

Future Internet: Fundamentals and Measurement

[Report of the COST Arcadia Future Internet Workshop]

Thrasylvoulos
Spyropoulos
INRIA, Sophia-Antipolis
2004, route des Lucioles
06902 Sophia Antipolis,
France
first.last@sophia.inria.fr

Serge Fdida
LIP6 Laboratory
University Pierre and Marie
Curie
75015 Paris, France
Serge.Fdida@lip6.fr

Scott Kirkpatrick
Benin School of Engineering
and Computer Science
Hebrew University of
Jerusalem
91904 Jerusalem, Israel
kirk@cs.huji.ac.il

ABSTRACT

While the Internet is hardly “broken”, it has proved unable to integrate new ideas, new architectures, and provide paths for future integration of data, voice, rich media and higher reliability. The reason is that the basic concept of the Internet as an end-to-end packet delivery service has made its middle layer, networking services through TCP/IP, untouchable. If we wish to see any disruptive enhancements to security, routing flexibility and reliability, and robust quality of service guarantees in coming years, we will need to move towards an Internet in which networking environments offering differing strengths can coexist on a permanent basis. This view is gaining currency in the US, advocated by the FIND/GENI initiative [7, 6] and in Europe, where it forms the heart of the activities reviewed by ARCADIA. The ARCADIA activity, sponsored by COST [1] has been chartered to look at critical areas in which research on fundamentals in the Internet’s architecture and protocols, supported by accurate experiment, can unlock some of the Internet impasses. This paper attempts to describes the insight gained and conclusions drawn from the first ARCADIA workshop on the Future of the Internet, organized around the main themes of Virtualization, Federation and Monitoring/Measurement.

Categories and Subject Descriptors

C.2.1 [Computer Communication Networks]: Network Architecture and Design

General Terms

Design, Experimentation, Measurement, Security

Keywords

Future Internet, Virtualization, Federation, Testbeds, Monitoring, COST, Arcadia

1. INTRODUCTION

The Internet today is being used for a large variety of commercial and non-commercial purposes. For many people it is a crucial working tool, for others their business venue, and for a large majority of people a new and efficient means for communication, entertainment, and even education. However, its rapid growth, an integral part of its success, appears to be bringing it to a point of crisis. While the Internet is hardly “broken”, it has proved unable to integrate new ideas,

new architectures, and provide paths for future integration of data, voice, rich media and higher reliability because of the commitment to keeping the present structure intact and decentralized.

In fact, the Internet has become ossified due to its strong support of numerous activities of our society and the fact that it was not defined with the current vision of a business and societal infrastructure. For example, despite international support and some strong technical motivation to migrate, IPv6 has not been widely deployed yet. The reason is that the basic concept of the Internet as an end-to-end packet delivery service has made its middle layer, networking services through TCP/IP, untouchable. Both evolutionary and “disruptive” alternatives, if they attempt to propose a uniform “future Internet” solution which every participant will have to adopt, will meet the same fate as IPv6. If we wish to see any disruptive enhancements to security, routing flexibility and reliability, and robust quality of service guarantees in coming years, we will need to move towards an Internet in which networking environments offering differing strengths can coexist on a permanent basis. This view is gaining currency in the US, advocated by the FIND/GENI initiative [7, 6] and in Europe, where it forms the heart of the activities reviewed by ARCADIA.

The ARCADIA activity, sponsored by COST [1], has been chartered to look at critical areas in which research on fundamentals in the Internet’s architecture and protocols, supported by accurate experiment, can unlock some of these impasses, and to articulate the unique strengths that Europe brings to this opportunity.

