



**Fermi National Accelerator Laboratory**

**FERMILAB-Pub-98/026**

**What is the Internet Doing? Performance and Reliability  
Monitoring for the HEP Community**

R.L.A. Cottrell, Connie A. Logg and David E. Martin

*Fermi National Accelerator Laboratory  
P.O. Box 500, Batavia, Illinois 60510*

January 1998

Submitted to *Computer Physics Communications*

Operated by Universities Research Association Inc. under Contract No. DE-AC02-76CH03000 with the United States Department of Energy

## **Disclaimer**

*This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.*

## **Distribution**

*Approved for public release; further dissemination unlimited.*

# What is the Internet Doing? Performance and Reliability Monitoring for the HEP Community

R. L. A. Cottrell<sup>a</sup>, Connie A. Logg<sup>a</sup>, David E. Martin<sup>b</sup>

<sup>a</sup>*Stanford Linear Accelerator Center, Stanford, CA 94309, USA*

<sup>b</sup>*HEP Network Resource Center, Fermi National Accelerator Laboratory, Batavia,  
IL 60510, USA*

Collaborative HEP research is dependent on good Internet connectivity. Although most local- and wide-area networks are carefully watched, there is little monitoring of connections that cross many networks. Work is in progress at several sites to monitor Internet end-to-end performance between hundreds of HEP sites worldwide. At each collection site, ICMP ping packets are automatically sent periodically to sites of interest. The data is recorded and made available to analysis nodes, which collect the data from multiple collection sites and provide analysis and graphing. Future work includes improving the efficiency and accuracy of ping data collection.

*Key words:* Wide Area Networking; Internet; Monitoring; End-to-end; Performance, ICMP, IP, Ping

## 1 Introduction

High energy physics (HEP) research is characterized by large collaborations whose members are widely scattered at universities and laboratories throughout the world. Although rarely mentioned prominently in project plans, wide-area networking is critical to the success of most collaborations. Much of the day-to-day work of a collaboration is done over computer networks, from such simple tasks as reading electronic mail to such complex tasks as event reconstruction. Today, the Internet and the Internet Protocol (IP) are used for almost all HEP data exchange.

The current worldwide Internet consists of over 30,000 "networks" (both local-area and wide-area) interconnected at various points to provide a seemingly single network to end users. Each network is typically run by an organization

that monitors the network's physical links, routers and logical interconnections. For example, the Energy Sciences Network (ESnet) runs a network to interconnect major energy research facilities in the United States. ESnet provides 24-hour monitoring of all lines and equipment. Likewise, MCI runs and monitors the vBNS network that interconnects research institutions with National Science Foundation (NSF) supercomputer centers.

However, connections between users on different networks are rarely monitored. In order for an ESnet site to reach an NSF center, the traffic will cross a number of different networks. Monitoring performance and reliability of connections across many networks is difficult since no single organization has access to statistics stored on all intermediate nodes. Traditional network monitoring tools based on protocols like the simple network management protocol (SNMP) are unusable with such access.

In 1994, the Stanford Linear Accelerator Center (SLAC) embarked on a task to study connections to research sites collaborating with SLAC [1]. SLAC staff developed a system to collect network performance and reliability statistics and present them in both tabular and graphical formats. In 1996, the ESnet Site Coordinating Committee (ESCC) formed the Network Monitoring Task Force which chose to extend the SLAC work to allow for monitoring of connections between many different sites. The HEP Network Resource Center has been working to developing this new system while SLAC has been working to incorporate their analysis routines into the new system. This paper details the techniques for data collection and dissemination used in the new system and data analysis used in the present SLAC system and planned for the new system.

## 2 Technique Used

Since HEP traffic often crosses many different networks, it is impractical to try and gain access to statistics of transit nodes. Negotiating access rights to router statistics with even a few transit networks has proved to be impossible. The decision was made, therefore, to treat the entire network of intermediate nodes as a black box and monitor end-to-end performance only. Throughout this paper, such end-to-end connections will be referred to as links. Although this technique greatly simplifies data collection, it somewhat limits the utility of the data in diagnosing problems. Because Internet Control Message Protocol (ICMP) messages are almost universally supported, and because the ping command is ubiquitous, ICMP ECHO\_REQUEST messages as generated by the UNIX ping command were chosen as a basis for network monitoring.

