Computer
Sciences
Department

On the Precision of Active Probes for Packet Loss

Paul Barford
Joel Sommers

Technical Report #1442

September 2002 (revised)

UNIVERSITY OF
WISCONSIN
MADISON

On the Precision of Active Probes for Packet Loss

Paul Barford and Joel Sommers
Computer Science Department
University of Wisconsin–Madison
{pb, jsommers}@cs.wisc.edu

Abstract—**Active probing is one of the basic methods for measuring network properties and behavior. One of the most important aspects of these methods is that they provide information about end-to-end path characteristics without requiring direct access to individual links along a path. In this paper, we evaluate the *precision* of a standard type of active network measurement: probes for packet loss. We evaluate packet loss data gathered over a two week period in a widely deployed infrastructure with access to both backbone router interfaces and co-located hosts that send active probes in a full mesh. Our assessment of precision is based on comparing loss rate characteristics in each of these data sets. We find that there is generally low correlation between loss rate measurements from active probes and those reported from router interfaces. Specifically, we find that the correlation coefficients for time series of router measures of loss and active probes for all of the paths in our infrastructure to be low. We also compare the distributional characteristics of loss including lengths of loss free periods, loss rates during lossy periods, and measures of loss constancy. We find the degree of agreement between router measures and active measures for each of these characteristics to be quite low. Deeper evaluation of our data indicates that current methods for active probing for packet loss suffer from the coarse timescales over which they sample and from effects of end-host interface loss.**

I. Introduction

Packet loss due to congestion is a fundamental issue in wide area packet switched networks. Great effort has been expended to characterize and model this phenomenon and to design protocols and engineer networks that effectively avoid, control and recover from packet loss. While progress has certainly been made, packet loss and its effects on performance remain a significant problem for both network researchers and network operators.

Coupled with the evolution of network protocols and network systems is the process of deepening and broadening our basic understanding of packet loss behavior through empirical study. A number of significant studies of packet loss behavior have had important influence on current network systems. Two of the best examples of protocols that have benefited directly from empirical observations of packet loss behavior are the NewReno [1] and SACK [2] versions of TCP. However, the Internet is a constantly changing environment, and this makes continued evaluation of important phenomena, such as packet loss, critical.

There are two basic methods for measuring packet loss. The first is through passive monitors attached to network links or nodes. A standard example of passive monitoring capability is the set of Management Information Base (MIB) counters available on network nodes via the Simple Network Management Protocol (SNMP) [3]. These counters track a wide range of activity including packet losses due to congestion. The benefit of passive monitoring systems is that they accurately capture many of the important details of local traffic behavior. However, the cost for this detail is often high (*e.g.,* in terms of data storage requirements) and access to links or routers is frequently not possible.

The second means for measuring packet loss is active probing of the network. The most simple active probe which measures packet loss is the `ping` utility. Like all active probe tools, `ping` sends a series of packets into the network aimed at a target system and measures the response packets returned to the sending system. Lost packets are tracked by the sender through the use of sequence numbers. The benefit of active probes is that they can be run from virtually anywhere in the network and that they give an end-to-end perspective of network behavior. The difficulty is that the discrete nature of active probing limits the resolution of the measurements. If more frequent probes are sent into the network then resolution should increase, but if the frequency is too high then the probes themselves can skew the results (a so-called *Heisenberg effect*). Despite these difficulties, active probing remains one of the most important methods for gathering packet loss data.

In this paper we address a very simple question; "how precise are active measurements of packet loss?"[1] The simplest assumption is that precision is based on the sampling rate and that if we assume that the rate is sufficient, there should be good correlation between measured and true values. We address this question by measuring packet

[1] A definition of precision is the degree of agreement between a measurement and its true value.

loss over a two week period using SNMP at all backbone routers in the Abilene/Internet2 infrastructure. This environment enables us to gather SNMP loss data at 30 second intervals. We aggregate loss data from all interfaces along each path in the full mesh of paths to obtain end-to-end perspectives on loss behavior. We treat these measurements as the baseline from which we will compare a set of active probes for loss taken in the same infrastructure.

The active probing tool we use to measure loss is the `zing` utility [4] which sends probe packets at exponentially modulated intervals. This probing method should provide unbiased, time-average data for loss conditions along an end-to-end path. We take one-way measurements of packet loss by running `zing` between nodes in the Surveyor infrastructure [5] that are directly connected to the Abilene backbone routers. This enables us to probe paths in a full mesh in this backbone without the risk being unable to account for packet loss at intermediate routers. We set our average probe rate to 10Hz and then aggregate the measured loss rates into 30 second intervals to assess precision versus the SNMP data.

Instead of attempting to develop a single metric for precision, we evaluate the degree of agreement between the active measurements and the SNMP measurements along a number of dimensions. First, we compare the correlation coefficients for the time series of loss rates for each end-to-end path. Our results show that there is virtually no correlation between loss rates measured by active probes and loss rates measured by SNMP. Next, we compare distributional characteristics of loss measurements for a number of different loss properties including lengths of loss-free periods, loss rates during lossy periods, and measures of loss constancy as described in [6]. In each case we find a very low degree of agreement between the distributions. This leads to our overall conclusion that active probes for loss are generally quite imprecise.

It seems clear that there are a number of possible reasons for the lack of agreement between the two data sets. The first is that the sampling rate we employ in our active measurements is too coarse to enable typical loss episodes in this infrastructure to be measured accurately. We did not experiment with loss probes that sample more frequently and leave that for future work. We did experiment with the probe process by comparing Poisson modulated `zing` probes with simple `ping` probes sent at the same rate (10Hz). We found negligible difference between the two probes. We attribute this to the very low overall loss rates we observe in our data. Another possible reason for poor precision is that there may be artifacts in our measurements that bias the results. One such artifact is interface loss on the active probe systems. We see examples of interface loss in our data when there is a loss measured by the active probe but no associated loss measured by SNMP. We attribute these losses to the interface/end-host[2]. While occurrence of these losses is rare, we evaluate precision after censoring them from the data and still find very low precision in active probes.

Our work has implications in a number of different areas. First, it suggests that new active probe methods for loss may be necessary to get a more accurate picture of loss behavior due to congestion. Next, network operators and systems that monitor active probes for loss may need to consider other means for collecting this data. Another implication is for characterizations and models of loss processes which have been developed based on probe measurements. Our study suggests that these models may need to be revised.

