

Proceedings of USITS' 99: The 2nd USENIX Symposium on Internet Technologies & Systems

Boulder, Colorado, USA, October 11–14, 1999

STING: A TCP-BASED NETWORK MEASUREMENT TOOL

Stefan Savage



© 1999 by The USENIX Association

All Rights Reserved

For more information about the USENIX Association:

Phone: 1 510 528 8649

FAX: 1 510 548 5738

Email: office@usenix.org

WWW: <http://www.usenix.org>

Rights to individual papers remain with the author or the author's employer.

Permission is granted for noncommercial reproduction of the work for educational or research purposes.

This copyright notice must be included in the reproduced paper. USENIX acknowledges all trademarks herein.

Sting: a TCP-based Network Measurement Tool

Stefan Savage

Department of Computer Science and Engineering

University of Washington, Seattle

savage@cs.washington.edu

Abstract

Understanding wide-area network characteristics is critical for evaluating the performance of Internet applications. Unfortunately, measuring the end-to-end network behavior between two hosts can be problematic. Traditional ICMP-based tools, such as `ping`, are easy to use and work universally, but produce results that are limited and inaccurate. Measurement infrastructures, such as NIMI, can produce highly detailed and accurate results, but require specialized software to be deployed at both the sender and the receiver. In this paper we explore using the TCP protocol to provide more accurate network measurements than traditional tools, while still preserving their near-universal applicability. Our first prototype, a tool called *sting*, is able to accurately measure the packet loss rate on both the forward and reverse paths between a pair of hosts. We describe the techniques used to accomplish this, how they were validated, and present our preliminary experience measuring the packet loss rates to and from a variety of Web servers.

1 Introduction

Measuring the behavior between Internet hosts is critical for diagnosing current performance problems as well as for designing future distributed services. Unfortunately, the Internet architecture was not designed with performance measurement as a primary goal and therefore has few “built-in” services that support this need [Cla88]. Consequently, measurement tools must either “make do” with the services provided by the Internet, or deploy substantial new infrastructures geared towards measurement.

In this paper, we argue that the behavior of the commonly deployed Transmission Control Protocol (TCP) can be used as an implicit measurement service. We present a new tool, called *sting*, that uses TCP to measure the packet loss rates between a source host and some target host. Unlike traditional loss measurement tools, *sting* is able to precisely distinguish which losses occur in the forward direction on the path to the target and which occur in the reverse direction from the target back to the source. Moreover, the only requirement of the target host is that it run some TCP-based service, such as a Web server.

The remainder of this paper is organized as follows: In section 2 we review the current state of practice for measuring packet loss. Section 3 contains a description of the basic loss deduction algorithms used by *sting*, followed by extensions for variable packet size and inter-arrival times in section 4. We briefly discuss our implementation in section 5 and present some preliminary experiences using the tool in section 6.

2 Measuring packet loss

The rate at which packets are lost can have a dramatic impact on application performance. For example, it has been shown that for moderate loss rates (less than 15 percent) the bandwidth delivered by TCP is proportional to $1/\sqrt{\text{lossrate}}$ [MSM97]. Similarly, some streaming media applications only perform adequately under low loss conditions [CB97]. Not surprisingly, there has always been a long-standing operational need to measure packet loss; the popular `ping` tool was developed less than a year after the creation of the Internet. In the remainder of this section we'll discuss two dominant methods for measuring packet loss: tools based on the Internet Control Message Protocol (ICMP) [Pos81] and new measurement infrastructures.

