mail system so that the user can make sensible decisions based on the information?

Ports and application visibility

Today, the Internet names services (such as Web or email) using "well-known ports"—numerical indices that are statically assigned to each application and service at design time. Since these port numbers are included in each packet, this permits any observer in the network to determine what application is being used. And since these numbers are statically assigned, an attacker can easily launch an attack against an application on a given host. It need not be this way--an alternative would be to design a new mechanism for "service rendezvous", and to use random port numbers to identify connections. What would the implications be if packets did not reveal what applications they were associated with, and security tools could not associate applications with ports.

There are implications for economics and security if it is not possible to tell what application the users are running by looking at the packets "in the network"? ISPs observe port numbers to build models of what their users are doing. This information may help with capacity planning, and is the basis of discrimination among different classes of users that wish to use different services.

From a security perspective, well-known ports are a fundamental component of both attack and defense. Attackers combine host addresses and well-know ports to explore whether a host is running a version of the service with a known vulnerability. Firewalls respond by blocking access to specified port numbers. If port numbers were random, this could essentially eliminate the value of the attack known as port-scanning, but it would change the whole security landscape by limiting what firewalls can do based on packet inspection.

So the question: should a future Internet employ a different scheme for session rendezvous and connection identification.

Technical form of the question: Should the Internet move to a mechanism such as DNS-SERV to assign port numbers to services on hosts? Should the port field be increased, perhaps to 32 bits? Should we move to an application-specific mechanism for rendezvous and connection management? Should the port number be encrypted in some way?

Validating the authorship of web pages and other documents.

Today, we establish the authorship (or authority) of a piece of information indirectly. The way we confirm that a web page is (for example) created by CNN is that we download it from CNN. The confirmation of the authorship is associated with how we got it, not with the object itself. If A downloads a web page and then send it to B, there nothing about the web page itself that confirms to B that it is legitimate. Should the Internet include a way for authors to sign Web pages and other net documents, and how should this namespace be controlled?

There are lots of reasons why it would be important to validate the authorship of a piece of information independently of how one gets it. Today, bits of information are sent around from person to person as email attachments. The research community has proposed a number of creative schemes, such as peer-to-peer retrieval system, to support efficient and robust information dissemination. And as we more and more depend on search engines such as Google to find things, we have less and less confidence about the origin of anything we download. But today, none of these schemes can confirm the validity of the information they propagate. To mitigate these problems, what Is needed is a scheme that associates the authority with the object, not how it is fetched.

It is not difficult to conceive a scheme to sign documents with some sort of encryption scheme. The hard part of this problem is to devise the namespace in which the authors are named. Such a system has to resist forgery, be easy to manage, and easy for users to understand. Today we use DNS names and public key certificates to verify the source of data, and at first blush, we could just use those same names as names of authors. However, the ongoing debate about governance of the DNS suggests that we might consider designing a name space that sidesteps some of the tough political tangles that have snared the DNS.

One choice for naming is that there is some sort of registry of author names, similar to the registry of domain names. Provided the registry is trustworthy, this approach gives the receiver of an object some confidence about a page from a previously unknown author, but raises all the issues we have today with control and governance of the registry. Another approach is that names are not registered at all, but are just validated pairwise between sender and receiver. That is, a receiver comes to trust a sender by developing a relationship with that sender. The names provide continuity, but not the trust in the relationship. This concept is similar to the discussion about "bottom-up" signing of email.

Gated communities and virtual networks

The current Internet started out as "one big network", where everyone could send to anyone. Today, most of the Internet technology is not in the public Internet, but in private versions of the Internet, such as corporate networks, which are connected to the public Internet only in very constrained ways. The public Internet itself might go in this direction, fracturing into groups that self-define in some way, and which link to each other only in very constrained ways. In balance, would this be a good idea or a bad idea?

Today, the public Internet is "open", in that anyone can send to anyone. But different applications use this open basis in different ways. Email is fairly open, instant messaging is less so, and applications can impose very complex restrictions on which parties can communicate with each other. There is no reason to think that the "open by default" model is the one toward which the Internet will gravitate.

