

Comparing Probe- and Router-Based Packet-Loss Measurement

Empirical analysis of Internet traffic characteristics should not be biased by the measurement methodology used to gather data. This article compares probe- (active) and router-based (passive) methods for measuring packet loss both in the laboratory and in a wide-area network. The laboratory case study demonstrates the accuracy of passive Simple Network Measurement Protocol (SNMP) measurements at low loss rates; the wide-area experiments show that active-probe loss-rate measurements don't correlate with those measured by SNMP from routers in a live network. This case study's findings also reveal that common methods for active probing for packet loss suffer from high variance and from the effects of end-host interface loss.

**Paul Barford
and Joel Sommers**
University of Wisconsin, Madison

Packet loss due to congestion is a fundamental problem in wide-area packet-switched networks. Researchers have expended much effort in characterizing this phenomenon and designing protocols that operate effectively in lossy environments, but the Internet's dynamic nature requires continued empirical evaluation of packet-loss behavior.

We can divide the methods for measuring packet loss into two categories. The first uses passive monitors that are either attached to network links or available from network nodes. A standard means of passive monitoring is to use the set of Management Information

Base (MIB) counters available on network nodes via the Simple Network Management Protocol (SNMP). These counters track packet losses due to congestion in router subsystems, with the benefit being that they capture many important details about local traffic behavior. Unfortunately, the cost for this information can be high in terms of data-storage requirements, and SNMP access across administrative domains is usually impossible.

The second means for measuring packet loss is through active end-to-end probing with a tool such as the ubiquitous ping utility. Active probe tools

such as `ping` send a series of packets aimed at a target system and then measure the returned response packets; the sender can use sequence numbers to track lost packets. The benefits of active probes are twofold: they can run virtually anywhere in the network, and they give an end-to-end perspective of network conditions. However, like many sampling tools, the probes' discrete nature limits resolution. Increasing the probing rate can increase resolution, but the probes themselves can skew the results if the frequency is too high.

Our objective here is to address the question: "Do probe- and router-based measurements of packet loss correlate?" The work we describe in this article has three main implications. First, it exposes the limitations of probe-based loss measures in low-loss environments, the implication being that new probe-based methods might be necessary to get a more accurate picture. Second, our results demonstrate the accuracy of SNMP-reported loss measurements, suggesting that SNMP is an attractive alternative for measuring loss in low-loss environments. Finally, our study suggests that characterizations and models for packet loss based on active measurements might need to be reevaluated using data from new probing methodologies or passive measurements.

Setting up the Experiments

To get started, we examined the accuracy of SNMP loss measurements through a series of laboratory experiments. Using precise measurement systems, we found that SNMP accurately reported loss. We then took router-based SNMP measurements of packet loss over three collection periods at all backbone routers in Internet2 (www.internet2.edu) and aggregated data along each path to get end-to-end perspectives on loss behavior. Simultaneously, we took one-way probe-based measurements of packet loss using the `zing` utility between GPS-enabled nodes directly connected to each Internet2 backbone router. `zing` sends probe packets at exponentially modulated intervals, which should provide unbiased, time-averaged data for loss conditions along an end-to-end path. For our three measurement periods, we set the average probe rate at 10 Hz, 20 Hz, and 100 Hz, respectively, and then aggregated the measured loss rates to compare with SNMP data.

We evaluated the degree of agreement between the probe- and router-based measurements by comparing the correlation coefficients for the loss-

rate time series on each end-to-end path. The results showed little correlation between probe- and router-based data sets. Next, we compared distributional characteristics of loss measurements for different loss properties, including lengths of loss-free periods, loss rates during lossy periods, and loss constancy (periods without a trend in loss rate).¹ In each case, we found a low level of agreement between the distributions, so we concluded that probe- and router-based loss measures can give quite different perspectives.