2.1 ICMP-based tools

Common ICMP-based tools, such as `ping` and `traceroute`, send probe packets to a host, and measure loss by observing whether or not response packets arrive within some time period. There are two principle problems with this approach:

- *Loss asymmetry.* The packet loss rate on the forward path to a particular host is frequently quite different from the packet loss rate on the reverse path from that host. Without any additional information from the receiver, it is impossible for an ICMP-based tool to determine if its probe packet was lost or if the response was lost. Consequently, the loss rate reported by such tools is really:

$$1 - ((1 - \text{loss}_{fd}) \cdot (1 - \text{loss}_{rev}))$$

Where loss_{fd} is the loss rate the forward direction and loss_{rev} is the loss rate in the reverse di-

rection. Loss asymmetry is important, because for many protocols the relative importance of packets flowing in each direction is different. In TCP, for example, losses of acknowledgment packets are tolerated far better than losses of data packets. Similarly, for many streaming media protocols, packet losses in the opposite direction from the data stream have little or no impact on overall performance. The ability to measure loss asymmetry allows a network engineering to more precisely locate important network bottlenecks.

- *ICMP filtering.* ICMP-based tools rely on the near-universal deployment of the *ICMP Echo* or *ICMP Time Exceeded* services to coerce response packets from a host [Bra89]. Unfortunately, malicious use of ICMP services has led to mechanisms that restrict the efficacy of these tools. Several host operating systems (e.g. Solaris) now limit the rate of ICMP responses, thereby artificially inflating the packet loss rate reported by `ping`. For the same reasons many networks (e.g. microsoft.com) filter ICMP packets altogether. Some firewalls and load balancers respond to ICMP requests on behalf of the hosts they represent, a practice we call *ICMP spoofing*, thereby precluding real end-to-end measurements. Finally, at least one network has started to rate limit all ICMP traffic traversing it. It is increasingly clear that ICMP's future usefulness as a measurement protocol will be reduced [Rap98].

2.2 Measurement infrastructures

In contrast, wide-area measurement infrastructures, such as NIMI and Surveyor, deploy measurement software at both the sender and the receiver to correctly measure one-way network characteristics [Pax96, PMAM98, Alm97]. Such approaches are technically ideal for measuring packet loss because they can precisely observe the arrival and departure of packets in both directions. The obvious drawback is that the measurement software is not widely deployed and therefore measurements can only be taken between a restricted set of hosts. Our work does not eliminate the need for such infrastructures, but allows us to extend their measurements to include parts of the Internet that are not directly participating. For example, access links to Web servers can be highly congested, but they are not visible to current measurement infrastructures.

Finally, there is some promising work that attempts to derive per-link packet loss rates by correlating measurements of multicast traffic among many different hosts [CDH⁺99]. The principle benefit of this approach is that it allows the measurement of N^2 paths with $O(N)$ messages. The slow deployment of wide-area multicast

routing currently limits the scope of this technique, but this situation may change in the future. However, even with universal multicast routing, multicast tools require software to be deployed at many different hosts, so, like other measurement infrastructures, there will likely still be significant portions of the commercial Internet that can not be measured with them.

Our approach is similar to ICMP-based tools in that it only requires participation from the sender. However, unlike these tools, we exploit features of the TCP protocol to deduce the direction in which a packet was lost. In the next section we describe the algorithms used to accomplish this.

3 Loss deduction algorithm

To measure the packet loss rate along a particular path, it is necessary to know how many packets were sent from the source and how many were received at the destination. From these values the one-way loss rate can be derived as:

$$1 - (\text{packetsreceived}/\text{packetssent})$$

Unfortunately, from the standpoint of a single endpoint, we cannot observe both of these variables directly. The source host can measure how many packets it has sent to the target host, but it cannot know how many of those packets are successfully received. Similarly, the source host can observe the number of packets it has received from the target, but it cannot know how many more packets were originally sent. In the remainder of this section we will explain how TCP's error control mechanisms can be used to derive the unknown variable, and hence the loss rate, in each direction.