The question of "how open?" can be raised at several levels—at the packet level, at the application level, and at the "social" level. The Web and its possible futures may be a

case study that can stand for all the Internet. The original idea of the Web was that any page could link to any other page, that any user could be a producer and make fresh linkages in the Web, and that the power of the Web was its open nature. Of course, this simple picture is imperfect. Corporate information (which may be proprietary) is not open to this sort of linkage. Some ISPs are exploring the business opportunities that might arise from providing selective access to some content over other. There seems to be a trend for users to seek information from sources that they prefer. And some countries are demanding that their citizens be protected from seeing forbidden content.

Given that users only interact with the Internet through that applications that run on it, if applications control the degree to which the Internet is open or closed, has there been a loss of social value? Should application designers include features that can be used to control the reach of openness? Perhaps the loss of social value is related to which parties can control the degree of openness, and the issue is that if these controls are added at all, it is difficult to control which parties can control them. Should application designers concern themselves with the question of how their design influences which parties have control over the degree of openness. Should we be concerned with the design of search tools, and how they lead users to one or another answer to their queries?

Managing location information

There is no notion of physical location in the design of today's Internet. But with the increasing use of mobile and wireless devices, physical location will become of increasing importance. Today, the cellular system has been required to determine the location of every cell phone so they can respond properly to a 911 emergency call. Physical location is integral to some emerging on-line socialization, flirting and game-playing. It seems inevitable that physical location will be central to many future Internet applications. So now is the time to consider how a system for managing physical location should be designed.

How is location derived? Today, in the Internet, addresses are used as the starting point. Addresses are given out to Internet Service Providers, who give them in turn to customers. Most ISPs operate within a physical locale (which may be large in the case of a national ISP), and most ISPs know where their physical circuits go. If they know that a customer is at the end of a physical circuit, they can make a reasonable guess as to where that customer is. This sort of inference is used to prevent auction services from displaying Nazi memorabilia to citizens of France, for example. Services such as Google give different answers depending on what country you are in, and this requires some sort of guess as to physical location based on IP address.

In the future, there will be many more ways to determine location. More and more end-node devices have GPS receivers, so if they are outside (in view of the open sky and the GPS satellites) they can compute their own location. The question then is whether (and to what parties) they will pass on that information. GPS information is under the control of the end-node, not the ISP. But the wireless service providers are also getting better at

locating their end nodes. Wireless systems can attempt to geo-locate end-nodes using the radio signals, and triangulating.

The question, then, is what will operators do with this information? One obvious answer is to sell it. This may trigger calls for regulation of the use of this information. This will in turn raise question about what third-party players who attempt to guess a location will be allowed to do with their information. To help visualize what the issues might be, here are some possible scenarios of usage:

Case 1: Use of location in an ongoing interaction. In this situation, some number of parties are already in communication, and one of them attempts to determine the location of another. Since they are in communication, each end knows something about the other—at a minimum an IP address. There may be other information available. This is the situation with country-specific modification of services like Google or E-bay. In these cases, the question about location is asked about a specific node that is already known.

Case 2: Use of location to find possible candidates for interaction. In this case, the interaction starts with a request to determine a set of nodes that are within some locale, after which an attempt is made to communicate with them. Examples of this include sending a message to anyone on a beach that a tsunami is coming, or sending an advertisement to anyone nearby about a sale going on. The social value may vary widely, and there may be calls to allow this sort of location-based search only under certain circumstances, which will have to be debated.

Designing a new Internet—why FIND and GENI matters

Draft version 2.1 of Wednesday, February 15th 2006

The NSF FIND project pursues a bold vision, the vision of creating a new Internet suited for the needs of tomorrow. GENI is a research platform that will support experiments across a wide range of computer science, including the NSF FIND project. This document illustrates what the FIND research entails, what it will mean to create and evaluate a new Internet, the range of concepts and innovations that will go into this project, and how GENI will support the FIND program.