In a first meeting in late summer 2006, we decided to organize around the main themes of Virtualization, Federation and Monitoring/Measurement. A first step in this was a workshop on the Future of the Internet, with presenters from Europe, the US and Japan, as a one-day satellite to the CoNEXT meeting in Lisbon, Portugal, on Dec 4, 2006 [11]. The first three sessions of the workshop follow the three themes identified above. A fourth session explored both security issues and novel approaches to naming and addressing. There was also a final discussion session involving all speakers and additional ARCADIA members. This report summarizes the approaches, each grounded in existing projects, existing research, and existing measurement tools, that were advocated at the workshop.

2. VIRTUALIZATION

Proposing new technologies for the future Internet involves testing these technologies beforehand. Although theory, simulation, or emulation can be useful tools in the design process, experimentation with the “real” Internet is necessary to test a protocol’s reaction to unexpected, real events that evade simulation or emulation. The Internet is now an important production and business tool. ISPs would not risk taking it down to install new, untested routing software or risk running experimental software. Such an interruption or an unexpected failure that could take the whole network down could result in loss of a large number of existing business and resulting revenue. Thus, extensive experimentation is a crucial step for coming even close to convincing an ISP to deploy any new enhancement. On the other hand, deployment of the new technology might actually be the only way of testing it, under “realistic” conditions. These two facts create a vicious circle, where disruptive technologies are not deployed due to lack of enough experiments to prove that they would not harm deployed Internet services, and new protocols can never be tested to the extent necessary due to the ISPs’ averting to deploy them.

To overcome this deadlock, global experimentation platforms like PlanetLab [10] have been deployed with a goal of allowing researchers to (i) conduct large scale experiments in a “real” Internet topology and with “real” Internet traffic, while (ii) not disrupting the Internet’s regular behavior and performance. PlanetLab is an open, overlay network of “virtualized” computing and communication resources, that are connected through the Internet. By using virtualization software like VINI [13] it allows different experiments to share a large set of these resources by allocating dedicated “slices” or resources, while at the same time being exposed to real Internet events and traffic. Of course, when setting up testbeds and doing virtualization, both full realism in experimentation and reproducibility of experimental results may not always be achievable. Researchers will have to hit the right tradeoff between experiments that take into account real traffic and network conditions, and experiments that can be reproduced, or at least provide enough context to be meaningful for comparison.

Another effort towards that direction is XORP (eXtensible Open Router Platform) [15, 5]. XORP is an open software project that aims at allowing experimental networking software to run side by side with production software. It runs a virtual router stack on commodity PCs or servers, and is also used by the VINI platform. Through a carefully designed Forwarding Engine Machine (FEM) the forwarding tables are exposed to a large number of concurrently running routing protocols like OSPF, RIP, BGP, multicast protocols like PIM, as well as experimental protocols under testing. Each of these protocols runs as software on the user plane to enable better (sandbox) isolation and security. What is more, groups (or each) of these could run in a separate physical router. Thus, if multicast fails, for example, it would not at least bring down the unicast service. Furthermore, if one router is compromised by say a “router worm”, this software isolation could prevent its further spread (something that is not guaranteed by today’s Internet routers). Finally, this stack enables individual changes to the routing software to be implemented much more easily and quickly. In other words, XORP promises to deliver stability for the provider and flexibility for deployment.

As with any virtualization approach, XORP may face per-

formance considerations. It introduces an extra processing layer (FEM) and its respective processing overhead. A PC cannot be faster than specialized hardware. On the other hand, the technology of XORP could allow many cheap routers to be put together in a much more powerful router cluster. It has been demonstrated that XORP can “virtualize” up to 40Gbps connections. Nevertheless, it is still mostly envisioned to address “edge” routers rather than “backbone” ones. Finally, XORP may also be used inside generic virtualization software like VINI, to allow different experiments/ISPs to run different routing protocols. While XORP may sound like slow software, compared to the fast specialized hardware found inside today’s routers, it is claimed that in a high performance server it can support routing at bandwidths up to 40 Gbps.

