Fingerprinting Internet Paths using Packet Pair Dispersion

USC Computer Science Technical Report No. 06-876

Rishi Sinha
Department of Computer
Science
University of Southern
California
rsinha@netweb.usc.edu

Christos Papadopoulos
Department of Computer
Science
University of Southern
California
christos@isi.edu

John Heidemann
Department of Computer
Science
University of Southern
California
johnh@isi.edu

ABSTRACT

Path fingerprinting is an essential component of applications that distinguish among different network paths, including path selection in overlay networks, multi-path routing, monitoring and diagnosis of network problems, and developing a deeper understanding of network behavior. This paper proposes a new approach to Internet path fingerprinting based on the distribution of end-to-end packet-pair measurements. This approach allows detection of busy link sharing between two paths, even when those segments have low utilization and are not the paths' bottlenecks. While our fingerprints do not assure physically disjoint paths (since that requires information external to the network), they reflect the traffic and link characteristics of intermediate links. This methodology is therefore tolerant of opaque clouds such as VPNs, VLANs, or MPLS (unlike traceroute). Using analysis and simulation we explore the network factors that affect the fingerprints, and we introduce a simple method to compare them. Through measurements of up to a year over 15 Internet paths, we show that our fingerprints are both distinct and persistent over periods of several months, making their collection and use for path selection feasible.

1. INTRODUCTION

Path characterization is a process that measures the performance of a path, typically capacity, loss, delay. The goal of path characterization is to describe something about the path itself. *Path fingerprinting* is a process that generates a unique fingerprint for a path, typically aimed at distinguishing one path from another. The fingerprint may or may not be dependent on performance-related properties of the path. Both path characterization and fingerprinting are important to a wide range of applications such as overlay path selection, multi-path routing, monitoring and diagnosis of network problems, designing and testing protocols for realistic network conditions, and developing a deeper understanding of network behavior.

Path fingerprinting is especially useful for applications that require the selection of disjoint paths to improve performance or robustness. Path fingerprints allow such applications to determine if two paths share common links. A simple path fingerprint is the identity of all routers on the path. While tools such as traceroute aim to provide this information, these tools fail when the router topology becomes opaque due to layer-2 clouds (from ATM or optical switching), tunneling (such as MPLS or VPNs), or non-cooperative routers.

While an exact router-level topology is difficult to determine, users optimizing network performance are often interested in *performance isolation*. Here the goal is to identify paths that do not share highly utilized links. For such applications, the existence of shared bottlenecks must be identified, but over-provisioned links are of much less concern. Ideally, path fingerprints should also be relatively stable, allowing data collection costs to be amortized by reusing and sharing fingerprints.

In this paper we develop a new technique to fingerprint Internet paths based on the distribution of packet pair dispersions, which we call the *dispersion fingerprint* of a path (Section 2). Our approach has the following important properties: it generates *distinct* path fingerprints that can be reliably compared using simple techniques; fingerprints taken at the same time of day are *persistent* over periods of at least several months, and are thus reusable; and the fingerprints are shaped by traffic and link characteristics, especially links with higher utilization, and can be used to detect path overlap.

To understand how dispersion fingerprints are shaped by the network we first systematically study fingerprints through network simulation (Section 3), and compare these results with fingerprints from over 15 paths on Internet2 for periods of up to a year (Section 4). We corroborate these results with short-term measurements from commercial Internet paths (Section 5). Finally, (Section 6) we propose two applications of fingerprints, namely detecting sharing of busy links between two paths, and detection of traffic and link changes over time on a path, and discuss the advantages of our approach over existing solutions.

To our knowledge, the existence of distinct and persistent path fingerprints that capture link and traffic characteristics has not been demonstrated before. While there has been a great deal of work on path characterization [3, 5–7, 12, 26, 35], most has focused on determining specific network characteristics such as loss, delay, and throughput rather

than fingerprinting. The packet pair methodology has been widely used before in link capacity and available bandwidth estimation [9, 10, 13, 18, 20, 24, 30]. Our work differs because we aim to find an indicator of a path's unique identity rather than characterizing performance aspects.

Our contributions are as follows: (a) we define dispersion fingerprints as a new approach to identifying paths shaped by traffic and link characteristics; (b) we investigate the underlying physical basis for dispersion fingerprints and demonstrate through simulation, analysis and measurement how links and cross traffic affect dispersion; and (c) we demonstrate that that many Internet paths contain a distinct and persistent dispersion fingerprint. In addition, we begin to explore how our fingerprints can be used to identify paths with shared links and detect traffic changes on a path.

2. THE DISPERSION DISTRIBUTION AS A PATH FINGERPRINT

In this section we present evidence that the distribution of packet pair dispersions is a good candidate for fingerprinting Internet paths.

2.1 Packet Pair Basics

Packet pairs [18] and their variants (such as packet trains), are a useful building block commonly used in tools for link capacity and available bandwidth estimation [7, 9, 10, 13, 20, 24, 26, 30]. A packet pair typically consists of two of equal-sized packets sent with fixed initial *dispersion* δ_{init}. We use the term dispersion as it is commonly defined: the time interval between the first bit of the first packet and the first bit of the second packet of the pair. At the destination, the resulting dispersion δ_{final} is a function of both the physical capacity of the various links in the path as well as competing cross traffic on these links.

Figure 1 shows how network links and traffic affect dispersion. Observe the incoming and outgoing dispersions (δpre and δ_{post}) at the first router in each case. With no interference and equal capacity links, case (a) shows that there is no change in dispersion. In case (b), the packet pair experiences *expansion* (the dispersion increases) as it moves from a high-bandwidth to a low-bandwidth link. In case (c), the dispersion value again increases, but this time due to cross-traffic packets slipping between the packet pair, increasing their separation. Finally, in case (d), the packet pair experiences *compression* due to queuing at a busy link. The final dispersion δ_{final} at the end of the path is due to combinations of these effects along the entire path. We reevaluate these cases more precisely in Section 3.2.

Several other network factors can also affect dispersion values from packet pair measurements. Scheduling policies, multi-path routing and route changes could all affect measurements. Like most other prior work using packet pairs, we assume the common case of FIFO or RED queuing, single-path routing and stable routes between probes.

In this paper we explore the use of packet pairs to characterize network paths. From these examples we observe that the network can both increase or decrease the dispersion value unpredictably, making it an apparently difficult choice. We next explore how the the network forces regularities in dispersion values, allowing a modest number of probes to characterize a link.

(a) No change, $\delta_{post} = \delta pre$.

(b) Link expansion, $\delta_{post} > \delta pre$.

(c) Traffic expansion, $\delta_{post} > \delta pre$.

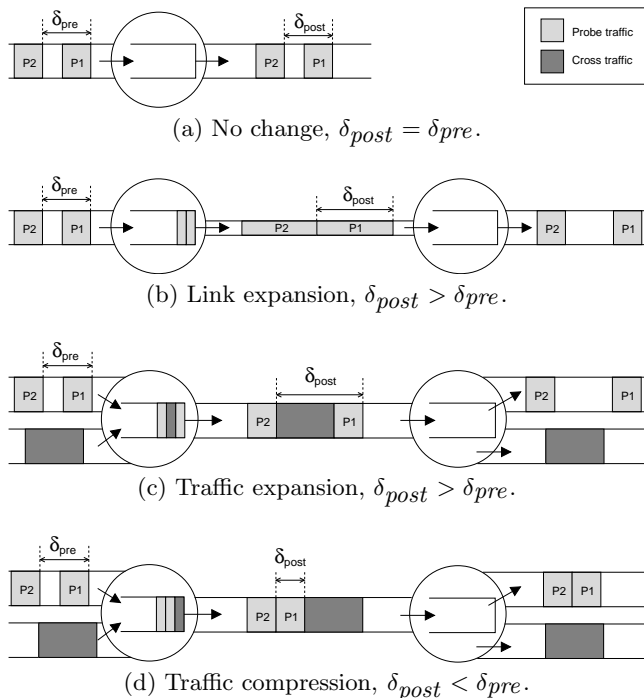(d) Traffic compression, $\delta_{post} < \delta pre$.

Figure 1: Examples of how packet pair dispersion may change.

2.2 Can Dispersion Values Provide a Path Fingerprint?

The discussion above suggests that packet pair dispersions reflect features of both the static (physical) and the dynamic (cross-traffic) properties of a path. However, fingerprints are most useful to the extent that they are *distinct* and *persistent* over at least moderate periods of time. We look into these questions next.

Figure 2(a) shows the results of an experiment where a sustained stream of UDP packet pairs was transmitted for a period of one hour between hosts in Germany and Massachusetts. We call this path DE–UM. The initial dispersion of each pair was 120 μs, which corresponds to back-to-back transmission of the two 1500 byte packets at 100 Mb/s. (Full details about our methodology are in Section 3.1.) The figure shows final dispersion δ_{final} on the y-axis with time on the x-axis. This scatter plot shows strong horizontal bands, indicating frequent dispersion modes around 135, 160, 200, 240, 280, 320, 360 μs, etc.

Although the figure shows strong trends in the data, it is not obvious how to interpret it. We would like to quantify the strength of given bands and understand their causes, so we switch to a cumulative distribution of the dataset in Figure 2(b). This representation makes quantification of the dispersions easier, since the strength of each band is proportional to the size of the step in the CDF. Figure 2(b) also adds the CDF of a similar experiment conducted about a month later. The strong similarity of the CDFs in Figure 2(b) suggests that the bands in Figure 2(a) are not transient but may be caused by the underlying network, indicating that *dispersion CDFs may provide persistent fingerprints.*

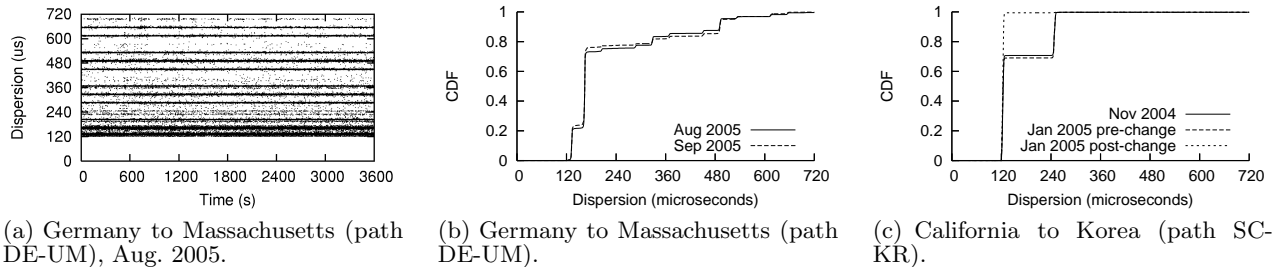Turning to the question of the distinctness, Figure 2(c)

(a) Germany to Massachusetts (path DE-UM), Aug. 2005.

