Quantifying the Effects of Routing Dynamics on End-to-End Internet Path Failures

Feng Wang
Department of Electrical and
Computer Engineering
Univ. of Mass., Amherst
Amherst, MA 01002, USA
fewang@ecs.umass.edu

Nick Feamster
Networks and Mobile Systems
Massachusetts Inst. of Tech.
Cambridge, MA 02139
feamster@lcs.mit.edu

Lixin Gao
Department of Electrical and
Computer Engineering
Univ. of Mass., Amherst
Amherst, MA 01002, USA
lgao@ecs.umass.edu

ABSTRACT

This paper quantifies the causes of end-to-end Internet path failures and the effects of routing instability on end-to-end performance. Studies have shown that end-to-end Internet failures (periods of prolonged packet loss) are widespread. These failures are typically attributed to either congestion or routing dynamics. Unfortunately, the extent to which congestion and routing dynamics cause end-to-end failures, and the effect of routing dynamics on end-to-end performance, are poorly understood.

This paper uses active measurements and routing data to characterize end-to-end failures observed over one month on a topologically diverse Internet testbed. We find that routing dynamics contribute significantly to end-to-end failures and, in particular, routing dynamics are responsible for most long-lasting path failures. We classify failures caused by routing dynamics into those that involve forwarding loops and those do not. Our results show that loop-free routing dynamics cause the majority of failures involving routing dynamics, and that failures involving loop-free routing dynamics typically last longer than those that involve loops. We find that failures caused by routing dynamics are widespread, and most of the long-lived end-to-end path failures that involve routing dynamics are caused by BGP convergence or instability. Our results provide new insights into the effects of routing instability end-to-end Internet path performance.

1. INTRODUCTION

The deployment of interactive applications, such as voice-over-IP (VoIP) applications, multiplayer games, and video conferencing, has made high availability of end-to-end Internet paths of paramount importance. Empirical studies have shown that end-to-end Internet path failures (periods of prolonged packet loss) are widespread and can last as long as 10 minutes [6, 11, 15]. These prolonged path failures can degrade the quality of interactive applications and even render them unusable. These failures are typically attributed to either congestion or routing dynamics; unfortunately, the extent to which congestion and routing dynamics cause prolonged end-to-end failures, and the effect of routing dynamics on end-to-end performance, are poorly understood. In particular, very little is known about (1) how routing instability affects end-to-end path performance (e.g., duration of reachability loss, packet delivery rates, delay, etc.), or (2) what causes the routing dynamics that result in these failures in the first place.

A better understanding of how routing dynamics affects end-to-end Internet path failures can help network engineers and protocol designers determine which aspects of routing dynamics have the most detrimental impact on end-to-end path performance. An understanding of how routing dynamics affect end-to-end path performance can also enable end-hosts to make informed reactive routing decisions using overlay networks.

Towards this end, this paper quantifies the causes of end-to-end Internet path failures and the effects of routing instability on end-to-end performance. To characterize path failures, we collect active measurements and routing data observed over one month on a topologically diverse Internet testbed. We find that routing dynamics contribute significantly to end-to-end failures, and that nearly all long-lasting path failures are caused by routing dynamics. We classify failures caused by routing dynamics into those that involve forwarding loops and those that are loop-free. Our results show that the majority of failures are caused by loop-free routing dynamics and that failures caused by loop-free routing dynamics typically last longer than those that involve loops.

To explore the feasibility of deploying reactive routing for masking various types of end-to-end path failures, we characterize the spatial diversity of these failures. We observe that most paths experience failures, and that failures caused by routing dynamics are spread across a large number of locations, which is consistent with observations from previous work [6]. We also observe that most path failures are independent (i.e., a single failure does not affect reachability across a large number of paths), and nearly all path failures caused by routing dynamics affected reachability between a single pair of end-hosts. (The same was not true for failures not caused by routing dynamics, which occasionally affected reachability to a large fraction of other end-hosts.) This finding suggests that reactive routing can potentially mask path failures caused by routing dynamics, although it may sometimes be less successful at masking other types of failures.

We find that many of the long-lived failures that are caused by routing dynamics are due to the behavior of today's interdomain routing protocol, Border Gateway Protocol (BGP) [16]. We observe that a significant portion of end-to-end failures caused by loop-free routing dynamics are the result of BGP's process of exploring alternate paths during convergence. We also observe that most path failures caused by routing loops can be attributed to BGP routing dynamics.

To the best of our knowledge, our work is the first to show that most long-lasting end-to-end Internet path failures are caused by routing dynamics (rather than persistent congestion or some other cause). It is also the first to measure the effects of actual observed routing dynamics on end-to-end performance (previous work has examined the effects of *injected* faults on end-to-end performance, but does not observe the effects of naturally occurring faults [10]). Our results have important implications for enhancing Internet reliability. They also underscore the necessity of enhancing today's interdomain routing architecture, should we ever hope to deploy interactive, mission-critical applications that cannot tolerate periods of prolonged end-to-end packet loss.

The rest of this paper is organized as follows. Section 2 describes our techniques to identify end-to-end failures caused by routing dynamics. We describe our measurement setup in Section 3. We characterize causes of end-to-end path failures at Section 4. In Section 5, we analyze the extent to which routing dynamics are caused by BGP. We list related works in Section 6. We end with conclusions in Section 7.

2. IDENTIFYING ROUTING DYNAMICS

There are two major causes of end-to-end path failures: *network congestion* and *routing dynamics.* During network congestion, packets are dropped due to buffer overflow. During routing convergence, packets might encounter a loop or blackhole and eventually be dropped. In this section, we discuss various types of routing dynamics; we then present techniques to identify failures that are caused by routing dynamics or network congestion.

2.1 Types of Routing Dynamics

When an event causes a set of routers to lose their current routing information, those routing changes will be propagated to other routers. We define routing changes following an event as *routing dynamics.* There are two classes of routing protocols in the Internet: Interior Gateway Protocol (IGP) and Border Gateway Protocol (BGP). Although they have different mechanisms to update routing information, BGP or IGP route changes can lead to routing dynamics.

IGP, such as OSPF and IS-IS, is a link-state routing protocol, and it requires each router to maintain a topology map of the network. When a network link changes state, a notification, called a link state advertisement (LSA) is flooded throughout the network. All routers note the change and recompute their routes accordingly. On the other hand, BGP is a path vector routing protocol, and it requires that each router simply advertise its best route for each destination to its neighbors. When there is an event affecting a router's best route to a destination, that router will compute new best route (if any) and advertise the routing change to its neighbors. If the router does not have any route to the destination, it will send a withdrawal messages to neighbors for that destination.

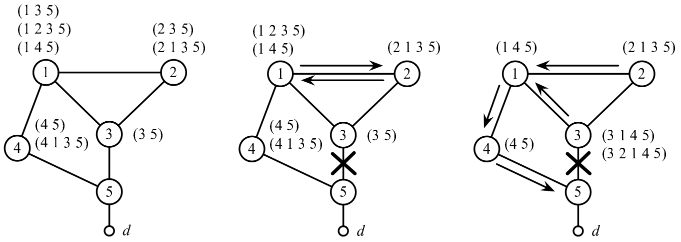

(a) Before a failure

(b) Route convergence period after the link between AS 3 and AS 5 is down

(c) After route converges

Figure 1: During BGP routes exchanging, looping is possible between AS 1 and AS 2. AS 1 and AS 2 will use transient routes, which are invalid, during the route from AS 4 converging

In the rest of this section, we explain how both BGP and IGP routing dynamics can cause end-to-end path failures.