Taking virtualization a step further, one could even envision different Internet architectures, rather than just experiments, to be running side by side. One such proposal, called CABO (Concurrent Architectures are Better than One) [14], advocates to separate the providers of the physical network infrastructure from the service providers themselves. Despite a large economy running over the “net”, its economic incentives are misaligned, stifling growth. Different ISPs control different, often small parts of the network, and are rarely on both ends of an end-to-end session. To change anything in the core functionality of the protocols, and offer a new service, a number of different ISPs need to come to agreement and coordinate to make anything happen. In addition to the inherent difficulty of such an endeavor, this also creates a disincentive for ISPs to compete by innovating. How can an ISP attract customers from another ISP by creating and offering a better service, when that service often requires the agreement of that second ISP in order to work? For example, imagine an ISP wishes to offer QoS guarantees to its customers. Even if it invests a lot of resources into developing its own network with QoS provisions, these would be useless as soon as its traffic would have to traverse other networks that do not comply.

CABO aims at better aligning these economic incentives. A service provider would now lease physical resources end-to-end from different infrastructure providers, and would be able to offer a better service to its customers. Furthermore, different service providers would not only have the ability to, but now also every incentive to evolve, improve, and innovate their (virtual) network, in order to attract customers to connect through them. Of course, this does not preclude an entity from being both the infrastructure and service provider in some cases, as for example when the national infrastructure provider (e.g. France Telecom) also acts as a service provided (e.g. for DSL).

However, this would also imply that different service providers would now have to share physical resources. Hence, some means is necessary (i) to be able to guarantee each the required resources to operate its services, and (ii) to isolate different service providers from inadvertently or maliciously harming each others’ services. A platform like VINI could once more be the tool to achieve both these goals. Although VINI was initially proposed to allow different experiment to co-exist over shared physical testbeds like PlanetLab, it could naturally be applied in this scenario, as well. In short, what experimenters are for VINI and PlanetLab, Service Providers would be for VINI and CABO. In fact, it is envisioned that a large scale, long-term experiment running on

a PlanetLab-like platform could mature to become a new service provided.

Despite the promise of this argument, a number of different issues need to be addressed. First, even though this shift appears attractive to the service providers, it is not as clear that it suits the infrastructure providers. Furthermore, although ISPs would now be able to control an end-to-end path, it would be very difficult for them to do so for all the end-to-end paths that the actual client might wish to establish. For example, even if an ISP leases enough resources to provide stable and reliable bandwidth for video streaming over the US backbone, what happens if the client wishes to connect to content that lies in Europe? Finally, a number of technical obstacles need to be overcome by the virtualization infrastructure, before service providers feel comfortable sharing physical resources with a competitor. Useful lessons from earlier efforts for such virtualization, like for example the CPLANE start-up, could probably be helpful to avoid making the same mistakes or re-inventing wheels in this process.

As a final note, in addition to virtualizing the “wired” Internet, considerable effort has been devoted towards the virtualization of wireless resources. Such endeavors are necessary to allow experimentation with a more diverse set of networks which are expected to be an integral part of a heterogeneous Internet of the future. A nearly completed effort has been undertaken at Rutgers University to provide virtualization over a 400-station 802.11 testbed, Orbit, which has been in operation since 2005 [18, 12]. They selected time-division multiplexing, and support 16 simultaneous “slices” operating on the entire testbed, with one slice at a time having access to all stations. They employ user mode Linux (UML) and require time slices to be greater than 100 ms to avoid problems with the skew that can occur when this many interfaces are simultaneously reset to the next time slice’s system and channel configurations. Performance studies under various application loads are presently underway.

3. FEDERATING TESTBEDS

PlanetLab had been designed as a global and shared experimental platform for researchers to be able to test new Internet architectures and protocols, before these get deployed, on a “real-like” environment without breaking the internet. It offers to researchers what simulation doesn’t, namely real nodes and networking links, and what emulation doesn’t either, namely spread of nodes, and “realistic” Internet conditions and traffic.