(b) Germany to Massachusetts (path DE-UM).

(c) California to Korea (path SC-KR).

Figure 2: Measurements motivating the use of dispersion distributions as path signatures.

shows dispersion CDFs from three measurements taken with the same methodology on a different path, California to Korea (SC–KR). We observe that the DE–UM CDFs were very similar to each other, while the SC–KR measurements appear quite different. This observation suggests that *dispersion CDFs may be distinct for a given path.* In Section 4 we measure 15 different Internet paths over periods of up to a year, and we find each one to have a distinctly different CDF.

While the DE–UM CDFs (Figure 2(b)) are very consistent, we see two different trends in SC–KR (Figure 2(c)). In January 2005, there was a significant and permanent change in the CDF of this path—before the change we see two modes (at 120 and 240 μs), but later only the first mode remains. We associate this change with a network outage and temporary route change, suggesting that the routing facilities where altered at this time. We explore this example in detail in Section 4.5. Here we observe that large changes in dispersion CDFs can be associated with changes in the underlying network; dispersion CDFs are not completely unchanging.

In this section we showed preliminary evidence that dispersion CDFs of Internet paths can persist for months, and can be distinct, suggesting that dispersion CDFs may provide a new type of path fingerprint. In the following sections, we investigate in detail the factors affecting our fingerprint and its properties and examine dispersion CDFs through simulation and a larger set of Internet measurements.

3. SYSTEMATIC STUDY OF THE DISPERSION CDF

We next explore factors that influence dispersion CDFs, beginning with simple one-router scenarios and case-by-case analysis of cross-traffic. We then generalize this discussion of single-hop dispersion, dividing the set of conditions into three regions of operation. Finally we turn to simulations of small scenarios with multiple routers.

Dispersion at a single hop has been studied before [13, 21, 25] when considering bandwidth estimation. Our goal here is to understand the factors most pertinent to dispersion CDFs used as path fingerprints.

3.1 Methodology for Computing the Dispersion CDF

Our simulations and experiments use the following methodology. We send UDP packet pairs with exponential interdeparture times (exploiting the PASTA principle [33]) at a mean rate of 100 pairs/s. Experiment lengths are either 100 s (simulation) or 1 hour (Internet). (We explore the effect of packet pair rate and measurement duration in detail in Section 5, and gets similar results with rates as low as 10 pairs/s and durations as short as 30 s.) Each packet is 1500 bytes long (including UDP and IP headers) We send the two packets back-to-back at the rate of the access link (typically 100 Mb/s), so δ_{init} is 120 μs. To factor out link speed we normalize the dispersion at the receiver by the initial dispersion and report $\tilde{\delta} = \delta_{final}/\delta_{init}$, the *normalized dispersion.*

3.2 Analysis of Dispersion After A Single Router

We begin with simple one-hop scenarios with dispersion from packet pairs interacting with cross traffic.

We use a trivial topology: one router with two incoming and one outgoing link (Figure 3(a)). One incoming link carries the packet pairs, the other carries cross traffic, and all packets exit the router through the outgoing link. Both incoming links have the same capacity C_{pre}, while the outgoing link may have the same or different speed C_{post}. We assume that a packet pair arrives at the router with arbitrary dispersion δ_{pre} and leaves with dispersion δ_{post}. We assume packet pairs and cross traffic packets have the same length, but write them L_P and L_X to allow this assumption to be relaxed later. Section 3.4.3 deals with the case of multiple packet sizes in cross traffic.

To describe how large a particular value of dispersion is, it is helpful to define *dispersion slack*, S, the duration the link would be idle between the two probes if no cross traffic were present (see Figure 3(a)). For any incoming pair, the input slack is $S_{pre} = \delta pre - L_P/C_{pre}$ and the output slack is $S_{post} = \delta_{post} - L_P/C_{post}$, where δ_{post} is the output dispersion that would result if no cross traffic contended for the output link. Note that $S_{post} = max(0, \delta pre - L_P/C_{post})$.

We want to consider each way cross traffic may alter probe dispersion. To do this, we consider when a potential cross-traffic packet arrives, and what the spacing of the packet pair is. We assume non-preemptive queueing, so we can discard all cases when cross traffic arrives after the head of the second probe packet.

3.2.1 Cross Traffic Arriving Between the Probes

First, consider when a cross-traffic packet arrives after the head of the first probe packet but before the head of the second. In this case, it will interfere, delaying the second packet and *increasing* dispersion, but the increase is dependent on the output slack S_{post}.

When there is no slack ($S_{post} = 0$), then any intervening packet will increase dispersion by exactly the service time of

(a) Dispersion slack.

(b) Dispersion CDFs for each region of operation.
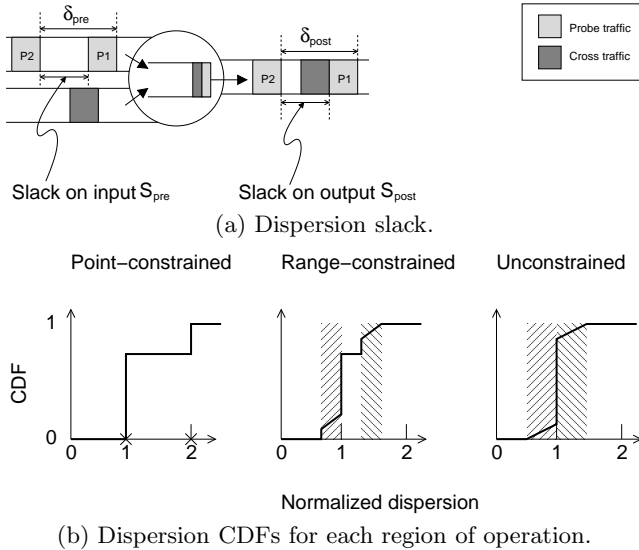
Figure 3: Dispersion slack is the key to the regions of operation.

that new packet. In other words, $\delta_{post} = \delta pre + L_X/C_{post}$ (if interference), or $\delta_{post} = \delta pre$ (if not). We call this region of operation *point-constrained*, because dispersion increases by a fixed amount.

Next, suppose there is a bit of slack, but less than a full packet's worth, that is, $0 < S_{post} < L_X/C_{post}$. Again, if no cross traffic arrives dispersion is unchanged, but the longer inter-packet gap makes this chance lower than when point-constrained. If cross traffic arrives immediately before the second probe packet, it will increase the dispersion by L_X/C_{post} (the same as point-constrained). However, if the packet enters a bit earlier, than part will be transmitted before the second probe packet enters, and the increase in dispersion will be slightly smaller: $L_X/C_{post} - \epsilon$. If the intervening packet enter immediately after the first probe packet, it can consume all the slack. Thus, an intervening packet may increase dispersion by any value between $L_X/C_{post} - S_{post}$ and L_X/C_{post}. We therefore call this region of operation *range-constrained*.

Finally, if there is more than a packet's worth of slack ($S_{post} \geq L_X/C_{post}$), then the dispersion can increase by any amount between 0 and L_X/C_{post}, since the slack can now absorb an entire cross-traffic packet. We call this region of operation *unconstrained*.

Figure 3(b) shows dispersion CDFs for these cases, showing when δpre is either L_P/C_{post}, $1.5L_P/C_{post}$ or $2L_P/C_{post}$.

3.2.2 Cross Traffic Arriving Before the First Probe

Next we consider when cross traffic arrives β before the head of the first probe. If it arrives well before the first probe, when $\beta \geq L_X/C_{post}$, then it can be completely serviced and does not affect dispersion. Otherwise the effects again depend on slack.

If $S_{post} = 0$, then both probes are delayed by β, but dispersion does not change because the probes remain back-to-back.

If $0 < S_{post} < \beta - L_X/C_{post}$, then the cross traffic delays the first probe's service time until after the second probe

arrives, decreasing δ_{post} to zero (point-constrained).

If $S_{post} \geq \beta - L_X/C_{post}$, the cross traffic again delays the first probe, decreasing dispersion, but the pairs do not leave back-to-back (range-constrained).

If $\beta = 0$ and $S_{post} \geq L_X/C_{post}$, then the dispersion is reduced by one full packet service time; $\delta_{post} = \delta pre - L_X/C_{post}$, and dispersion is point-constrained.

In these examples, the reduction in dispersion occurs because packets queue behind cross traffic. This cause is the same cause of ACK compression [34], but in our case, with data packets.

3.2.3 Faster Output Link

Next consider different link speeds, $C_{pre} \neq C_{post}$. First consider the case when the output link is faster than the input link, $C_{post} > C_{pre}$. Assume the packets arrive back-to-back on the input link, with no idle time between them, so $\delta pre = L_P/C_{pre}$ and $S_{pre} = 0$. This still implies non-zero output slack. Specifically, the slack $S_{post} = L_P/C_{pre} - L_P/C_{post}$, a value greater than zero, since $C_{post} > C_{pre}$. This slack on the output link creates the potential for the faster link to silently "absorb" some cross traffic, provided the cross traffic can slip into the output link's slack. We will explore this issue in more detail when we consider simulation results in the next section.

3.2.4 Multiple or Faster Input Links

With multiple sources of cross traffic or a faster single source of cross traffic, dispersion values can increase by integral multiples of each of the above cases. (We assume round-robin servicing of all input links.) So if the probes arrive with no slack, then they may exit with $\delta_{post} = L_P/C_{post} + n(L_X/C_{post})$ for any integer $n > 0$ in the point-constrained cases. If there is more slack when the probes enter, they will be constrained into range "stripes" with the maximum of each range at some multiple $n(L_X/C_{post})$, and the minimum at some multiple $n(L_X/C_{post} - S_{post})$. (Note that the width of the range is wider at higher multiples.)