2.1.1 Routing Loops

When a set of routers make inconsistent routing decisions a routing loop can arise. We present several examples to describe how routing loops can arise due to BGP, IGP, interactions BGP and IGP, and static default routing. In Section 5, we present techniques to identify those causes based on observations of IP-level path information alone and attribute routing loops to the following various causes:

- **Routing loops due to BGP.**

  In Figure 1(a), there are five ASes. The text beside each AS indicate the routes to d in the order of most to least preferred. Each AS uses the direct path to *d* as the best route. Suppose the link between AS 3 and AS 5 fails. If AS 1 and AS 2 receive a withdrawal message from AS 3 at the same time, these two ASes will each select the path via the other to reach *d*. As a result, there is a routing loop, as shown in Figure 1(b). After AS 1 and AS 2 exchange their new routes, AS 1 will delete the path from AS 2 and select the path from AS 4 as the best path. Finally, all ASes will use the path via AS 4 en route to *d*.

- **Routing loops due to IGP.**

  Figure 2, shows three routers running IGP; the number beside each path is the weight of the path. As shown in Figure 2(a), router 1 selects the path via router 2 as its shortest path. Suppose the link between router 2 and router 3 fails, and a LSA is broadcast in the network, and that router 2 is the first one to realize the link failure. The new shortest path for router 2 is via router 1. As a result, all traffic via router 2

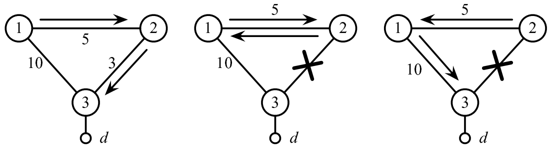

(a) Before a failure

(b) Route convergence period after the link between 2 and 3 is down

(c) After route converges

Figure 2: During link states exchanging, looping is possible between router 1 and 2. Router 1 receives the LSA later than router 2 does so that router 2 changes it route to forward packets to router 1, but router 1 still uses an out-of-date path to forward packets to router 2.

will be forwarded to router 1. However, during this period, router 1 has not received the LSA yet, so it still forwards traffic to router 2. In this case, there is a forwarding loop between router 1 and router 2, as shown in Figure 2(b). After router 1 receives the LSA, the routing loop stops, and traffic is forwarded to the destination, as shown in Figure 2(c).

- **Routing loops due to both BGP and IGP.**[1]

  Figure 3 has the same topology as Figure 2, with the only difference being that all routers are running both iBGP and IGP. Initially, every router uses the shortest path to forward traffic. Suppose that the IGP link weight changes from 3 to 20. If router 2 runs its BGP decision process first and updates its forwarding table, router 2 will forward traffic to router 1. However, router 1 still forwards traffic toward router 2 until it re-runs the BGP decision process (due to timers, this may be as much as 60 seconds later).

2.1.2 Loop-free Routing Dynamics

Now we consider the impact of *loop-free routing dynamics*—those that do not lead to loops. During the period that loop-free routing dynamics occur, packets can be dropped due to lack of up-to-date routing information. One possible reason is the time lag in obtaining routing information. We define loop-free routing dynamics which involve delay in obtaining routing information, as *routing lags*.

BGP-speaking routers can experience routing lags as follows. In Figure 4, each node represents an AS, while each link represents the connectivity between a pair of ASes. AS 2 selects the direct path as its best path and AS 1 prefers the path from AS 2 to its direct path. BGP as a path vector protocol applies "poison reverse" to avoid routing loops. In this example, AS 1 uses AS 2 to reach a destination, AS 1 does not announce its best route to AS 2. As a result, AS 2 does not have routing information about AS 1's direct path to the destination. Suppose the link between AS 2 and AS 0 fails. When AS 2 detects the failure, it will send a withdrawal message to AS 1. Subsequently, AS 1 will advertise its direct path to AS 2. Until AS 2 learns the alternate path from AS 1, all packets from AS 2 to AS 1 are dropped. Routing lags can also occur within an AS.

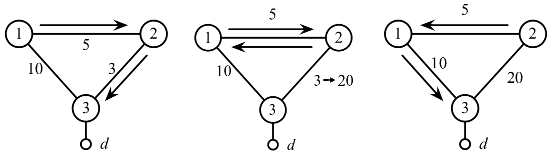

(a) Before the link weight changes

(b) Route convergence period after the link weight changes

(c) After route converges

Figure 3: After the link weight for link (2 3) changes, looping is possible between router 1 and 2. Router 2 runs BGP routing decision process first and forward traffic toward router 1, but router 1 still forward packets toward router 2.

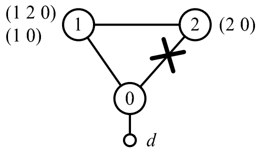

Figure 4: An example of routing lag occurring at AS 2. Before the failure between AS 0 and AS 2, the alternate path (1 0) in AS 1 is invisible to AS 2 due to BGP "poison reverse". Before AS 2 obtains the alternate path (1 0) from AS 1, packets are dropped at AS 2.

The time lag of obtaining alternate routes depends on the distance to the AS that provides an alternate path and the *rate limiting timers*, which are used to control the frequency of route announcements. The latency of obtaining an alternate route from an AS is about $2|d|MRAI$, where $|d|$ is the the number of hops to the AS, and MRAI is the rate limiting timer: it takes $|d|$ hops for the withdrawal message to arrive at the AS that has the alternate path, and each hop will apply MRAI timer to the withdrawal. Similarly, when the alternate path is sent from the AS, every router along the path applies MRAI timer before re-advertising the alternate path. Therefore, during routing lags may significantly delay the propagation of an alternate path.

2.2 Techniques to Identify Routing Dynamics

An ideal method to identify whether a failure can be attributed to routing dynamics, is to correlate the failure with

[1] This example has also been described in previous work [18].

routing changes, including BGP and IGP routing information, from all routers involving in that failure. Unfortunately, studying routing dynamics on end-to-end paths requires obtaining such a large set of routing information from multiple ISPs and multiple routers, which is extremely difficult (if not impossible). Instead, we use *IP-level path* changes, as measured by traceroute, to identify routing dynamics.

2.2.1 Heuristic to Identify Routing Dynamics

Assume that we have a set of IP-level forwarding paths from a source to a destination. Those paths include (1) an IP-level path before the failure, (2) a set of IP-level paths during the failure, (3) an IP-level path after the failure. From the set of IP-level paths, we can identify a routing loop if a traceroute shows the same sequence of routers multiple times. Therefore, we first identify those routing dynamics that lead to routing loops.[2] Then, we identify loop-free routing dynamics by IP-level path changes.

Ideally, we could attribute all observed IP-level path changes to loop-free routing dynamics, but this approach is not always correct. For example, load balancing within an AS can lead to IP-level path changes. An AS may split traffic load between multiple links based on link state protocol [8]. In the scenario shown in Figure 5, suppose the load balancing scheduler sends the first three packets along the path (1 2 3), and next three packets are along the path (1 3). If link (2 3) is congested, the first three packets may be dropped, while the other three still can be forwarded via another IP paths. Thus, even though the packet loss coincides with an IP-level path change, it would be incorrect to say that the failure was caused by routing dynamics.

To resolve this potential ambiguity, we continue to examine if those routing dynamics show either different *AS-level forwarding paths* before and after failures or different *egress points* within an AS. Here, we distinguish AS-level forwarding path that is derived from traceroutes, from AS path that shows in BGP table. If those IP-level path changes are also AS-level forwarding path changes, they are due to BGP routing dynamics. On the other hand, if those IP-level path changes have the same AS-level forwarding path, we examine their egress points within an AS. We select the first AS where the IP-level paths change. If, within that AS, the IP paths before and after the failure have the different egress points, we conclude that the failure is caused by routing dynamics, because load balancing typically splits traffic to the same egress point within an AS.