3.1 TCP basics

Every TCP packet contains a 32 bit sequence number and a 32 bit acknowledgment number. The sequence number identifies the bytes in each packet so they may be ordered into a reliable data stream. The acknowledgment number is used by the receiving host to indicate which bytes it has received, and indirectly, which it has not. When in-sequence data is received, the receiver sends an acknowledgment specifying the next sequence number that it expects and implicitly acknowledging all sequence numbers preceding it. Since packets may be lost, or re-ordered in flight, the acknowledgment number is only incremented in response to the arrival of an in-sequence packet. Consequently, out-of-order or lost packets will cause a receiver to issue duplicate acknowledgments for the packet it was expecting.

<pre> for i := 1 to n send packet w/seq# i dataSent++ wait for long time </pre>	<pre> for each ack received ackReceived++ </pre>
---	--

Figure 1: Data seeding phase of basic loss deduction algorithm.

<pre> lastAck := 0 while lastAck = 0 send packet w/seq# n+1 while lastAck < n + 1 dataLost++ retransPkt := lastAck while lastAck = retransPkt send packet w/seq# retransPkt dataReceived := dataSent - dataLost ackSent := dataReceived </pre>	<pre> for each ack received w/seq# j lastAck = MAX(lastAck, j) </pre>
---	---

Figure 2: Hole filling phase of basic loss deduction algorithm.

3.2 Forward loss

Deriving the loss rate in the forward direction, from source to target, is straightforward. The source host can observe how many data packets it has sent, and then can use TCP's error control mechanisms to query the target host about which packets were received. Accordingly, we divide our algorithm into two phases:

- *Data-seeding.* During this phase, the source host sends a series of in-sequence TCP data packets to the target. Each packet sent represents a binary sample of the loss rate, although the value of each sample is not known at this point. At the end of the data-seeding phase, the measurement period is concluded and any packets lost after this point are not counted in the loss measurement.
- *Hole-filling.* The hole-filling phase is about discovering which of the packets sent in the previous phase have been lost. This phase starts by sending a TCP data packet with a sequence number *one greater* than the last packet sent in the data-seeding phase. If the target responds by acknowledging this packet, then no packets have been lost. However, if any packets have been lost there there will be a “hole” in the sequence space and the target will respond with an acknowledgment indicating exactly where the hole is. For each such acknowledgment, the source host retransmits the corresponding packet,

thereby “filling the hole”, and records that a packet was lost. We repeat this procedure until the last packet sent in the data-seeding phase has been acknowledged. Unlike data-seeding, hole-filling must be reliable and so the implementation must timeout and retransmit its packets when expected acknowledgments do not arrive.

3.3 Reverse Loss

Deriving the loss rate in the reverse direction, from target to source, is somewhat more problematic. While the source host can count the number of acknowledgments it receives, it is difficult to be certain how many acknowledgments were sent. The ideal condition, which we refer to as *ack parity* is that the target sends a single acknowledgment for every data packet it receives. Unfortunately, most TCP implementations use a *delayed acknowledgment* scheme that does not provide this guarantee. In these implementations, the receiver of a data packet does not respond immediately, but instead waits for an additional packet in the hopes that the cost of sending an acknowledgment can be amortized [Bra89]. If a second packet has not arrived within some small timeout (the standard limits this delay to 500ms, but 100-200ms is a common value) then the receiver will issue an acknowledgment. If a second packet does arrive before the timeout, then the receiver will issue an acknowledgment immediately. Consequently, the source host cannot reliably

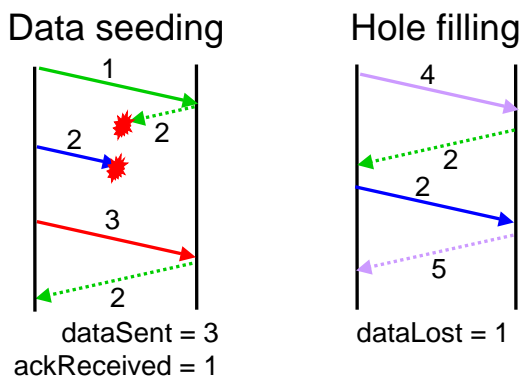


Figure 3: Example of basic loss deduction algorithm. In each time-line the left-hand side represents the source host and the right-hand side represents the target host. Right-pointing arrows are labeled with their sequence number and left-pointing arrows with their acknowledgment number.

differentiate between acknowledgments that are lost and those which are simply suppressed by this mechanism.