Although PlanetLab has been largely successful in terms of attracting researchers and experiments, it is currently oversubscribed and faces a number of important limitations. First, PlanetLab is a wired testbed environment with most nodes connected to high-performance academic backbones, like Abilene and Geant, etc.; this often provides a less than representative picture of the Internet; what is more, it leaves out a growing demand for experimenting with new wireless technologies and networks. Second, PlanetLab does not expose the underlying network, and thus does not offer enough control to the experimenter. Routing between PlanetLab nodes is currently performed solely by the Internet, and events like disconnections or failures cannot occur on demand. Finally, PlanetLab is currently under one single administration (Princeton University). This makes it

rather difficult to enlarge the network. Further, PlanetLab is a homogeneous structure that is neither flexible enough to accommodate different group goals and/or heterogeneity in the networks, nor can it provide guaranteed resources (e.g. for Grid computing) or authenticated /authorized access to “special” resources (e.g. a wireless sub-testbed).

OneLab [8] (in Europe) and GENI [7] are two projects that aim to address these crucial issues. Two of their design goals are to “deepen” PlanetLab by providing more/full access to the underlying network (e.g. user-defined routing protocols for inter-node routing, emulation of link/node failures, etc.) as well as to “widen” it by connecting to it a number of wireless platforms including sensor networks, mesh (ad-hoc) networks, vehicular networks, etc.

In addition to these technical improvements, one of the most important goals of OneLab is Federation. OneLab aims to be PlanetLab-Europe which will federate with PlanetLabs in US, Japan, and other places, as well as private PlanetLabs (e.g. EverLab [4]). Federation is necessary not only to allow this global testbed to scale smoothly, but also to accommodate differences in goals and perspective (e.g. Europe might have a different research agenda or policy than US). Furthermore, federation will encourage a number of private PlanetLabs to share their resources while respecting their special requirements. Federation should be an incentive for sharing testbeds and enhance their sustainability. For example, a rare resource like a wireless testbed would require different access policies than general PlanetLab nodes, to ensure the semantics of the wireless virtualization are respected and no over-subscription occurs. The Panlab [9] Specific Support Action will enable the trial and evaluation of service concepts, technologies, system solutions and business models in federated testbeds and, as such, develops a complementary vision of this concept. The ultimate goal of these initiatives is a global, federated network of heterogeneous testbeds.

However, a number of issues need to be addressed for federation to be successful. First, the right API abstractions are necessary that will hit a balance between flexibility for the individual PlanetLab(s) and inter-operability between all. Additionally, private PlanetLabs may need to provide predictable and reliable resources for experimentation, authentication and authorization of users, and monitoring of resources and usage. As a result, different policies may be enforced for internal and external users in terms of access to resources or even monitored data. Usage models, namely, how a user will specify its objectives and data to be collected from the experiment, are important to understand and specify. In the end, this global testbed will provide authorized and monitored access to a large set of heterogeneous assets. However, this might also imply that an experiment will have to define in advance all the resources that it will need throughout its run. Slice-monitoring software based on PlanetLab’s CoMo infrastructure was described by the EverLab project, and provides a first step towards making such privatization of usage policies possible. Finally, federation will also raise critical issues related to governance, privacy, business models and policies.

4. MONITORING: MEASURING THE FUTURE

As the Internet has evolved rapidly in the last decade, so has the interest in measuring and studying its structure,

performance and unacceptable behaviors. Measuring and understanding the Internet's properties is necessary: (i) to troubleshoot and debug failures faster and more efficiently, (ii) to design future protocols that are better adapted to its special characteristics, (iii) to monitor performance and to react appropriately to changes in real-time, and (iv) to better understand complex socio-economical phenomena as the Internet is currently expanding into a large number of developing economies. However, measuring the Internet is a daunting task. Its growth has been following a frantic pace, with tens to hundreds of millions of nodes connected currently. Its structure is heterogeneous with routing policies over long distances that reflect economic considerations, not solely performance needs. Finally, a growing number of applications have resulted in very complex traffic and usage dynamics. Hence, to measure the Internet of the future a number of obstacles need to be overcome.