Figure 6 shows the heuristic we have developed to identify routing dynamics. The function $ASpath()$ maps an IP-level path to an AS-level forwarding path, while $Egress()$ derives the egress router within an AS.

2.2.2 Heuristic to Identify Routing Lags

In order to understand how loop-free routing dynamics can lead to failures, we design a heuristic to identify failures that are caused by routing lags. The heuristic is based on the changes of failure points. We define the last router in a traceroute that drops packets to a destination as the *failure point.* If packets are dropped at *multiple failure points*, and failure points get progressively closer to the source, we conclude this failure is caused by routing lag. The reason is that the sequence of failure points, which has decreased distance to the source, which corresponds to the propagation direction of BGP withdrawals.

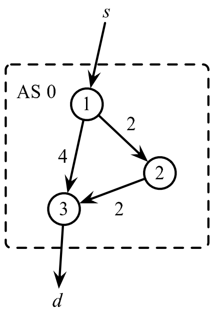

Figure 5: Load balancing using link state within an AS can result in IP-level path changes that are not associated with failures.

```
Input:
  (1)$P_0$,$P_t$: IP-level paths before and after a path failure.
  (2)$F$: a set of IP-level paths during the path failure.

 1. If there is a loop in $F$ then
        the failure is due to routing loop.
 2. If $P_0 \neq P_t$, then
        If $ASpath(P_0) \neq ASpath(P_t)$, then
          the failure is due to loop-free routing dynamics.
        If $Egress(P_0) \neq Egress(P_t)$, then
          the failure is due to loop-free routing dynamics.
```

Figure 6: Heuristic to identify routing dynamics.

For example, in Figure 7, there are multiple IP-level paths from the source s to destination d. Suppose that link between router 4 and d fails at time $t = 0$. The solid path indicates the path used during the path failure, while the dashed path represents the path after the path failure. When an upstream router receives a withdrawal message, for example, router 3, it will delete the entry to the destination in its routing table. Because router 3 cannot find any alternate paths, it sends a withdrawal to its upstream hop, router 2. At time $t = 1, 2, 3$, we will observe failure points at router 4, 3, 2, respectively. At $t = 4$, router 1 receives the withdrawal, and it has an alternate path, so all packets are forwarded along the path (s 1 5 6 d). From this example, we find that IP-level path changes with multiple failure points reflect routing lags.

2.3 Heuristic to Identify Network Congestion

After identifying a failure that cannot be attributed to routing dynamics or routing loops, we continue to decide if the

[2] Note that a loop in one traceroute run may be caused by an upstream routing change. The upstream router adds more hops so that the sequence of hops is the same as previously. We use the same method as described in [15] to consider the same sequence or routers shown at least three times as a forwarding loop.

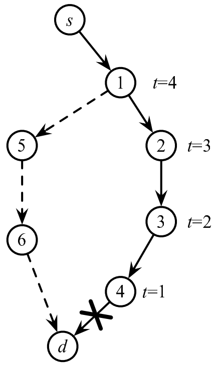

Figure 7: Routing lags with multiple failure points.

failure is due to network congestion. When congestion occurs, only packets are dropped, and it does not involve any routing change. Based on the location of failure point, which is defined as we identify routing lags, we design a heuristic to identify failures that are caused by congestion. If all packets are dropped at the *same failure point*, and before and after the failure there is no IP-level path changes, we conclude this failure is caused by congestion. If not all packets during the failure are dropped, i.e., *IP-level* forwarding paths sometimes can reach the destination, we consider this failure is also caused by congestion. However, our heuristic may misidentify routing dynamics with the same property as congestion because such transient routing dynamics might not be captured by traceroutes.

3. MEASUREMENT SETUP

To study the effects of routing instability on data plane behavior, we collected traceroutes and active probing data from the RON testbed on three separate occasions: from September 9, 2004 to October 7, 2004 (trace T_1) between 23 pairwise hosts, from November 28, 2004 to December 8, 2004 between 19 pairwise hosts (trace T_2), and from March 11, 2005 to March 21, 2005 between 9 pairwise hosts (trace T_3). We collected trace T_2 and T_3 to increase our confidence in the results; T_2 and T_3 include considerably more frequent traceroutes, to increase our ability to correlate our traceroute measurements with the failures observed in the data plane. Table 1 summarizes the hosts we used for each experiment. The traceroute data consists of both periodic "snapshots" of the testbed topology and traceroutes that were triggered by failures that were detected by the active probes. The dataset contains over 430 million active probes and 4 million traceroutes. These hosts are geographically and topologically diverse: the connections of these testbed hosts included low-bandwidth upstream connections such as cable modem and DSL, as well as higher bandwidth connections to both research networks (e.g., Internet2) and commercial ISPs. We collected BGP routing updates at the sites where 7 of these sites were located, as shown in boldface in Table 1. The remainder of this section describes our measurement techniques. We first describe our active probing techniques (i.e., the data plane measurements); we then describe our control plane measurements, which consist of IP-level traceroutes and BGP routing data.

Name	**Location**	T_1	T_2	T_3
Aros	Salt Lake City, UT	●	●	●
AT&T	Florham Park, NJ	●	●	
* Brown	Providence, RI	●	●	
CA-DSL	Foster City, CA	●	●	●
CCI	Salt Lake City, UT	●		
Coloco	Laurel, MD	●	●	●
* CMU	Pittsburgh, PA	●	●	●
* Cornell	Ithaca, NY	●		
Cybermesa	Santa Fe, NM	●	●	●
Digitalwest	San Luis Obispo, CA	●	●	●
GBLX-ANA	Anaheim, CA	●		
GBLX-CHI	Chicago, IL	●	●	
GBLX-LON	London, England	●	●	●
Intel	Palo Alto, CA	●	●	
MA-Cable	Cambridge, MA	●	●	
* MIT	Cambridge, MA	●		
* **MIT-BGP**	Cambridge, MA	●	●	
NC-Cable	Durham, NC	●	●	
* NYU	New York, NY	●	●	●
PSG	Bainbridge Island, WA	●		
PWH	San Jose, CA	●	●	
* UCSD	San Diego, CA	●	●	●
* Utah	Salt Lake City, UT		●	
Speakeasy	Cambridge, MA	●	●	

Table 1: The hosts between which we measured network connectivity. Asterisks indicate U.S. universities on the Internet2 backbone. Hosts where we also collect BGP data are shown in boldface. All hosts listed except for Utah were used in trace T_1.

3.1 Data Plane Measurements

The data plane measurements consist of active probes between pairs of testbed hosts. The active probes allow us to continuously monitor packet loss and delay characteristics of the end-to-end paths in the testbed topology. The active probes allow us to determine when various paths are experiencing outages or periods of high delay.

Because we are interested in witnessing end-to-end path properties that result from routing protocol behavior that occurs on the order of tens of seconds, we probe each end-to-end path once every five seconds. Each probing packet is assigned a unique identifier. When a host transmits a packet, it logs the time when the packet was sent; the host that receives the packet (1) logs the time when the packet was received and (2) sends a reply probe to the sender, logging the time at which it sent the reply. The initial sender then logs the time when it receives the reply packet. All of the testbed hosts are synchronized to within 1 millisecond, which allows us to measure the one-way delay of every transmission.

Limitations. Our data plane measurements only test end-to-end reachability of each path once per five seconds. Therefore, we are not guaranteed to capture any failures that last shorter than 5 seconds and we are not equipped to characterize those we do see; that is, we can only determine whether a failure lasted at least 5 seconds if we see two (or more) lost probes. Thus, not only will we fail to observe short failures (i.e., those that last less than 5 seconds), but we also do not have the ability to study the effects of routing dynamics on these short-lived failures. The results of our study should thus be interpreted as results for failures that last longer than five seconds.