All nodes running IP are required to respond to ICMP messages, a family of

packet types used to perform various low-level IP routing maintenance and network diagnostics [2]. An ICMP ECHO\_REQUEST packet (also known as a *ping*) has an IP and ICMP header (which contains a sequence number), followed by an 8-byte timestamp, and then a number of "pad" bytes used to fill out the packet to a specified length. When an Internet node receives such a packet, it responds with an ICMP ECHO\_RESPONSE packet with the same timestamp, sequence number and pad bytes. Since this is a datagram protocol, either the ECHO\_REQUEST or ECHO\_RESPONSE packet may be lost or duplicated.

A very common application of this protocol is the UNIX *ping* command which (by default) sends a single 64-byte ECHO\_REQUEST packet to the host specified and reports whether a resulting ECHO\_RESPONSE packet was received within twenty seconds [3]. Typical options to the ping command allow control of the number of ECHO\_REQUESTs sent, the interval between each request, the number of pad bytes, and the time to wait for an ECHO\_RESPONSE. When used in batch mode, the ping command gives the percentage of packets lost and the minimum, maximum and average response times over all responses received.

ICMP messages are not usually available at the user level and, in fact, on a UNIX system a normal user is forbidden from sending or receiving any type of ICMP packets. On UNIX, therefore, the ping command runs at the root level by doing a *setuid* upon invocation. Receipt of ECHO\_REQUEST packets and response with ECHO\_RESPONSE packets are performed at a low level in the operating system without user-level intervention, making it a good probe of network response time rather than system response time. Unless a system is very heavily loaded, ping packets should be received and responded to without significant delay. An exception to this is some brands of routers, which give low priority to ICMP messages. They may ignore ICMP messages even during relatively light load.

At first anecdotal evidence was used to verify ping as a good measure of user-perceived network performance and reliability. User complaints about a link could often be matched to large packet loss or high response time on that link. Similarly, user reports of improved performance were often matched to reductions in packet loss or response time. Although the correlation was not perfect, data from ping studies was successfully used to choose Internet service providers for SLAC telecommuters, among other uses. In order to provide a more rigorous validation, a study was done to compare times of Hypertext Transport Protocol (HTTP) transfers with ping response times. The study showed that the response time seen by ping is a good predictor of application-level network performance [4]. This correlation was also shown in [5].

Ping, though, is not a perfect measure. It is more likely to give a false posi-

tive, indicating a problem where there is none, rather than a false negative, indicating the network is fine when it is not. Also, pings only give an instantaneous view of the state of the network so periodic pings may miss transient problems.

### 3 Data Collection and Distribution

Data collection is performed on UNIX workstations by running a Perl script that is scheduled by the cron facility. This Perl script is called pingtime. Every thirty minutes, a list of hosts is scanned sequentially. For each entry in the list, one ping packet with a 100-byte payload is sent to the host, then ten such packets are sent to the host, then ten ping packets with a 1000-byte payload are sent to the host. All pings are sent at intervals of at least one second. The first ping is used to prime router caches and address resolution tables and its results are discarded. 100-byte pings were chosen to represent interactive traffic, 1000-byte pings to represent batch transfers. The results of the 100-byte and 1000-byte bursts are stored in a single line that contains:

- IP name of destination node;
- IP address of destination node;
- date and time of beginning of batch job (in long format);
- percentage of loss, minimum, maximum and average of response times (for 100-byte pings);
- percentage of loss, minimum, maximum and average of response times (for 1000-byte pings);
- date and time of first ping to this particular node (in both UNIX ctime and long format).

The line is written into a file containing the entire month of data collected by the source site. This format is based on the original SLAC system. Note that it does not record the source node since all data collection was done from a single node at SLAC.

A new Perl script is being phased into use. It does a number of pings in parallel to increase the number of nodes that can be pinged. In addition it uses a more compact format:

- IP address of source node;
- IP address of destination node;
- size of pings;
- date and time of first ping (in UNIX ctime format);
- percentage of loss, minimum, maximum and average of response times. (If percentage of loss is 100%, minimum, maximum and average are omitted.)