An obvious method for guaranteeing ack parity is to insert a long delay after each data packet sent. This will ensure that a second data packet never arrives before the delayed ack timer forces an acknowledgment to be sent. If the delay is long enough, then this approach is quite robust. However, the same delay limits the technique to measuring packet losses over long time scales. If we wish to investigate shorter time scales, or the correlation between the sending rate and observed losses, then this mechanism is insufficient. We will discuss alternative mechanisms for enforcing ack parity in section 4.

3.4 Putting it all together

Figures 1 and 2 contain simplified pseudo-code for the algorithm as we've described it. Without loss of generality, we assume that the sequence space for the TCP connection starts at 0, each data packet contains a single byte (and therefore consumes a single sequence number), and data packets are sent according to a periodic distribution. When the algorithm completes, we can calculate the packet loss rate in each direction as follows:

$$Loss_{fwd} = 1 - (dataReceived/dataSent)$$

$$Loss_{rev} = 1 - (ackReceived/ackSent)$$

Figure 3 illustrates a simple example. Here, the first data packet is received, but its acknowledgment is lost. Subsequently, the second data packet is lost. When the third data packet is successfully received, the target responds with an acknowledgment indicating that it is still

waiting to receive packet number two. At the end of the data seeding phase, we know that we've sent three data packets and received one acknowledgment.

In the hole filling phase, we send a fourth packet and receive an acknowledgment indicating that the second packet was lost. We record the loss and then retransmit the missing packet. The subsequent acknowledgment for our fourth packet indicates that the other two data packets were successfully received.

$$Loss_{fwd} = 1 - (2/3) = 33\%$$

$$Loss_{rev} = 1 - (1/2) = 50\%$$

4 Extending the algorithm

The algorithm we described is fully functional, however it has several unfortunate limitations, which we now remedy.

4.1 Fast ack parity

First, the long timeout used to guarantee ack parity restricts the tool to examining background packet loss over relatively large time scales. If we are interested in examining losses over shorter time scales, or exploring correlations between packet losses and packet bursts sent from the source, then we must eliminate the long delay requirement.

An alternative technique for forcing ack parity is to take advantage of the *fast retransmit* algorithm contained in most modern TCP implementations [Ste94]. This algorithm is based on the premise that since TCP always acknowledges the last in-sequence packet it has received, a sender can infer a packet loss by observing duplicate acknowledgments. To make this algorithm efficient, the delayed acknowledgment mechanism is *suspended* when an out-of-sequence packet arrives. This rule leads to a simple mechanism for guaranteeing ack parity: during the data seeding phase we skip the first sequence number and thereby ensure that all data packets are sent, and received, out-of-sequence. Consequently, the receiver will immediately respond with an acknowledgment for each data packet received. The hole filling phase is then modified to transmit this first sequence number instead of the next in-sequence packet.

4.2 Sending data bursts

The second limitation is that we cannot send large packets. The reason for this is that the amount of buffer space provided by the receiver is limited. Many TCP implementations default to 8KB receiver buffers. Consequently, the receiver can accommodate no more than five

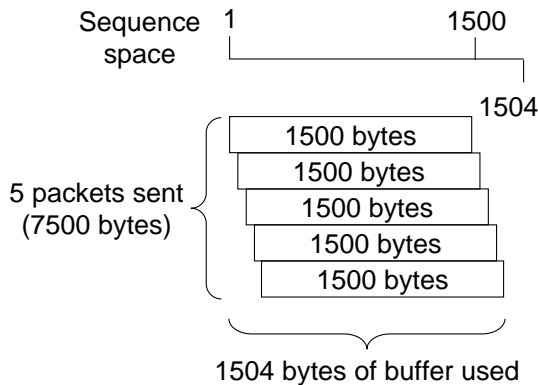


Figure 4: Mapping packets into sequence numbers by overlapping sequence numbers.