There are several feasible explanations for the lack of agreement between the data sets, one possibility being that artifacts in our measurements, such as interface loss, bias the results. (We attributed loss measured by active probes to host interfaces when we didn't observe simultaneous loss in the SNMP data.) We found these losses to be relatively rare, and censoring them from the data didn't improve the correlation between data sets. The most plausible explanation for the overall lack of correlation is that the sampling rates we used in probe-based measurements were too coarse to let us accurately measure typical loss episodes. Although we used three different probe rates, correlation didn't significantly improve with faster probe rates: the overall SNMP loss-rate measurements were extremely low and would have required sampling for very long periods of time before close correlations could have been established.

Statistical Issues in Probe-Based Measurements

A standard statistical technique for getting an unbiased estimate of a random process's average state is to sample at exponentially distributed intervals. An extension of this approach led to the well-known Poisson Arrivals See Time Averages (PASTA) theorem,² which states that exponentially distributed arrivals at a queue will "see" the system's average state. Because this theorem expresses an asymptotic result, it must be considered carefully in practice.

Let X_t be a binary process with states describing whether a packet is lost due to congestion (1) or not (0). We are interested in estimating loss rate p as the probability of loss due to congestion, that is, $p = \Pr(X_t = 1)$. Sampling n times at Poisson intervals, we get \bar{X}_n , the average of the n samples; thus, the expected value $E(\bar{X}_n) = \Pr(X_t = 1) = p$. As the number of samples $n \rightarrow \infty$, $\bar{X}_n \rightarrow p$, but note that this estimate might have a very large variance — namely, $\text{Var}(\bar{X}_n) \approx p/n$. For the standard devia-

Related Work in Packet-Loss Behavior

Many researchers have studied packet-loss behavior over the Internet. Bolot¹ and Paxson² used active probe measurements to establish much of the baseline for understanding packet-loss characteristics in wide-area networks, including correlation structures on fine time scales and typical loss rates. Another probe-based method for measuring packet loss is to use tomography (coordinated end-to-end measurements to infer link-level loss rates).^{3,4} Yajnik and colleagues evaluated correlation structure on longer time scales and developed Markov models for temporal dependence structures.⁵ Zhang and colleagues assessed three different aspects of constancy in loss rates based on active probing in Internet^{2,6}. The future promises interesting results from several ongoing passive and active measurement projects (see <http://moat.nlanr.net/AMP> for a look at the

National Laboratory for Applied Network Research's recent work with the Active Measurement Program).^{7,8} While data from those projects will provide an important empirical perspective on packet-loss behavior, results from this article suggest that data from active measurements be considered carefully.

References

1. J. Bolot, "End-to-End Packet Delay and Loss Behavior in the Internet," Proc. ACM Special Interest Group on Data Communications (SIGCOMM '93), ACM Press, 1993, pp. 289–298.
2. V. Paxson, "End-to-End Internet Packet Dynamics," Proc. ACM Special Interest Group on Data Communications (SIGCOMM '97), ACM Press, 1997, pp. 139–152.
3. M. Coates and R. Nowak, "Network Loss Inference Using Unicast End-to-End Measurement," Proc. Int'l Teletraffic Congress Seminar on IP Traffic, Measurement, and Modeling, //publisher?/, 2000, pp. 281–289.
4. N. Duffield et al., "Inferring Link Loss Using Striped Unicast Probes," Proc. IEEE Conf. on Computer Comm. (INFOCOM '01), IEEE CS Press, 2001, //pages cited?/.
5. M. Yajnik et al., "Measurement and Modeling of Temporal Dependence in Packet Loss," Proc. IEEE Conf. on Computer Comm. (INFOCOM '99), IEEE CS Press, 1999, pp. 915–923.
6. Y. Zhang et al., "On the Constancy of Internet Path Properties," Proc. ACM Special Interest Group on Data Communications (SIGCOMM), ACM Press, 2001, pp. 197–212.
7. C. Fraleigh et al., "Packet-Level Traffic Measurements from the Sprint IP Backbone," IEEE Network, vol. 17, no. 6, 2003, pp. 6–16.
8. W. Matthews and L. Cottrell, "The PINGer Project: Active Internet Performance Monitoring for the HENP Community," IEEE Comm., vol. 38, May 2000, pp. 130–136.

tion to be approximately $0.1 p$, we need $n \approx 10/p$; for average loss rates on the order of 10^{-4} , we therefore need $n \approx 10^5$ samples.