The nature of the contemplated research

We envision a program of research that combines scientific discovery, invention of new mechanism, and the integration of new concepts into coherent overarching proposals for an Internet of tomorrow. We call these integrated proposals *architecture*, and the intended outcome of the FIND research agenda is the creation and validation of candidate architectures for an Internet of tomorrow. These will be tested on GENI.

The Internet has been with us long enough that it is easy to forget how recently its basic ideas were conceived, and how novel they were at the time. The central concept of the Internet, that data is broken into small units called packets that are statistically multiplexed over circuits, is less than 50 years old, and when it was first proposed it was viewed as a radical and unworkable idea. There were fears of instability and of uncontrolled queuing. The validity of the idea had to be proven by a combination of mathematics, simulation and real-world trials. In fact, the first workable approach to controlling queue length took over 10 years to conceive from the start of the Internet, and only with real world experience did the designers come to understand that part of the problem was system-wide oscillation caused by the network being driven at the fundamental frequency of the flow-control window size. Simpler models had simply failed to capture this effect.

Designing a new Internet is not just the discovery of one or two new, breakthrough ideas. Even the singular idea of packet switching is only part of what defines the Internet. Creation of an architecture for a system is different from scientific discovery, and fits within the general domain of *systems engineering*. Systems engineering is a process that includes specification of requirements, invention of new approaches, and a complex process of trade-off and balance among different functional objectives, and as well among different stake-holders and constituents. Validation involves a process that first tests specific proposals for new mechanisms to better understand their limitations and relative advantages, and second tests more complete architectures to determine fitness of purpose within the multi-dimensional space of requirements.

Designing a new Internet is perhaps like designing a new airplane. Lots of new innovations may be proposed, and the process of design must validate these innovations.

But the success of the overall design is the integration of concepts to produce a design that balances a number of disparate considerations that include fuel efficiency, noise abatement, minimum required runway length, capacity, air safety and cost. In aircraft design, individual innovations become useful when a new plane is designed. The airframe market sees a steady introduction of new planes, which provides a platform for new ideas to enter the market. Similarly, there are some innovations in networking that can only enter the market if we contemplate a new Internet.

Some examples: a critique of the current Internet

One way to contemplate what a new Internet might look like is to catalog some of the key components of the current Internet, note what is wrong with each part, and list some of the proposals that have been put forward to improve them. This approach has the risk that it can lock us too much into the current set of parts, but it has the merit that it permits a concrete example of what the experiments on GENI might look like. So with that warning, we can look at the current Internet.

Packets and multiplexing

A basic assumption of the Internet is that data is broken into packets, which are then multiplexed statistically along communications paths. Most (though not all) researchers conclude that the concept of packets is a good one that should be a part of a future Internet. But in the center of the network, there is an increasing view that individual packets are too small to be handled individually. Instead, we need a design that can process aggregates of packets. Today, this is done outside the architecture, using a separate mechanism (such as MPLS). If **routing and management of aggregates** were included into the architecture of the Internet, it would allow both packets and aggregates of packets to be handled in a unified way. In particular, the concepts of routing, traffic engineering and topology management should be unified in a future Internet. Fault recovery should be unified across layers. The inclusion of switched optical components in GENI will allow researchers to experiment with algorithms for rapid reconfiguration of aggregates.

While statistical multiplexing of paths leads to good link utilization and cost-effective design, it is also a security risk, in that an attacker may be able to flood certain links to the point where good users are squeezed out. There are several approaches that have been proposed to solve this problem. One is **Quality of Service**, which is now being used in private networks, but only partially in the public Internet. Another approach is **virtualized resources**, in which simple statistical multiplexing is replaced with a more complex layered approach to sharing in which classes of users or activities are given static shares, and only within these classes is there dynamic sharing. GENI will be used to test the concept of virtualization. Another approach to controlling abuse is **diffusion routing**, discussed below

Addressing and forwarding

Once we agree that the network will carry packets, the next step is to design the mechanism that allows packets to be forwarded across the network. The Internet contains elements called routers, which look at the *address* in packets to determine how to forward them. The original Internet assigned a global address to every destination, and allowed any computer to send a packet to any place. This open pattern of communication was critical to the early success of the Internet, but has caused a number of serious problems, which only became apparent over time. For this one topic of packet addressing and forwarding, we have cataloged over 24 proposals for alternative addressing and forwarding schemes, most of which have gone nowhere, because there is no way to validate them. GENI will allow us to try out alternatives to today's scheme that might provide better security, better management, and better functionality.