A faster single source of cross traffic causes changes similar to multiple links. In the next section we will simulate cross traffic entering in link speeds ten times that of the probe traffic.

3.2.5 Multiple Hops

Cross traffic over multiple hops causes increase or decrease in dispersion at every hop. Moreover, an increase in dispersion at one hop increases the slack available at the next, so interactions quickly become complicated. We consider them in simulation in the next section.

3.3 Generalized Regions of Operation on a Single Hop

In Section 3.2 we saw specific examples of how packet interactions give rise to three regions of operation at a single hop. Now we generalize the conditions underlying each region. for arbitrary combinations of the packet arrival and link speed scenarios separated above. There are no assumptions other than constant cross traffic packet size. Let n be the number of intervening packets of cross traffic, which may be zero. Each packet pair falls into one of the three regions (Figure 3(b)) depending on its value of slack S_{post}, which is determined by δpre. When $\delta pre \leq L_P/C_{post}$, the slack S_{post} is zero, and this is the point-constrained region,
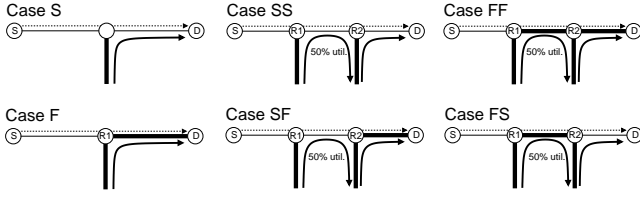
Figure 4: The six possible topologies involving one or two busy links and two link capacities. Thick lines represent fast (F) links, while thin lines represent slow (S) links.

where each value of n results in exactly one value of δ_{post}. When $L_P/C_{post} < \delta pre < (L_X + L_P)/C_{post}$, the slack S_{post} can absorb only part of a cross traffic packet, and this is the range-constrained region, where each value of n gives rise to non-overlapping "stripes" of possible δ_{out} values. When $\delta pre \geq (L_X + L_P)/C_{post}$, the slack S_{post} can absorb one or more cross traffic packets, and the "stripes" now overlap, making any δ_{out} value above L_P/C_{post} possible.

3.4 Simulation of Dispersion With Multiple Routers

To explore more complex scenarios, we now use simulation and the six topologies shown in Figure 4. We send packet pairs from source S to the destination D through routers R1 and, when present, R2. All links except the access link (S–R1) carry cross traffic. We use either slow links (100 Mb/s, thin lines), or fast links (1 Gb/s, bold lines). We label the topologies using S (for Slow) and F (for Fast) based on the speeds of the busy links, so SF indicates two busy links, a slow link followed by a fast link. (Links feeding cross traffic are always fast links to allow single or multiple packet arrivals between probes.)

In each topology we vary cross traffic such that utilization at the last link is 30%, 50% or 90%. (Values from 10% to 90% were similar.) With two busy links (SS, SF, FF and FS), we fix cross traffic on R1–R2 at 50% and vary utilization on the R2–D link. Initially, we assume cross traffic is only 1500 byte packets with Poisson arrivals; in Section 3.4.3 we consider multiple packet sizes. Cross traffic is sent using exponential interarrival times. An open issue is to explore non-Poisson distributions of cross traffic.

Probe traffic is injected on the 100 Mb/s link S–R1; the same speed as slow intermediate links, but one-tenth the speed of fast intermediate links. The probes are sent as described in Section 3.1 back-to-back ($S_{pre} = 0$).

We use the ns-2 simulator [8]. Each simulation run lasts 100 s, but we discard the first 5 s to avoid start-up transients. (The simulation is stable after 10 s, so the duration is ample.) We repeated each simulation 5 times with different randomization but found no significant variation, so we present the results of a single instance here.

3.4.1 The Dispersion CDF with One Busy Link

We begin with the simplest case: one busy link that is either the same speed or faster than the access link of the sender. We will find that these two cases are in the point-constrained and unconstrained regions of operation, respectively. Range-constrained operation is an intermediate region with characteristics of the other two regions, and is not

illustrated in these simulations.

Case S. Figure 5(a) shows the dispersion CDFs with homogeneous link speeds. In this topology, all pairs arrive at the busy link with a dispersion L_P/C (where C is the capacity of R1–D) and have zero slack. This scenario is point-constrained, so we see steps in the CDF only at integer multiples of L_X/C. Considering the case of light load (30% utilization), we see that about 75% of pairs show no dispersion change—they enter and leave the congestion link without change in their spacing. About 20% of pairs show a dispersion of $2\delta_{init} = L_P/C + L_X/C$, corresponding to one cross-traffic packet arriving at R1 after the first probe of the pair but before the second. Finally, we occasionally (about 5% of the time) see dispersion factors higher than 2, corresponding to capture of multiple packets of cross traffic. Since the packet size is fixed, we always see CDF steps at multiples of the packet size, and we characterize this kind of CDF as *stairstep*.

With 90% load, again we see a stairstep dispersion CDF, but with more cross traffic, large dispersions are more common since multiple packet insertions between probes are more common.

These simulations confirm the analysis of Section 3.2 and 3.2.4. With Poisson cross traffic we can compute the expected CDF. The fraction with no dispersion change correspond to the probability that no cross traffic arrives in duration L_P/C. The expected number of arrivals in this time is $(L_P/C) \times (\lambda_X/L_X)$, where λ_X is the average data rate of cross traffic. The observed CDFs in case S are in agreement with this distribution.

Case F. Next we turn to case F, where a fast link carries cross traffic. Since the capacity C of the busy link R1–D is 10 times that of the access link, the incoming dispersion at R1 is $10(L_P/C) = 5(L_P + L_X)/C$, which puts all pairs in this case in the unconstrained region. Figure 5(b) shows its CDF. which we describe as *smooth*. Rather than the large, discrete steps of the stairstep CDF we obtained in the point-constrained case, here we see a spread of dispersion values, centered at δ_{in}. Like case S, normalized dispersions greater than δ_{init} correspond to cross traffic arriving between two probes and spreading them apart. Unlike case S, the amount of increase is unconstrained, and thus the resulting CDF shows weight at continuous values rather than discrete steps. Another difference from case S is that we see output dispersions δ_{final} that are *smaller* than input dispersion δ_{init}, since, as we argued analytically in Section 3.2.2, pairs can be pushed together if they must queue behind cross traffic. However, the lower bound on δ_{final} is L_P/C.

Comparing the amounts of cross traffic for case F, we see that when there is more cross traffic, we get higher variation in the dispersion (compare the 90% to 30% utilizations). Again, this result is due to a higher probability of probes capturing multiple cross packets. At 90% load, there are steps at multiples of L_X/C (which is $0.1\delta_{init}$), because both probes in a pair are often queued along with intervening cross traffic, resulting in zero output slack for many pairs.

3.4.2 The Dispersion CDF with Multiple Busy Links

We now turn to cases SS, FF, SF and FS, where there are two busy links, R1–R2 and R2–D. In each of the previous cases S and F, the incoming dispersion at the busy link was the same for all pairs, so all pairs were in the same region of operation. In the multiple-busy-link cases, the in-
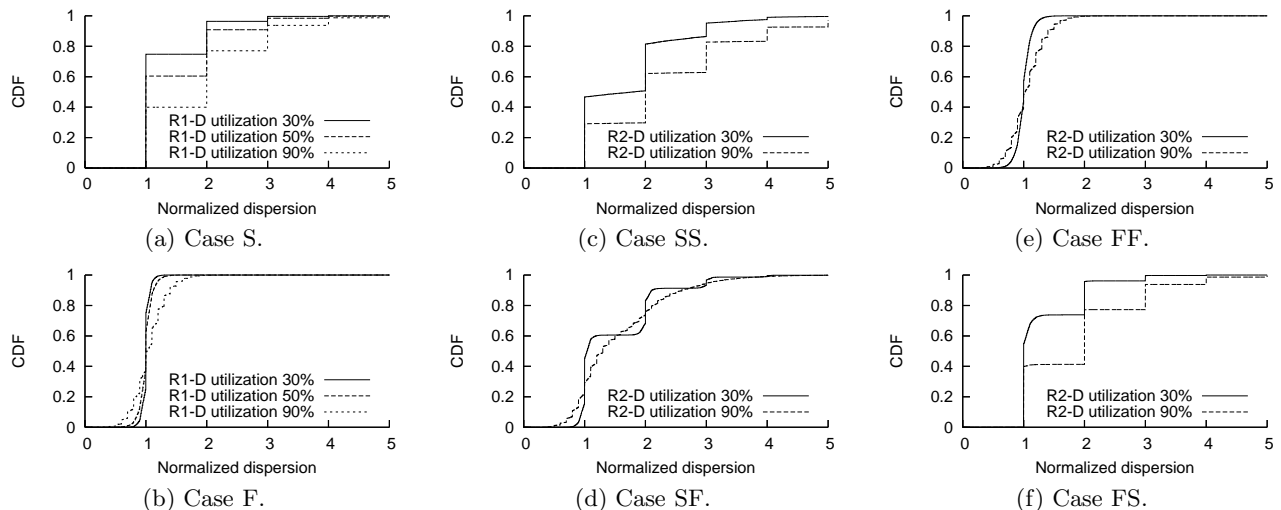
Figure 5: Dispersion CDFs from the six simulation cases.

coming dispersions at R1 are still the same for all pairs, but pairs entering R2 may have different dispersions because of interference at R1. Consequently, all pairs do not operate in the same region on R2–D. For multiple link cases we fix the utilization of R1–R2 at 50%, so the 50% load curves in Figures 5(a) and 5(b) show the distribution of dispersions approaching R2 for cases SS and SF, or FS and FF, respectively.

Case SS. When both busy links are slow links (Figure 5(c)), the incoming dispersions at R2 are either in the point-constrained or unconstrained region. This is seen from the 50% curve in Figure 5(a)—approximately 60% of dispersions stay unchanged at δ_{init}, and these are point-constrained with respect to R2–D; the rest of the dispersions are too large to be in the range-constrained region, and are all in the unconstrained region with respect to R2–D. The 60% of pairs that are point-constrained produce jumps at integer values in Figure 5(c), while the rest of the pairs produce both integer and non-integer values. Although the regions of operation don't change at 90% utilization of R2–D, the higher load in the last link dominates the CDF, forcing the dispersion CDF into a nearly stairstep appearance.