This simple analysis has several important implications. For loss rates of 10^{-7} (not uncommon in a highly engineered network such as Internet2), we must send probes for nearly two weeks to get an accurate projection of loss rates, even if we use a relatively fast probe rate of 100 Hz. Furthermore, if we simply decide to increase our probe rate to reduce the time required, we're inevitably forced to make trade-offs because of increasing bandwidth consumption from probe traffic and the potential for skewing the loss measurements.

Data Collection

We collected our wide-area measurement data in the Internet2/Abilene backbone, which researchers frequently use in network measurement and characterization projects. We sent active probes across a full mesh of end hosts directly connected to the routers; we then sent SNMP queries every 30 seconds to collect packet count and loss data from router interfaces (we selected this period as a compromise between increased load on routers and sufficiently detailed data). The SNMP measurements included packet drops due to full queues, packet corruption, and interface errors. In total, we took

probe-based measurements on 56 distinct paths and router-based measurements from 30 router interfaces. Further details on the measurement infrastructure and methods appear elsewhere.³

We collected our active measurement data by using the *zing* utility.⁴ We sent 256-byte UDP probes at exponentially distributed intervals with means of 100 ms, 50 ms, and 10 ms (for 10 Hz, 20 Hz, and 100 Hz probes, respectively). The literature does not treat the problem of probe-size choice for loss measurements, so we selected our probe size to be consistent with another study¹ and to ensure modest bandwidth consumption. Because Cisco GSR12000 series routers in Internet2 use *buffer carving* (to queue by packets, not bytes),⁵ probe-packet size had no impact on loss measurements for the buffer configurations on all routers during our tests. In our analysis, we called these "*zing* traces." In parallel with the *zing* probes, we sent 256-byte UDP probes with uniform spacing, the methodology of which is essentially the same as the *ping* tool (although in our case, probe packets flow in only one direction). We sent probes continuously over each data-collection period, with each type of probe following the same forwarding path as normal packet traffic through the routers. To compare *zing/ping* traces with the SNMP data, we aggregated the probe traces in 30-second inter-

vals, which made our analysis more conservative when reporting loss events. Even if these events were measured both by SNMP and an active probe in the same interval, the active probe appeared to detect the congestion-loss event.

We also took `traceroute` measurements across the full mesh of end hosts every 10 minutes to determine the sets of router interfaces encountered along each end-to-end path. We used the physical interface layout information (openly available from the Abilene network operations center (www.abilene.iu.edu) to complete path details; we found the routes to be extremely stable during the course of our study.

We collected data over three periods: 24 April 2002 to 8 May 2002 (10-Hz probes), 24 July 2002 to 31 July 2002 (20-Hz probes), and 8 August 2002 to 9 August 2002 (100-Hz probes). Due to the immense amount of data generated from the 100-Hz measurements, we collected that data for just two days. Link utilizations over our study period averaged 12 percent, 8 percent, and 7 percent, with standard deviations of 11 percent, 5 percent, and 4 percent for each measurement period, respectively.

Using the router-based measurements, we calculated the loss rates for paths with multiple hops using *union-of-loss probabilities*, which means taking the product of loss rates at each hop. Specifically, we calculated loss-rate L for a multi-hop path p of length n interfaces 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 interface i , and t_i is the sum of packets transmitted and packets lost at the same interface during the same period.

This calculation assumes *independence-of-loss events* (meaning loss events unrelated) at each hop in the path. We feel this is reasonable in highly engineered networks; it's unlikely that a single flow or even a small group of flows can cause correlated congestion losses at two points on a path. We calculated correlation coefficients for both loss periods and loss rates on all multihop paths, and found all coefficients very tightly bunched around zero. Although this result doesn't prove independence, it's consistent with our assumption of it.