One shortcoming of current measurement projects is that they only involve a small number of measurement nodes. Further, these nodes tend to be placed close to the core. This provides a less than representative view of the Internet often missing many "horizontal links" between ASs. DIMES [20, 3] is a highly distributed measurement infrastructure that addresses this problem by using a large set of interacting measurement agents. These agents are light weight, low signature agents that run on host nodes doing active measuring as a background process, and provide the location diversity necessary to map the Internet down to PoP (point-of-presence) or even router level. Currently, there are a few thousands of DIMES nodes running in more than 600 ASs. Ninety per cent of the agents are outside academic networks. The project has already revealed a much richer connectivity picture than existing projects based on BGP routing messages (e.g. RouteView). An example of this is a twofold increase in the average degree of connectivity, partly due to discovering many horizontal links between peering ASs, which can only be discovered by monitors installed inside at least one of them (due to BGP policies). The end goal of DIMES is to produce a realistic and predictive model of Internet evolution and dynamics.

DIMES is an active measurement tool that runs mostly on non-privileged hosts. In addition to such active measurements near the edge, backbone measurements are necessary to provide a more complete view of the network as well as to provide the information necessary to respond to failures and changes in performance. In an Internet that seems to be integrating more and more of the existing and new communication services, it is not clear how reliable this infrastructure is. Imagine for example a network failure, which needs to be resolved by calling the system administrator. What happens if this phone service runs as VoIP over the exact same failed network?

By tapping onto a backbone interface one is able to collect almost all relevant information, process it, identify events of interest, and act accordingly. However, an important and difficult decision to be made first is the granularity of the monitored data. Packet monitors provide almost all possible information about a session, but the amount of data that needs to be filtered is immense. Keeping only flow statistics improves this situation slightly, but still results in a flood of data. SNMP monitoring provides information about the actual devices, but these may often be too coarse grained. In addition to the processing and filtering overhead there are

also some deployment issues with these data collection methods, when it comes to commercial ISPs. Looking at routing information might be easier, but also requires extensive filtering to find the "useful" data. It is important therefore not only to reduce the amount of data to be analyzed significantly but also to be able to identify what to measure when the purpose is unclear. Finally, adding metadata to measurements, which will provide the necessary context, is crucial to be able to compare and validate different datasets.

As a final note, installing a *global monitoring infrastructure* doesn't face only technical challenges. There are important legal issues (e.g. privacy) involved in accessing other people's data. No matter how anonymous the data are made, there is often some way to reverse-engineer it back to the content or user. Consequently, a legal framework needs to be established that will bind the entities involved (e.g. monitoring agent, academic institutions, commercial ISPs, etc.). Experience from such a monitoring infrastructure running over a UK backbone network (SJ5) shows that a legal framework is a crucial enabler for participation. A second interesting question is where the monitored data is processed and who has access to it. One needs to backhaul measured data (in or out of band) to somewhere useful. Furthermore, law may prevent sharing of these data. As a result, one may be forced instead to ship the code to monitor to the ISP itself.

5. SECURITY, APPLICATION ROUTING, AND NAMING

The Internet today is not only a major production tool for numerous companies and individuals, but also supports a rapidly expanding online economy. Security is both of increasing priority as well as a growing concern. Internet threats include attacks that exploit software vulnerabilities (e.g. Internet worms), malware that the user may download and invoke (e.g. spyware, viruses), and distributed denial of services (DDoS) attacks. The arms race between the attackers and the defense systems is intense and is expected to stay so, as both become more and more sophisticated and adapt to new technologies (e.g. component architectures).