Case SF. Now consider case SF, where the second busy link is a fast link (Figure 5(d)). The incoming dispersions at R2 are still distributed as the output of case S with 50% load. However, because R2–D is now faster than the probe access link, router R2 operates in the unconstrained region for all pairs. We expect a stairstep distribution to be induced at router R1, and at 30% utilization these steps are each smoothed as in case F. At 90% utilization, the cross traffic at R2 dominates the dispersion CDF giving it an overall smooths shape, but actually with many small bumps at the 1 Gb/s back-to-back spacing.

Case FF. For case FF (Figure 5(e)), the incoming dispersions at R2 are distributed as in the 50% curve in Figure 5(b). The vast majority of these dispersions, those greater than $0.2\delta_{init}$, are in the unconstrained region when they arrive at R2. Thus, the dispersion CDF for case FF is like case F itself, although with slightly longer tails and smoother edges.

Case FS. Finally, in case FS (Figure 5(f)), pairs arriv-

ing at R2 again have dispersions distributed as the output of case F with 50% load, but in this case dispersions up to δ_{init} are point-constrained, and dispersions between δ_{init} and $2\delta_{init}$ are range-constrained. A very small number of pairs with dispersion more than $2\delta_{init}$ are unconstrained. Since most of the dispersions are point-constrained, Figure 5(f) shows steps at integer values in the 30%-load curve. The effect of range-constrained dispersions is also evident between normalized dispersions 1 and 2, though the higher-valued band repetitions are too weak to be visible. It is interesting that cases FS and SF are similar in topology, but there are differences in their dispersion CDFs. From this observation we conclude that the *order* of busy links affects the dispersion CDF.

These results demonstrate how the dispersion CDF is shaped by link capacities, utilizations and order of busy links when there is traffic on multiple links.

3.4.3 The Dispersion CDF with Multiple Packet Sizes

Cross traffic in all the above simulations consists of a single packet size. To consider multiple packet sizes, we must revisit the classification into regions of operation, as well as the simulations.

In general, the three regions still exist for multiple cross traffic packet sizes, using the smallest packet size for L_X in the expressions in Section 3.3. In particular cases, however, it is possible for the range-constrained region to merge with the unconstrained region. This happens when bands due to different packet sizes overlap to cover an unbounded range of δ_{final} values.

The typical Internet packet size distribution is believed to consist mainly of a few strong modes, around 40, 576 and 1500 bytes [23, 29, 32]. We recently reported [28] that this distribution appears to have shifted to a mostly bimodal distribution with modes around 40 and 1500 bytes. More details are in Appendix A.

Based on this distribution, we repeated the simulations in the six cases of Figure 4 with a bimodal distribution of packet sizes in cross traffic on the last link in each case. The mix was 30% 1500 byte packets and 70% 40 byte packets. As expected, we observed that the effect of introducing the
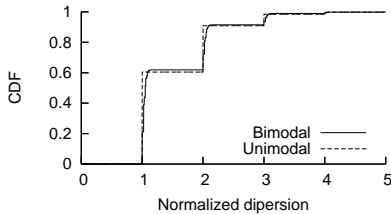
Figure 6: Dispersion CDF from simulation case S with bimodal packet size distribution.

Table 1: Measurement sites

Site	Domain	Abbr.
U. of Southern California	usc.edu	SC
UC Santa Barbara	ucsb.edu	SB
UC San Diego	ucsd.edu	SD
U. of Mass.	umass.edu	UM
U. of Maryland	umd.edu	MD
Inha U., Korea	inha.ac.kr	KR
Nat. Tech. U. of Athens, Greece	ntua.gr	GR
T.U. Braunschweig, Germany	tu-bs.de	DE

smaller packets was to remove sharp step transitions in the CDFs, where they existed. For example, Figure 6 shows simulations for case S (with 50% utilization on R1–D) with both unimodal and bimodal cross traffic. All pairs in both curves were in the point-constrained region, but a greater number of possible δ_{final} points exists in the bimodal case. However, because a 40 byte packet is much smaller compared to a 1500 byte packet, the CDF is shaped largely by 1500 byte packets.

This evidence on packet size distribution and the effect of smaller packets indicates that dispersion CDFs on the Internet will be shaped largely by packets near 1500 bytes long.

4. MEASUREMENT OF DISPERSION IN THE INTERNET

In this section we study dispersion measurements taken from paths on the Internet. Our measurements span a period between October 2004 and October 2005 and cover 15 Internet paths between the 8 sites shown in Table 1. We do not have measurements for all 56 paths between the 8 sites due to unstable hosts. Our hosts were located on academic sites, so the paths we measured are over Internet2 and other research networks with a capacity of at least 100 Mb/s. We complement our measurements with a commercial site in Section 5.1 to provide short-term validation of our primary survey; they generally confirm our long-term results.

4.1 Experiment Methodology

Basic experiment parameters are described in Section 3.1. A measurement lasts for 24 hours in segments of one hour each. Before and after each segment, we record traceroute results for each path. All but one of the machines used were on 100 Mb/s networks. Pairs from these machines were transmitted back-to-back, with initial dispersion being 120 μs. For the path (MD-SC) the sender was on a 1 Gb/s LAN during the measurement, but we maintained the same initial dispersion of 120 μs. We did not use the 1 Gb/s

host as a receiver because interrupt coalescing tainted our measurements.

We did attempt to use PlanetLab [4] for our experiments, but we discovered that high machine load interfered with our measurements.

Overall we made about 56 measurements over twelve months. Since most measurements consist of 24 one-hour segments, we have close to 1300 hours of measurements. Figure 7 shows a small sample of these measurements, one graph for each path we measured. In each graph, we show several dispersion CDFs taken one or more months apart. Within each path, all CDFs are from the same time of day. We do not have measurements for all paths for all months due to unstable hosts.

4.2 Understanding Internet Dispersion CDFs

The data in Figure 7 is taken from the "wild" Internet, thus we do not have ground truth about the link and traffic characteristics. We therefore turn to the understanding of small scale network effects on dispersion CDFs developed in Section 3 to interpret our wide-area data.

The captions of Figure 7 identify the topology class (S, F, SF, FS) from Section 3 we believe most closely matches the dispersion CDF. We omit classes SS and FF since they are subsumed by classes S and F, respectively.

To map between our experimental data and the basic analysis/simulation in Section 3 we must make some assumptions about link speeds and packet size distributions in the Internet. The dispersion CDFs show that all links on our paths have at least 100 Mb/s capacity. Prior studies of distributions of Internet packet sizes have consistently shown strong modes at 1500 and 40 bytes, with occasional weaker modes at 576 or 1200 bytes [23, 28, 29, 32]. In Section 3.4.3 we demonstrated through simulation that bimodal traffic results in slightly rounder steps than unimodal traffic. We consider that effect when relating topology classes to wide-area data. We also label each based on how well it matches our expectations: "expected modes", "unexpected modes", or "no modes". We describe these next.

Expected modes. The Jan 2005 CDF from path DE-SB and all CDFs from paths SB-KR and SC-KR directly match our expectations of a link with cross-traffic (classes S and FS) and point-constrained operation. The strongest modes are at close-to-integer multiples of initial dispersion, consistent with mostly 1500 byte cross traffic packets on a 100 Mb/s link. Closer inspection shows smaller steps, consistent with bimodal cross traffic that includes 40 byte packets.

Unexpected modes. Many of the CDFs show strong modes, but not at values consistent with 100 Mb/s links. We characterize these CDFs (MD-SC, UM-SC, KR-SC, SB-UM, SC-UM, DE-UM, GR-SC and DE-SC) as having *unexpected* modes. We can break these into three groups with similar features. The first group (MD-SC, UM-SC, KR-SC, and DE-SC) all terminate at USC, and show strong modes at about 1.6 to 1.65× the initial dispersion (final δ of 190–200 μs). The second group all terminate at UMass (SB-UM, SC-UM, DE-UM), and all show modes at 1.3 to 1.4× initial dispersion (final spacing of 155–170 μs). Figure 2(a) shows the raw distribution of a portion of a DE-UM measurement, showing the very consistent output dispersion at certain multiples. Finally the third group contains the singleton GR–SC, with modes at about 1.2 and 2.3 times initial dispersion.
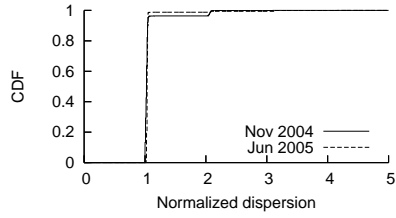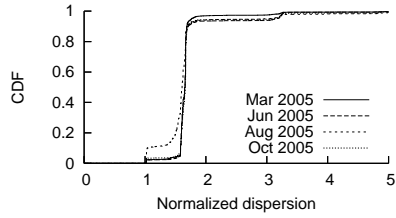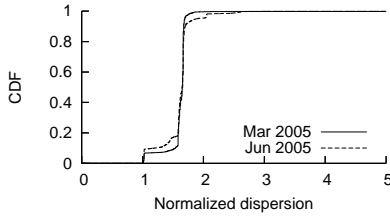
(a) UCSB to Korea (SB–KR): S, expected modes.
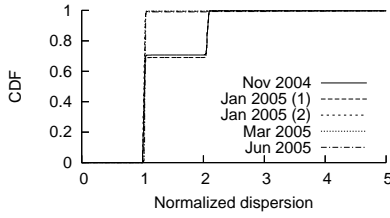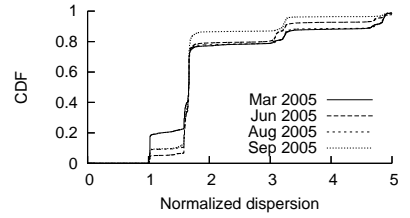
(b) USC to Korea (SC–KR): S, expected modes.
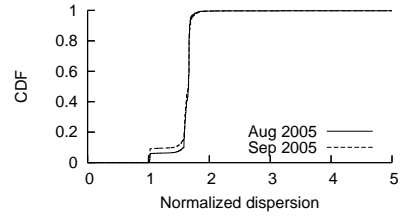
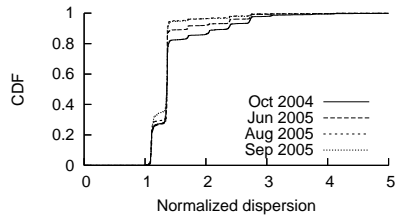(c) UMd to USC (MD–SC): S, unexpected modes.
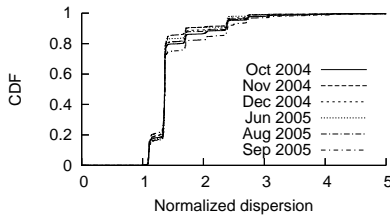
(d) UMass to USC (UM–SC): S, unexpected modes.