Laboratory Evaluation of SNMP Loss Counters

To effectively compare probe- and router-based packet-loss measurements, we experimentally evaluated the packet-loss counters on the Cisco GSR router. We used a Spirent AX4000 to generate traf-

fic on an OC-12 interface, which terminated at a Cisco GSR; this traffic was then routed back to the AX4000 over an OC-3, forming the bottleneck over which packet loss was generated. In each direction, we used optical splitters to Endace DAG3.5 capture cards (the same cards used in SprintLab's IPMON environment⁶), which have precise packet-measurement capabilities. By tuning the packet-emission parameters at the AX4000, we generated varying degrees of packet loss at the router.

Our three experiments, each of which lasted for more than two hours, consisted of loss regimes created with the AX4000 to generate approximately 0.1 percent, 0.01 percent, and 0.001 percent packet loss, respectively. We uniformly used 256-byte packets and generated packet bursts of varying sizes such that the combination of the average inter-burst time and the average burst length created the desired loss rate. The correlation coefficients between SNMP and DAG traces were 0.87, 0.94, and 0.96, thus demonstrating precision in the SNMP loss measurements. Lower correlation at higher loss rates is primarily an artifact of edge conditions due to the 30-second sampling interval.

Loss-Rate Comparison

We compared results for all loss measures (SNMP, `zing`, and `ping`, at 20 Hz) along an arbitrarily chosen canonical path (from Indianapolis, Indiana, to Los Angeles, California) representative of many paths in our study. A detailed description of all our loss measurements appears elsewhere.³

We quantified the degree of distributional agreement between `zing/ping` and router counters using the χ^2 goodness-of-fit test parameterized to conservatively favor finding agreement between two distributions. χ^2 is a parametric goodness-of-fit hypothesis test that is extremely robust to underlying distributional characteristics. Although we could use other comparison methods, our objective was to make a straightforward quantitative assessment.

Loss-Rate Time Series

Figure 1a (next page) shows time-series graphs of loss measurements for the canonical path. Clearly, active probes largely overestimate loss rate compared to router interfaces, but it's important to note the lower bound of the probe-measured rate when we group measurements into sample intervals for time-series analysis (in which the lower bound is defined as measuring a single loss event within a specified time interval). This bound is a

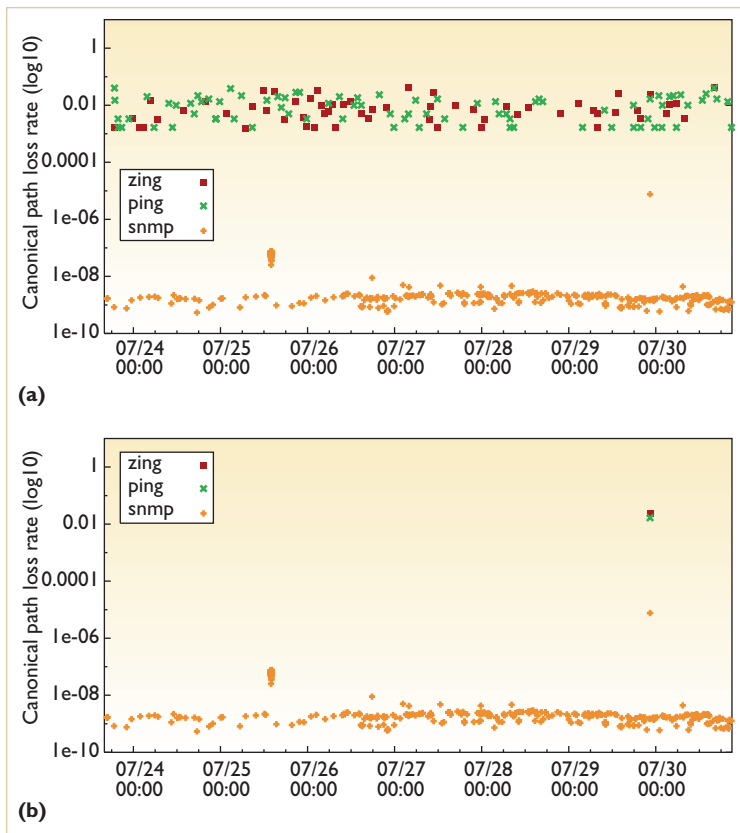


Figure 1. Loss rates. We compared loss rates on the Indianapolis, Indiana, to Los Angeles, California, canonical path for a 20-Hz probe rate. Simple Network Measurement Protocol (SNMP) measurements were identical for each data set, but filtered data sets contained only active probe events when we measured a loss event at a router in the corresponding sample interval.