One example of this fierce arms race is that of Internet worms. State-of-the-art systems (e.g. Honeycomb, Auto-graph) currently identify and detect a signature for a worm, and then count incidences of such a signature. A high incidence number implies a worm that might be replicating itself. However, tools already exist to create polymorphic worms. Polymorphic worms try to minimize the invariant content between replicas (e.g. encrypted payload, obfuscated decryption routine) in order avoid detection. To deal with this "smarter" adversary it is still possible to identify some invariant content between worm replicas. This however has to be combined into a common, identifiable string. Polygraph [17] is such a system that uses learning-based techniques (e.g. Bayesian classifiers), among other things, to construct and detect signatures for polymorphic worms. However, even such a system could be tricked by a yet more adaptive adversary, which injects obscure sub-sequences to falsely train the classifier to look for a longer signature than the actual one. It is foreseeable that both host-based and network-based detection systems need to be built and coordinate to successfully deal with the adaptive adversary of the Future Internet.

With a number of new and diverse, overlay applications appearing regularly in the Internet, application-level routing has emerged as another important component of the current infrastructure. The classic Internet perspective of routing was based on a unique 32-bit address and its sole purpose was to *bring the message closer to the destination*. However, application-level routing breaks away from the classic view in order to achieve the desired application functionality over the "legacy" infrastructure. Applications use routing not just in its traditional sense, but also to *pass a message through one or more application servers*. In fact, "middle-boxes" running as intermediaries is often the application (e.g. protocol conversion, web page caching, building multi-point connections out of point-to-point ones).

Today, application-level routing is largely performed in an ad-hoc manner, and on a per-application basis. Furthermore, the source of a session needs to (i) be aware of the middlebox (e.g. web proxy) and address its traffic to it rather than the destination, and (ii) needs to cooperate to include the application server in the path. A method proposed to overcome these problems and provide a common framework for application-level routing is that of "Source-subscription Routing". A source subscribes first to the application server(s) and then addresses its subsequent messages to the destination itself. The application server then transparently receives all relevant traffic from the subscribed sources and performs the in-network processing necessary before forwarding them further to the destination. Source-subscription routing can address many problematic situations, as for example node mobility. However, it requires to assume symmetry between the source and the destination, and also that the application has control over related messages.

Another example where a type of application-level routing might be necessary is Delay Tolerant Networks [2, 19, 16]. A decade ago, the Internet was essentially a set of interconnected wired networks. However, in the future the Internet is expected to integrate and interconnect a large number of wireless network technologies, most of them near the edge (e.g. WiFi networks, sensor networks, mesh networks, vehicular networks, etc.). Delay Tolerant Networks is an architecture that aims at providing connectivity in a number of "challenged" wireless environments. In these environments connectivity disconnections or disruptions may be the rule rather than the exception, either due to inherent characteristics of the application (e.g. Interplanetary Networks, Underwater Networks) or for cost and efficiency reasons (e.g. low-cost Internet provision to developing or remote communities). In such networks both ends may not be present at the same time and transmissions/connectivity opportunities may be short, sporadic, and unpredictable, resulting in long(er) delivery delays. To enable data delivery, it is proposed that messages are stored as bundles (application-level messages) and carried until an appropriate communication opportunity arises. Then, they get forwarded on a hop-by-hop basis, without often being aware of the final destination(s). In this context, addressing is to be performed using names, and late binding is necessary. These networks need to also be addressable by the current or core Internet, and addresses are *Endpoint IDs* (URIs) to which nodes register (with the possibility of multi-homing). Finally, the lack of contemporaneous connectivity and long delay feedback loop raises the security bar even higher.