The rest of this paper is organized as follows. In Section II we discuss work related to this study. In Section III we discuss some of the problems inherent in active probes of loss. Section IV presents the details of the data that we collected and evaluated in this work. In Section V we compare the active probe loss measurements with the SNMP loss measurements to assess the degree of agreement between the two. We summarize our study and discuss future work in Section VI.

II. Related Work

To our knowledge there has been no prior work which attempts to assess the precision of active probes for packet loss. However, recent work by Pasztor and Veitch identifies limitations in active measurements, and proposes an infrastructure using the Global Positioning System (GPS) (quite similar to Surveyor [5]) as a means for improving accuracy of active probes [7]. Their work is validated by comparing passive measurements of packet delays *at end-hosts* to delay measured by active probes but does not address the precision of loss measurements from the perspective of nodes in the network.

There have been many studies of packet loss behavior in the Internet. Work by Bolot [8] and Paxson [9] used active probe measurements to establish much of the baseline for understanding packet loss characteristics in the wide area. These characteristics include correlation structures on fine time scales and typical loss rates. Yajnik *et al.* evaluated correlation structure on longer timescales and developed Markov models for temporal dependence structures [10]. Recent work by Zhang *et al.* assesses three different aspects of *constancy* in loss rates in an infrastructure which

[2] The other possible cause would be imprecision in the SNMP measurements. However, these hardware counters are verified by their vendor to be extremely precise.

has many similarities to our own (many of the links traversed by their active probes were in Internet2/Abilene). That work evaluates an important notion of a loss process called a "change free period", which is a period of time during which a loss rate appears well-modeled as steady. We evaluate change free periods in our data.

We use `zing` to measure packet loss in one direction in this study. This relies on coordinated end-hosts which are not always available when taking active measurements. Savage developed the *Sting* [11] tool as a means for solving this problem. Sting uses a clever scheme for manipulating a TCP stream to measure packet loss *in both the forward and reverse direction* from a single end host.

There are a number of widely deployed measurement infrastructures which actively measure wide area network characteristics [5], [12], [13]. These infrastructures use a variety of active probe tools to measure loss, delay, connectivity and routing from an end-to-end perspective. Of these systems, only Surveyor can monitor individual nodes *within* the network.

III. Issues in Active Probes for Loss

The appealing features of active probing are offset by a number of issues related to network behavior. One issue already described is that of the *Heisenberg effects*. Another that is central to this study is that of measurement time scale.

In an attempt to develop intuition about how precisely loss is measured by active probes, we performed a number of simulation studies using the ns-2 simulator [14]. The topology we used is a simple "dumbbell", with two routers (A and B) connected by a bottleneck link. For the experiments we describe here, the bottleneck bandwidth between A and B is set at 1.5Mbps with a delay of 25 milliseconds. The probes are sourced from A to B at Poisson intervals with a mean of 100ms in one experiment, and with a mean of 50ms in the second. An infinite TCP source from A to B runs for the duration of the simulation, which is 10 simulation minutes. Background traffic in the direction of the probes and the TCP source is generated through a fixed number of ON/OFF sources. We used three regimes of background traffic for our studies in order to generate different levels of congestion at the bottleneck router. The first regime was no background traffic, so any congestion generated at the router is a result of the TCP source. The second and third regimes were tuned to generate loss averages over the duration of the simulation of 1% and 10%.

Figure 1 depicts the difference in loss measured by the probes versus the loss recorded at router A with relatively low background traffic (i.e. an average of 1% loss is generated) for two different probe sampling rates. In Figure 1(a), the probes are sent at an average of 10Hz, while Figure 1(b) shows probes which are sent at an average rate of 20Hz. Vertical lines indicate the arrival of a probe at the bottleneck router. Hashes indicate when any packet is dropped at the router, and boxes indicate when probes are dropped. The figures are representative portions of each simulation.

The key feature to note is that relative to the loss recorded at the bottleneck router, very few probe packets are lost. Over the course of these 10 minute simulation runs, only 8 probes are lost out of nearly 6000 for the 10Hz case and only 17 probes are lost out of around 12000 for the 20Hz case. Since overall packet loss for each simulation run is close to 1%, we conclude that the measurement time scale of the probes is essentially mismatched with the time scale granularity of loss events at the bottleneck router.

Further experiments involving higher probe rates and different levels of background traffic also show that the active probes either continue to miss many loss episodes, or that they begin to induce congestion themselves, thus measuring more loss than would have taken place otherwise. Even at rates slightly higher than 20Hz, the probes begin to induce unnecessary loss.

IV. Data

A. Measurement Infrastructure

Our measurement infrastructure is unique in that it consists of widely dispersed end-host measurement stations as is typical in Internet measurement projects, and also includes production routers in the Abilene backbone of Internet2. We send active probes across the full mesh of end hosts and collect one way loss measurements at each host. We also periodically query backbone routers via SNMP to collect router interface counters.

In typical active measurement studies, the network under study must in most cases be viewed as a *gray box* [15]. Inferences are made based on information known about algorithms employed in the network (*e.g.,* drop-tail queueing behavior) or known topological information. Because of this limited visibility into the system, we have a gray box instead of an opaque, *black box*. The ability to capture router interface counters exposes many details of the network links that make up the paths under study, and permits us to compare our limited view from end-hosts with a more "omniscient" view from inside the network. We are, however, limited in our knowledge of internal (*i.e.,* router) loss. We discuss these limitations below.

Figure 2 depicts the topology of our infrastructure. Each circle represents a backbone router in Abilene (the only