function of the considered probe rate and the time interval. With a mean probe rate of 20 Hz, we sent an average of 600 packets per 30-second sample interval, which set the effective lower bound on loss at a rate of 0.15 percent. However, the effective lower bound for SNMP was much lower. Assuming an average packet size of 300 bytes, the minimum loss rate over a 30-second period for an OC-48 (2.4×10^9 bps) was roughly 3.0×10^{-8} . We noted minimum SNMP-measured loss rates on the order of 10^{-9} , which is consistent with the average packet sizes computed with other MIB variables.

To estimate the effect of transmission loss due to network interface drops, we compared our raw data with a filtered set. We created this set from raw data by retaining all SNMP measurements and removing any losses reported by zing/ping that SNMP didn't report during each 30-second sample interval. Figure 1b shows the filtered results of the same path. Although congestion loss could have

occurred at the router measured by the active probe instead of the router counters, we didn't consider this to be a significant possibility based on our previous experiments. From filtered data, we see that probes appear to miss many of the loss events the router recorded. Rows 1 and 2 in Table 1 further quantify the effects of transmission loss due to host interface drops. Although the overall loss rate was low for both raw and filtered data sets, the filtered data's loss rate was often an order of magnitude lower – occasionally even zero. This highlights the shortcomings of active probing for loss even if host interface drops could be completely avoided.

Next, we calculated the time-series correlation coefficients for each path between router-based measures and each of the probe traces. Correlation coefficients were very low for both the raw and filtered traces: more than 50 percent of paths had zero correlation.

Loss-Free Periods

A loss-free period is defined as the number of consecutive 30-second sample intervals during which no loss is measured. Evaluation of loss-free periods as measured by SNMP over all paths showed a wide range of values; Figure 2 shows the cumulative distributions of loss-free periods for each measurement method along the canonical path. Router-based measures clearly show loss events more closely spaced in time than probe-based measures. χ_2 values rejected the fit hypothesis even at the 1 percent acceptance level, indicating that the probe-based marginal distributions (marginals) of loss-free periods are not good approximations of router-based marginals.

Loss Periods

We assessed the loss rates measured only during the 30-second intervals over which we detected packet loss, but we observed a wide range of SNMP-measured loss rates over all paths. For the canonical path, Figure 3 (next page) shows that probes experienced vastly different loss rates than router-based measures did. The lower bound on the probe-measurable loss rate is obvious from the curves; for this path, zing/ping measured similar loss rates. Results from χ_2 tests indicated that the probe-measured loss periods reasonably fit the SNMP data. This good fit is based on our conservative choice of χ_2 parameters and breaks down quickly when we use more degrees of freedom.

Table 1. Summary statistics for a canonical path from Indianapolis, Indiana, to Los Angeles, California.