3.2 Control Plane Measurements

To study control plane dynamics, we measure the IP-level forwarding paths with pairwise traceroutes between testbed hosts. We also collected BGP routing data at seven of the testbed sites to allow us to monitor BGP routing instability and route changes.

3.2.1 Traceroute Measurements

If a host sends two consecutive packets without receiving a reply from the destination host, then the sender initiates a series of traceroutes to the destination. The sender immediately sends a traceroute to the destination, and subsequently sends one traceroute to the destination every ten seconds for ten minutes or until the destination becomes reachable again, whichever occurs first. The traceroutes allow us to study the properties of the IP-level path once we have ascertained the existence of a problem in the data plane. Previous work has also used the combination of active probes and traceroutes to study IP-level path properties during failures, albeit on a much coarser timescale [6]. Due to the large number of testbed paths and the frequency with which we ran traceroutes, it was necessary to rate-limit our traceroute probes; as such, we do not capture traceroutes that correlate to all path failures, but we believe that the sample for which we do measure path performance is representative, especially for longer failures: we capture traceroutes for roughly 80% of failures for which at least three probes were lost and roughly 75% of failures for which at least two probes were lost.

To discover IP-level path *changes*, we must also collect periodic "snapshots" of the IP-level paths in the testbed. Accordingly, in addition to the failure-triggered traceroutes, each host initiates a traceroute to every other testbed destination every five minutes. (In dataset T_2 and T_3, we increased this frequency to once per minute.) These measurements provide us with a view of the likely IP-level path *before* the end-to-end path failure occurred, thus allowing us to ascertain whether a particular path failure resulted in an IP-level path change, and what the nature of that path change was. To study data plane characteristics on shorter timescales, we also ran a brief 10-hour experiment on one of the testbed paths, sending traceroutes once per second and active probes twice per second in each direction. To estimate the AS of each IP-level hop in the traceroute, we use the origin AS of a routeviews routing table in combination with the routing registry data.

After we have identified that a sequence of traceroutes corresponds to some type of routing dynamics, we conclude that any data path failure event that starts within 30 seconds of the start of that sequence of traceroutes was likely caused by the observed routing dynamics.

Limitations. We emphasize several important subtleties and limitations of our traceroute measurements:

1. We do not use the traceroute measurements are to detect the presence of a failure on a forward path (as traceroute alone cannot ascertain such information), but rather to *provide information* about the IP-level path characteristics of a failure on a forward path that was detected using active probes as described in Section 3.1. This additional information helps us classify which failures are due to routing dynamics versus those that are likely not.
2. As our traceroute measurements are relatively infrequent (once per ten seconds, in the best case), we may be unable to attribute some failures to routing dynamics (i.e., if the routing dynamics in between our traceroutes). As such, the number of failures that we have attributed to routing dynamics should be interpreted as a *lower bound*. In other words, routing dynamics may be an even greater contributor to end-to-end failures than our results suggest.

Host	*BGP Peers*
MIT (AS 3)	Genuity, Cogent, Comcast, Internet2
PSG (AS 3130)	Genuity, Verio
GBLX-LON (AS 3549)	Many ISPs
Aros (AS 6521)	UUNet, Electric Lightwave
PWH (AS 6549)	7 ISPs
Cybermesa (AS 14818)	Global Crossing, Xspedius

Table 2: Hosts from which we collected BGP data.

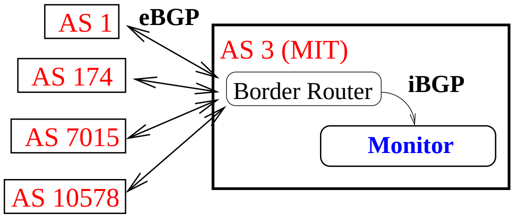

Figure 8: At each collection host, we collect BGP messages from the network's border router. The figure shows the configuration for MIT, which obtains upstream connectivity from Genuity (AS 1) Cogent (AS 174), Comcast (AS 7015), and the Northeast Exchange (via AS 10578).

3.2.2 BGP Routing Measurements

Table 2 shows the hosts where we collected BGP messages. These hosts ran Zebra, an open source software router [20], configured to log all BGP updates.

Figure 8 shows where the MIT collection host sits in relation to the border router of the hosting network and the rest of the Internet; other monitors sit in similar positions relative to their border routers. MIT's border router has four upstream feeds: a commercial feed via Genuity/Level3 (AS 1), Cogent (AS 174), Comcast (AS 7015), and to Internet2 via the Northeast Exchange (AS 10578).

The testbed host receives BGP updates from the border router. Because of the configuration, the monitors will not see all BGP messages heard by the border router; they see only BGP messages that cause a change in the border router's choice of *best* route to a prefix. Nevertheless, the monitor sees all BGP changes that would be relevant to transient routing failures, because it sees all BGP updates

Failure Type	*Number*	*Lost Packets*	*Fraction*
Routing Loops	50	5,249	0.0975
Loop-Free Dynamics	113	28,848	0.5360
Congestion	1,033	19,720	0.3664

Table 3: Packets lost due to each type of failure for trace T_1 (for which more than two consecutive probes were lost).

Failure Type	*Number*	*Lost Packets*	*Fraction*
Routing Loops	12	281	0.0157
Loop-Free Dynamics	788	12,053	0.6747
Congestion	1307	5,530	0.3096

Table 4: Packets lost due to each type of failure for trace T_2 (for which more than two consecutive probes were lost).

Failure Type	*Number*	*Lost Packets*	*Fraction*
Routing Loops	6	165	0.0230
Loop-Free Dynamics	182	5,854	0.8176
Congestion	156	1,141	0.1594

Table 5: Packets lost due to each type of failure for trace T_3 (for which more than two consecutive probes were lost).

that could possibly cause the end-to-end paths between a host and its destination to change.

4. PATH FAILURE CHARACTERISTICS

In this section, we analyze various characteristics of the end-to-end path failures that we observed on the data plane. We study data plane behavior of various types of path failures based on the following three categories: (1) those that involve a routing loop; (2) those that involve routing dynamics but not a loop (i.e., loop-free routing dynamics); and (3) failures that we cannot reliably attribute to a control plane. It is probably reasonable to assume that failures in the third category are likely due to other phenomena such as congestion, as they correspond to neither changes in the IP-level path before, during, or after the failure nor any visible routing updates. or any visible routing updates. Although we can reliably infer when a path failure is caused by routing dynamics, we unfortunately cannot attribute the third class of failures to congestion with absolute certainty, because the control plane failure might not have been observable with our traceroute-based measurements: it may have lasted less than five seconds, not involved an IP-level path change, or both. Therefore, we emphasize that routing dynamics may be an even *greater* contributor to end-to-end path failures than we are able to ascertain by our measurements alone.

In this section, we study the characteristics of these three types of failures involving at least one lost packet. Table 3 summarizes the number of packets lost due to each type of failure, for all failures longer than 1 minutes, for trace T_1. Tables 4 and 5 reflect the same statistics for trace T_2 and T_3, respectively.

Although we observe that most lost packets are caused by congestion or some other change that is not related to routing dynamics, when failures involving routing dynamics *do* occur, they are responsible for considerably longer failures.

4.1 Failure Duration

We first study the effects of various types of failures on the duration of end-to-end path failures. Routing dynamics do not account for the majority of lost packets, but when these types of failures *do* occur, they are responsible for failures that last considerably longer than failures that are not caused by routing instability. Figure 9 shows a cumulative distribution function (CDF) of the duration of each class of end-to-end path failures. (Figure 10 shows the CDF of the number of packets lost due to each type of routing event. Because each host probed each path once every five seconds, the number of probes lost is simply another way to look at failure duration.) While 90% of failures that do not involve routing instability last less than 10 seconds, half of all failures involving routing dynamics last longer than 60 seconds. Failures involving loop-free routing dynamics often last quite long: 20% of these failures last longer than 15 minutes. We also note that almost no failures last longer than 30 minutes. The order of magnitude of these routing failures is consistent with BGP-related failures in previous work [6, 11].