1500 byte packets, a number far too small to be statistically significant. While we could simply create a new connection and restart the tool, this limitation prevents us from exploring larger packet bursts.

Luckily, we observe that TCP implementations *trim* packets that overlap the sequence space that has already been received. Consequently, if a packet arrives that overlaps a previously received packet, then the receiver will only buffer the portion that occupies “new” sequence space. By explicitly overlapping the sequence numbers of our probe packets we can map each large packet into a single byte of sequence space, and hence only a single byte of buffer at the receiver.

Figure 4 illustrates this technique. The first 1500 byte packet is sent with sequence number 1500, and when it arrives at the target it occupies 1500 bytes of buffer space. However, the next 1500 byte packet is sent with sequence number 1501. The target will note that the first 1499 bytes of this packet have already be received, and will only use a single byte of buffer space. Using this technique we can map every additional packet into a single sequence number, eliminating much of the buffering limitation. This technique only allows us to send bursts of data in one direction – towards the target host. Coercing the target host to send arbitrarily sized bursts of data back to the source is more problematic since TCP’s congestion control mechanisms normally control the rate at which the target may send data. We have investigated techniques to remotely bypass TCP’s congestion control [SCWA99] but we believe they represent a security risk and aren’t suited for common measurement tasks.

4.3 Delaying connection termination

One final problem is that some TCP servers do not close their connections in a graceful fashion. TCP connections are full-duplex – data flows along a connection in both directions. Under normal conditions, each “half” of the connection may only be closed by the sending side (by sending a FIN packet). Our algorithms implicitly assume this is true, since it is necessary that the target host respond with acknowledgments until the testing period is complete. While most TCP-based servers follow this termination protocol, we’ve found that some Web servers simply terminate the entire connection by sending a RST packet – sometimes called an *abortive release*. Once the connection has been reset, the sender discards any related state so any further probing is useless and our measurement algorithms will fail.

To ensure that our algorithms have sufficient time to execute, we’ve developed two ad hoc techniques for delaying premature connection termination. First, we ensure that the data sent during the data seeding phase contains a valid HTTP request. Some Web servers (and even some “smart” firewalls and load balancers) will reset the connection as soon as the HTTP parser fails. Second, we use TCP’s *flow control* protocol to prevent the target from actually delivering its HTTP response back to the source. TCP receivers implement flow control by advertising the number of bytes they have available for buffering new data (called the *receiver window*). A TCP sender is forbidden from sending more data than the receiver claims it can buffer. By setting the source’s receiver window to zero bytes we can keep the HTTP response “trapped” at the target host until we have completed our measurements. The target will not reset the connection until its response has been sent, so this technique allows us to inter-operate with such “ill-behaved” servers.

5 Implementation

In principle, it should be straightforward to implement the loss deduction algorithms we have described. However, in most systems it is quite difficult to do so without modifying the kernel and developing a portable application-level solution is quite a challenge. We observe that the same problem is true for any user-level implementation of TCP. The principle difficulty is that most operating systems do not provide a mechanism for redirecting packets to a user application and consequently the application is forced to coordinate its actions with the host operating system’s TCP implementation. In this section we will briefly describe the implementation difficulties and explain how our current prototype functions.

5.1 Building a user-level TCP

Most operating systems provide two mechanisms for low-level network access: *raw sockets* and *packet filters*. A raw socket allows an application to directly format and send packets with few modifications by the underlying system. Using raw sockets it is possible to create our own TCP segments and send them into the network. Packet filters allow an application to acquire *copies* of raw network packets as they arrive in the system. This mechanism can be used to receive acknowledgments and other control messages from the network. Unfortunately, another copy of each packet is also relayed to the TCP stack of the host operating system; this can cause some difficulties. For example, if *sting* sends a TCP SYN request to the target, the target responds with a SYN of its own. When the host operating system receives this SYN it will respond with a RST because it is unaware that a TCP connection is in progress.