	Data set	10 Hz		20 Hz		100 Hz	
		μ	σ	μ	σ	μ	σ
Loss rate (raw)	SNMP	4.1×10^{-8}	2.4×10^{-6}	4.2×10^{-10}	2.9×10^{-8}	5.2×10^{-10}	2.4×10^{-9}
	ZING	5.9×10^{-5}	3.5×10^{-3}	2.8×10^{-5}	6.7×10^{-4}	9.9×10^{-6}	1.2×10^{-4}
	PING	5.9×10^{-5}	3.8×10^{-3}	3.4×10^{-5}	7.5×10^{-4}	1.1×10^{-5}	1.3×10^{-5}
Loss rate (filtered)	SNMP	4.1×10^{-8}	2.4×10^{-6}	4.2×10^{-10}	2.9×10^{-8}	5.2×10^{-10}	2.4×10^{-9}
	ZING	0	0	8.2×10^{-7}	1.2×10^{-4}	0	0
	PING	0	0	3.0×10^{-6}	2.1×10^{-4}	3.0×10^{-6}	7.2×10^{-5}
Loss-free periods (raw)	SNMP	1.4×10^2	1.8×10^5	1.9×10^1	1.5×10^3	6.2×10^0	4.4×10^1
	ZING	5.0×10^2	6.4×10^5	3.2×10^2	6.2×10^4	1.4×10^2	2.9×10^3
	PING	4.6×10^2	6.5×10^5	2.6×10^2	4.8×10^4	1.5×10^2	5.0×10^3
Loss periods (raw)	SNMP	3.6×10^{-3}	3.4×10^{-3}	1.6×10^{-8}	6.7×10^{-14}	7.6×10^{-9}	1.4×10^{-16}
	ZING	2.8×10^{-2}	5.3×10^{-3}	9.0×10^{-3}	6.5×10^{-5}	1.4×10^{-3}	2.0×10^{-7}
	PING	2.6×10^{-2}	5.9×10^{-3}	9.0×10^{-3}	6.6×10^{-5}	1.7×10^{-3}	3.2×10^{-8}
Change-free period (raw)	SNMP	2.4×10^5	2.7×10^{11}	1.2×10^6	0	8.6×10^4	0
	ZING	1.2×10^6	0	2.4×10^3	3.1×10^7	1.2×10^3	4.6×10^6
	PING	1.2×10^6	0	4.1×10^3	1.3×10^9	1.3×10^3	7.1×10^6
Number of change-free periods (raw)	SNMP	5		1		1	
	ZING	1		511		75	
	PING	1		299		65	

Change-Free Periods

Finally, we compared measures of loss constancy. A previous work defined a time series as “a series of piece-wise steady regions delineated by change points,”¹ that is, periods without a statistically discernable trend in loss rate. Thus, the task of identifying change-free periods (CFPs) reduces to identifying change points. We used the bootstrapping method to identify change points.¹

CFP duration distributions for all paths using router-based measures showed a wide range of durations, including several paths for which conditions didn’t change for days and several others for which loss conditions changed with much higher frequency. Figure 4 (next page) shows the cumulative distributions of CFP duration for the canonical path, indicating that `zing` and `ping` both experienced high proportions of short durations of steady loss rates. Seen through the router interfaces for this particular path, however, the loss rate remained constant over the entire collection period. χ_2 statistics rejected the fit hypothesis even at the 1 percent acceptance level, indicating that probe-based CFP marginals are not good approximations for router-based CFP marginals.

Future Work

Our results highlight both the need for great care

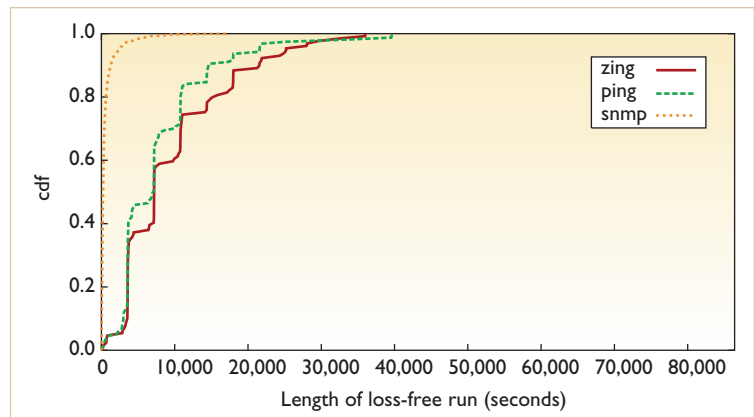


Figure 2. Loss-free periods on the canonical path. During the 20-Hz probe period, the cumulative distributions of consecutive 30-second sample intervals for all paths show that router-based measures display loss events more closely spaced in time than probe-based measures. Thus, router-based measures indicate loss events are much more common.

in the use of active probes for loss characterization and the potential for the use of router-based measures. Our results also suggest that prior published results based on active probes require scrutiny in their interpretation, and that future active probe tools that sample network characteristics must be designed with traffic variability in mind. Our next




Figure 3. Loss rates during loss intervals on the canonical path. During the 20-Hz probe period, the cumulative distributions of loss rates for sample intervals for all paths using probe-based measurements differ vastly. Router-based measures indicate loss rates during loss intervals are much lower.

