Measuring the Contributions of Routing Dynamics to Prolonged End-to-End Internet Path Failures

Feng Wang
CCIT
Liberty University
Lynchburg, VA 24502
fwang@liberty.edu

Nick Feamster
College of Computing
Georgia Tech
Atlanta, GA 30332
feamster@cc.gatech.edu

Lixin Gao
Department of Electrical & Computer Engineering
University of Massachusetts, Amherst
Amherst, MA 01002
lgao@ecs.umass.edu

Abstract— **This paper studies the contributions of routing dynamics to the duration of long-lived end-to-end Internet path failures. 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 contribute to long-lasting path failures, and the effect of routing dynamics on end-to-end performance, are not well understood. This paper uses a joint analysis of 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 coincide with most prolonged end-to-end failures, suggesting that routing dynamics contribute significantly to the duration of these failures. We also find that most long-lived end-to-end path failures that coincide with routing dynamics are caused by BGP convergence or instability. Our results provide new insights into the effects of routing instability on end-to-end Internet path performance.**

I. Introduction

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 [7], [14], [17]. These failures can degrade the quality of Internet applications. Although most failures initially result from either congestion or changes in the underlying network topology (*i.e.*, node and link failure and recovery), the *duration* of the failure comprises not only the time that a link is down but also the time that the routing protocol takes to react to the failure and discover alternative paths. During the period, routing protocols propagate routing changes among sets of routers, which are called routing dynamics. This paper seeks to understand the extent to which routing dynamics contribute to prolonged end-to-end Internet path failure, a property which, until now, has been poorly understood. More specifically, 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 types of routing dynamics can prolong path failures in the first place.

Using a joint analysis of active measurements and routing data over one month on a topologically diverse Internet testbed, this paper studies how routing dynamics contribute to prolonged end-to-end Internet path failures. Building on previous work that studied the *correlation* between BGP routing dynamics and Internet path failures [7], we study the *types* of BGP routing dynamics that can cause end-to-end path failures, and we observe how often specific types of routing dynamics coincide with prolonged path failures. As opposed to actively injecting exogenous routing faults into the network and observing the effects on end-to-end performance [13], [21], we observe *endogenous* faults and the passive response of the routing protocols to these faults.

We find that routing dynamics contribute significantly to end-to-end failures and that nearly all long-lasting path failures coincide with routing dynamics. Additionally, 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) [18]. We observe that a significant portion of end-to-end failures are caused by a specific class of routing dynamics that BGP experiences during convergence. Although it is well known that BGP experiences long convergence delay, BGP convergence does not necessarily cause path outages per se, and long convergence delay does not necessarily lead to prolonged path failures.

To the best of our knowledge, our work is the first to show that routing dynamics contribute to most prolonged end-to-end Internet path failures. Our results have important implications for enhancing Internet reliability: We believe that the results in this paper will help both network engineers and protocol designers determine which aspects of routing dynamics have the most detrimental impact on end-to-end path performance. Our results also underscore the necessity of enhancing today's interdomain routing architecture and explicitly point out the kind of of routing dynamics to avoid in the future design of interdomain routing protocol.

The rest of this paper is organized as follows. Section II describes our measurement methods, as well as important caveats and limitations. Section III describes the data used in our experiments. Section IV explains the techniques we use to classify routing dynamics that lead to prolonged end-to-end path failures and also shows the extent that those routing dynamics are caused by BGP instability (as opposed to other types of routing dynamics). Section V describes related work, and Section VI concludes.

II. Measurement Techniques

In this section, we describe our measurement techniques. We use active probes to detect a failure on the forward path and traceroutes to *provide information* about the IP-level path characteristics of a failure on a forward path. This additional information helps us classify which failures are due to routing dynamics versus those that are likely not.