Like the original system, all data for the month is stored in a single file. Since current data analysis and presentation software is still based on the original format, a Perl script is available to convert from the new format to the old. In the original SLAC system, all pinging originated at a single node. This made analysis and reporting simple, but provided only a limited view of the network. The current system improves the breadth of sites examined by providing multiple collecting sites. A collecting site is one that has agreed to compile a list of nodes (called remote sites) to ping and has agreed to run pingtime on a local node. In addition, this node must run the *ping\_data* CGI/Perl script and an HTTP server. The *ping\_data* program allows a remote site to retrieve data for links over a specific time period. The goal is to have several collecting sites per collaboration or other affinity group.

In the original SLAC system, all data analysis was done on a single node at SLAC. The current system makes use of a number of analysis sites. Analysis sites run the *ping\_collect* Perl program which collects data from the collecting sites and store a copy locally. Analysis sites run the SAS environment which provides for graphical and tabular analysis of the data. This analysis is discussed in detail in the next section. The results of the analysis are made available through the world-wide web in pre-packaged overnight reports and dynamically generated reports.

## 4 Reporting

Data is analyzed and presented in reports accessible via the WWW. The examples given in this section are from the original SLAC system, but the plan is to migrate these reports to the new system. The reports provide both short term (last few hours or days) and longer term (last fortnight, last 10 weeks, last 180 days and going back several years) information. The short term reports are mainly used for trouble-shooting and understanding the current state of connectivity. The longer term reports are mainly for looking at trends and planning.

### 4.1 Short Term Reports

Short term reports provide information on the performance (response time, packet loss, reachability) of remote hosts measured so far this day, for yesterday, for the last 14 days, and for the last 30 days. Figure 1 shows an example of such a report for one day. The day versus night (at RAL) response time differences are striking in this figure and indicative of congested links during the English daytime.

RAL (England) – Jan. 30th, 1997

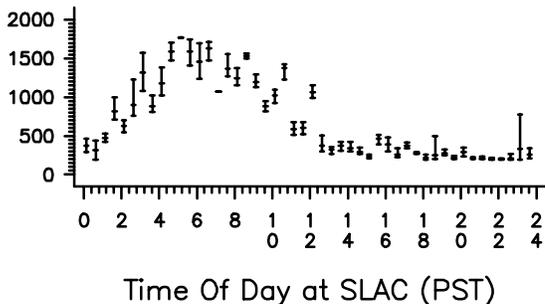


Fig. 1. Average, minimum and maximum (of ten 1000 byte payload pings) response time between SLAC and the Rutherford Appleton Laboratory (RAL) in England, for one day. The time of day is in Pacific time, so midnight on the graph is 8 am in England.

To generate alerts we calculate the averages and standard deviations of the ping response and packet loss to each remote host for the previous 10 weeks. This is repeated for the last 7 days and for yesterday. Alerts are raised if the last 7 days or yesterday’s averages are over 3 standard deviations greater than for the last 10 weeks. These alerts may be provided via email or active links in WWW reports.

#### 4.2 Medium Term Reports

Medium term trends are provided by plots showing the average daily ping response and packet loss for the last 180 days for each remote host. Figure 2 shows a typical example measured between SLAC and the University of California at Davis (UCD). Immediately visible in this figure is a degradation in weekday response by almost a factor of 4 in this 180 day period. The packet losses (not shown here) for this period are more variable and increase by about 50%. The big difference between weekday and weekend performance is again indicative of congestion.

To provide a measure of performance predictability (in particular the variability between day and night time performance), we calculate for each set of 30 minute interval measurements the 100-byte payload packets the ping  $success = 1 - \text{packets lost}^1 / (\text{total number of packets})$ , and for 1000-byte payload packets the ping  $data\_rate = 2000 \text{ bytes} / (\text{average response time of 10 consecutive ping packets})$ . Then for all  $successes$  and  $data\_rates$  to a given host in one day we calculate the dimensionless ratios:  $s = avg(success) / max(success)$

<sup>1</sup> This excludes measurements with 100% packet loss, these are accounted for in the *unreachability* analysis.