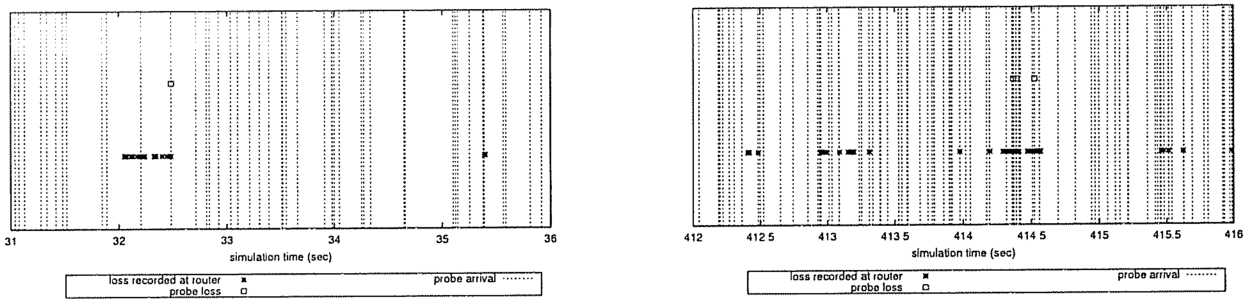

(a) 10Hz Probe Rate

(b) 20Hz Probe Rate

Fig. 1. Loss Measured by Poisson Probes in Simulation

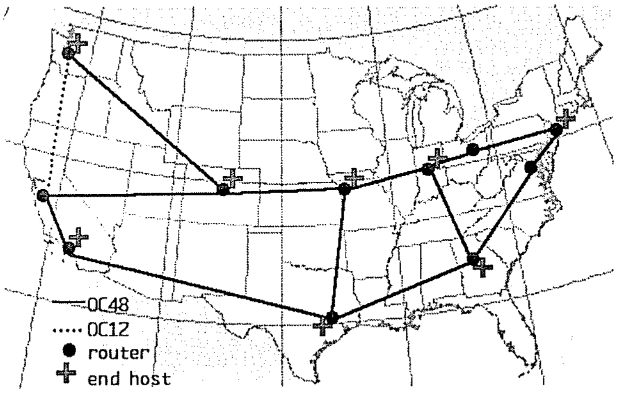

Fig. 2. Map of Abilene Backbone and Measurement Stations

Abilene backbone router absent from the picture is in Chicago, which we exclude from our study). Each hash represents the location of a measurement host. In all cases but one (New York), these hosts are directly connected to a backbone router. In three cases (Sunnyvale, Cleveland, Washington D.C.), we do not have measurement hosts. In the case of New York, there are two hops from the measurement host to the backbone router in New York. In total, we take end-to-end probe measurements from eight hosts (comprising 56 distinct paths) and collect SNMP interface data from eleven routers (roughly 30 interfaces.)

The end hosts each run BSD/OS version 3.1. Six of the eight hosts have ATM OC-3 interfaces directly connected to the backbone routers using Fore (Marconi) 200e ATM network interface cards. The remaining two hosts (New York and Houston) are connected to the backbone routers via 100Mbps Ethernet. The reason for distinguishing these two types of connections is that the routers under study process incoming packets differently in each case. Data arriving on an ATM interface may take a "fast path" through the router, while data arriving on an Ethernet interface takes the "slow path" in all cases. This difference is discussed further below.

The routers are Cisco 12008 Gigabit Switch Routers (GSRs). The GSRs run a variant of IOS version 12[3]. The backbone links are all OC-48 (2.4Gbps), except for the link between Seattle and Sunnyvale, which is OC-12 (622Mbps.) Link utilization over the period of our study averaged 12% with a standard deviation of 10%, indicating that utilization rates vary widely across the Abilene backbone.

B. Data Collection

The data we present and analyze in this paper was collected over the two week period from April 24, 2002 to May 8, 2002. In this section we describe the specifics of data collection for the active measurements and for the routers.

B.1 SNMP Router Interface Data

Our router interface data was collected through a process which queried backbone link interfaces MIBs every 30 seconds. Ingress and egress packet counts, interface drop counts and error counts were collected from counters in the MIB-II `ifTable` and `ifXTable`. In addition, a Cisco enterprise MIB that gives more complete information on interface drop counts was polled. For each measurement, we additionally noted the last interface change time stamp available in the MIB-II `ifTable` and the operational and administrative statuses to ensure we did not

[3] The specific versions of IOS on the GSRs are a mix 12.0S and 12.0ST. Build revisions are mostly the same for each subversion S and ST at 21 and 19, respectively.

collect invalid data, and to assist in detection of counter wrap-around.

The reason we must poll the Cisco-specific MIB is that the `ifInDiscards` entry in the MIB-II `ifTable` only counts one type of packet discard which can happen on input[4]. Inexplicably, output counts do not have this limitation. While we do not have detailed categorization of why packets are dropped, such as can be obtained from the IOS `show interface` command, we have complete information on packets which are dropped at a given router interface.

Polling more often than 30 seconds yields diminishing returns. Besides increasing router CPU load, the router MIBs are not updated in real time. Individual interfaces propagate local counters to the main processor module approximately every 10 seconds. We decided on 30 seconds as a compromise between increased load on routers and sufficiently detailed data.

It is important to provide some detail on how packets are actually lost inside the GSRs from an operational perspective, and for understanding the meaning and limitations of our measurements. Tracing the lifetime of a packet through a GSR, the packet may be dropped in the following areas:

Burst buffer Upon arrival at an interface, the packet is copied into a "burst buffer" of size 2 × MTU where it awaits input buffer allocation. The primary cause for this type of drop is the inability of the physical interface module to allocate buffer space in a timely manner. This situation can occur with extremely heavy volume of small packets. Actual buffer space may exist, but it cannot be acquired fast enough.

Input queue drop In deciding the output interface for a packet, the GSR attempts to make a routing decision in an interface interrupt handler using a cached exact match. This fast path routing decision is the most common path packets take through a Cisco GSR. For packets which cannot be routed using this fast path logic, they are queued on input awaiting slower processing by the main router processing module. If this input queue exceeds the configured size, packets are dropped. All packets bound for the router itself must take this slow path. Additionally, packets arriving on some interfaces invariably take the slow path. Notably, this slow path is taken for packets arriving on 100Mbps Ethernet interfaces, and other interfaces with relatively slow line rates.

No input buffers Lack of a properly sized input buffer can cause a packet to be dropped. This drop can occur either on the slow path or fast path of routing decision.

Switching fabric Internal switching fabric congestion can result in packet loss internal to the router.

Output queue drop This loss situation is the archetypical situation of congestion in a statistically multiplexed packet switched network. The output line rate is less than the aggregated input source rates. Packets are queued awaiting transmission and are dropped when the output queue is full according to an algorithm such as drop-tail or RED.

Of the above, the only type of drop we cannot measure through our SNMP polling is loss due to congestion in the internal router switching fabric. This type of loss detection requires debugging capability to the router and cannot be gathered from SNMP. Losses of this type are thought to be very rare, although we were not able to find a means for quantifying this phenomenon.

Using the SNMP-based loss rates measured at each router, we calculate the loss rates for paths with multiple hops using a union of loss probabilities. Specifically, we calculate loss rate L for a multi-hop path p for a given 30 second period as $L_p = 1 - \prod_{i=1}^{n}(1 - l_i/t_i)$ where l_i is the sum of packets lost during a 30 second period at router i and t_i is the sum of packets transmitted and packets lost at the same router during the same period.