One problem with global addressing is that it allowed the Internet to be a vector to deliver security attacks, since any machine, including a malicious one, can send traffic to a security target. A future Internet must provide what has been called **trust-modulated transparency**: trusting nodes should be able to communicate at will, as in the original conception of the Internet, but nodes should be protected from nodes they do not want to communicate with. There are several approaches to achieving this balance, which we expect to validate using GENI. One is **address indirection**, in which senders do not know the address of the receiver, but only a name. A protected element in the network (which would have to be invented for the purpose) would check the identity of the sender, and decide whether to forward the packet to the address of the named recipient. A second approach is the **permit** approach, in which the receiver gives to the sender a special token, which is then included in the packets from the sender. Again, a protected node in the network would check the token to decide whether to forward the packet. These schemes, in general, are examples of taking the concept of a **firewall**, which is an afterthought in the current design, and considering from scratch how to integrate this component into the Internet as a first-class element.

A second problem with the original addressing scheme is that it did not take into account mobile devices, which are becoming the norm, and may dominate the Internet in 10 years. Today, Internet addresses are used to convey a weak form of identity as well as location on the net. Since the IP address captures the notion of identity, it is not possible to change the IP address in an ongoing conversation, which means that a node that is mobile cannot change its address as it changes it location. A future Internet must have a scheme for **dynamic address reassignment** and a scheme for automatic **connection persistence** for mobile end nodes.

On the other hand, as we better deal with mobility, the value of an address as a signal of identity may erode. This raises the question of whether there needs to be some explicit form of identity that is visible to an observer of packets in the network. Different answers have very different implications for privacy, for accountability and for policing. One response to this question is that there will be different needs for identity in different parts

of the network, so the packet header should include an identity field, but not a rigid specification of what that field should contain. One proposition for an experiment on GENI is a **semantics-free identity field** in the packet header.

A final example of a problem with the current Internet addressing scheme is that IP addresses are normally bound to specific physical machines, but in many cases a message needs to be sent to a more abstract entity—a *service* rather than a *machine.* A scheme called **anycast** has been proposed to solve this problem; this scheme allows the same address to be assigned to multiple machines, and the particular machine to receive the packet is determined by the routing protocol at run time. Anycasting may solve a number of security problems as well as problems of service location and session initiation, but it has never been fully elaborated or tested. A new Internet may contain a mechanism like this, which will have to be evaluated on GENI.

Routing

Routing is the process of computing the best path from sources to destinations. Routing is not the same as forwarding—routing computes the best paths, forwarding uses the results computed by routing to take the correct action as each packet arrives at the router.

Today, the Internet uses a two-level routing scheme, with a top-level mechanism called Border Gateway Protocol, or BGP, to connect different administrative regions, and a second level of protocol inside each region. The **region structure** of the Internet seems like a fundamental, and in fact may be more explicitly expressed in a future Internet design. This means that we will have to set up experiments on GENI to capture the idea that different parts of the Internet are run by different organizations.

The BGP of today is flawed: it limits the business relationships that ISPs can negotiate, it recovers from some failures much too slowly, it is not sufficiently secure, and under some circumstances it can be unstable and lead to routing oscillations. None of these issues were fully understood until the protocol was put into use on a large scale Internet. Alternatives to BGP are being developed that provide better **convergence after equipment failures**. A platform such as GENI is critical to testing these schemes. Evaluating a route-computation service in GENI would enable experiments that measure routing-protocol convergence delay and the effects on end-to-end performance when topology changes occur. This would involve "injecting" link failures under today's Internet routing architecture and under the new design. This experiment would be difficult to conduct without GENI because simulations do not accurately capture the overheads and delays on the routing software, and operational networks would not permit researchers to intentionally inject failures.