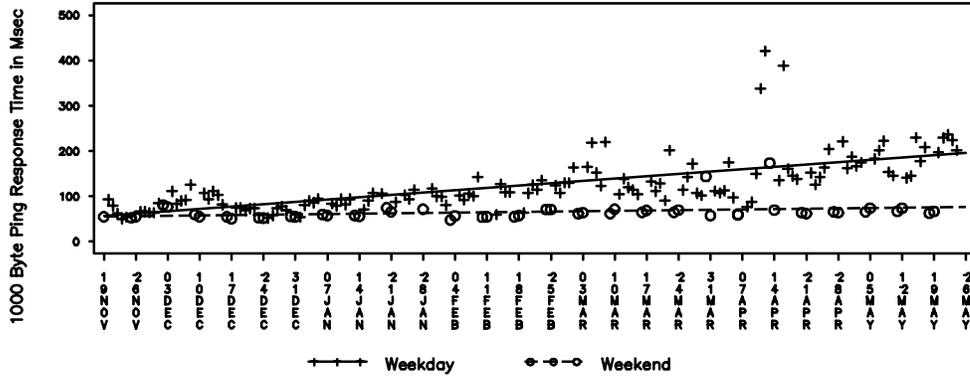


Fig. 2. 180-day trend plot of the ping response (for 1000 byte ping payload) between SLAC and UCD starting 19 November 1995. The lines are linear regression fits to aid the eye.

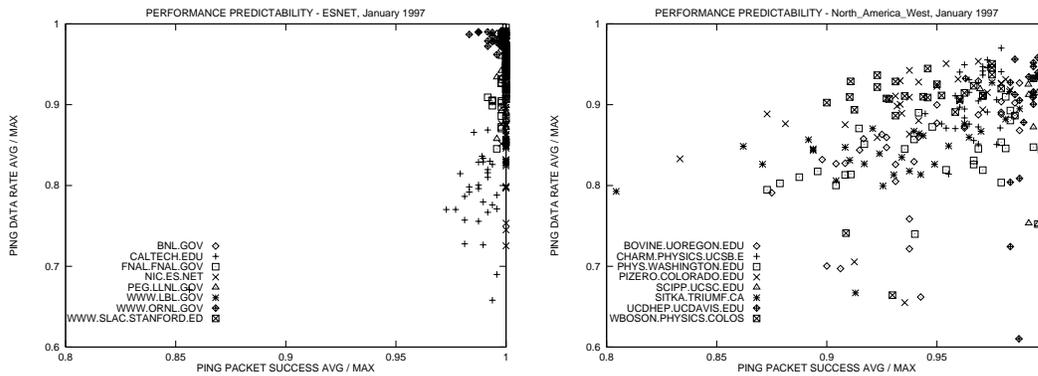


Fig. 3. Performance predictability between SLAC and ESnet hosts and between SLAC and western N. American hosts monitored by SLAC for January 1997. Each point on the plots represents  $(s, r)$  for one host for one day.

and  $r = avg(data\_rate)/max(data\_rate)$ . We then scatter plot the daily success ratio  $(s)$  versus the rate ratio  $(r)$  for a given month for a set of hosts. Values of the ratios close to one indicate the average performance is close to the optimal performance. Ratios much less than one occur particularly on links which are congested during prime time. Examples of such plots are seen in Figure 3. which indicate that performance predictability was much better between SLAC and ESnet hosts than between SLAC and western North American hosts, presumably because SLAC has a better connection to ESnet hosts without passing through the commercial Internet.

### 4.3 Long Term Reports

To provide trends going back over longer periods, we calculate the average response time and the average ping loss for each month for each host. Since

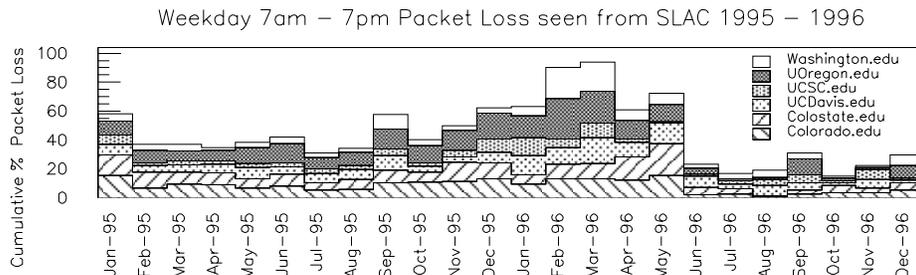


Fig. 4. Average SLAC prime time (7am - 7pm, weekday) monthly packet loss between SLAC and some western N. American HEP hosts. Note the general improvement Apr-Jun '96 following the improvement in ESnet connections to the Internet. Also it is seen that there is considerable variability from month to month and host to host.

most of the interest concerns the performance during working hours, we include only weekday ping data measured between 7am and 7pm (SLAC time). An example of such a plot is seen in Figure 4.