Another way to consider this loss rate calculation is from the perspective of a total loss rate for each path. Using this measure we would sum the number of packets lost at all hops along a given path, dividing this value by the total number of packets transmitted plus total lost at all hops. This calculation would result in lower path loss rates than the formula we used. An argument could be made that this rate is important as well since it reflects a notion of total loss in the network. We chose not to employ this method in our evaluation of precision because end-to-end measurements would not be able to infer this rate unless tomography methods were employed to identify loss rates at individual hops [16].

B.2 Active Probe Data

The active measurement data was collected by using a modified version of the `zing` utility installed at our eight end hosts. We sent 256 byte probes at exponentially distributed intervals with a mean of 100ms. In our data analysis, we refer to these traces as `zing` traces. In parallel, we sent 256 byte probes with a uniform spacing of 100ms. This uniform probing methodology is essentially the same as the ubiquitous tool `ping` and in our data analysis, we refer to this data set with the same name. The probes were sent continuously over the two week period of our study.

We had to modify `zing` because we were unable to use

[4] The `ifInDiscards` counter in MIB-II counts input drops due to lack of buffers, which is distinctly different than lack of input queue space. We have to consult a Cisco interface table in order to obtain input queue drops.

the packet filter capability of the utility due to practical limitations with the kernels installed on the measurement hosts with ATM interface cards. We also modified `zing` in order to facilitate data storage in reasonably sized files. Retrieval of the data was initiated from a host at the University of Wisconsin-Madison and was done every four hours.

In addition to running probes for packet loss, we took `traceroute` measurements across the full mesh of end hosts every 10 minutes. This data enabled us to determine the sets of router interfaces that were encountered along each end-to-end path in our mesh. The loss data from specific sets of interfaces was then compared with active measurement traces between end hosts using the method described above. Since our study was conducted in the backbone of Internet2, routes were extremely stable. There were 122 unique paths observed, thus 66 route changes during the period of our study. These changes were confined to a three day period - all happened at around 4am UTC at the Denver router. The regularity and specificity of the changes led us to believe that this was a standard maintenance activity on the Denver router.

Because we cannot use packet filters at our end hosts, we had no way to determine whether measured loss was due to some event internal to the network (uncounted in our router measurements), or whether it occurred at an endpoint. We could not differentiate measurement host OS buffer overruns or interface drops from network congestion. We can (and did), however, detect interface errors by periodically running `netstat`. We return to this issue in our data analysis.

In order to compare our `zing` and `ping` traces with the SNMP data, we aggregated the probe traces in intervals of 30 seconds to match the SNMP query frequency. The result is that we have comparable time series at the possible cost of lost insight into events on smaller time scales for the active measurements. This aggregation causes our analysis to be conservative in the sense that even if loss events are measured both by SNMP and by an active probe in the same interval, it appears that the active probe has detected the router loss event. The aggregation of data is therefore favorable for making the case that active probes can indeed precisely measure end-to-end loss and unfavorable for making the counter argument.

V. Results

Our analysis of the precision of active probes is based on the assumption that measurements of loss from SNMP represent a "true value" to which the active probe measurements can be compared. We noted above that this measure does not, in fact, account for *all* losses within a router. It does, however, provide a more refined measure of loss and should serve well as a baseline.

The first step in our analysis was a qualitative comparison of loss rates for the two different measures. We followed this assessment by comparing four distributional characteristics of our data: loss rates, lengths of loss-free periods, loss rates during loss periods, and loss constancy (based on the notion of change free periods).

In each of our analysis, we first explore the characteristics of the router measured loss over all our paths to provide an understanding of the baseline to which the probe data would be compared. Next, we look at the distributional characteristic for all loss measures (SNMP, `zing` and `ping`) along a "canonical path." Finally, we quantify the degree of distributional agreement between `zing` and the router counters and between `ping` and the router counters. We additionally provide tables giving summary statistics for each area of analysis.

The standard definition for precision *validity* is "the deviance of a measure from the true value." Since we take the SNMP data as the true measure of loss, we come to a slightly different formulation of this definition: the validity of precision is the deviance of a measure from a measure *as close as possible to the true value.* To compare deviation from the true value from a distributional perspective, we use the χ^2 goodness-of-fit test with with 9 degrees of freedom. We arbitrarily chose 9 degrees of freedom as a level which conservatively favors finding agreement between two distributions. Other measures of agreement between distributions such as relative entropy could have been employed however our object was to make more simple quantitative comparisons while at the same time demonstrating details of the distributional characteristics.

We chose the path from Denver to Indianapolis as our canonical path. The route from Denver to Indianapolis was normally measured as traversing a single intermediate router in Kansas City. On April 25, we detected two route changes. Around 4:15 UTC, the route switched to a path traversing Sunnyvale, Los Angeles, Houston, Atlanta, ending at Indianapolis. Later, at 4:25, the route changed to Sunnyvale, Los Angeles, Houston, Kansas City, and Indianapolis. By 4:35, the route had returned to its original path, with one hop through Kansas City. We could see from our SNMP-level traces that the interface between Denver and Kansas City had been administratively taken down, causing the route fluctuation. Our choice of canonical path was arbitrary, but is qualitatively representative of other paths under study.

We also note that the end points of our canonical path have direct ATM interfaces to routers. This choice is also arbitrary since we do not see fundamental differences between loss measurements taken between hosts connected by ATM or by Ethernet.

A. Qualitative Comparison

In Figure 3 we show time series graphs for the router, zing, and ping data for the canonical path. The left graph shows the entire measurement period, while the right graph gives a detailed view of 6 hours on 6 May 2002. Note that the y-axis is log scale.

Qualitatively, zing and ping largely overestimate the lost packets counted by the router interfaces. What is important to note in these graphs is the lower bound of loss rate measured by active probes. This bound is a function of the probe rate and the time interval considered. For example, with our mean probe rate of 10Hz, we send an average of 300 packets per 30 seconds. These parameters set the effective lower bound on loss at a rate of 0.003.

In order to estimate the effect of interface drops on our data we compared the raw data with a "filtered" set. In this data set we removed the losses reported by zing or ping but not recorded by router interfaces during each 30 second interval. While it could be the case that loss episodes occurred at the router which were measured by the active probe and not by the router counters, we do not consider this to be an significant possibility. Filtered results of the same path are shown in Figure 4. After filtering the raw data, we notice that the active probes appear to miss many of the loss events recorded by the router which corresponds to our intuition about probe rates. However, the time scale over which the active probes are taken still effectively overestimates the loss rate during intervals of loss.