One of the concerns with BGP is that it does not provide adequate levels of security. GENI can be used to evaluate the route-computation service of a **security-enhanced alternative to BGP**. For example, upon learning a BGP route announcement from a neighboring domain, the service can classify the route as "suspicious" if the Autonomous System originating the route does not agree with past history. The service can prefer

non-suspect routes over suspect routes. The experiment could evaluate this new variant of the BGP path-selection process and see if it effectively protects the experimental network from attackers injecting false routing information.

Today, the user has little choice over the route his packets take. There is no equivalent of "picking your long-distance provider" in the Internet, and little support for a host or edge network that wants to have multiple paths into the network. This lack of support for multi-homing is a major contributor to poor availability. It has been proposed that Internet routing should be redone to support **end-node route selection** so that the user has more control over how his packets are carried, both to support multi-homing and to impose the discipline of competition on the service providers. We need to experiment with this to see if we can design tools that make end-node route selection practical and usable, and to see if this approach actually improves availability, since the approach solves some problems and raises others.

Routing algorithms today attempt to find the optimal path for traffic, given the overall traffic pattern. As the traffic pattern changes, routes must be constantly recomputed. An alternate idea is to take traffic and diffuse it across all possible paths from source to destination. It can be shown that **traffic diffusion** provides stable traffic allocation to links for all feasible traffic patterns. In other words, it eliminates the need for traffic engineering. It also may improve security by eliminating the ability of an attacker to concentrate his traffic onto one circuit in order to overload it. In order to test this idea, what is needed is a network with a high degree of route diversity, which GENI can provide by means of virtualization.

In today's Internet, the route computation is performed in the same physical devices (the routers) that also forward packets. One proposal for a future Internet moves the route computation out of the individual routers, and into a separate **route computation service**. This approach offers great advantages in consistency, manageability, and scale. It allows competing route computation services to offer alternative algorithms for different customers. However, it raises new challenges for robustness and resilience in the face of arbitrary failure. This breaking up of the router function will also shift the industry landscape and create new opportunities for innovation and competition. We need to experiment with this sort of scheme in a real network with rich connectivity and real-world failure modes. In particular, since GENI provides the option of interconnection with operational ISPs, it can be used to test new routing regimes in the real world.

Today's routing protocols have some severe flaws with respect to management. One example is that they are not designed to take a component out of service gracefully. They can deal with a *failed* component, but there is always a transient glitch as they compute new routes. If network operators know they are going to take a component out of service, it should be possible for the routing algorithm to plan for this so the users never see a momentary outage. A demonstration on GENI could evaluate the effects on end-to-end performance of a **graceful maintenance protocol,** which would make step-by-step changes in the routing tables to prepare for removing equipment from the network.

Putting it all together—architecture

The list of issues above, and the examples of approaches to deal with them, are only a very partial catalog of what the research community is preparing do using the GENI test facility. It is important, as we consider this list, to look at the whole and not the parts. Each one of the ideas above, and the many others that have been suggested by the research community, may be interesting in their own right, but it is when they are put together, their interactions explored, their joint implications worked out, that the real payoff occurs. It is through the combination and harmonization of many ideas like this that new architecture emerges. GENI can be used to support initial experiments to explore individual ideas, but the most important experiments on GENI will support the testing of these new architectures—combinations of these new ideas that greatly improve the fitness for purpose of the Internet.

It would be nice if we could discover simple, independent fixes that each improved one problem with the current Internet. It would be much simpler if we could find "the security fix", "the management fix", "the availability fix", and so on. But in fact, most of the new ideas cataloged above shift the landscape in all of these areas. This is why "putting the parts together" is such an important part of the project. As a further illustration of how complex the relationship is between mechanism and requirement, here is one final example about new Internet design. The proposed change seems small, but it affects almost every requirement we have posed for a future Internet.