Our findings involving failure duration confirm the commonly held view that congestion-related failures (and other failures that do not involve the control plane) are typically short, while failures that involve control plane instability last considerably longer. These findings make sense: while congestion-related failures are typically caused by short-lived events (e.g., full queues), failures that involve control plane phenomena such routing protocol convergence are likely to last considerably longer.

4.2 Interarrival Time

We hypothesized that failures involving the control plane would be less frequent than other types of failures, since control plane-induced failures likely reflect more serious problems that occur less frequently (e.g., BGP session reset, failure of router interface, fiber cut, misconfiguration, etc.), whereas other types of failures such as congestion are likely caused by burst events that tend to occur more frequently and are dictated by events that occur on a much faster timescale, such as packet arrivals.

Figure 11 shows that our hypothesis holds in the case of very long failures: nearly 25% of failures involving routing dynamics are not followed by another control plane-related failure within the same day. On the other hand, other types of failures tend to occur much more frequently.

4.3 Delay Characteristics

Failures that involve routing dynamics should not be reflected by problems in the data plane before the failure occurs; on the other hand, we expected that a path that experiences a congestion-related failure might experience an increase in delay before the actual packet loss occurs. To test this hypothesis, we examined the difference between average one-way delay along a path over 10 seconds before the failure occurred and the delay after the failure ended.

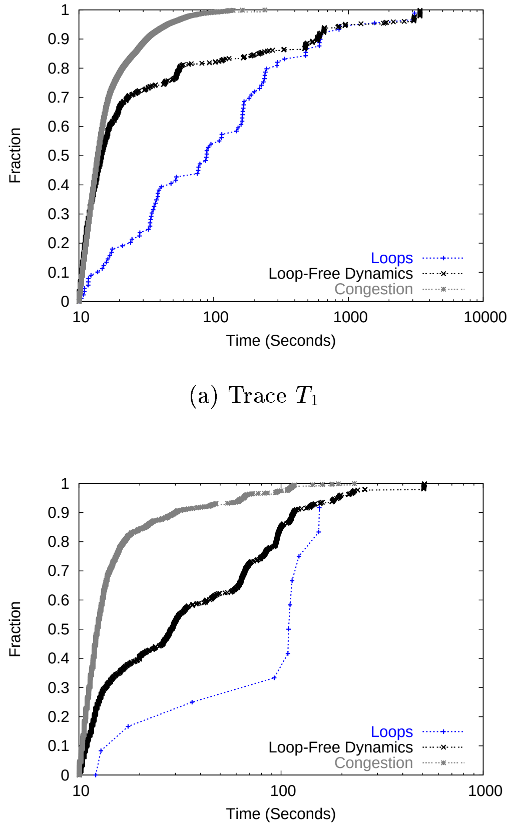

(a) Trace T_1

(b) Trace T_2

Figure 9: CDF of failure duration for failures of each type.

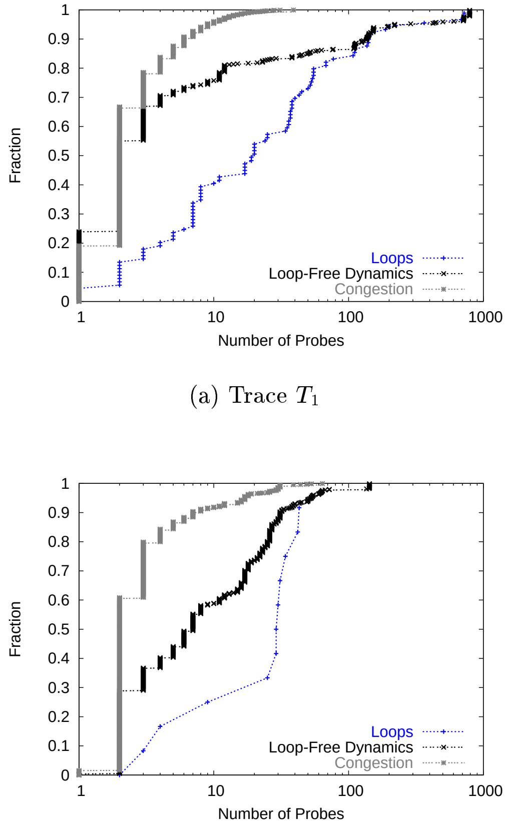

(a) Trace T_1

(b) Trace T_2

Figure 10: CDF of number of probes lost for failure of each type.

Only half of all failures of each type experienced an improvement in one-way delay after a failure; thus, it seems that control plane failures also experience degradation of path quality before the actual packet loss occurs. This degradation may be caused by routing instability that causes packets to take circuitous or otherwise suboptimal routes but nevertheless does not prevent packets from reaching their destination. Alternatively, when a control plane failure ends, the IP-level path may be different and *shorter* than that which was being used before the failure; in this case, we are not witnessing a degradation before the control plane failure, but rather a path improvement after the failure.

When a path does experience an improvement in one-way delay after a failure ends, the improvements in delay are often significantly larger when the failure was related to congestion. Figure 12 shows a CDF of these delay characteristics for the failures that experienced an *improvement* in one-way delay after the path failure ended. Although 90% of failures of *any* type experience a decrease in average delay of 10 milliseconds or less, about 2% of failures that do not involve the control plane are associated with a one-way delay improvement of more than 100 milliseconds (this characteristic is particularly true in trace T_1. Nearly all control plane failures do not involve significant changes in delay characteristics of the path. The same can be said of most other failures, but a small fraction of other failures involve significant changes in the delay characteristics of the path.

4.4 Spatial Characteristics

In this section, we study the spatial characteristics of path failures, addressing the following questions: Where do failures appear? How are failures distributed across paths? and Do paths fail independently (i.e., how likely is it that other paths from the *same* source to other destinations also fail)? These questions are important because reactive routing techniques (e.g., RON [17]) are most successful at masking path failures that occur both independently and further from the edge of the network [6]. Answering these questions for different types of failures lends insight into what types of failures reactive routing will be most successful at masking.

4.4.1 Where do failures appear?

To determine whether failures caused by routing dynamics occur close to the network edge, we compute the following metric for each failure:

$$I_d = \frac{|P_{su}|}{|P_{sd}|}$$

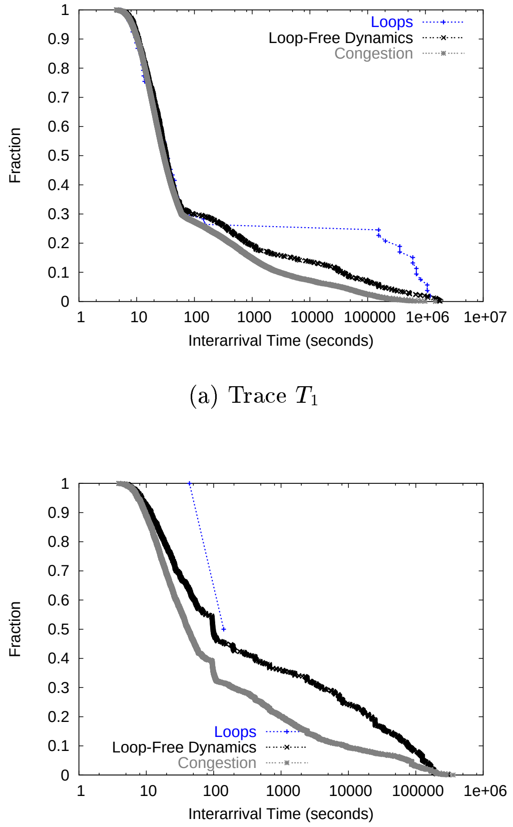

(a) Trace T_1

(b) Trace T_2

Figure 11: CDF of failure interarrival times.