In analysis to follow, we compare the SNMP data with the raw zing and ping traces. Our reasoning for continuing analysis with the raw data is that many active probing studies have suffered the same restriction of inability to use packet filters to measure interface or operating system buffer drops. A typical reason for disallowing use of packet filters is that administrative ("root") privileges are normally required in order to perform this activity. Active measurements are often taken at remote sites which have volunteered to assist in a particular study.

Rows 1 and 2 of Table II help further quantify the effect due to interface loss. While the overall loss rate is very low for both raw and filtered data sets, the loss rate of the filtered data is an order-of-magnitude lower.

B. Loss Rates

We now consider loss rate distributions for each measurement, including all intervals regardless whether loss has occurred in a given interval or not. Figure 5(a) shows log-log complementary distribution functions for all 56 paths for the SNMP data. Table I also gives minimum and maximum mean loss rates, and the minimum and maximum standard deviations for all SNMP path traces.

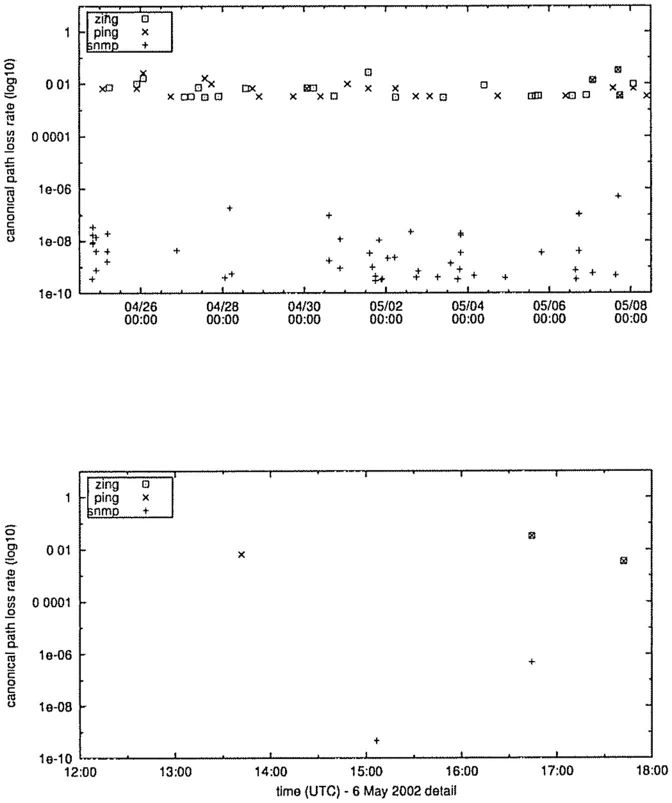

Fig. 3. Qualitative Comparison - Raw Data

Figure 5(b) shows the log-log complementary distribution of loss rates for SNMP, zing and ping over the canonical path. As stated above, the raw zing and ping data sets are used. It is interesting to note that SNMP measured a higher loss rate than the probes. Referring to Figure 3, this SNMP loss measure lies in the early morning hours of 25 April. We previously described routing changes that occurred during this time. While the loss rate reported by the interfaces is quite high, it occurs in such a way that zing and ping are not able to detect the event.

Next, we calculated the correlation coefficients for each path between SNMP and each of the probe traces, and constructed corresponding cumulative distribution functions. These functions are shown in Figure 6. Note that for both the raw and filtered traces, correlation is poor (zero or only slightly above zero). Another feature to note is that neither zing nor ping have distinct correlational advantages.

TABLE I

SUMMARY STATISTICS - SNMP FOR ALL PATHS

	min μ	**max μ**	**min σ**	**max σ**
Loss Rate	0.0	4.6×10^{-5}	0.0	5.9×10^{-3}
Duration of Loss-Free Periods (sec)	141	13439	134	23084
Loss Rate in Loss Periods	1.1×10^{-9}	2.4×10^{-2}	3.8×10^{-12}	1.3×10^{-1}
Duration of Change-Free Periods (sec)	5191	1209600	0	667242

TABLE II

SUMMARY STATISTICS FOR CANONICAL PATH

	Data Set	μ	σ
Loss Rate (raw)	SNMP	1.2×10^{-6}	2.4×10^{-4}
	ZING	4.9×10^{-6}	2.7×10^{-4}
	PING	4.9×10^{-6}	2.7×10^{-4}
Loss Rate (filtered)	SNMP	1.2×10^{-6}	2.4×10^{-4}
	ZING	8.2×10^{-7}	1.6×10^{-4}
	PING	8.3×10^{-7}	1.7×10^{-4}
Loss-Free Periods (raw)	SNMP	9.2×10^{2}	2.0×10^{6}
	ZING	1.5×10^{3}	1.3×10^{6}
	PING	1.6×10^{3}	1.4×10^{6}
Loss Periods (raw)	SNMP	9.8×10^{-4}	5.8×10^{-5}
	ZING	7.9×10^{-3}	5.9×10^{-5}
	PING	8.2×10^{-3}	5.8×10^{-5}
Change-Free Period Duration (raw)	SNMP	4.0×10^{5}	4.5×10^{10}
	ZING	1.3×10^{4}	8.5×10^{8}
	PING	1.2×10^{4}	7.8×10^{8}
Number of Change-Free Periods (raw)	SNMP	3	
	ZING	91	
	PING	97	