Well-known ports and service initiation

Today, the Internet names services (such as Web or email) using "well-known ports"—numerical indices that are statically assigned to each application at design time. Since these port numbers are included in each packet, this permits any observer in the network to determine what application is being used. And since these numbers are statically assigned, an attacker can easily launch an attack against an application on a given host, just by combining a host address with the port number, and using that destination as the target of an attack. An alternative would be to design a new mechanism for "**service rendezvous**", and to use random port numbers to identify connections. This change, perhaps combined with an increase in the range of port numbers, would essentially eliminate the value of the attack known as port-scanning, and would provide more privacy from observers in the network. However, a sparse port-space would change the whole security landscape by changing what firewalls can do based on packet inspection. In fact, this change would force a complete re-conception of what a firewall does and the balance of power in the space of attack and defense. The alternative would also change the economic landscape by making it harder for Internet Service Providers to discriminate among customers based on what applications they want to run. Presumably, they would respond by inventing some other form of discrimination. The change would make the network more useful to consumers, by eliminating some of the restrictions that are imposed by the invention of Network Address Translation units as network attachment devices.

A broader set of experiments on GENI

GENI is not just about experiments to support FIND and the creation of Future Internet architectures. It will be used as well to support a broad set of experiments that range from evaluation of new technology options to new designs for distributed systems and applications. Here are some further examples of experiments to be run on GENI, which illustrate the range of experiments that it can support.

1. **A new end-to-end protocol model** for mobile data: The ubiquitous TCP/IP protocol used for most Internet services has several known weaknesses when applied to mobile data scenarios. In particular, the IP network layer framework does not support disconnected operation or caching/storage within the network while the window flow-control mechanism in the TCP transport layer performs poorly over wireless access links with high error rate. Numerous solutions to these problems have been investigated by the wireless networking research community, including mobility service overlays and modified TCP or all-new transport layer protocols, but none of these solutions have migrated to general use due to legacy staying power and the difficulty faced by innovators in deploying their protocols on a a large-scale network to test performance and end-user acceptance.

A GENI experiment of near-term interest to the mobile networking community would be to deploy one or more alternative protocol solutions on an end-to-end basis with a significant user population. Experimental measurements of interest include short-term numerical performance measures such as packet delay, packet loss rate and user throughput for specified applications and mobility patterns. In addition, longer-term service quality measures such as % dropped connections and availability will be measured.

2**. Delay Tolerant Networking** (DTN): DTN is a very different model for an Internet architecture, which is suited to disconnected operation; this can include wireless devices that go in and out of range, as well as communication among planets. This concept has been under development for a number of years, and is an obvious candidate to deploy and test on GENI. (The current design of GENI does not include inter-planetary links, however.)

3. Evaluation **of location-aware networking techniques** for mobile and sensor applications: Location (defined in terms of geographic coordinates) is being recognized as an increasingly important aspect that needs to be integrated into mobile and sensor network applications. For example, mobile users seek geographically relevant information about people, products and services in the immediate vicinity. Sensor applications require addressing of sensors by location rather than by network or MAC address. Vehicular safety applications require multicasting to nearby cars within a certain bounded region. In all these instances, techniques for naming, addressing and routing in the network need to be extended to account for geographic location. Techniques such as location service overlays and geographic routing have been proposed but never validated at sufficient scale or realism.

A location-aware network experiment to be run on GENI involves instrumenting one or more wireless subnetwork with location determination services based on signal triangulation or other methods, along with implementations of overall or new network layer protocols for location service, georouting, etc. This experiment would start with a bottom-up validation of the accuracy with which location can be tracked by the network along with an estimate of protocol overheads and latencies associated with providing location information to higher layer protocols or applications. Once the protocol is validated and performance/overhead measured, it is anticipated that GENI would be used to offer long-running location services to new mobile and sensor applications with real end-users, leading to identification of one or more viable protocol designs.

Location management is an excellent example of a design problem that will benefit from a multi-discipline approach, since an architecture for location management must take into account issues of privacy, ownership of personal location information, and rights of third parties and the state to gain access to this information.