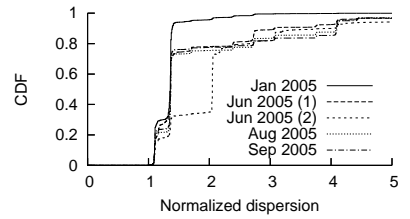(e) Korea to USC (KR–SC): S, unexpected modes.
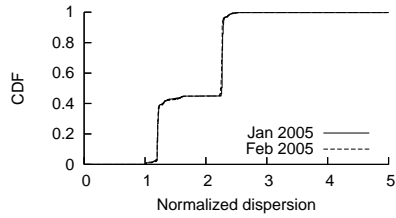
(f) Germany to USC (DE–SC): S, unexpected modes.

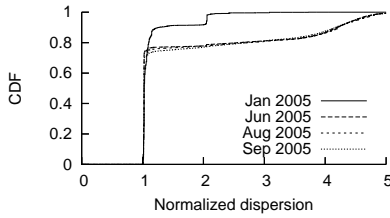(g) UCSB to UMass (SB–UM): S, unexpected modes.

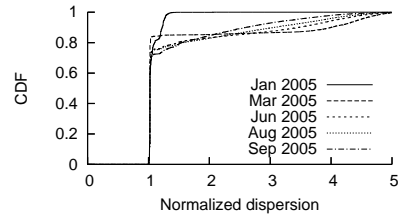(h) USC to UMass (SC–UM): S, unexpected modes.
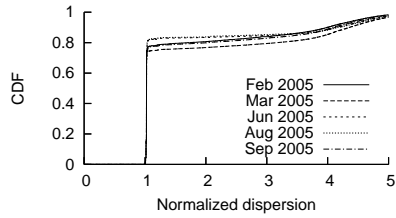
(i) Germany to UMass (DE–UM): S, unexpected modes.

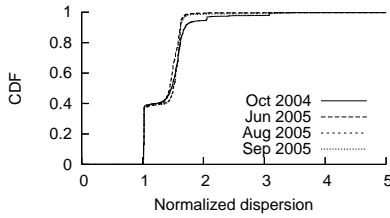(j) Greece to USC (GR–SC): S, unexpected modes.
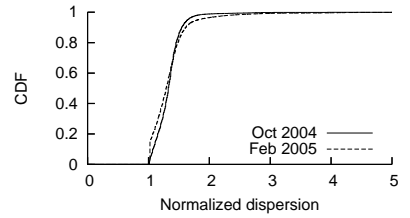
(k) Germany to UCSB (DE–SB): FS/SF, no modes.

(l) UCSD to Germany (SD–DE): F/SF, no modes.

(m) USC to Germany (SC–DE): SF, no modes.

(n) USC to UCSB (SC–SB): SF, no modes.

(o) UMass to UMd (UM–MD): SF, no modes.

Figure 7: Dispersion CDFs from Internet experiments.

Since we do not have ground truth for these paths we cannot fully explain them. However, our simulations provide some basis for a hypothesis. Taking the UMass traces as an example (specifically the SC–UM trace), we first observe that about 60% of pairs show dispersions around 170 μs or slightly less. This spacing is consistent with a half allocation of an OC-3 link (155 Mb/s halves to 77 Mb/s, consistent with 17 μs spacing); we believe this link represents the path bottleneck. The third strongest mode (about 10% of packets) is at twice this value ($2.4\delta_{init}$ or 34 μs), consistent with the pairs trapping a 1500-byte packet on this bottleneck.

Several remaining modes for this path occur at 1.1, 1.7, 2.05 and 2.75 times δ_{init}. In fact, the mode at $1.1\delta_{init}$ is the second strongest mode with about 18% of pairs. These weaker modes are each about $\pm0.3\delta_{init}$ off of the strongest mode and its double. These modes are consistent with traffic passing through a 300 Mb/s link.

The group of paths terminating at USC shows a primary mode around $1.65\delta_{init}$ (198 μs), with about 80% of the pairs making a smooth, S-shape between 1.6 and $1.7\delta_{init}$. The main secondary mode at about $1.02\delta_{init}$. The primary mode is consistent with traffic passing through a bottleneck link around 62 Mb/s. We believe the secondary mode is caused by compression: the pairs gain a large amount of slack from the bottleneck link, then a few queue behind a router at a 100 Mb/s link.

While these interpretations are consistent with our analysis, the non-standard bottleneck speeds imply that we do not yet have a complete understanding of this network. We are working on obtaining ground truth on these paths.

No modes. We consider Jun–Sep 2005 CDFs from path DE–SB, and all CDFs from paths SD–DE, SC–DE, SC–SB and UM–MD to have "no modes", with relatively smooth CDFs and no clear modes (except possibly at 1, which represents unchanged final dispersion.) We believe that the continuous shape of these CDFs is caused by cross traffic at a fast link, consistent with case SF in Figure 3.4.2. Our hypothesis is that a busy high-speed link (much faster than 100 Mb/s) causes operation in the unconstrained region in this class, and that any 100 Mb/s links causing significant dispersion are upstream of this fast link.

We now move on to overall observations drawn from the ensemble of dispersion CDFs.

4.3 Observation 1: Persistence

The first observation from Figure 7 is that the dispersion CDFs remain fairly persistent over periods of months, as can be seen by very similar modes and shapes across CDFs that are months apart. For example, path SC–DE (Figure 7(m)) had consistent shape (80% of packets have their initial dispersion and the rest have a smooth range of other values) for seven months.

Explanation: Persistence of dispersion CDFs indicates relatively stable underlying traffic conditions. As we have seen in Section 3, background traffic intensity mainly affects the magnitude of the CDF (the y-axis) and packet size distribution affects the location of the modes. The stability in our results indicates that neither of those changes drastically in timescales of months.

Caveats: For a few paths (DE–SB, SD–DE, and SC–DE) we see a diurnal cycle as shown in Figure 8. This is well known phenomenon and our fingerprints are able to catch it. We expand on these diurnal changes in Section 4.5 and
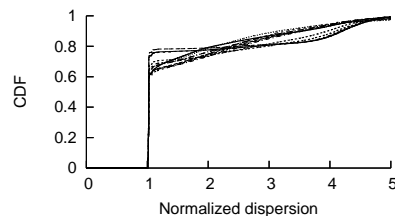

Figure 8: Dispersion CDFs every two hours for path DE–SB on Friday, June 24, 2005.

quantify their differences in Section 6.2.

4.4 Observation 2: Distinctness

The dispersion CDFs in Figure 7 show another interesting property: often signatures for different paths are quite distinct. For example, DE–SC, MD–SC and GR–SC, all of which terminate at the same destination, and DE–SC and SC–DE are quite different. The difference between DE–SC and SC–DE is not surprising as they do not have any traceroute hops in common.

Several groups, however, show strong similarities. For example, MD–SC, UM–SC, KR–SC, DE-SC form one class (terminating at USC); SB–UM, SC–UM, DE–UM form a second class (terminating at UMass); and DE–SB (omitting Jan. 2005), SD–DE (omitting Jan. 2005), SC–DE form a third class. We believe these similarities indicate that the paths share common busy links. We can confirm this with traceroute. We see four common hops at the destination for the first group and four for the second. The third group has paths in both directions involving Germany so we see common hops (at least 9) only on the two forward paths. We speculate that common hops exist on the reverse path also, and may be detectable by alias resolution.

Explanation: Dispersion CDFs are influenced by link speeds and cross traffic. Paths with different link speeds and traffic generate different dispersion CDFs, while paths that share many links (particularly bottleneck or busy links) generate similar CDFs.

Caveat: In some cases we have observed fingerprint changes that we associate with traffic changes (rather than routing changes). We examine these next.

4.5 Characterizing Fingerprint Changes

Our statements about persistence and distinctness are based on the claim that dispersion CDFs reflect underlying link and traffic characteristics. Thus, it is not surprising that significant routing or traffic changes can result in corresponding changes in the dispersion CDFs. We have seen three kinds of changes: on three paths we observe diurnal changes of fingerprints, on three cases we observed significant changes between measurements, and in two cases we captured significant changes during our 24 hour observations. We describe each of these below, focusing on the two cases that we captured in action.

Although most paths are very stable regardless of observation time, we observe diurnal changes on on paths DE–SB, SD–DE and SC–DE. Figure 8 shows 12 measurements for DE–SB. This change in CDF is consistent with a change in cross traffic volume consistent with our simulation cases F and SF. In Section 6.2 we quantify these changes. The presence of diurnal effects suggests that, if dispersion CDFs are
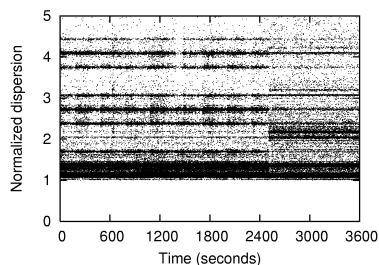
Figure 9: One-hour dispersion time series on path DE-UM.

used to identify paths, than they must either be taken over the entire day (to smooth out traffic variation), or taken at a consistent time of day. The plots in Figure 7 confirm long-duration persistence once diurnal effects are factored out by measuring at a consistent time of day.