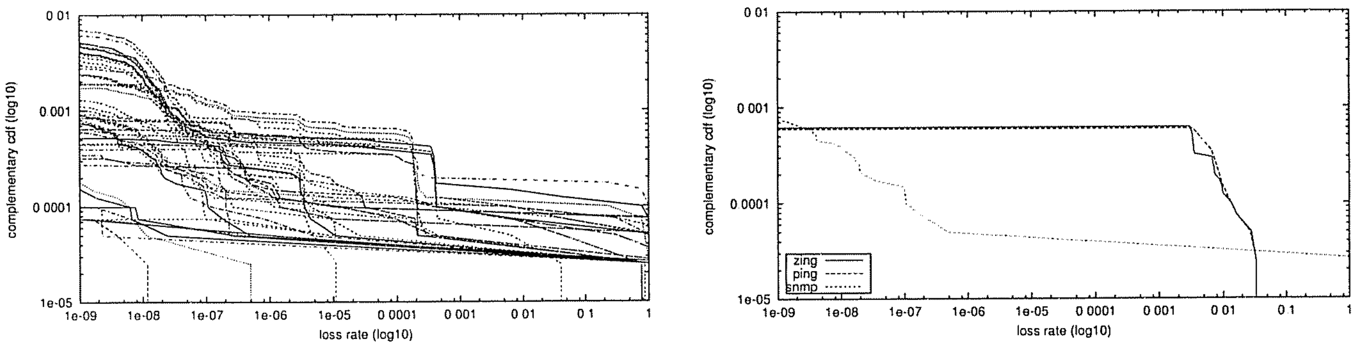

(a) All Paths for SNMP

(b) Canonical Path

Fig. 5. Loss Rates

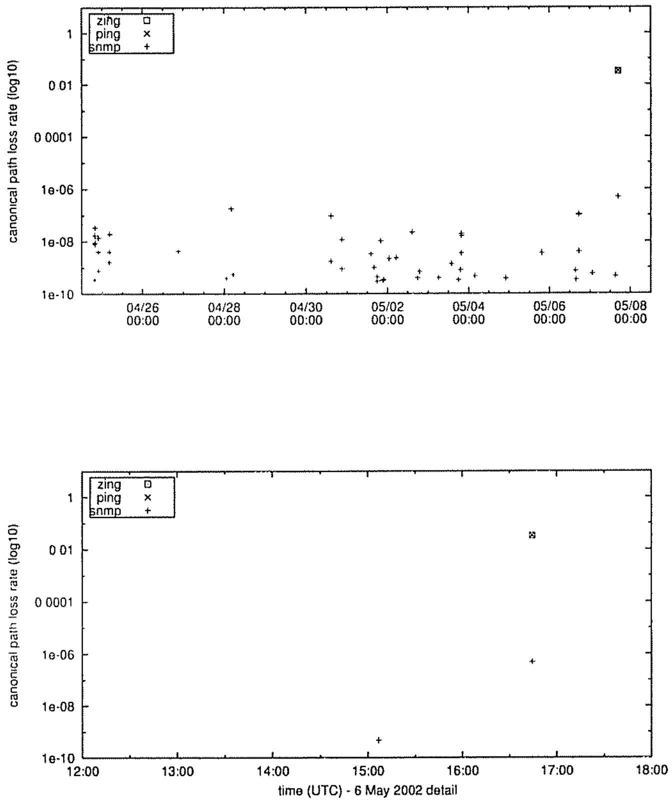

Fig. 4. Qualitative Comparison - Filtered Data

C. Loss-Free Periods

In this area of analysis, we compare the distribution of loss-free period durations. A loss-free period is defined as the maximum number of consecutive 30 second bins which do not measure any loss. Another way of understanding this measure is to think in terms of loss event interarrival times.

Figure 7(a) shows the cumulative distribution function of loss-free periods for all paths measured by the router interface counters. It is clear that for some paths, losses occur frequently, and apparently quite regularly. For other paths, however, losses occur infrequently. This wide variety of loss interarrival times poses a challenge for determining how to best conduct active network measurements for loss.

Figure 7(b) plots the cumulative distribution functions of loss-free periods for each measurement method along the canonical path. The key feature to notice is that losses are more closely spaced in time as measured by the routers than by `zing` or `ping`.

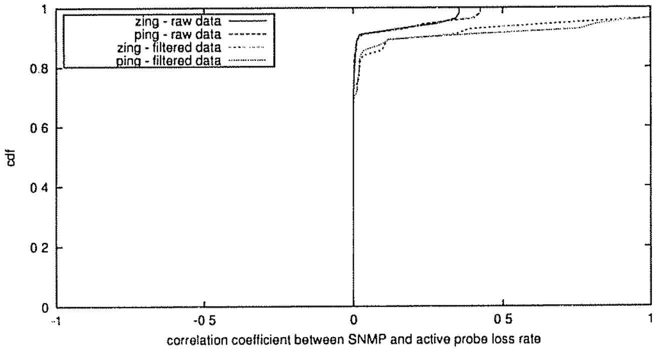

Fig. 6. Distribution of Loss Rate Correlation Coefficient

Figure 8 gives the cumulative distribution functions of the χ^2 goodness-of-fit statistic for `zing` and SNMP, and for `ping` and SNMP. We also plot vertical lines indicating the 95% and 1% acceptance levels[5]. Note that the x-axis is plotted on a log scale. It is immediately clear that even at the 1% acceptance level, we must find that `zing` and `ping` are not good fits to the distribution of loss-free periods measured by SNMP.

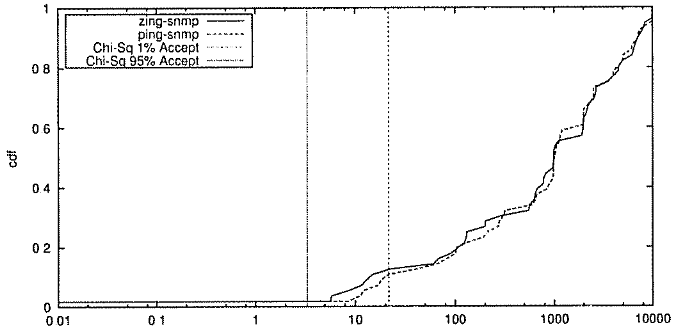

Fig. 8. Loss-Free Periods - χ^2 Distribution

D. Loss Periods

We now assess the loss rates measured during the 30 second intervals over which packet loss is detected. Figure 9(a) plots cumulative distribution functions of loss rates during these loss periods for all paths measured by the SNMP traces. From the figure, we see that most loss measured by router interfaces is on the order of 1×10^{-8} or

[5] The χ^2 goodness-of-fit test is a hypothesis testing procedure. A fit hypothesis is accepted at a given confidence level if the χ^2 metric is less than the χ^2 distribution value with specified degrees of freedom