6. CONCLUSIONS

The evolution of the current Internet will continue as a support of many valuable applications of our society, economy and digital life. Nevertheless, the internet principles were not designed to address the future challenges raised by mobility, security, management or scalability. It is the right time to start exploring new research ideas to enable the design of the future digital infrastructure. At the same time, it is fundamental to develop a disruptive research on methodologies to assess and understand the properties, laws and performance of the future systems. Indeed, nowadays, and despite the availability of useful tools such as emulab, planetlab or ns2, we lack benchmarking environments as well as trustable tools that provide reliable and reproducible results. These exist in most other scientific areas, e.g. physics, life sciences, data mining, image processing . . . As a scientific community and despite the fact that, in our context, technologies are developed within short timescales, we should be able to explore various methodologies appropriate to assess innovation.

It was reassuring to learn how much has been accomplished in virtualization in software and the fact that routing function and computing function are already beginning to merge. But the present approaches solve individual problems under many limitations. Their performance and scalability needs to be explored and understood. Probably considerable added invention is needed, but this work is important because it is exposing the basic problems to be dealt with.

In addition to looking at the "traditional" Internet structure, fresh ideas may also come from seemingly different networking paradigms. The wide range of activities and problem targets that delay tolerant networking for challenged network environments is attacking comes as a bit of a surprise. We should try to see what lessons are emerging from this work, like for example in the area of addressing, for heterogeneous networks of the future, which are only challenged in parts. Furthermore, how do you federate or even integrate a DTN with the rest of the world? What's the next step up in performance, and are some of these ideas still relevant to alleviating the congestion that today's protocols engender in the intermediate levels of the Internet?

Furthermore, in terms of experimentation and testbeds, there will continuously be a myriad of deployment in the future, as there is a need to experiment with technological platforms, service testing, management testbeds, applications and living labs as well as networks to support e-Science. Federation appears as an emerging concept to ease the sustainability of these testbeds and provide incentives for sharing. They are immediate opportunities for federation in the European private planet labs and beyond, such as work that will incorporate new heterogeneous worlds and technologies, and deepen our measurement and simulation capability. There is also considerable incentive into federating with and integrating efforts going on in other parts of the world.

Finally, in thinking about designing the future and experimenting with new architectures, it is expected that researchers will have to deal with conflicting goals in the process. First, when setting up testbeds and doing virtualization, both full realism in experimentation and reproducibility of experimental results may not always be achievable. Researchers will have to hit the right tradeoff between ex-

periments that take into account real traffic and network conditions, and experiments that can be reproduced, or at least provide enough context to be meaningful for comparison. Furthermore, the issue was raised regarding the extent to which measurements and designs based on the "current" Internet will be relevant or useful for the Internet of the future. It will, thus, be important for researchers to be able to identify the aspects of the Internet that are expected to remain important and invariant, and focus their attention close to them. The data to be collected on the future experimental facility have their own value for the community and should be made available at large. Research agencies are always instrumental in motivating the research community to address new challenges and areas. They should enforce funding projects with experimental objectives in mind, motivating these projects to understand how to federate their work. The design of a new system will certainly require skills from various areas: technologies, economics or social sciences. It is also a challenge to drive researchers from these communities to join efforts to address this common goal.

