Bench–style Network Research in an Internet Instance Laboratory

Paul Barford and Larry Landweber

Abstract— **Network researchers employ a variety of experimental methods and tools including analytic modeling techniques, simulators, and widely deployed measurement infrastructures. It is natural to assume that the overall scope of network research may be limited by the type and capability of the tools and test systems that are available. In this paper we describe a new, bench–style approach for conducting network research that we argue is essential for effectively investigating different classes of important problems. We describe the architecture for the workbench environment which enables this approach - what we call the Internet Instance Laboratory (IIL). The conceptual model for an IIL is a highly configurable laboratory environment containing commercial networking equipment typical of any end–to–end path in the Internet. An IIL would also have the capability to create accurately a broad range of conditions across all networking layers. The two most important advantages of an IIL are the ability to instrument entire end–to–end paths and the ability to install new equipment or protocols in any location in the environment. Clearly, neither of these opportunities is available in the live Internet. While an IIL offers significant capabilities, developing such a testbed is not without significant challenges. We describe these challenges and approaches for addressing them. Finally, we discuss different classes of research questions which would become tractable if an IIL were available.**

I. Introduction

There is an old saying: "If the only tool you have is a hammer, then everything looks like a nail". The tools that network researchers employ to investigate problems can be generally classified in three categories: analytical, simulation, and *in situ* measurement. While these approaches have obviously enabled great expansion in the performance, reliability and capability of networks, it is critical to evaluate limitations on future research arising from the tools and methods available to the network research community.

A recent National Research Council report on the state of research in networking came to a number of important conclusions related to tools and methods [1]. Principal among these was the danger of ossification in networking research due to past success and the strictures of the systems that have been developed to date. To address this risk the report advocates improved measurement capability, improved modeling capability and the need for "disruptive" prototypes. In this paper we present a new tool and approach for network research that would greatly expand measurement, modeling and prototype deployment capability. We call this tool the Internet Instance Laboratory (IIL).

An IIL is an environment for performing bench–style experiments to explore a wide variety of network research questions. The vision of an IIL is a *large* set of commercial systems (end hosts, switches, routers, etc.) which would typically be found along any end–to–end path in the Internet. The systems in an IIL would be interconnected in such a way as to enable many different topological configurations, *i.e.* different instances of the Internet. The environment would include services available in the Internet such as the Domain Name System (DNS), and would also be able to recreate faithfully traffic conditions that occur at any point in the Internet. We define an IIL as a bench–style enviornment since it would exist in a single laboratory (or small collection of laboratories connected via dedicated channels) where researchers have open access to all of the physical equipment. We argue that this environment is essential for addressing many questions currently facing the network research community.

The implementation of an IIL would allow effective investigation of many open problems. For example, network path property inference tools such as bandwidth estimators [2], [3] and loss probes [4], [5] would be much easier to evaluate in an IIL. Tools like these suffer from the difficulties of validating their accuracy which would be trivial in a fully instrumented IIL. In addition, development of new network behavior models would benefit greatly from the availability of an IIL. In [6], Willinger argues for development of "physical–based" network models that relate directly to network systems, protocols and structure. The essence of an IIL is creating instances of all aspects of the network. This would provide insight and enable validation and assessment of sensitivity of models to changing parameters. An IIL would also enable the evaluation of

P. Barford is a member of the Computer Science Department at the University of Wisconsin, Madison. E-mail: pb@cs.wisc.edu.

L. Landweber is an emeritus member of the Computer Science Department at the University of Wisconsin, Madison. E-mail: lhl@cs.wisc.edu.

prototype systems such as routers [7] under much more realistic conditions. While proving absolute accuracy of traffic in an IIL or the frequency of invocation of protocols such as routing updates will be difficult, it is clear that the ability to test disruptive prototype systems with commercial implementations of protocols would greatly improve their evaluation. We discuss research areas enabled by an IIL in greater detail in Section IV.