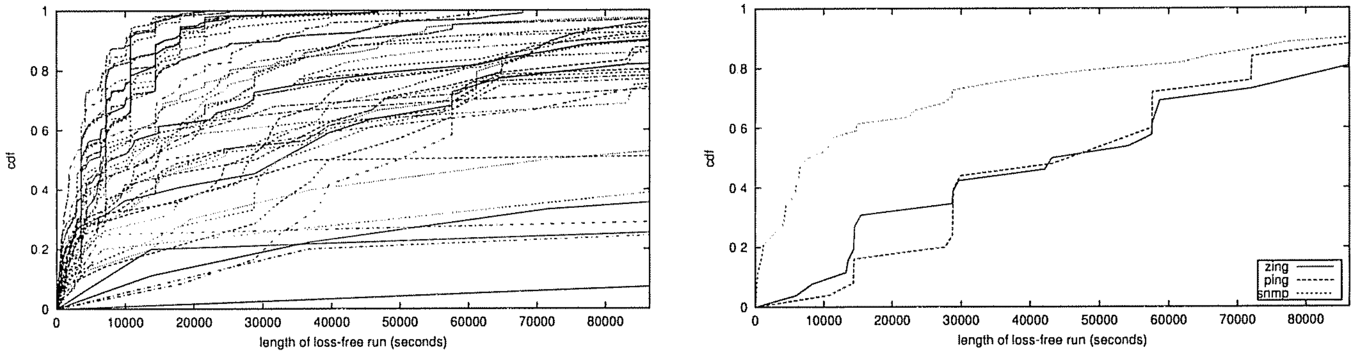

(a) All Paths for SNMP

(b) Canonical Path

Fig. 7. Loss-Free Periods

less. We also note that a non-neglible percentage of the traces experienced much higher loss rates, even surpassing 0.01.

For the canonical path, Figure 9(b) shows that zing and ping experience vastly different loss rates than are measured by SNMP. The lower bound on loss rate measurable by the probes because of the sampling rate is obvious from the curves, and for this path zing measures lower loss rates than ping during these periods of loss.

We do not plot the results for the χ^2 test on loss period distributions. The reason is that the test falsely indicates that the loss periods measured by zing and by ping are good fits to the SNMP measurement. The reason for this is simple if we consider the effect of binning when computing the statistic: if we use 10 bins and the maximum loss rate measured is more than 1%, almost all of the measured values for all three data sets will fall in the lowest bins (recall Figure 3 and Table II), thus giving a (false) positive indication for goodness of fit with high confidence.

E. Change Free Periods

Finally, we compare how loss constancy is measured by probes versus the loss constancy measured along a path of router interfaces. We use the notion of *change free periods*, as described in [6]. In our study, we used the bootstrapping method for generating change points. As noted in [6], this method is conservative in the sense that it is more likely to produce false change points than to miss them. An area for future work would be to explore other methods for finding change points in our SNMP data.

Figure 10(a) shows cumulative distribution functions of the duration of change free periods for all paths measured using router SNMP traces. Analogous to Figure 7(a), it indicates that there is a wide range of durations over which path loss is steady. There are a number of paths for which conditions do not change for days, and there are also a number of paths on which loss conditions change with higher frequency.

Figure 10(b), showing cumulative distribution functions of the duration of change free periods for the canonical path, indicates that zing and ping both experience high proportions of short durations of steady loss rates. The view of constancy seen through the router interfaces for this particular path, however, is that change free periods are quite long.

Figure 11 plots the cumulative distribution function of the χ^2 statistic for comparing change free periods seen by zing and SNMP and ping and SNMP across all paths. Vertical lines are plotted indicating the 95% and 1% acceptance levels. Clearly, neither zing nor ping are good fits to the SNMP data.

Finally, we plot the cumulative distribution function for the number of change free periods for all the SNMP traces, all the zing traces, and all the ping traces in Figure 12. Immediately, we notice that there are many fewer change points measured by the router interfaces across all paths. Comparing Figure 12 with Figure 10(a), we infer that there are fewer numbers of change free regions of short duration which are recorded by the routers, and more, rather long regions of constancy. Our data indicates that zing and ping, in contrast to the routers, tend to measure more change free periods of short durations.

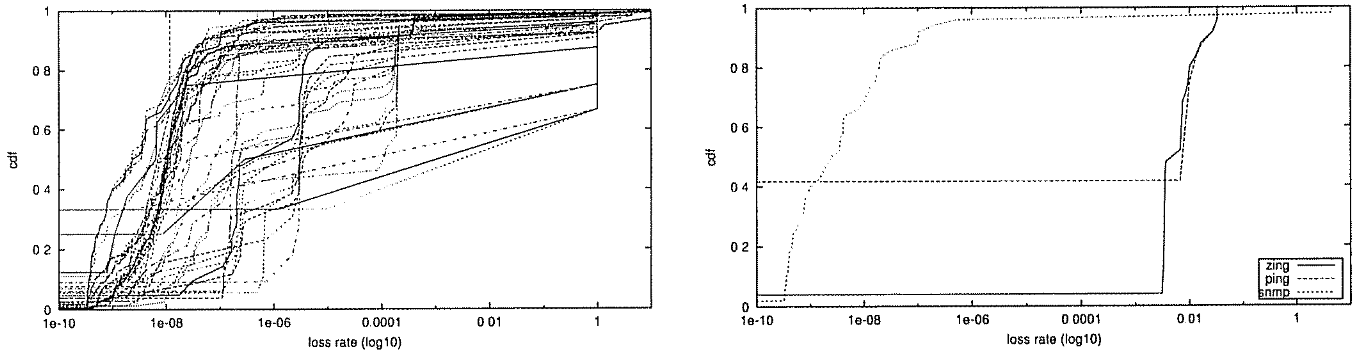

(a) All Paths for SNMP

(b) Canonical Path

Fig. 9. Loss Periods

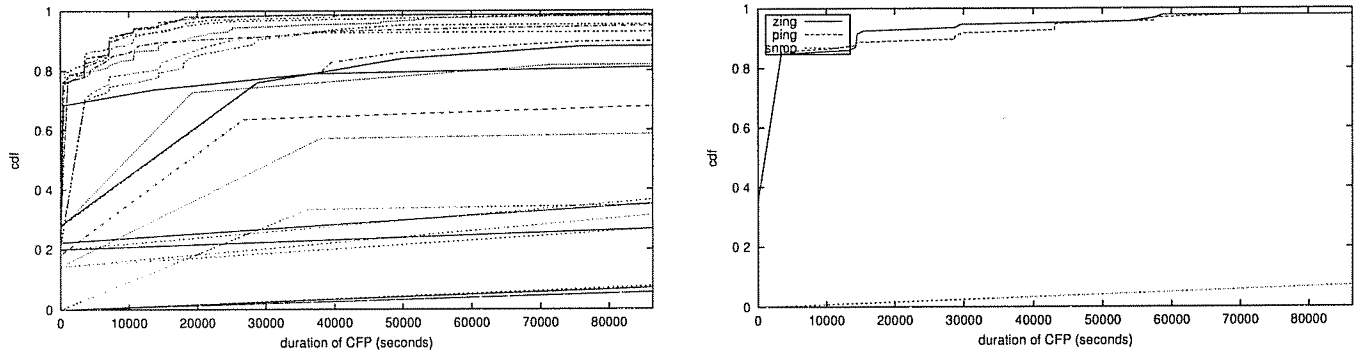

(a) All Paths for SNMP

(b) Canonical Path

Fig. 10. Change Free Periods

VI. Conclusions and Future Work

We present an evaluation of the precision of active probes for loss. Our study is based on measurements taken over a two week period using the Surveyor infrastructure and the Internet2/Abilene backbone. We gather packet loss data from Abilene backbone routers via SNMP in 30 second intervals. We use `zing` to actively probe for loss at a rate of 10Hz and then aggregate these values into 30 second intervals in order to assess how well the two measures correlate with each other.