7. REFERENCES

- [1] COST: European Cooperation in the field of Scientific and Technical research. <http://www.cost.esf.org>.
- [2] Delay tolerant networking research group. <http://www.dtnrg.org>.
- [3] The dimes project. <http://www.netdimes.org/new/>.
- [4] EverLab: Next Generation PlanetLab Network. <http://www.everlab.org/>.
- [5] The eXtensible Open Router Platform. <http://www.xorp.org/>.
- [6] Future Internet Network Design. <http://find.isi.edu/>.
- [7] Global Environment for Network Innovations. <http://www.geni.net/>.
- [8] OneLab. <http://www.fp6-ist-onelab.eu/>.
- [9] Pan European Laboratory. <http://www.panlab.net/>.
- [10] PlanetLab: an open platform for developing, deploying, and accessing planetary-scale services. <http://www.planet-lab.org/>.
- [11] Future Internet Workshop, collocated with CONEXT'06, Dec. 2006. <http://www.adetti.pt/events/CONEXT06/Main.php?contents=fiw.htm>.
- [12] S. Banerjee. Time-based virtualization of an 802.11 wireless facility. GENI Development Effort, 2006. <http://www.geni.net/dev.php>.
- [13] A. Bavier, N. Feamster, M. Huang, L. Peterson, and J. Rexford. In vini veritas: realistic and controlled network experimentation. In *Proc. of ACM SIGCOMM '06*, Sept. 2006.
- [14] N. Feamster, L. Gao, and J. Rexford. How to lease the internet in your spare time. Technical report, Georgia Institute of Technology (GaTech), 2006. <http://www-static.cc.gatech.edu/~feamster/papers/cabo-tr.pdf>.
- [15] A. Ghosh, M. Handley, O. Hodson, E. Kohler, and P. Radoslavov. Designing extensible ip router software. In *Proceedings of the 2nd Symposium on Networked Systems Design and Implementation (NSDI)*, May 2005.
- [16] S. Keshav. Naming, addressing, and forwarding reconsidered. Manuscript. <http://blizzard.cs.uwaterloo.ca/keshav/home/Papers/data/05/naming.pdf>, year = 2005,.
- [17] J. Newsome, B. Karp, and D. Song. Polygraph: Automatically generating signatures for polymorphic worms. In *Proceedings of the IEEE Symposium on Security and Privacy*, May 2005.
- [18] S. Paul and S. Seshan. Virtualization and slicing of wireless networks. GENI Design Document 06-17, Wireless Working Group, Sept. 2006. <http://www.geni.net/GDD/GDD-06-17.pdf>.
- [19] A. Seth, D. Kroeker, M. Zaharia, S. Guo, and S. Keshav. Low-cost communication for rural internet kiosks using mechanical backhaul. In *Proc. MobiCom '06*, 2006.
- [20] Y. Shavitt and E. Shir. Dimes: let the internet measure itself. *ACM SIGCOMM Computer Communication Review*, 35(5):71–74, 2005.

APPENDIX

A. WORKSHOP PROGRAM

4th December 2006	
9:00	Virtualization (<i>session chair: Jennifer Rexford</i>)
(20 min)	“Cabo: Concurrent Architectures are Better than One” Nick Feamster, Ga Tech
(20 min)	“Experimental testbeds and XORP” Mark Handley, UCL
(20 min)	“Towards Wireless Network Virtualization on Commodity Hardware” Suman Bannerjee, Wisconsin
(30 min)	Comments and Discussion
11:00	Federation (<i>session chair: Serge Fdida</i>)
(20 min)	“OneLab and PlanetLab,” Timur Friedman, Universite Pierre et Marie Curie
(20 min)	“Federating a European Private PlanetLab,” Elliot Jaffe, HUJI
(30 min)	Comments and Discussion
13:30	Monitoring (<i>session chair: Scott Kirkpatrick</i>)
(20 min)	“Active Measurements – where we are and where we should go,” Yuval Shavitt, TAU
(20 min)	“Passive Measurement in the Backbone” Anja Feldmann, DT and TU Berlin
(20 min)	“Monitoring in Experimental Testbeds” Andrew Moore, Queen Mary College, UK
(30 min)	Comments and Discussion
15:30	Security and naming (<i>session chair: Kenjiro Cho</i>)
(20 min)	“Naming, Addressing, and Routing for Disconnection-Tolerant Networks” S. Keshav, Waterloo
(20 min)	“Networked Systems: Vulnerabilities and Adaptive Adversaries” Brad Karp, UCL
(20 min)	“Naming, Addressing, Routing and Forwarding from an Application Perspective,” Pamela Zave, ATT Labs
(30 min)	Comments and Discussion
17:00	Panel (<i>panel chair: Jon Crowcroft</i>)
	Jon Crowcroft, Christophe Diot, Serge Fdida
	Kenjiro Cho, ETP strategist, speakers, session chairs