We identify a host as being down or unreachable when no ping response is obtained in the set of 21 pings made each 30 minute period. Using this identification, we calculate the ping *unreachability* = (# periods host is down) / (total number of periods), the number of down periods per month, the Mean Time Between Failures for each remote host. The unreachability is plotted (one point per host/month) and the other information is provided in tabular form.

To provide a broader overview of the performance, we average the various indicators (response, packet loss, unreachability, unpredictability (defined as the distance of the coordinate  $(s, r)$  from  $(1,1)$ )) for each month over groups of hosts. Typically the grouping is geographical or by service provider. The main host groupings we use are: ESnet, western N. America, eastern N. America, international (non ESnet and non N. American), and local Internet Service Providers (ISPs) in the San Francisco Bay Area. Figure 5 shows examples of group ping unreachability and unpredictability.

One of the more easily understood and critical metrics, for the end user, is the packet loss, since packet loss results in timeouts which have a large impact on the performance of network applications. To provide an upper level view of the packet loss we arbitrarily divide the losses into 5 *quality* categories:

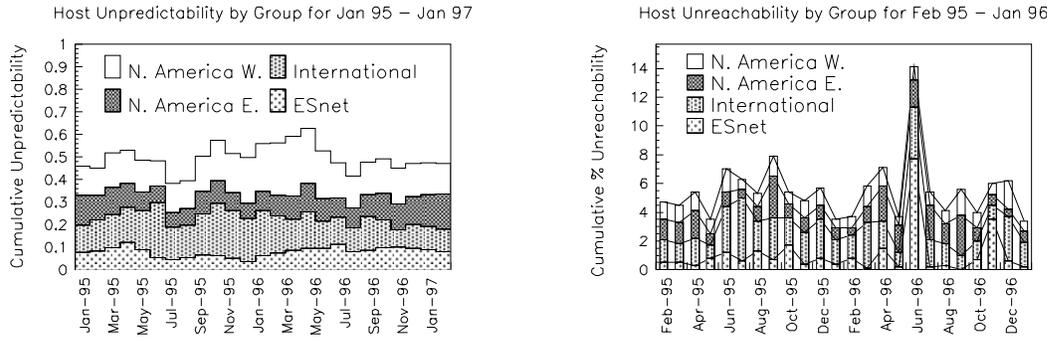


Fig. 5. Group ping unpredictability and unreachability. It can be seen that the unpredictability is worse for international hosts and western N. American hosts. Western N. American hosts became particularly unpredictable in Spring 1996 before the ESnet links to the Internet were improved. The unreachability peak in June 1996 was due mainly to a host which was shutdown while it was relocated.

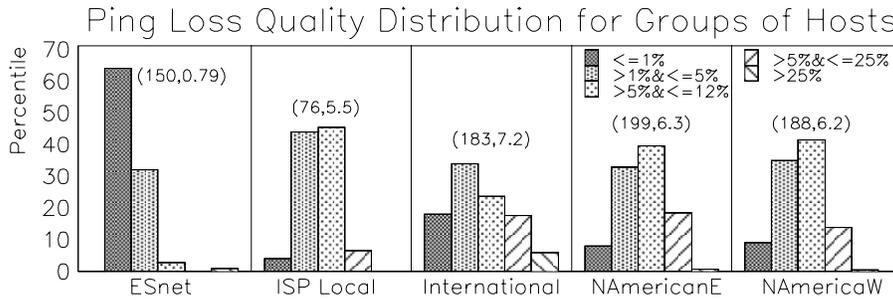


Fig. 6. SLAC prime time ping loss quality distributions for groups of hosts from January 1995 through December 1996. The numbers in parentheses are the number of host-months and the median packet loss.

- $\leq 1\%$ <sup>2</sup> packet loss  $\equiv$  Good WAN performance.
- $> 1\%$  &  $\leq 5\%$  packet loss  $\equiv$  Acceptable WAN performance.
- $> 5\%$  &  $\leq 12\%$  packet loss  $\equiv$  Poor WAN performance.
- $> 12\%$ <sup>3</sup> &  $\leq 25\%$  packet loss  $\equiv$  Bad WAN performance.
- $> 25\%$  packet loss  $\equiv$  Unusable WAN Performance.