We saw three cases where there were significant changes in the month between our observations. The three paths are DE–SB, SD–DE and DE–UM (Figures 7(k), 7(l) and 7(i)), all involving a source or destination in Germany, with the change occurring between January 2005 and later traces. (Our one other path involving Germany (DE–SC) does not show this change because we lack data for January 2005.) Examination of traceroutes indicate that this change in CDF corresponds to a change in path: an extra hop appears in all of the January paths. In January, each path showed some level of traffic, while later traces show more traffic but a smooth CDF. This change is consistent with an increase in bandwidth of some intermediate link with some utilization (for example, perhaps a change in the LAN used at a peering point).

Of the two major changes *during* an observation, the first was on path SC–KR, during a 24-hour observation in January 2005. Dispersions CDFs before the change, from November 2004 and January 2005 are similar, and have a large mode at $2\delta_{init}$. CDFs after the change are also consistent with each other, with 99% of dispersions nearly unchanged (within δ_{init} and $1.06\delta_{init}$). During January 2005 we observed a 2.5 hour outage, with a new signature after the outage. The route reported by traceroute did not change (but both routes show four anonymous hops). The outage may indicate a hardware change and the change in CDF indicates that fewer pairs experienced large increases in dispersion after the change. This change in CDF suggests presence of traffic before and the reduction of traffic afterwords, or perhaps an increase in link capacity.

Our last example of a major signature change *during* an observation was on path DE–UM in June 2005. We observe a stable path (call it R1), followed by an alternate route (R2) on Thu June 23 2005 between hours 12 and 13 of the 24-hour observation period, reverting to the original route after hour 13. The R1–R2 change keeps the same dispersion CDF, but the return of route R1 (after hour 13) is concurrent with a very different dispersion CDF than before hour 12 (see Figure 9 around time 2500 s). This condition persists through the rest of June, but is gone in our August and later measurements. This case is very puzzling, since the route change does not correspond to the CDF change. Considering Figure 9, we see that the modes before and after are similar, but the strengths of each band vary greatly. Perhaps this

indicates that the outage corresponds to a change in the peering point that results in changed cross traffic.

Finally, we also observed the converse of these examples: long-term stable paths and stable dispersion CDFs. For example, for path SC–UM, the dispersion CDFs are quite stable. The router-level path (from traceroute) shows some variation over that time (16 to 18 hops), but its structure remains unchanged with five ISPs (USC, CENIC, Internet2, Northern Crossroads, and UMass) for the twelve months of observation. We observed a similar level of stability for other paths not mentioned above.

These examples indicate that that significant changes in network are often associated with changes in the dispersion CDF, but that long-term measurements indicate generally stable and distinct dispersion CDFs. Based on these observations, we propose to use the dispersion CDF as a path signature. In the next sections we briefly review other factors that might affect dispersion CDFs (Section 5) and then turn to applications in Section 6.

5. SENSITIVITY ANALYSIS

Our approach to collection dispersion CDFs requires a number of measurements. We next evaluate the sensitivity of these measurements to measurement parameters, such as probe rate and measurement duration, and to environmental factors such as system load and clock drift. We first validate that our experimental results are not biased by high-speed academic networks, then we look at varying measurement parameters on a specific path. (We use the SC–UM path for these measurements, chosen because we have the longest observations for this path.)

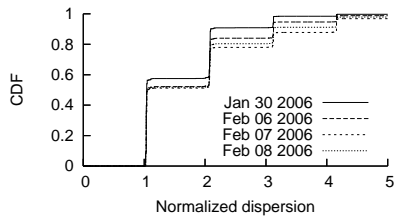### 5.1 Paths in the Commercial Internet

The dispersion measurements presented in Section 4 begin and end at hosts in academic networks, and so the paths often exploit research networks such as Internet2. Internet2 is designed to be very high speed, and perhaps business concerns force commercial links to be run at higher utilizations. To investigate if our observations hold on commercial networks as well as academic networks we took took short-term measurements from a host placed in in the commercial address space of Los Nettos, a regional network, labeled here "LN".

We measured paths UM–LN, MD–LN, SB–LN, DE–LN and SD–LN, and the reverse path LN–UM. We verified that each of these paths goes over non-Internet2, commercial links. We do not reproduce all this data here, but in Figure 10 we show dispersion CDFs for paths DE–LN, and LN–UM.

During a week of observation, we found that the dispersion CDFs on these paths had similar behavior as those on the academic paths—CDFs are persistent and in general different on different paths. These experiments serve as preliminary verification of dispersion CDF properties on commercial paths.

5.2 Effect of Packet Pair Rate

Throughout this paper we use the relatively high probe rate of 100 pairs/s (with exponential interarrivals); a data rate of 2.4 Mb/s. We chose this rate to get good accuracy without great overhead relative to the 100 Mb/s link capacity or more. However, a lower probe rate would be attractive to reduce the overhead. Here we test the affect of
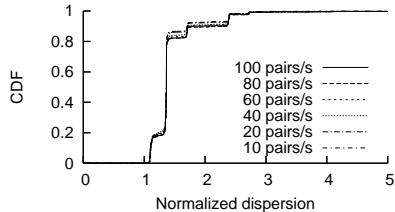
(a) Germany to Los Nettos (DE–LN): S, expected modes.
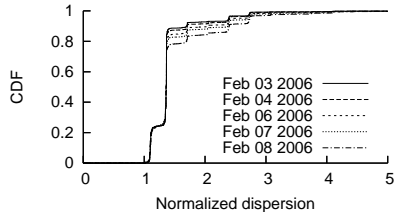
(b) Los Nettos to UMass (LN–UM): S, unexpected modes.

Figure 10: Dispersion CDFs from verification experiments on commercial Internet paths.


Figure 11: Dispersion CDFs with various packet pair rates over the SC-UM path.

lower probe rates on the dispersion CDF.

We computed dispersion CDFs for average probe rates of 100, 80, 60, 40, 20 and 10 pairs/s on the SC–UM path (Figure 11). We found that dispersion CDFs look reasonably similar for most packet-pair rates. When we compared them with the approach described in Section 6.1, finding the largest difference of 0.0023, well less than our threshold indicating difference. Thus, we conclude that our measurement process is not particularly sensitive to packet-pair rate, which can be reduced by a factor of 10 and still produce acceptable results. We plan to more systematically look at reduced probe rates across all paths.

5.3 Effect of Measurement Duration
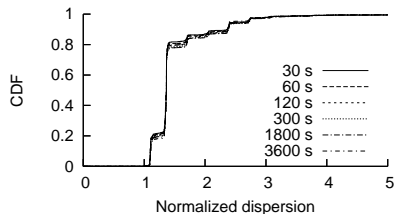
Figure 12: The effect of varying measurement duration on the SC-UM path.

In Section 4 we presented the dispersion CDFs of paths with measurements that lasted one hour (3600 s). An important question is how long should we measure a path before we can get a reliable representation of its fingerprint? In Figure 12 we show the CDFs produced from data for the first 30, 60, 120, 300 and 1800 s of the hour-long experiment for the path SC–UM, sampled at 100 pairs/s. We find that the CDFs change very little for different measurement duration, which shows that the path fingerprint is quite stable even at shorter time scales (seconds vs. minutes or hours).

5.4 Other Parameters

We ran experiments to check the sensitivity of our results to several additional factors.

Different end-hosts: To verify that the fingerprints are not dependent on the destination host we carried out experiments in sites where we had access to different source and destination machines on the same subnet. Our results showed that the fingerprints were virtually the same and did not depend on the specific destination machine.

System load: We discovered that high system CPU utilization *can* affect dispersion CDFs by distorting the spacing of probes at the sender (by changing δ_{init}) and the accuracy of timestamps at the receiver (by changing δ_{final}). During our experiments we periodically examine the CPU utilization to ensure that the load is acceptable during our experiments. An area of future work is to understand if sender-side kernel-level timestamping can permit accurate dispersion CDFs from hosts under load.

6. FINGERPRINT COMPARISON AND APPLICATIONS

In previous sections we have shown that the dispersion CDFs of different paths are visually distinct. In this section we explore two applications of dispersion CDFs: identifying paths that share common components, and identifying changes in traffic on a given path.

Our technique relies on active probing, which raises questions about scalability to many concurrent users if it is widely adopted. However, we expect each user to have a relatively light impact on the network for the following reasons. First, collection of path fingerprints does not involve heavy measurement traffic for most paths; as we have shown in Section 5 packet pair rates as low as 10 pair/s for durations as low as 30 s could be sufficient. Second, we have shown that path fingerprints are persistent and thus reusable, so users do not need to measure paths frequently. Finally, path stability allows multiple users at the same site to share path fingerprints, should user density be high enough to warrant development of such infrastructure. Thus, we believe that widespread use of these tools is feasible.

We next describe the applications in detail. We begin by defining an approach to quantify differences in dispersion CDFs.

6.1 Comparing Dispersion CDFs

To compare CDFs, we measure the area enclosed between two normalized dispersion CDFs (Figure 13) and compare to a threshold. We adopted this approach because of its simplicity; as future work we plan to investigate alternate techniques such as by Bin Tariq et al. [31].

To compute areas we must first normalize the x-axis by
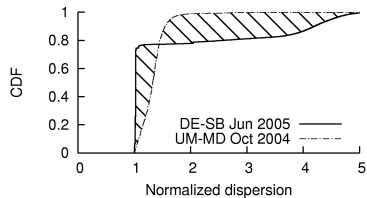
Figure 13: The area enclosed by two CDFs is used to quantify difference between them. The normalized enclosed area is 0.0278 for this example.

scaling it across a fixed range from 0 to some maximum value $m\delta_{init}$ (we assume both measurements have the same δ_{init}). We then compute the area between the CDFs, normalized by the maximum area that can be enclosed between any two CDFs. More specifically, if D_X and D_Y are two dispersion CDFs to be compared, the comparison metric $C(P_X, P_Y)$ between them is:

$$C(D_X, D_Y) = \frac{1}{\lceil m/b \rceil} \sum_{i=1}^{\lceil m/b \rceil} |D_X(i) - D_Y(i)|,$$

where b is the bin size for the discrete CDF. For our work below we select $m = 25$ to capture the vast majority of dispersion values. This choice of m was guided by the maximum normalized dispersion we observed, which ranged from 4.16 to 24.99 for the CDFs shown in Figure 7. This metric defines a value between 0 and 1, with higher values indicating greater difference.