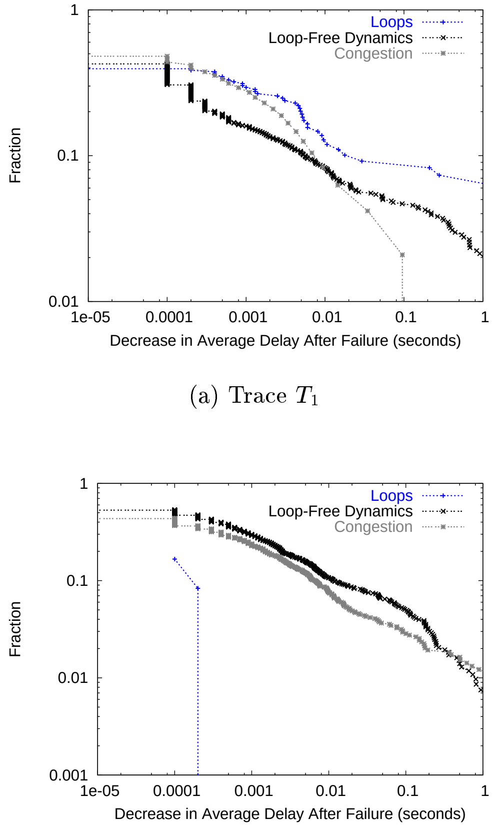

(a) Trace T_1

(b) Trace T_2

Figure 12: CDF of one-way delay characteristics for failures of each type.

where $|P_{sd}|$ denotes the length of AS-level forwarding path from a source s to a destination d and $|P_{su}|$ denotes the length of AS-level forwarding path from the source to the AS u where the failure appears. If I_d value is close 0 or 1, the failure is close the source or the destination, respectively.

Figure 13 shows the location of failures due to loop-free routing dynamics, and Figure 14 shows the location of failures caused by routing loops for trace T_1. In these figures, we show only failures lasting less than 10 minutes or less because only 10% of failures last more than 10 minutes. Although path failures can appear both close to end hosts and in the network "core", we observe that a significant fraction of failures due to routing dynamics occur in the core: 43% of failures involving loop-free routing dynamics and 63% of routing loops have $0.2 \leq I_d \leq 0.8$.

Loops that last longer than 100 seconds are extremely rare: we witnessed only four such cases; we observed two loops with $I_d = 0$ and a duration of 400 seconds. When these long-lived routing loops occurred, they all appeared in the same AS as the source host. (In Section 5, we will show that these routing loops are caused by static default routes.)

4.4.2 *How are failures distributed across paths?*

Figure 15 shows the number of failures caused by loop-free routing dynamics experienced by pairs of end hosts, and Figure 16 shows the number for routing loops for trace T_1. Most end-to-end paths failures. Furthermore, paths that experience failures due to loop-free routing dynamics are much more likely to experience routing loops than those that do not.

4.4.3 *Do paths fail independently?*

We wanted to determine whether failures related to routing dynamics were independent; that is, when a particular path from a source to destination fails, how likely is it that other paths from the *same* source to other destinations also fail? Figure 17 shows these characteristics for each type of failure. Most failures of each type involve only one source-destination pair, but control plane-related failures tend to involve a smaller number of destinations than other types of failures. This result suggests that failures that do not involve the control plane, such as congestion, typically occur near the edge of the network (i.e., close to an end host), thus preventing a single source from reaching a large number of destinations. On the other hand, control plane failures tend to occur further from the network edge.

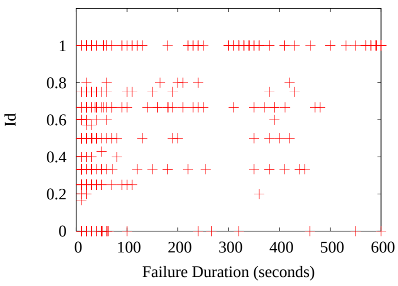

Figure 13: Spatial characteristics of failures caused by loop-free routing dynamics.

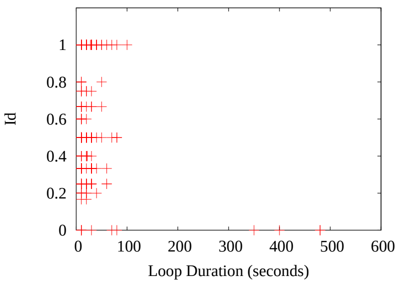

Figure 14: Spatial characteristics of failures involving routing loops.

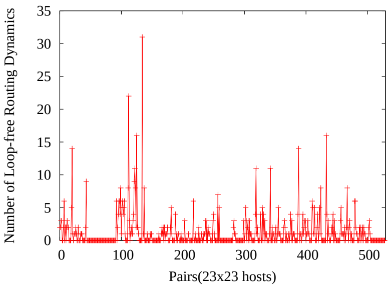

Figure 15: The number of failures due to loop-free routing dynamics experienced by each pair of end hosts.

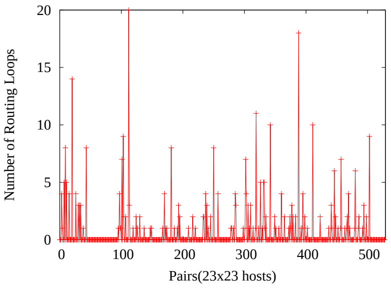

Figure 16: The number of failures due to routing loops experienced by each pair of end hosts.

Because, in all cases, more than 90% of failures involve only one path, reactive routing techniques should, in theory, be able to avert the majority of incidents that result in lost packets (assuming that these techniques could react to the failure events quickly enough). Our results also suggest that these techniques will be most successful at masking failures involving routing dynamics, which more commonly involve only a single path than other types of path failures.

5. ROUTING DYNAMICS CAUSED BY BGP

In this section, we aim to understand the extent to which routing dynamics are caused by BGP. We first study loop-free routing dynamics caused by BGP; we then discuss routing loops that appear to caused by BGP.

5.1 Loop-free Routing Dynamics

We use IP-level path changes to help us identify loop-free routing dynamics caused by BGP. Even though we cannot verify that all loop-free routing dynamics, we support our results from three observations, as shown in Table 6. The first observation is that failures occur at multiple failure points. As we mention above, IP-level paths with multiple failure points are due to the propagation of BGP withdrawals. We find that for both datasets, T_1 and T_2, about 17% of all failures caused by loop-free routing dynamics can be attributed to routing lags.

Our heuristic cannot identify routing lags with the same failure point. For example, in Figure 18, before a failure, router 2 has only one available routes to d because of BGP's "poisoned reverse", i.e., router 5 uses the path via router 2 to reach the destination. When the link between router 2 and 3 fails, packets are dropped at router 2 until it receives an alternate route from router 5. We know that this case should be attributed to routing lags. However, it is difficult to attribute this to routing lag from IP-level information since there are no visible *multiple failure points*. For example, for dataset T_1, we find about 111 failures that belong to this case, and we suspect that some of them belong to routing lags.

The second observation is that about 83% of failures involving loop-free routing dynamics experience a change in AS-level forwarding path, which indicates that those failures must indeed involve BGP.