The examples above point to the principal advantages of an IIL: namely the ability to measure and monitor the entire system and the ability to deploy and evaluate any prototype in the testbed. While measurement infrastructures in the Internet such as [8], [9], [10] offer excellent opportunities for *in situ* measurement of the Internet, they are limited either to end hosts or to within a single network. The vision of an IIL is to enable configurations which create multiple networks from end–to–end–through–core, *and* the ability to measure all points along the paths. While the ability to fully instrument an IIL alone is significant, prototype deployment and evaluation in an IIL may be even more important. Consider the difficulty in trying to convince a Internet Service Provider (ISP) to place a prototype system in their operational network!

Despite these and other benefits, there are many significant challenges in developing an IIL. Challenges fall into two general classes: technical and logistical. Examples of technical challenges include developing representative traffic generation capability, developing accurate propagation delay emulation capability, developing a configuration and test management environment and developing a measurement database system. Each one of these problems represent research opportunities in and of themselves and will likely require on–going enhancement as lab hardware is upgraded with the latest technology. Logistical challenges include the identification and acquisition of equipment for the IIL, maintenance of equipment, validation of results and overall management of the IIL. While these challenges are significant, aspects of many which have been addresses in other work that IIL development can use as a foundation.

The system which is perhaps the most similar conceptually to an IIL is the University of Utah's EMULAB [11]. This environment consists of a large number of general purpose computers which are configured to emulate the behavior of different types of network systems. While EMULAB was developed to provide a network research testbed, it is clear that IIL's use of actual networking equipment make it fundamentally different from EMULAB and enables research which is not possible in EMULAB. However, the work done by the EMULAB team in developing configuration management tools for general purpose computers in network research testbeds could be used to manage the systems in an IIL used for traffic generation.

The rest of this paper is organized as follows. In Section II we present the details of our position advocating an IIL. In Section III, we outline the design considerations for an IIL, factors which are critical for its success and challenges in its development. In Section IV we further support our case for development of an IIL by presenting research areas facilitated by an IIL. In Section V we summarize and conclude.

II. ARGUMENTS FOR AN IIL

Our argument in support of an IIL is based on the limitations of current networking research tools and on the functions which are enabled by an IIL. These have already been touched upon in the prior section and we expand on these ideas in this section. We feel that the list of limitations and new capabilities make a compelling case for the development of an IIL.

A. *Limitations of Current Tools*

Analytic modeling and characterization are classic means for evaluating systems. There have certainly been many successes in model development and characterizations in network research that have provided significant insight into network behavior. Examples include descriptions of self–similar network traffic behavior [12], [13], [14] and characterizations of heavy–tailed distributions [15], [16], [17]. Perhaps the most successful constructive model for Internet behavior is the ON/OFF model described in [18]. However, attempts to expand beyond this have been relatively incremental. While there are a number of purely analytic models that accurately fit observed distributions (*e.g.* time series or Markov models), at best these provide limited insight into the underlying physical systems [6] and are thus unlikely to enable significant improvements to the network. These problems by no means imply that the application of new modeling methods will not enable systems to be improved. They simply highlight the fact that the modeling process is very difficult and could certainly be aided by an IIL which gives immediate access to all of the physical systems in a network.

Another research tool that has been used heavily in network research is simulation. The most popular of these is the *ns* simulator which has been used in hundreds of studies over the past decade [19]. While simulation has been successfully employed to solve some classes of problems, its effective use is typically limited to simple topological configurations and focused on transport layer issues. Paxson and Floyd discuss the many challenges to simulating accurately a network as complex as the Internet in [20].

These include but are not limited to size, complexity, and rapidly changing conditions. The fact that simulation has been used so heavily is something that sets networking research apart from many other scientific disciplines.

Bench–style research environments are fundamental to many other scientific disciplines. Typical examples are wet labs for chemical and biological research and exotic examples include wind tunnels for aeronautical research and particle accelerators in high energy physics research. The major benefit of these testbeds is that they enable systems with highly complex characteristics and behaviors to be effectively evaluated without having to create abstractions of important features of the environment. Without these environments investigations in these areas would be virtually impossible. It is also interesting to note that in other disciplines bench–style environments have been developed long before simulation models. An example is wind tunnels which have been used for decades while only recently have simulation capabilities become mature enough to use in state–of–the–art airplane design. An IIL would provide the similar benefits to the on–going development of network simulation technology.