An alternative implementation would be to use a secondary IP address for the *sting* application, and implement a user-level proxy ARP service. This would be simple and straightforward, but has the disadvantage that users of *sting* would need to request a second IP address from their network administrator. For this reason, we have resisted this approach.

Finally, many operating systems are starting to provide proprietary firewall interfaces (e.g. Linux, FreeBSD) that allow the user to filter outgoing or incoming packets. The former ability could be used to intercept packets arriving from the target host, while the later ability could be used to suppress the responses of the host operating system. We are investigating this approach for a future version.

5.2 The Sting prototype

Our current implementation is based on raw sockets and packet filters running on FreeBSD 3.x and Digital Unix 3.2. As a work-around to the SYN/RST problem mentioned previously, we use the standard Unix `connect()` service to create the connection, and then hijack the session in progress using the packet filter and raw socket mechanisms. Unfortunately, this solution is not always sufficient as the host system can also become confused by acknowledgments for packets it has never sent. In our current implementation we have been forced to change one line in the kernel to control such unwanted interactions.¹ We are currently unsure if a completely portable user-level implementation is possible on today's Unix systems.

Figure 5 shows the output presented by *sting*. From the command line the user can select the inter-arrival dis-

¹We modify the ACK processing in `tcp_input.c` so the response to an acknowledgment entirely above `snd_max` is to drop the packet instead of acknowledging it.

```
# sting www.audiofind.com
```

```
Source = 128.95.2.93
Target = 207.138.37.3:80
dataSent = 100
dataReceived = 98
acksSent = 98
acksReceived = 97
Forward drop rate = 0.020000
Reverse drop rate = 0.010204
```

Figure 5: Sample output from the *sting* tool. By default, *sting* sends 100 probe packets according to a uniform inter-arrival distribution with a mean of 100ms.

tribution between probe packets (periodic, uniform, or exponential), the distribution mean, the number of total packets sent, as well as the target host and port. Our implementation verifies that the wire time distribution conforms to the expected distribution according to the tests provided in [PAMM98].

We have tested our implementation in several ways. First, we have synthetically dropped packets in the tool and using an emulated network [Riz97] and verified that *sting* reports the correct loss rate. Second, we have compared the results of *sting* to results obtained from `ping`. Using the derivation for `ping`'s loss rate presented in section 2 we have verified that the the results returned by each tool are compatible. Finally, we have tested *sting* with a large number of different host operating systems, including Windows 95, Windows NT, Solaris, Linux, FreeBSD, NetBSD, AIX, IRIX, Digital Unix, and MacOS. While we occasionally encounter problems with very poor TCP implementations (e.g. laser printers) and Network Address Translation boxes, the tool is generally quite stable.

6 Experiences

Anecdotally, our experience in using *sting* has been very positive. We've had considerable luck using it to debug network performance problems on asymmetric access technologies (e.g. cable modems). We've also used it as a day-to-day diagnostic tool to understand the source of Web latency. In the remainder of this section we present some preliminary results from a broad experiment to quantify the character of the loss seen from our site to the rest of the Internet.

For a twenty four hour period, we used *sting* to record loss rates from the University of Washington to a collection of 50 remote web servers. Choosing a reasonably-sized, yet representative, set of server sites is a difficult task due to the diversity of connectivity and load expe-




Figure 6: Forward loss measured across a twenty four hour period. Each point on this scatter-plot represents a measurement to one of 50 Web servers.