Our third observation is the correlation between BGP instability and IP-level path changes. We correlate IP-level path changes observed from 7 end hosts with BGP instability collected from the same end hosts. Because of BGP's convergence delay, we use a 60-minute time window to correlate IP level path changes with BGP instability. Suppose at time t there is a failure with an IP-level path change; we examine if there is any BGP update for the destination during the time $[t-30, t+30]$. We observe about 56% of IP-level path changes for dataset T_1 and about 48% for dataset T_2, are correlated with BGP updates from those end hosts. Note that if an IP-level path change occurs within an AS, this change may be invisible at end hosts. That is, this intra-AS IP-level path change does not change AS-level path so this failure may not be visible from the end hosts. (Note that this observation is consistent with previous work, which observed that end-to-end path outages coincide with BGP

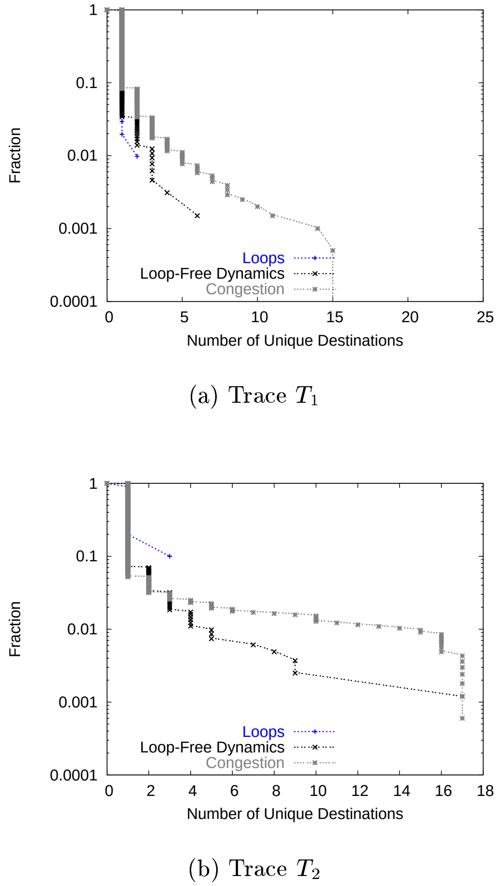

(a) Trace T_1

(b) Trace T_2

Figure 17: Complementary CDF of number of destinations involved for failures of each type.

instability roughly half of the time [6].)

5.2 Routing Loops

In this section, we introduce heuristics to identify the various causes of routing loops and estimate the prevalence of routing loops due to different causes. Here, we focus on routing loops shown in all datasets.

Table 7 shows the number of forwarding loops correlated with four routing loops. Over the course of our study observe 114 routing loops that were responsible for 63 distinct end-to-end path failures. Roughly 78% of loops are due to BGP protocol (iBGP or eBGP), which is consistent with previous observations [3]. Our inference techniques could not identify the cause of 37 of these loops.

The rest of this section describes our techniques for identifying the causes of forwarding loops and explores the properties of these different types of loops in more detail. We first study loops caused by default routes; then, we discuss loops that appear to be caused by BGP.

Observation	*Trace T_1*	*Trace T_2*
routing lags	17%	17.5%
AS-level forwarding path change	83%	82.5%
BGP instability	56%	48%

Table 6: Loop-free routing dynamics caused by BGP for trace T_1 and T_2.

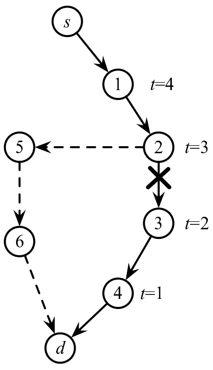

Figure 18: Routing lags with the same failure point.

5.2.1 Loops Caused by Default Routes

Routing tables in transit ASes are typically default-free, but stub ASes often use default routes for a significant number of Internet destinations. We use BGP routing tables from hosts to examine forwarding loops occurring at source hosts and destination hosts and determine whether the routing loop was likely caused by a default route. If a network receives a BGP withdrawal for some destination, the network has no BGP route for the destination. However, if the host continues to forward packets after the withdrawal occurs, and the packets are caught in a loop, we conclude that the loop is caused by a static default route. We observed four routing loops caused by default routes at one source host, all of which lasted more than 100 seconds.

5.2.2 Loops Caused by BGP

To distinguish loops caused by BGP routing events from those caused by the IGP, we exploit the fact that the intermediate paths look different during BGP convergence than they do during IGP convergence. When an IGP event occurs, routers recompute new paths according to a complete view of network topology; on the other hand, a router selects a BGP route among the route advertisements it receives from neighboring routers. When a single IGP routing event occurs, a router will only change paths once, and other routers on the new path should select consistent shortest paths. All IGP routers involved in loops ultimately switch to a shortest path that is consistent with the final shortest path assignment when IGP finishes converging. On the other hand, due to path exploration, the routers involved in a BGP-induced forwarding loop may temporarily use paths that are not consistent with the final path assignment.

Figure 19 shows an example of an IGP-induced forwarding loop: there are 5 IGP routers; each link has an IGP weight. Before the link between router 4 and 5 fails, all

Cause	*Number of routing loops*
default route	4
eBGP	35
iBGP	54
unknown (iBGP, IGP, or both)	21

Table 7: Routing Loops due to BGP, or static default route for all datasets.

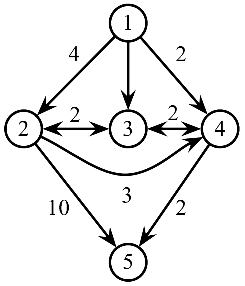

Figure 19: An example of IGP routing loop.

routers forward traffic via router 4. If that link fails, router 4 will switch to a new shortest path via router 3. Due to propagation delay of link-state advertisements (LSAs) and shortest-path computation, router 3 may not yet have computed a new path before router 4 switches paths. Before router 3 receives the LSA, there is a loop between routers 3 and 4. When router 3 receives the LSA, it derives a new shortest path via router 2, which is consistent with the path derived by router 4. Similarly, before router 2 receives the LSA, there is a loop between router 2 and router 4. Finally, the loop is resolved when router 2 receives the LSA. The new shortest path derived by router 2 is still consistent with that derived by router 4. Note that this behavior can also result from routing loops caused by interactions between iBGP and the IGP, although the convergence process will typically be much slower than with IGP alone.

This insight allows us to distinguish forwarding loops caused by BGP from those that are caused by IGP. If a forwarding loop traverses multiple ASes, we consider the loop due to eBGP loop. If the forwarding loop remains within a single AS, we then determine whether the routers involving in a loop make consistent routing decisions during convergence. We compare a set of IP-level paths during a loop with the IP-level path after the loop. If the path after the loop ends still fails to reach the destination, we conclude that the loop is due to iBGP loop. Otherwise, we examine if the first router in the loop has different next hops during and after the loop (i.e., whether its routing decisions were consistent). If the next hop that the first router uses during the loop is different than its next hop when the loop ends, we conclude the loop is caused by iBGP. Using this technique, we were able to identify 114 forwarding loops caused by either eBGP or iBGP convergence, as shown in Table 7; iBGP loops accounted for more than two-thirds of these forwarding loops.

Because of the probing granularity of our experiments (i.e., one probe per path every 5 seconds, and one traceroute every 10 seconds), we do not expect to observe IGP convergence, which typically occurs on the order of milliseconds to seconds. Thus, it is very likely that the 37 forwarding loops that we could not attribute to either iBGP or eBGP were caused by the interactions between iBGP and IGP, as described in previous work [18].

5.2.3 Correlating Routing Loops and BGP Instability

We correlate routing loops from 7 end hosts with BGP instability collected from the same end hosts. We use the same 60-minute time window to correlate them. We observe about 44% of routing loops are correlated with BGP instability.

6. RELATED WORK

In this section, we survey related work that has studied: causes of routing instability, the characteristics of packet loss, and properties of end-to-end Internet reachability.