The final tool in our taxonomy is measurement infrastructures deployed in the Internet. These environments enable *in situ* measurement of network behavior on a limited basis and have been used effectively to understand many different aspects of Internet behavior. The largest of these infrastructures today are commercial systems such as Keynote [21]. There are also a number of well–known research infrastructures including NIMI [9] and Surveyor [8]. These consist of general purpose computers (less than 100) deployed in different networks, principally at universities. Due to their limited size and difficulties in managing widely distributed systems, these measurement infrastructures are not typically available to the general network research community. Representativeness of measurements taken in this kind of infrastructure will be a topic of continued debate. There is also no way to explore a wide variety of conditions or repeat conditions which have been observed. Paxson describes other difficulties in dealing with Internet measurements in [22]. Perhaps their most significant shortcoming is their reliance on probe–based inference tools to gather data, without measurement capability at each hop along entire end–to–end paths.

B. New Capabilities in an IIL

As mentioned in Section I, two of the most significant capabilities of an IIL are the ability to instrument fully the entire system and to deploy prototypes at any place in the testbed. These are not the only new capabilities provided by an IIL. Another important capability of an IIL is the fact that actual protocol implementations will be used in all aspects of experiments. This is a true departure from simulation environments where even if a protocol is implemented following all details of an Internet Engineering Task Force Request For Comment (IETF RFC), there is no reason to believe that it will behave like a protocol implementation from a vendor. This is because RFC's do not specify *every last* detail of a protocol; they only specify what is required. There is the further issue that some implementations of protocols by major vendors are in fact *not* exactly what is specified in RFC's. Evaluating Internet behavior without this consideration could lead to incorrect results. Running the actual protocols in an IIL will improve the representativeness of traffic conditions versus simulated or emulated environments such as EMULAB.

An additional benefit of running the actual protocols and hardware configurations in an IIL is that failure modes will be more accurately reproduced. Failure modes are rarely considered in modeling or simulation although clearly they are encountered in wide area measurement studies. Any work on protocols or systems which must be robust in the face of failures will be more accurately evaluated in an IIL.

Configuration, measurement and management of an IIL will require development of a significant software infrastructure. The introduction of new equipment and expansion of an IIL will mean that this software will have to be maintained on an on–going basis. Because an IIL is built using commercial networking hardware, software developed for the IIL could be deployed in real networks. One might imagine, for example, that development of a router configuration management system for IIL might be very valuable in a production network. If this were the case then the transition from IIL to the production network would likely be trivial.

Once developed, an IIL will enable a wide variety of different and new technologies to be quickly deployed and evaluated. Unlike simulation or modeling environments, in which new systems and protocols require potentially large efforts from the investigators, an IIL simply requires one or a number of new units of the equipment which can then be easily deployed for evaluation in the testbed. Our assumption here is that, with industrial support, the task of acquiring new hardware will be much less than developing new models for new equipment.

III. Design Details

The goal designing an IIL is to have the ability to create a wide variety of "instances of the Internet" where an instance is defined along multiple dimensions including equipment, system configurations, interconnect structure, and traffic conditions. We will treat each of these issues in this section. We will also discuss other factors which are important to consider in an IIL's implementation and challenges which must be overcome.

A. Lab Layout

The guiding notion for lab equipment selection and layout is to enable configuration of end–to–end–through–core instances of the Internet paths using both typical and state–of–the–art commercial networking gear. This definition enables identification of components based on their typical deployment in the Internet. A useful taxonomy of network nodes can be found in [23] which defines nodes by their position in the Internet. We use similar definitions to classify nodes as either end hosts, distribution routers, access routers, and core routers. A conceptual diagram of IIL equipment is shown in Figure 1[1]. On the edge are general purpose computers which will be used for load generation. Moving from end host toward the core means moving from lower cost/speed/capability systems to higher cost/speed/capability systems. It also means moving from a larger number of systems toward the edges to a smaller number in the core.

While we do not propose absolute quantities for systems in an IIL, it seems clear that there are a number of consideration in this regard. First, the number and capability of systems selected for the core will influence the numbers in other parts of the lab. If core nodes are deployed with the highest speed line cards, then there must be fan–out sufficient to load the core. Second, sufficient number of systems must be available to recreate a variety of instances of the Internet. An example instance might be a single Internet service provider with multiple peers. There is a minimum number of systems required to fully exercise an interior gateway routing protocol.