We define precision as the degree of agreement between the active probe data and the SNMP-based data which we consider to be the *true* loss value. We assess precision by comparing loss rates along all paths in the full mesh in our measurement infrastructure. We consider the degree of correlation between loss rate time series and the degree of agreement between distributions of loss characteristics including lengths of loss free periods, loss rates during loss periods, and the duration of change free regions.

Our results show that loss rates as measured by active probes are not well correlated with those measured by SNMP. We also show that the distributions of values for loss free periods, loss rates during loss periods, and the duration of change free regions as seen by active probes do not align closely with the distributions of the same values as seen by SNMP. While we do not assess differences in

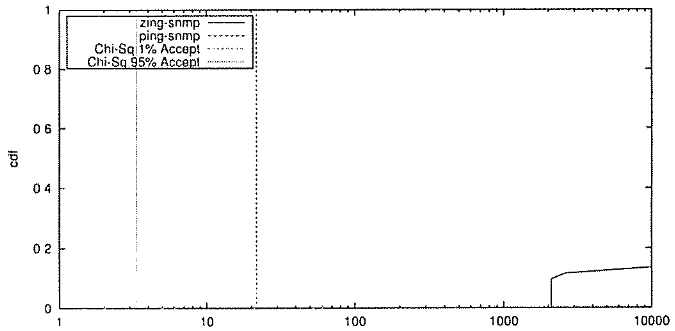

Fig. 11. Duration of Change Free Periods - χ^2 Distribution

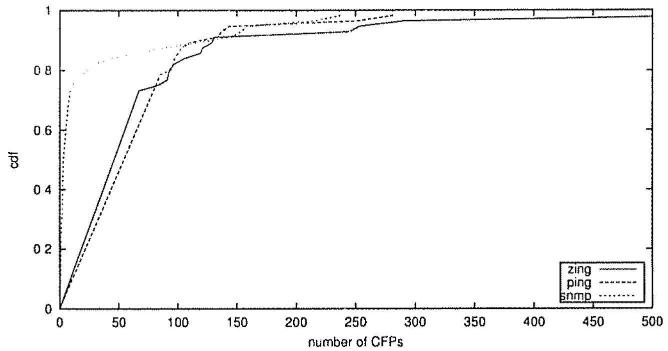

Fig. 12. Number of Change Free Periods

sampling rates in our work, it is clear that even the fairly fast probe rate of 10Hz is inadequate for getting a good picture of the true status of loss conditions in wide area high performance networks.

We also evaluate the differences between loss rates as measured by the Poisson modulated zing tool and the more simple ping utility which sends out probes at constant intervals. At least for the low loss rates seen in our measurement infrastructure, ping provides qualitatively the same level of precision as zing.

Our work has implications in a number of areas including network operations, loss modeling, loss probing and loss characterizations. Our next steps in this work will be to investigate new methods for probing that are both lightweight and provide more precise measurements of loss.

Acknowledgements

The authors would like to thank Matt Zekauskas, Guy Almes, Chris Robb and the people at Internet2/Abilene for their help in the data collection effort. Without their hard efforts, this study would not have been possible. The authors would also like to thank Larry Landweber for fruitful discussions.

References

[1] J. Hoe, "Improving the start-up behavior of a congestion control scheme for TCP," in *Proceedings of ACM SIGCOMM '96*, Palo Alto, CA, August 1996.

[2] M. Mathis, J. Mahdavi, S. Floyd, and A. Romanow, "TCP selective acknowledgement options," IETF RFC 2018, 1996.

[3] J. Case, M. Fedor, M. Schoffstall, and J. Davin, "A simple network management protocol (SNMP)," IETF RFC 11157, 1990.

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

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

[6] Y. Zhang, N. Duffield, V. Paxson, and S. Shenker, "On the constancy of Internet path properties," in *Proceedings of ACM SIGCOMM Internet Measurement Workshop '01*, San Francisco, November 2001.

[7] A. Pasztor and D. Veitch, "A precision infrastructure for active probing," in *PAM2001, Workshop on Passive and Active Networking*, Amsterdam, Holland, April 2001.

[8] J. Bolot, "End-to-end packet delay and loss behavior in the Internet," in *Proceedings of ACM SIGCOMM '93*, San Francisco, Setpember 1993.

[9] V. Paxson, "End-to-end Internet packet dynamics," in *Proceedings of ACM SIGCOMM '97*, Cannes, France, September 1997.

[10] M. Yajnik, S. Moon, J. Kurose, and D. Towsley, "Measurement and modeling of temporal dependence in packet loss," in *Proceedings of IEEE INFOCOM '99*, New York, NY, March 1999.

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

[12] NLANR Acitve Measurement Program - AMP, ," http://moat.nlanr.net/AMP.

[13] W. Matthews and L. Cottrell, "The PINGer Project: Active Internet Performance Monitoring for the HENP Community," *IEEE Communications Magazine*, May 2000.

[14] UCB/LBNL/VINT Network Simulator - ns (version 2), ," http://www.isi.edu/nsnam/ns/, 2000.

[15] A. Arpaci-Dusseau and R. Arpaci-Dusseau, "Information and control in gray-box systems," in *Proceedings of the Eighteenth Symposium on Operating Systems Principles (SOSP'18)*, Banff, Canada, October 2001.

[16] N. Duffield, J. Horowitz, F. Lo Presti, and D. Towsley, "Network delay tomography from end-to-end unicast measurements," in *Proceedings of 2001 International Workshop on Digital Communications*, Taormina, Italy, September 2001.