We then find the percentage of months for which each host fell in each of the above categories. We average these percents for each group of hosts and plot the distributions (i.e. host group vs. category vs. percentile in each category) which is shown in Figure 6. ESnet ping loss performance is seen to be good or acceptable over 95% of the host-months. The other groups, however, typically have packet losses which are poor or worse over half the host-months.

<sup>2</sup> 1% is the threshold we use on the SLAC local area network for generating alerts.

<sup>3</sup> The "Internet Weather Report" (<http://www.internetweather.com/>) marks networks as "RED" if the packet loss is  $> 12\%$ .

## 5 Conclusions and Futures

Ping is an excellent tool for end-to-end network performance monitoring. It provides almost universal coverage. Administrators at the monitored remote hosts do not have to install any software. It has low network impact if used wisely. It provides useful short- and long-term measures of bottleneck bandwidth, available bandwidth [5], response time, packet loss, reachability, and predictability which can be related to user applications.

The distributed architecture of the new system provides the ability to scale the system to many collection sites and thousands of links monitored. The amount of work for remote and collection sites is very low. Providing a central repository of ping data at an analysis site allows for trend analysis across all links monitored.

A major challenge has been coming up with simple, intelligible ways to characterize and visualize the enormous amounts of data. The results indicate that by most measures, performance within ESnet is acceptable to good. However packet loss performance between ESnet and the Internet at large is, on average, poor or worse for the hosts monitored. Packet loss seen from SLAC for non-ESnet hosts improved dramatically between April and June 1996, and the improvement has been sustained. In general performance is very variable in both the short and long-term, particularly for international hosts. From SLAC, average monthly response times by host groups are typically 300-500 ms. for international hosts, 150-220 ms. for eastern N. American hosts, 80-140 ms. for western N. American hosts, and 40-50 ms. for ESnet hosts.

The methodology is also being utilized to: select ISPs and monitor their performance possibly with a view to writing a service contract; help decide which universities to connect directly to ESnet; and, to identify bottlenecks in order to decide where to focus efforts.

Possible future work includes: performing the measurements with Poisson sampling, which, in principle results in unbiased measurements, even if the sample rate varies [6]; looking at better definitions of prime hours so the definition is less monitoring-site-specific and more realistic for non-U.S. links; increasing the number of monitoring sites by distributing the monitoring software (in particular add monitoring hosts in non-U.S. countries); responding to requests by a number of HEP-related organizations to add more remotely-monitored hosts; more carefully choose the remote hosts to monitor; and, install a range of fixed size WWW pages at various sites to look at long-term WWW responsiveness.

## 6 Acknowledgements

We would like to thank Bill Wing for help and encouragement in this work, Lois White and Charley Granieri for help with producing graphs, John Halperin for expert criticism, and Mike Wendling for much help with reformatting plots for embedding in LaTeX.

This work was supported by the United States Department of Energy under contracts DE-AC03-76FO0515 and DE-AC02-76CH03000.

## References

- [1] C. A. Logg and R. L. A. Cottrell *Network Monitoring and Performance Management at SLAC*. SLAC-PUB-95-6744, Stanford Linear Accelerator Center, CA. Presented at Networld+Interop Engineers' Conference, Las Vegas, NV, March 1995.
- [2] J. Postel, *Internet Control Message Protocol*, RFC-792, DARPA Internet Program, September 1981.
- [3] SunSoft, *SunOS 5.4 Reference Manual, Section 1M*, SunSoft, Mountain View, CA, 1994.
- [4] R. L. A. Cottrell, C. D. Granieri, J. H. Halperin, G. Haney, C. A. Logg, D. E. Martin, W. R. Wing, *Internet Monitoring in the Energy Research Community*, submitted for publication in IEEE Communications Magazine, 1997.
- [5] R. L. Carter and M. E. Crovella *Measuring Bottleneck Link Speed in Packet-Switched Networks*. Tech. Report BU-CS-96-006, Computer Science Department, Boston University, March 1996.
- [6] V. Paxson *End-to-End Routing Behavior in the Internet*. Proc. SIGCOMM '96, August 1996.