Our general design from end host through core provides a blueprint for an IIL's design, however there are many ways in which is could be realized. We do not advocate any specific pieces of equipment but argue for heterogeneity in all aspects of the equipment used in the lab since there is a wide variety of commercial equipment used in production networks. Many aspects of design details will be determined by logistical considerations.

Once a basic IIL has been developed it can serve as a substrate for testbeds focused in specific areas. Examples of test environments that could easily be built on top of an IIL include a network attached storage testbed, a Internet security testbed, a peer–to–peer file sharing testbed or a grid computing testbed. Some of these testbeds would require additional equipment specific to the area of study. The lab could also easily be extended to include emerging network technologies such as wireless or optical routing.

The final aspect of an IIL's layout is it's physical realization in a room. Considerations for power and cooling are typical to any lab. However, the quantity and diversity of equipment coupled with the needs for easy and frequent reconfiguration make the task of rack layout important. In the end layouts will be determined by the equipment selected, growth plans and the need to keep interconnection complexity to a minimum.

B. Factors critical to an IIL's Success

Development of the lab itself is not sufficient to insure that an IIL will be able to broaden successfully the way in which networking research is done. There are many factors beyond the lab design which must be addressed and many of these will quite likely require study and development. One key issue will be the development of configuration and management systems which will enable the IIL to be quickly, easily and accurately reconfigured for different experiments. While there are certainly commercial network management systems, the need for possibly frequent physical and operational reconfigurations will quite likely place demands on these systems that they are not designed to handle. This is also coupled with the need to configure and manage end hosts which will principally be used to generate load in the IIL.

Load generators in an IIL must be able to create a wide variety of traffic conditions. While some work has been done in this area such as [24], there is significant need to be able to generate traffic from a mixture of applications and to keep this mixture current as new applications are introduced. Traffic load must also be highly scalable so that high bandwidth line cards in core routers can be evaluated under heavy load conditions. While we advocate the use of general purpose computers for load generation, there are also commercial load generation systems available which could be used effectively in an IIL.

Another difficulty which must be addressed is the means for measuring and monitoring all aspects of the IIL. Many devices have built in monitoring capabilities (*e.g.* IP flow monitors and SNMP counters in routers) which can be used, but these may or may not be sufficient for all experiments. It seems clear that deploying additional monitoring hardware such as the OC3MON's [25] or even developing new monitoring methods for an IIL will be required.

Widely deployed measurement and monitoring capability in an IIL enable the potential generation of *gigantic* datasets which must be archived and analyzed. The scale

[1] The figure is meant to show only the most general aspects of an IIL's design. The interconnections between components are not meant to be representative of any particular topology and significant components such as switches and last mile gear have been omitted