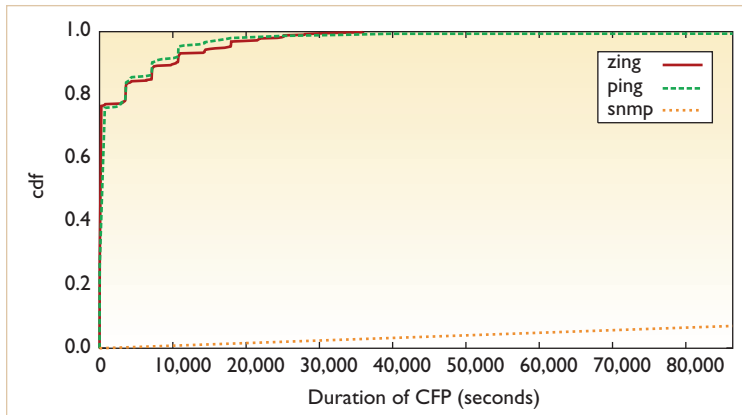


Figure 4. Change-free periods on the canonical path. During the 20-Hz probe period, the cumulative distributions of change-free periods (CFPs) for all paths don't match exactly, indicating that probe-based CFP marginals are not good approximations for router-based ones that suggest much more steady loss rates.

step will be to investigate new lightweight probing methods and tools that have the ability to detect loss over shorter time frames.

Acknowledgments

The authors thank Matt Zekauskas, Guy Almes, Chris Robb, Larry Landweber, Dave Donoho, Jay Rosenbloom, Nevil Brownlee, Mark Crovella, Sue Moon, Yin Zhang, Cisco Systems, and Spirent Communications.

References

1. Y. Zhang et al., "On the Constancy of Internet Path Properties," *Proc. ACM Special Interest Group on Data Communications (SIGCOMM)*, ACM Press, 2001, pp. 197–212.

2. R. Wolff, "Poisson Arrivals See Time Averages," *Operations Research*, vol. 30, no. 2, 1982, pp. 223–231.
3. P. Barford and J. Sommers, *A Comparison of Probe-Based and Router-Based Methods for Measuring Packet Loss*, tech. report TR-1498, Computer Science Dept., Univ. Wisconsin–Madison, 2003.
4. V. Paxson et al., "An Architecture for Large-Scale Internet Measurement," *IEEE Comm.*, vol. 36, no. 8, 1998, pp. 48–54.
5. "How to Read the Output of the Show Controller frfab—tofab Queue Commands on a Cisco 12000 Series Internet Router," Cisco Systems, 2004; www.cisco.com/warp/public/63/fabqueues_18002.html.
6. C. Fraleigh et al., "Packet-Level Traffic Measurements from the Sprint IP Backbone," *IEEE Network*, vol. 17, no. 6, 2003, pp. 6–16.

Paul Barford is an assistant professor of computer science at the University of Wisconsin at Madison; he's also the founder and director of the Wisconsin Advanced Internet Laboratory. His research interests are in the design, measurement, and modeling of wide area networked systems and network protocols. Barford received his BS in electrical engineering from the University of Illinois at Champaign-Urbana and his PhD in computer science from Boston University. Contact him at pb@cs.wisc.edu.

Joel Sommers is currently pursuing his PhD in computer science at the University of Wisconsin-Madison. His research focuses on Internet measurement and analysis of network dynamics. He received a BS in math and computer science from Atlantic Union College, South Lancaster, and an MS in computer science from Worcester Polytechnic Institute, Worcester. Contact him at jsommers@cs.wisc.edu.