As examples of the comparison metric, we consider path DE–SB $C(\text{Jan.'05}, \text{Jun.'05}) = 0.02274$, a large change; and $C(\text{Jun.'05}, \text{Aug.'05}) = 0.0007$, a small change.

Choosing a threshold: Given the above metric we can compare two dispersion CDFs and quantify their difference, but it is not clear how large a numeric difference is truly "different".

To select an initial threshold we examined our Internet data, considering over 13,000 pairs of dispersion CDFs, where each pair of CDFs (D_X, D_Y) corresponds to two different hours in the same 24-hour period for the same path. Assuming these represent very similar paths (based on our observations in Section 4.3 that dispersion CDFs are quite stable), we measure the 95th percentile of $C(D_X, D_Y)$ for all pairs. We use the result $C_{95\%} = 0.0105$ as a threshold, while C values beneath it considered similar, and above it different. Thus, we suggest that paths X and Y with dispersion CDFs D_X and D_Y are likely to have common links if $C(D_X, D_Y) < C_{95\%}$.

6.2 Fingerprint Variations for a Single Path

We now show how our comparison metric D can be used to monitor a path for changes in traffic or link characteristics. In Section 4 we argued that significant changes in these characteristics are reflected in changes in their dispersion CDFs.

To illustrate this application we consider for three sample paths: (a) a stable path, (b) with a sudden change in the fingerprint during measurement, and (c) with a diurnal pattern. For each path, we take dispersion measurements for 24

hours, and split this dataset into hour-long segments. We then compute $C(D_i, D_j)$ for all 276 combination of hours i and j for that path.

Figure 14 shows graphical representations of each type of path, each with a separate gray-scale showing the range of values of $C(D_i, D_j)$. (Note that each graph has a different range of gray-scale.) Hour labels are time since measurement beginning, not time of day.

Figure 14(a) shows a very stable path (SC–KR), which shows low variation over the course of a day. We see no particular pattern in the similarity or difference of different hour-long measurements. The maximum difference between any two hours is $C(16, 6) = 0.000612$, a very low value compared to the 95th percentile stated above.

In contrast, Figure 14(b) shows a path (DE–UM) that experienced a sudden change in hour 13, with otherwise consistent behavior before and after. Visually, the event is clearly represented as a bright rectangular region in the graph, and the maximum comparison is $C(x, y) = 0.035$, fifty times larger than the stable graph and well over the 95th percentile of 0.0105. We conclude that our comparison approach clearly captures this change, allowing quantitative comparison of the CDFs shown in Figure 7(i).

Finally, Figure 14(c) shows a path (DE–SB) where a moderate diurnal cycle is visible. Visually, the cycle is represented as a progression of bright/dark regions. The maximum difference in the area metric is 0.014, a twenty-fold increase over the stable example and slightly over the 95th percentile. Diurnal variation of the CDF is not common in our measurements; we found it exists on three paths (DE–SB, DE–SD and SC–DE). The variation is sufficiently large that the two extreme CDF shapes are from virtually different paths with respect to busy link characteristics, as confirmed by the large values of C in Figure 14(c). On most paths, variations are short-lived and aperiodic.

6.3 Detecting Busy Links Common to Two Paths

An important application of our fingerprints is detecting when two paths share common busy links. Detection requires data collection and then comparison. First, we produce fingerprints for paths X and Y to be compared, measuring the paths over a short period of time. (In Section 5 we determined that measuring the path for as little as 30 s is sufficient to get a reliable fingerprint.) Then we use the comparison test with the threshold defined above to determine whether the difference $C(D_X, D_Y)$ is statistically "large" or "small." (Recall that in Section 4.4 we observed that dispersion CDFs tend to be distinct.) If the difference is small, the two paths are likely to have shared busy links.

To demonstrate this approach we compared all one-hour segments of paths SC–UM and SB–UM for a common 24 hour period, computing $C(D_{\text{SC–UM},i}, D_{\text{SB–UM},j})$, for each hour i and j. We found that $C(D_{\text{SC–UM},i}, D_{\text{SB–UM},j})$, ranged from 0.002 to 0.008 over these 576 comparisons. The maximum value is less than our $C_{95\%}$ threshold, implying that all combinations are similar. In fact, this range corresponds to the 67th through 91st percentile of our empirical distribution from Section 6.1. Thus we conclude any busy links are likely shared between these paths.

This evaluation is only preliminary; more work is needed to demonstrate we can correctly find the presence and absence of common busy links. Our method is attractive if it withstands more detailed verification, because it promises

(a) Stable dispersion CDF on one path (SC-KR).

(b) Sudden change in dispersion CDF one one path (DE-UM).

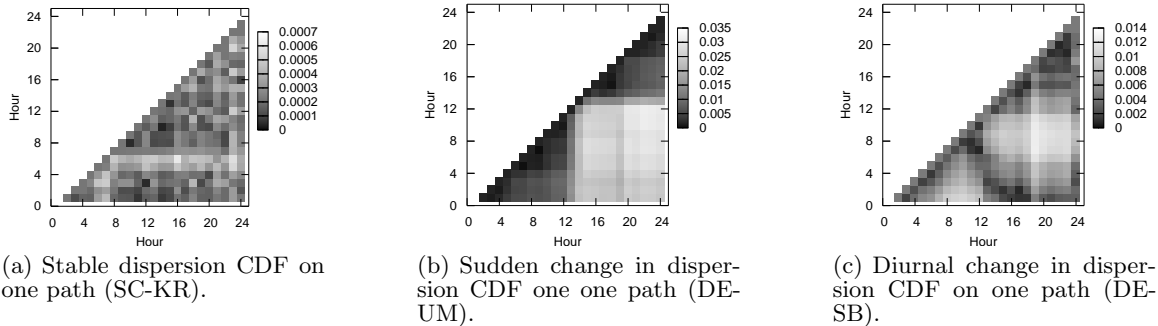(c) Diurnal change in dispersion CDF on one path (DE-SB).

Figure 14: CDF comparisons within one path.

several advantages over current methods of shared-link detection (see Section 7 for a description of existing methods). First, we do not need to measure the paths at same time. As we showed earlier, our path fingerprints are stable over long periods of time, so they can be reused long after they were generated. Second, we do not require the presence of a common congested link between paths. Our fingerprints can detect sharing even if the links are not highly utilized. Finally, our fingerprints do not rely on strongly correlated events such as shared packet losses or delay patterns, which imply both the need for synchronous measurement and the presence of a common congested link.

7. RELATED WORK

Our work builds on prior work to identify link or path characteristics, and some work to identify shared links.

Prior work in Internet path characterization includes Paxson's comprehensive study of delays, loss rates, bottleneck link capacities, normalized available bandwidth and the time-scales for queue length variation [26]. Zhang et al. [35] and Bolot [7] measured path loss and delay characteristics and the PingER project [3] uses RTT and loss measurements to characterize paths on a global scale. Pathchar and other related tools make point estimates of per-hop properties such as capacity and delay [14]. Our goal of fingerprinting differs from path characterization in *identifying* a path, rather than quantifying specific path characteristics.

The packet-pair technique was first used by Keshav [18] to estimate available bandwidth in the context of flow control. We build on this fundamental work. In recent stochastic analysis of multi-hop dispersion, Liu et al. [22] attempt to generalize the treatment of packet pairs. We complement this type of work with case-by-case analysis and network measurements.

Bolot [7] and the Bprobe [9] tool were among the first to use the principle of inter-packet dispersion to measure narrow link capacity. Paxson [26] introduced measurements with active receivers, eliminating errors due to the return path. Nettimer [20] improved estimation by filtering measurement noise using statistical methods. The filtering mechanism used by Pathrate [10] uses dispersion measurements from both packet pairs and packet trains. For available bandwith estimation, TOPP [24], IGI [13] and Pathload [15] use packet pair or packet train methods that essentially search for the data rate at which the path is saturated. Spruce [30] assumes that the tight and narrow links are the same and of known capacity, and uses Poisson sampling to

ensure it measures the average cross traffic rate after it forces queuing on the tight link. Our work differs from this prior work by using packet pairs to fingerprint paths rather than to estimate available bandwidth or capacity.

Pásztor et al. use dispersion distributions as a source of signatures [25], as we do, but their work focuses on isolation of signatures of particular links to detect bottlenecks, while we develop different signatures that identify the path. The MultiQ tool [17] also analyzes inter-arrival time distributions, to find capacity bottlenecks on a path. We use the distribution to characterize the entire path rather than just specific links or bottlenecks.

Prior work in detecting link sharing focuses on congested links. Rubenstein et al. [27] proposed delay and loss correlation techniques for detection of shared congestion. Kim et al. [19] improved the technique by adding wavelet denoising to remove dependence on common endpoints and relax the synchronization requirement to 1 s. Katabi et al. [16] used entropy minimization between inter-arrival time distributions when the bottleneck link is saturated. The technique by Harfoush et al. [11] correlates packet losses to detect shared congestion. Unlike this work, our method of detecting busy link sharing does not assume losses, congestion or simultaneous probing.

Finally, Bin Tariq et al. [31] use the Kullback-Leibler distance to compare CDFs. That approach is likely more statistically rigerous than our proposed comparision metric and a good direction for future work.

8. CONCLUSIONS

We have introduced the dispersion CDF as a new way to fingerprint Internet paths, and shown that it can persist for months, and distinctly identify unique paths. Applications include detecting shared busy links between two paths and detection of traffic changes on a path. In the future, we plan to make longer-duration measurements on commercial paths and expand simulations and experiments to better understand what affects dispersion CDFs.

9. ACKNOWLEDGMENTS

10. REFERENCES

[1] National laboratory for applied network research: Passive measurement and analysis. `http://pma.nlanr.net`.

[2] Nlanr traffic traces for september 13, 2005. `http://pma.nlanr.net/Traces/Traces/daily/20050913/`.

[3] The PingER project. `http://www-iepm.slac.stanford.edu/pinger/`.

[4] Planetlab: An open platform for developing, deploying, and accessing planetary-scale services. `http://www.planet-lab.org`.

[5] A. Akella, S. Seshan, and A. Shaikh. An empirical evaluation of wide-area Internet bottlenecks. In *Proc. of SIGMETRICS*, pages 316–317. ACM, 2003.