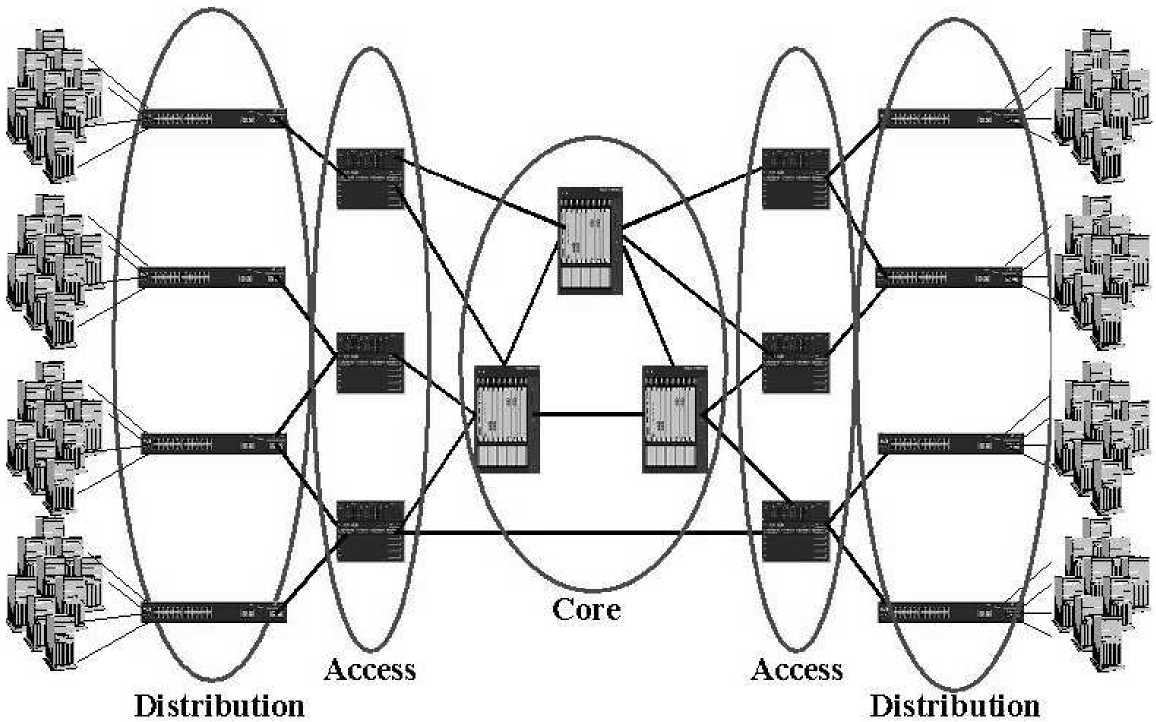

Fig. 1. Conceptual Layout for an IIL

and objectives of experiments will determine the amount of data generated and the types of analysis which will need to be performed. As network technology advances this problem will only increase. This will quite likely result in the development of new means and methods for sampling, summarizing, compressing, archiving and accessing network measurement data.

While the idea of this bench–style environment is to create paths and conditions that are exactly the same as the Internet, validation of results, especially in early stages of IIL development, will be critical. We propose developing partnerships with research staff at ISP's and with researchers managing measurement infrastructures. This will enable access to data which can help in validation efforts.

The final factor critical to the success of an IIL is the support and guidance of the networking community. An IIL is a resource that will be a failure if it is not used by the community. One of the strengths of the *ns* simulator is that the protocols that have been implemented in it have been reviewed by literally thousands of people. Similar attention is important for an IIL.

C. Challenges in IIL Development

There are many technical challenges in realizing the goals of an IIL. Each aspect of an IIL's design and implementation is a challenge many will be on–going. While it is certainly the case that large network equipment manufacturers and ISP have significant network lab facilities, they are typically available for customers and for testing their own equipment and not for general networking research. Other technical challenges include precise packet delay emulation, introduction of new network technologies and visualization of results.

The logistics involved in developing an IIL are perhaps even more significant than the technical challenges. Foremost among these is paying for the equipment and then maintaining it and upgrading it with the latest technology on an on–going basis. The only way to develop and maintain this kind of a system is through significant partnerships between researchers, industry and federal funding agencies. If these entities work together, the vision of an IIL can easily be realized.

IV. Research Enabled by an IIL

In this section we present a partial list of general research areas that are either enabled by or could benefit greatly if an IIL were available. It is clear that almost any network systems problem could be evaluated in an IIL. This list is presented to support our case for an IIL.

- **Network Inference:** Accurate network probes will continue to be important until widely deployed measurement capability becomes available. The challenge of current research in probing methods is that validation is very difficult. The ability to monitor entire paths between endpoints in an IIL would facilitate validation of network probes under a variety of conditions.