4. **Integrating adaptive cognitive radio networks** with the global Internet: Cognitive radio technology is expected to emerge in the next 5-10 years and offers the promise of programmable radios which can adapt to their environment, share spectrum efficiently and organize themselves into collaborative multi-hop networks. The wireless research community has several hardware and protocol design projects in this area, and would like to validate these designs further and study issues involved with integrating these networks with the global Internet. Protocol design topics of particular interest include naming and addressing for dynamically changing constellations of radios, as well as routing between ad-hoc radio networks and the wired Internet infrastructure. Of particular interest is the level of cross-layer protocol support required on an end-to-end basis.

The wireless research community plans to conduct several experiments with cognitive radio networks using the GENI system. In particular, the cognitive radio subnetwork planned for GENI will be instrumented to support dynamic spectrum measurements needed to evaluate the benefits of shared spectrum, agility and etiquettes. In addition, the testbed will be used to conduct large-scale validation of discovery, ad hoc collaborative network formation and routing in cognitive networks. Both short-term network performance (throughput, delay) and longer-term system behavior (spectrum utilization, availability) will be measured using GENI. Once a cognitive radio network is validated at the subnetwork level, end-to-end experiments will be carried out to investigate issues such as naming, addressing and cross-layer routing support.

The importance of building

Here is a final example of a specific experimental program. It is a story based on work that was carried out on Planet Lab, which does not support the range of experiments that GENI will. But this example illustrates the power of experimentation, innovation and

discovery when the community is able to carry out large scale experiments with real users.

A researcher designed a new system for Content Distribution that he believed scales better under load, yet has response time that's comparable to the best-known techniques. Using the best methodology of the day, he simulated the system and quantified the potential improvement in aggregate throughput. He published a paper that reports 60-91% improvement over the state-of-the-art.

Then Planet Lab became available, which allowed the researcher to deploy the system in a realistic setting, at scale, with real user traffic. The researcher took advantage of the facility, and within days of deploying the system (v1), learned an important lesson: unanticipated traffic compromises the security of the system, making it unusable. The researcher designed and deployed a redesigned system (v2) that took this lesson into account.

Within weeks of deploying the new system, the researcher discovered that performance is compromised by failures of the Domain Name System. Based on additional observations, the researchers discovered the root cause, and in response, demonstrateed how the Content Distribution Network (which was designed to make web content more available) could be adapted to also make DNS resolution more robust. The researcher modifiee his system (v3) and deployed it on Planet Lab.

Based on instrumentation of this system, the researcher discovered that the best known models of DNS behavior were all wrong, and produced a new model that others can use by other researchers.

Based on other data collected by instrumenting the system, the researcher discovered that he is able to observe two orders of magnitude more Internet failures than any existing observation platform has yielded. This results in a more accurate model of Internet behavior that other researchers are able to incorporate into their research.

The researcher also recognized that he can augment his original system (v4) to also diagnose Internet failures in real-time, and use this system to build adaptive applications that are able to route around failures, resulting in an even more robust service.

After gaining further experience, the researcher discoverd that his system performs poorly when distributing big files, especially to a large set of clients, but that by including new mechanisms, he was able to redesign the system (v5) to yield large-object throughput that scaled with the number of clients that request the object. One of the more interesting lessons of this exercise is that the algorithms proposed by others to solve this problem do not work well in practice, and it is only by a thorough evaluation of various engineering tradeoffs that he was able to design a system (v6) with robust performance under a wide-range of conditions.

Epilogue 1: Researcher never bothers to return to the issue of the specific algorithms used

in his original system, as they are in the noise relative to the other factors that influence an Internet service.

Epilogue 2: Researcher understands factors that influence network robustness at a deep level, and sets out to create a clean-slate network architecture that incorporates these lessons in the core of the design. The new architecture is dissemination-oriented rather than client/server oriented, but must include completely new approaches to security because content is now decoupled from specific points (servers) in the network.

GENI will be the test platform to evaluate this new architecture.