[6] M. Allman, W. M. Eddy, and S. Ostermann. Estimating loss rates with tcp. *SIGMETRICS Perform. Eval. Rev.*, 31(3):12–24, 2003.

[7] J.-C. Bolot. End-to-end packet delay and loss behavior in the Internet. In *Proc. of SIGCOMM*, pages 289–298. ACM, 1993.

[8] L. Breslau, D. Estrin, K. Fall, S. Floyd, J. Heidemann, A. Helmy, P. Huang, S. McCanne, K. Varadhan, Y. Xu, and H. Yu. Advances in network simulation. *IEEE Computer*, 33(5):59–67, May 2000.

[9] R. L. Carter and M. E. Crovella. Measuring bottleneck link speed in packet-switched networks. *Perform. Eval.*, 27-28:297–318, 1996.

[10] C. Dovrolis, P. Ramanathan, and D. Moore. What do packet dispersion techniques measure? In *Proc. of INFOCOM*, pages 905–914. IEEE, 2001.

[11] K. Harfoush, A. Bestavros, and J. Byers. Robust identification of shared losses using end-to-end unicast probes. In *Proc. of the Eighth IEEE Inter. Conf. on Network Protocols*. IEEE, 2000.

[12] N. Hu, L. E. Li, Z. M. Mao, P. Steenkiste, and J. Wang. Locating Internet bottlenecks: algorithms, measurements, and implications. In *Proc. of SIGCOMM*, pages 41–54. ACM, 2004.

[13] N. Hu and P. Steenkiste. Evaluation and characterization of available bandwidth probing techniques. *IEEE J. of Selected Areas in Communication*, 21(6):879–894, Aug. 2003.

[14] V. Jacobson. Pathchar; a tool to infer characteristics of Internet paths. `ftp://ftp.ee.lbl.gov/pathchar/`.

[15] M. Jain and C. Dovrolis. End-to-end available bandwidth: measurement methodology, dynamics, and relation with tcp throughput. In *Proc. of SIGCOMM*, pages 295–308. ACM, 2002.

[16] D. Katabi, I. Bazzi, and X. Yang. A passive approach for detecting shared bottlenecks. In *Proc. of the Tenth Inter. Conf. on Computer Communications and Networks*, pages 174–181, Oct. 2001.

[17] S. Katti, D. Katabi, C. Blake, E. Kohler, and J. Strauss. Multiq: automated detection of multiple bottleneck capacities along a path. In *Proc. of Internet Measurements Conference*, pages 245–250. ACM, 2004.

[18] S. Keshav. A control-theoretic approach to flow control. In *Proc. of SIGCOMM*, pages 3–15. ACM, 1991.

[19] M. S. Kim, T. Kim, Y. Shin, S. S. Lam, and E. J. Powers. A wavelet-based approach to detect shared congestion. In *Proc. of SIGCOMM*, pages 293–306. ACM, 2004.

[20] K. Lai and M. Baker. Measuring bandwidth. In *Proc. of INFOCOM*. IEEE, 1999.

[21] X. Liu, K. Ravindran, B. Liu, and D. Loguinov. Single-hop probing asymptotics in available bandwidth estimation: sample-path analysis. In *IMC '04: Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*, pages 300–313. ACM, 2004.

[22] X. Liu, K. Ravindran, and D. Loguinov. Multi-hop probing asymptotics in available bandwidth estimation: Stochastic analysis. In *IMC '05: Proceedings of the 2005 Internet Measurement Conference*, 2005.

[23] S. McCreary and k claffy. Trends in wide area IP traffic patterns: A view from Ames Internet Exchange. In *Proc. of 13th ITC Specialist Seminar on Measurement and Modeling of IP Traffic,*, 2000.

[24] B. Melander, M. Björkman, and P. Gunningberg. A new end-to-end probing and analysis method for estimating bandwidth bottlenecks. In *Proc. of GLOBECOM*, pages 415–420. IEEE, 2000.

[25] A. Pásztor and D. Veitch. The packet size dependence of packet pair like methods. In *Proc. of the IEEE/IFIP Inter. Workshop on Quality of Service*. IEEE, 2002.

[26] V. Paxson. End-to-end Internet packet dynamics. In *Proc. of SIGCOMM*, pages 139–152. ACM, 1997.

[27] D. Rubenstein, J. Kurose, and D. Towsley. Detecting shared congestion of flows via end-to-end measurement. In *Proc. of SIGMETRICS*, pages 145–155, New York, NY, USA, 2000. ACM.

[28] R. Sinha, C. Papadopoulos, and J. Heidemann. Internet packet size distributions: Some observations. `http://netweb.usc.edu/~rsinha/pkt-sizes/`.

[29] Sprint. Sprint IPMON DMS, packet trace analysis. `http://ipmon.sprint.com/packstat/packetoverview.php`.

[30] J. Strauss, D. Katabi, and F. Kaashoek. A measurement study of available bandwidth estimation tools. In *Proc. of Internet Measurements Conference*, pages 39–44. ACM, 2003.

[31] M. M. B. Tariq, A. Dhamdhere, C. Dovrolis, and M. Ammar. Poisson versus periodic path probing (or, does PASTA matter?). In *Proc. of Internet Measurements Conference*, pages 119–124, Berkeley, CA, USA, Oct. 2005.

[32] K. Thompson, G. Miller, and R. Wilder. Wide-area Internet traffic patterns and characteristics. *IEEE Network*, pages 10–23, 1997.

[33] R. Wolff. Poisson arrivals see time averages. *Operations Research*, 30(2):223–231, 1982.

[34] L. Zhang, S. Shenker, and D. D. Clark. Observations on the dynamics of a congestion control algorithm: the effects of two-way traffic. In *Proc. of ACM*

SIGCOMM Conference, pages 133–147, Zurich, Switzerland, Sept. 1990. ACM.

[35] Y. Zhang and N. Duffield. On the constancy of Internet path properties. In *Proc. of the Internet Measurements Workshop*, pages 197–211. ACM, 2001.

APPENDIX

A. PACKET SIZE DISTRIBUTION OF INTERNET TRAFFIC

In measurements of packet size distributions in Internet traffic [28], we observed some shifts in packet sizes compared to previously observed results [23, 29, 32]. We have two surprising observations.

First, current packet sizes seem mostly bimodal at 40 bytes and 1500 bytes (at approximately 40% and 20% of packets, respectively). This observation represents a change from earlier measurements such as the pre-2000 data that reports tri-modal packet sizes around 40, 576, and 1500 bytes.

Second, in some cases we observe a strong mode around 1300 bytes. This represents a new phenomenon.

The first observation holds across all measurements at 5 different network points, including Los Nettos (our regional ISP, carrying a mix of academic and commercial traffic), a USC Internet2 connection, and three connections monitored by NLANR [1]. The second observation does not hold universally, but is very strong at Los Nettos and USC Internet2, and is noticeable in all traces.

Figure 15 shows packet size distributions computed from tcpdump traces collected at Los Nettos and USC's Internet2 link (Figure 15(a)) and from publicly available traces collected by NLANR (Figure 15(b)). The Los Nettos traces are from October 10, 2005 and the USC Internet2 traces are from December 2, 2004. The NLANR traces are from September 13, 2005 and were obtained from the Web repository for this date [2].

By "packet size" we mean the byte-count in the length field of the IP header.

The shift away from 576-byte packets is not suprising, since it is consistent with evolution of operating systems and widespread use of Ethernet with a 1500-byte MTU.

The growth at 1300-byte packets (seen at Los Nettos and the USC Internet2 link) was suprising to us. We have tentatively identified 1300-byte packets as stemming from widespread use of VPN software, and possibly from recommendations from DSL providers.

Our observations do not point to wide use of end-to-end VPN over WANs, but to VPN use at the edge network, since the 1300 byte size noted is presumably that of packets that have exited a VPN tunnel. This edge-network use is certainly true at USC, where most wireless traffic traverses a VPN over the wireless hop and then proceeds unencrypted over the rest of the Internet. This behavior explains why Los Nettos and USC Internet2 traffic show the strongest 1300-byte modes of the sites we observe.
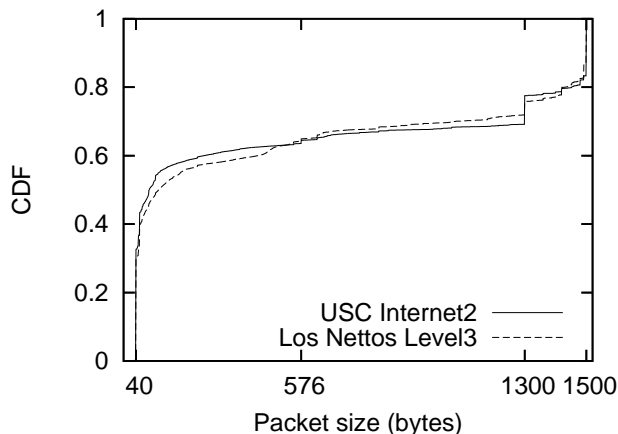
(a) CDFs from five-minute USC Internet2 and 43-second Los Nettos Level3 traces.
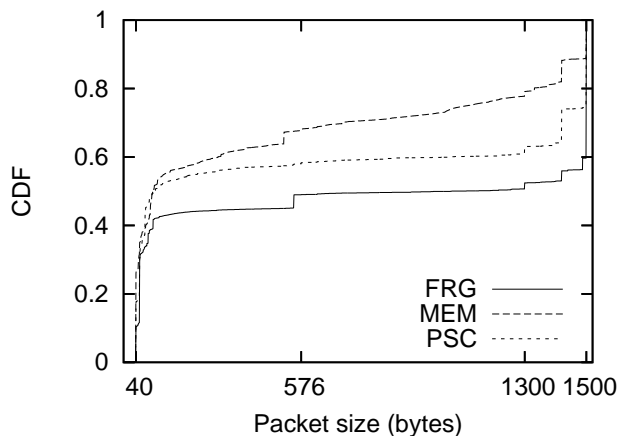
(b) CDFs from 90-second traces from NLANR sites FRG (Front Range GigaPOP), MEM (University of Memphis) and PSC (Pittsburgh Supercomputing Center).

Figure 15: Packet size distributions in Internet traffic.