- **Protocol Interaction** Quite little is known about how protocols at different layers in the network interact with each other. For example, a change in Boarder Gateway Protocol route advertisements can result in changes of the paths of TCP flows. How this and other kinds of protocol interactions affect network performance is not well understood. The ability to change dynamically many aspects of a network environment would enable this kind of cause–effect relationship to be explored in an IIL.
- **Packet Loss Behavior** For some time Sally Floyd has had a set of research questions regarding packet loss posted on a Web page [26]. Many are quite simply stated such as, "where is the congestion in the Internet?" and "what is the duration of congestion at a router?" however, they remain unanswered. The ability to take fine–grained measurements at all routers along all end–to–end path in an IIL would enable these kinds of questions to be addressed.
- **Modeling and Simulation** While we argue the need for an IIL in this paper we are most certainly not advocating that it replace prior methods for conducting networking research. Quite the contrary. We argue that an IIL would provide a critical means for developing, assessing and validating both analytic and simulation models for networks.
- **Resource Deployment** Determining where to place resources such as caches [27] or measurement systems [28] in the Internet has been an area of investigation for some time. An IIL would enable different kinds of resources to be physically located in different places in an Internet instance network and their performance evaluated under a variety of conditions.
- **Adaptive Systems** Many kinds of network systems from RED routers to applications use dynamic values to govern their behavior. An IIL would make it possible to evaluate the behavior of commercial systems that use feedback under a variety of conditions. This would enable both their stability and performance to be effectively evaluated.
- **Network management** Network managers are responsible for insuring that their networks run efficiently and reliably. Their task is hindered by the size and complexity and heterogeneity of their network and the lack of precise and timely information about it's behavior. Researchers using an IIL will be faced with some of the same challenges. However, the fact that the entire network is self–contained should enable more efficient and effective problem diagnosis. As stated earlier, all software developed for the management of an IIL should also be easy to transition to production networks.
- **Security** An IIL offers a unique environment for evaluating attacks, trace–back methods and intrusion detection. For example virus propagation could be evaluated by releasing the virus in the lab. Similarly, groups of nodes could be configured to participate in distributed denial of service attacks or stealthy port scans which would enable evaluation of both detection and trace–back methods.

V. Summary and Conclusion

In this position paper we present and argue for the development of a bench–style environment for networking research. We call this environment an Internet Instance Laboratory. It would consist of commercial equipment typically found along Internet paths from end–to–end–through–core. The equipment would reside in a single laboratory or possibly in a small number of distributed laboratories connected over dedicated channels. The equipment would be installed to facilitate a wide variety of physical configurations. An IIL would have the capability to generate traffic from common applications and to place varying loads on the network. An IIL would also be fully instrumented enabling measurement data to be collected anywhere in the system. The environment would have significant system management, data archival and data analysis capabilities.

We compare an IIL to current tools and methods for networking research and discuss how an IIL can significantly broadens the way in network research could be done. We argue that limitations of current tools used in networking research, combined with an IIL's capability to measure all aspects of a network and the ability to deploy disruptive prototypes make an IIL's development compelling.

We present general aspects of an IIL's design and outline the capabilities required for its success. The design is based on a simple topological model for selecting equipment and provides great flexibility. Required capabilities include configuration management, workload generation, delay emulation, and access to measurements taken at any point in the network. We also discuss the challenges facing an IIL's development which consist of developing the required technical capabilities and solving the logistical problems of funding and on–going support. While these challenges are significant, support from industry, federal agencies and the network research community will enable realization of an IIL.

Acknowledgements

The authors would like to thank the following people for their input on IIL ideas and design: Fred Baker, Paul Beebe, Jim Gast, Bill Jensen, Robert Olsen, David Parter, David Plonka, Jay Rosenbloom, David Rossetti, Mike Sanders, and Peter Smale.

REFERENCES

[1] Committee on Research Horizons in Networking, *Looking Over the Fence at Networks: A Neighbor's View of Networking Research*, National Academy Press, Washington D.C., 2001.

[2] A. Downey, "Experiences with pathchar," in *Proceedings of ACM SIGCOMM '99*, Cambridge, MA, August 1999.

[3] K. Lai and M. Baker, "Nettimer: A tool for measuring bottleneck link bandwidth," in *Proceedings of the USENIX Symposium on Internet Technologies and Systems*, March 2001.

[4] S. Savage, "Sting: A tool for measuring one way packet loss," in *Proceedings of IEEE INFOCOM '00*, Tel Aviv, Israel, April 2000.

[5] V. Paxson, G. Almes, J. Mahdavi, and M. Mathis, "Framework for IP performance metrics," IETF RFC 2330, 1998.