Figure 7: Reverse loss measured across a twenty four hour period. Each point on this scatter-plot represents a measurement to one of 50 Web servers.

rienced at different points in the Internet. However, it is well established that the distribution of Web accesses is heavy-tailed; a small number of popular sites constitute a large fraction of overall requests, but the remainder of requests are distributed among a very large number of sites [BCF⁺99]. Consequently, we have constructed our target set to mirror this structural property – popular servers and random servers. Half of the 50 servers in our set are chosen from a list of the top 100 Web sites as advertised by `www.top100.com` in May of 1999. This list is generated from a collection of proxy logs and trace files. The remaining 25 servers were selected randomly using an interface provided by Yahoo! Inc. to select pages at random from its on-line database [Yah].

For our experiment we used a single centralized data collection machine, a 200Mhz Pentium Pro running FreeBSD 3.1. We probed each server roughly once every 10 minutes.

Figures 6 and 7 are scatter-plots showing the overall distribution of loss rates, forward and reverse re-

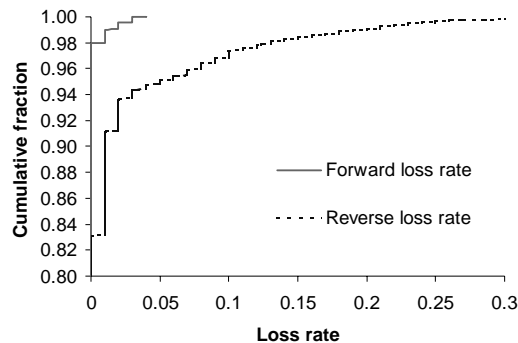


Figure 8: CDF of the loss rates measured to and from a set of 25 popular Web servers across a twenty-four hour period.

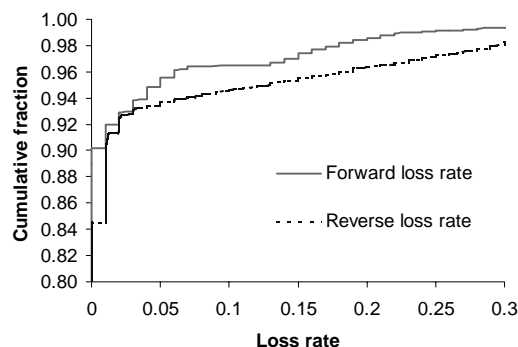


Figure 9: CDF of the loss rates measured to and from a set of 25 random Web servers across a twenty-four hour period.

spectively, during our measurement period. Not surprisingly, overall loss rates increase during business hours and wane during off-peak hours. However, it is also quite clear that forward and reverse loss rates vary independently. Overall the average reverse loss rate (1.5%) is more than twice the forward loss rate (0.7%) and at many times of the day this ratio is significantly larger.

This reverse-dominant loss asymmetry is particularly prevalent among the popular Web servers. Figure 8 graphs a discrete cumulative distribution function (CDF) of the loss rates measured to and from the 25 popular servers. Here we can see that less than 2 percent of the measurements to these servers ever record a lost packet in the forward direction. In contrast, 5 percent of the measurements see a reverse loss rate of 5 percent or more, and almost 3 percent of measurements lose more than a tenth of these packets. On average, the reverse loss rate is more than 10 times greater than the forward loss rate in this population. One explanation for this phenomenon is that Web servers generally send much more traffic than they receive, yet bandwidth is provisioned in a full-duplex fashion. Consequently, bottlenecks are much more likely to form on paths leaving popular Web

servers and packets are much more likely to be dropped in this direction.

We see similar, although somewhat different results when we examine the random server population. Figure 8 graphs the corresponding CDF for these servers. Overall the loss rate is increased in both directions, but the forward loss rate has increased disproportionately. We suspect that this effect is strongly related to the lack of dedicated network infrastructure at these sites. Many of the random servers obtain network access from third-tier ISP's that serve large user populations. Consequently, unrelated Web traffic being delivered to other ISP customers directly competes with the packets we send to these servers.