Previous work has studied routing instability and end-to-end performance separately but has not examined the effects of routing instability on end-to-end performance. Labovitz *et al.* studied BGP route instability, focusing on the stability of paths between Internet Service Providers and artificially injected routing failures to discover their effects on Internet path performance [10]; we extend this work by quantifying the effects of real-world routing instability on end-to-end performance. Recent work attempts to identify the cause and origin of routing dynamics but does not study the effects of routing dynamics on end-to-end performance [4, 5, 7]. Other work has characterized failures that are correlated with IS-IS routing updates [3]. They classify failures according to their underlying causes such as maintenance activities, router-related and optical layer problems. Teixeira *et al.* measure the effects of intradomain routing on BGP routing stability but do not examine how this instability affects end-to-end performance [18]. Other work has also examined the effects of various routing protocol artifacts (e.g., timers, route flap damping parameters) on convergence time but does not explore the effects of this slow convergence on end-to-end performance [9, 13].

Conversely, other studies have examined the correlation between packet delay and packet loss and model congestion-induced packet loss [14, 19], but these studies do not examine the effects of routing dynamics on packet loss. Our work extends these previous studies by quantifying the effects of these instabilities on end-to-end performance and the extent to which routing instability degrades end-to-end performance.

Measurement studies have *correlated* routing instability and end-to-end performance, without identifying to what extent routing instability actually *causes* end-to-end performance degradation. Paxson identified Internet failures, routing loops, and routing pathologies using end-to-end traceroutes collected in 1994 and 1995 [15] and discovered that routing instability can disrupt end-to-end connectivity. We build on this work by examining the extent to which various types of routing instability are responsible for packet loss and degradations in end-to-end performance. Feamster *et al.* studied the location and duration of end-to-end path failures and correlated end-to-end path failures with BGP routing instability [6]. Their results show that most path failures last less than 15 minutes and most failures that coincide with

BGP instability appear in the network core. Our paper extends this study by examining the effects of various types of routing instability on end-to-end performance, rather than simply the correlation between instability and end-to-end performance.

Recent work has also examined the effects of routing instability data plane performance within a single AS. Agarwal *et al.* correlated BGP routing changes with packet traces from a large backbone ISP and found that BGP routing instability usually has little effect on shifts in traffic within a single AS [1]. Boutremans *et al.* use active and passive measurements to study the impact of network congestion, link failures and IS-IS routing instability on voice over IP service on a tier-1 backbone network [2]. Markopoulou *et al.* also observed failure durations and interarrival times within a single AS, but do not observe how these failures affect end-to-end path performance [12]. Our work focuses on how routing dynamics affect *end-to-end* packet loss and performance, rather than traffic shifts within a single AS.

7. CONCLUSION

Despite the fact that increasingly many Internet applications depend on high availability of end-to-end paths, our understanding of (1) how routing dynamics affect end-to-end path performance and (2) what types of routing events are responsible for dynamics that result in long-lived has been extremely limited to date. This paper explores how routing dynamics affect end-to-end path reachability and performance; we believe that this paper presents the first in-depth study of the effects of routing dynamics on end-to-end paths. Our combines measurements from both the data and control planes (i.e., active probes, traceroutes, and BGP routing data) and employs new techniques to identify the causes of end-to-end path failures using only information the IP-level path as measured from end-hosts.

Our findings suggest that while most packet losses are caused by phenomena other than routing dynamics (e.g., congestion), when routing dynamics *do* cause path failures, these path failures can last significantly longer than other types of failures. Path failures involving loop-free routing dynamics tend to last longer than those that involve loops. Most path failures due to routing dynamics occur closer to the network "core" involve only a small number of end hosts, while failures due to other causes (e.g., congestion) often involve a larger number of paths. This result suggests that reactive routing can be successful at masking the types of failures that result from routing dynamics and that it may occasionally have trouble masking long-lived failures caused by other factors. Finally, we note that most long-lived failures that are caused by routing dynamics can be attributed to the interdomain routing protocol, BGP. BGP is the source of many cases of routing dynamics that result in long-lived end-to-end path failures; redesigning some of BGP's artifacts that result in slow convergence may eliminate the vast majority of end-to-end path failures caused by routing dynamics.

8. REFERENCES

[1] S. Agarwal, C. Chuah, S. Bhattacharyya, C. Diot, The impact of BGP dynamics on Intra-Domain Traffic. *ACM Sigmetrics, Performance Evaluation Review Special Issue*, vol. 32, no. 1, pp. 319-330, June 2004.

[2] C. Boutremans, G. Iannaccone, and C. Diot, Impact of link failures on VoIP performance, in *ACM NOSSDAV*, May 2002.

[3] C. Boutremans, G. Iannaccone, S. Bhattacharyya, C. Chuah, and C. Diot, Characterization of Failures in an IP Backbone, in *Proc. ACM SIGCOMM Internet Measurement Workshop*, November 2002.

[4] M. Caesar, L, Subramanian, and R. H. Katz, Root cause analysis of Internet dynamics. NANOG presentation, http://www.nanog.org/mtg-0402/caesar.html. Feburary, 2004

[5] D. F. Chang, R. Govindan, and J. Heidemann, The temporal and topological characteristics of BGP path changes, in *Proc. International Conference on Network Protocols*, November 2003.

[6] N. Feamster, D. Andersen, H. Balakrishnan, M. Kaashoek. Measuring the Effects of Internet Path Faults on Reactive Routing. *ACM SIGMETRICS*, San Diego, CA, June 2003.

[7] A. Feldmann, O. Maennel, Z. M. Mao, A. Berger, B. Maggs, Locating Internet Routing Instabilities, In *Proceedings of ACM SIGCOMM*, 2004.

[8] B. Fortz, J. Rexford, and M. Thorup, Traffic engineering with traditional IP routing protocols, IEEE Communication Magazine, October 2002.

[9] Timothy G. Griffin and Brian J. Premore. An experimental analysis of BGP convergence time. In *IEEE International Conference on Network Protocols (ICNP)*, Riverside, CA, November 2001.

[10] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian. Delayed Internet Routing Convergence. *IEEE/ACM Transactions on Networking*, 9(3):293–306, June 2001.

[11] C. Labovitz, A. Ahuja, F. Jahanian. Experimental study of Internet stability and wide-area network failures. Proc. of Fault Tolerant Computing Symposium, June 1999.

[12] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C. Diot. Characterization of Failures in an IP Backbone. *Proc. Infocom*, March 2004.

[13] Zhuoqing Morley Mao, Ramesh Govindan, George Varghese, and Randy Katz. Route Flap Damping Exacerbates Internet Routing Convergence. In *Proc. ACM SIGCOMM*, Pittsburgh, PA, August 2002.

[14] S. B. Moon, J. Kurose, P. Skelly, D. Towsley, Correlation of Packet Delay and loss in the Internet, Technical Report 98-11.

[15] V. Paxson. End-to-end routing Behavior in the Internet, *IEEE/ACM Transactions on Networking* 5,5(1997), 601-615.

[16] Y. Rekhter, T. Li, and S. Hares. A Border Gateway Protocol 4 (BGP-4). Internet Draft draft-ietf-idr-bgp4-25.txt, September 2004.

[17] MIT RON Project. http://nms.lcs.mit.edu/ron/.

[18] R. Teixeira, A. Shaikh, T. Griffin, and J. Rexford, Dynamics of Hot-Potato Routing in IP Networks, In *Proc. ACM SIGMETRICS*, June 2004.

[19] M. Yajnik, S. Moon, J. Kurose, and D. Towsley, Measurement and Modelling of the Temporal Dependence in Packet Loss. in INFOCOM, 1999, vol. 1, pp. 345–52.

[20] Gnu Zebra, http://www.zebra.org/.