[6] W. Willinger, "A critical assessment of the current state of internet traffic modeling and characterization," Invited talk at BU/NSF Workshop on Internet Measurement Instrumentation and Characterization, August 1999.

[7] E. Kohler, R. Morris, B. Chen, J. Jannotti, and F. Kaashoek, "The click modular router," *ACM Transactions on Computer Systems*, vol. 18, no. 3, pp. 263–297, August 2000.

[8] The Surveyor Project, "http://www.advanced.org/csgippm/," 1998.

[9] A. Adams, J. Mahdavi, M. Mathis, and V. Paxson, "Creating a scalable architecture for Internet measurement," *IEEE Network*, 1998.

[10] N. Duffield and M. Grossglauser, "Trajectory sampling for direct traffic observation," in *Proceedings of ACM SIGCOMM '00*, Stockholm, Sweden, September 2000.

[11] J. Lepreau, C. Alfeld, D. Andersen, and K. VanMaren, "A large–scale network testbed," in *Proceedings of IEEE SIGCOMM '99 - New Research Session*, Boston, MA, September 1999.

[12] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of Ethernet traffic (extended version)," *IEEE/ACM Transactions on Networking*, pp. 2:1–15, 1994.

[13] V. Paxson and S. Floyd, "Wide-area traffic: The failure of poisson modeling," *IEEE/ACM Transactions on Networking*, vol. 3(3), pp. 226–244, June 1995.

[14] M. Crovella and A. Bestavros, "Self-similarity in World Wide Web traffic: Evidence and possible causes," *IEEE/ACM Transactions on Networking*, vol. 5, no. 6, pp. 835–846, December 1997.

[15] M. Arlitt and C. Williamson, "Internet Web servers: workload characterization and performance implications," *IEEE/ACM Transactions on Networking*, vol. 5(5), pp. 631–645, October 1997.

[16] P. Barford, A. Bestavros, A. Bradley, and M. Crovella, "Changes in Web client access patterns," *Special Issue on World Wide Web Characterization and Performance Evaluation; World Wide Web Journal*, December 1998.

[17] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web caching and Zipf-like distributions: Evidence and implications," in *Proceedings of IEEE INFOCOM '99*, New York, NY, March 1999.

[18] W. Willinger, M. Taqqu, R. Sherman, and D. Wilson, "Self-similarity through high-variability: Statistical analysis of Ethernet LAN traffic at the source level," *IEEE/ACM Transactions on Networking*, vol. 5, no. 1, pp. 71–86, February 1997.

[19] UCB/LBNL/ISI, "Network simulator - ns (version 2)," http://www.isi.edu/nsnam/ns/, 2000.

[20] S. Floyd and V. Paxson, "Difficulties in simulating the Internet," *IEEE/ACM Transactions on Networking*, vol. 9(4), pp. 392–403, August 2001.

[21] Keynote Systems Inc., "http://www.keynote.com," 1998.

[22] V. Paxson, ," Invited talk, Workshop on NetworkRelated Data Management, May 2001.

[23] J.-J. Pansiot and D. Grad, "On Routes and Multicast Trees in the Internet," *Computer Communications Review*, vol. 28, no. 1, January 1998.

[24] P. Barford and M. Crovella, "Generating representative workloads for network and server performance evaluation," in *Proceedings of ACM SIGMETRICS '98*, Madison, WI, June 1998, pp. 151–160.

[25] J. Apisdorf, K. Claffy, K. Thompson, and R. Wilder, "OC3MON: flexible, affordable, high performance statistics collection," in *Proceedings of the USENIX Tenth System Administration Conference (LISA X)*, Chicago, IL, September 1996.

[26] S. Floyd, "Questions," http://www.icir.org/sally/questions.html, 2002.

[27] B. Li, M. Golin, G. Italiano, X. Deng, and K. Sohraby, "On the optimal placement of Web proxies in the Internet," in *Proceedings of IEEE INFOCOM '99*, New York, New York, March 1999.

[28] P. Francis, S. Jamin, C. Jin, Y. Jin, D. Raz, Y. Shavitt, and L. Zhang, "IDMaps: a global internet host distance estimation service," *IEEE/ACM Transactions on Networking*, vol. 9(5), October 2001.