7 Conclusion

In this paper, we have described techniques for measuring packet loss using TCP. Using these methods the sting tool can separately derive measures for packet loss along the forward and reverse paths to a host, and can be used to probe any TCP-based server. In the future we plan to develop TCP-based techniques for estimating bandwidth [LB99], and round-trip time, and for deriving queue length distributions.

Acknowledgments

We would like to thank a number of people for their contributions to this project. Neal Cardwell, Geoff Voelker, Alec Wolman, Matt Zekauskas, David Wetherall and Tom Anderson delivered valuable feedback on this paper and on the development of sting. Vern Paxson provided the impetus for this project by suggesting that it would be hard to do. USENIX has graciously provided a student research grant to fund ongoing work in this area. Finally, Tami Roberts was exceptionally generous and accommodating in supporting the author during the paper submission process.

References

- [Alm97] Guy Almes. Metrics and Infrastructure for IP Performance. <http://www.advanced.org/surveyor/presentations.html>, 1997.
- [BCF⁺99] Lee Breslau, Pei Cao, Li Fan, Graham Phillips, and Scott Shenker. Web Caching and Zipf-like Distributions: Evidence and Implications. In *Proceedings of the IEEE INFOCOM '99*, March 1999.
- [Bra89] R. Braden. Requirements for Internet Hosts – Communication Layers. RFC-1122, October 1989.
- [CB97] Georg Carle and Ernst W. Biersack. Survey of Error Recovery Techniques for IP-Based Audio-Visual Multicast Applications. *IEEE Network Magazine*, 11(6):24–36, November 1997.
- [CDH⁺99] R. Caceres, N.G. Duffield, J. Horowitz, D. Towsley, and T. Bu. Multicast-Based Inference of Network-Internal Characteristics: Accuracy of Packet Loss Estimation. In *Proceedings of the IEEE INFOCOM '99*, New York, NY, March 1999.
- [Cla88] D. Clark. The Design Philosophy of the DARPA Internet Protocols. In *Proceedings of the ACM SIGCOMM '88*, pages 106–114, Palo Alto, CA, September 1988.
- [LB99] Kevin Lai and Mary Baker. Measuring Bandwidth. In *Proceedings of the IEEE INFOCOM '99*, New York, NY, March 1999.
- [MSM97] Matthew Mathis, Jeffrey Semke, and Jamshid Mahdavi. The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm. *ACM Computer Communications Review*, 27(3):67–82, July 1997.
- [PAMM98] V. Paxson, G. Almes, J. Mahdavi, and M. Mathis. Framework for IP performance metrics. RFC-2230, May 1998.
- [Pax96] Vern Paxson. End-to-End Routing Behavior in the Internet. In *Proceedings of the ACM SIGCOMM '96*, pages 25–38, Stanford, CA, August 1996.
- [PMAM98] Vern Paxson, Jamshid Mahdavi, Andrew Adams, and Matthew Mathis. An Architecture for Large-Scale Internet Measurement. *IEEE Communications*, 36(8):48–54, August 1998.
- [Pos81] J. Postel. Internet Control Message Protocol. RFC-792, September 1981.
- [Rap98] Chris Ravier. ICMP and future testing (IPPM mailing list). <http://www.advanced.org/IPPM/archive/0606.html>, December 1998.
- [Riz97] L. Rizzo. Dummynet: A Simple Approach to the Evaluation of Network Protocols.

ACM Computer Communications Review,
27(1), January 1997.

- [SCWA99] Stefan Savage, Neal Cardwell, David Wetherall, and Tom Anderson. TCP Congestion Control with a Misbehaving Receiver. Draft, in review, 1999.
- [Ste94] W. Richard Stevens. *TCP/IP Illustrated*, volume 1. Addison Wesley, 1994.
- [Yah] Yahoo! Inc. Random yahoo! link url. <http://random.yahoo.com/bin/ryl>.