A. One-Way Failures

We use active probes between pairs of testbed hosts to identify failures. The active probes allow us to continuously monitor packet loss and delay characteristics of the end-to-end paths in the testbed topology and determine when various paths are experiencing failures. Each host generates a unique packet identifier for each probe packet before sending the probe. 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. Because each testbed host logs every time it sends or receives a packet, we can merge these logs to identify one-way path failures between host pairs (this process differs from traditional "pings", which cannot differentiate between failures on the forward path vs. failures on the reverse path).

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 are shorter than 5 seconds. In this paper, because we are only concerned with *prolonged* failures, we only attempt to characterize failures that last longer than 10 seconds (*i.e.*, failures for which at least two consecutive probes were lost). The results in this paper apply only to failures that last longer than ten seconds; determining the causes of transient, short-lived failures is beyond the scope of this paper.

B. Identifying Routing Dynamics

We study how routing protocol dynamics contribute to the end-to-end path failures observed in Section II-A. The Internet has two types of routing protocols: intradomain routing protocols, or Interior Gateway Protocols (IGPs); and interdomain routing protocols, the primary example of which is the Border Gateway Protocol (BGP) [18]. When some event (*e.g.*, a link or node failure) causes an end-to-end path failure, these routing protocols react to the failure by propagating routing changes among sets of routers. Studies have shown that link state intradomain protocols such as OSPF and ISIS can converge in a few hundred milliseconds [1], [2], [3], [9]. On the contrary, BGP may converge very slowly [12], [13], during which time packets may be dropped [20], [21].

IGP or BGP routing messages may indicate routing dynamics, but obtaining these messages for all routers across the Internet is not possible. Additionally, routing failures may not be reflected in BGP updates [7], particularly if the observation point is far from the source of the failure. To address these limitations, we use traceroute to measure *IP-level path* changes that occur around the time of path failures to identify routing dynamics.

First, we perform traceroutes *immediately after the failure occurs*. If a host sends two consecutive probe packets without receiving a reply from the destination host, 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. These 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. Due to the large number of testbed paths and the frequency with which we ran traceroutes, it is 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.

Second, to gain information about the behavior of the routing system *before* the failures, we collect periodic "snapshots" of the IP-level paths between pairs of hosts in the testbed. In addition to the failure-triggered traceroutes, each host initiates a traceroute to every other testbed destination every minute.

From these sets of traceroute measurements, we can identify the IP-level paths around the time of a failure: (1) an IP-level path before the failure, P_0 (taken from our snapshots); (2) a set of IP-level paths during the failure, denoted as P (taken from the triggered traceroutes); and (3) an IP-level path that reaches the destination after the failure, P_t (also taken from the triggered traceroutes, or from the snapshots if the triggered traceroutes do not have such a path). Some IP-level paths in P may not reach the destination, but P_0 and P_t must reach the destination. We use the following guidelines to help us classify path failures caused by routing dynamics:

1) We say that the failure coincided with routing dynamics if some IP-level path in P can reach the destination if the IP-level path changes between P_0 and the path in P or if the IP-level path changes between P_0 and P_t.
2) We attribute the failure to a routing loop if any paths in P have duplicate routers.
3) If the paths in P fail at different routers, and these routers get progressively closer to the source, we attribute the failure to the propagation of BGP withdrawals.

This method may misclassify a failure if the failure does not induce a path change. For example, suppose that a router has only one route to a destination. When the router experiences a failure, the traceroute may not reflect any IP-level path change, since no alternate path exists. Our triggered traceroute measurements may also fail to attribute some failures to routing dynamics if we do not observe the IP-level path change (*e.g.*, if routing dynamics occur in between our traceroutes). As such, the number of failures that we have attributed to routing dynamics is a *lower bound*. That means, routing dynamics may be an even greater contributor to end-to-end failures than our results suggest.

III. Data

This section summarizes the data sets that we collected for our experiments.

1) Traceroutes and Active Probes: To study the effects of routing instability on data plane behavior, we collected traceroutes and active probes from the RON testbed [19] on two separate occasions: from November 28, 2004 to December 8, 2004 between 19 pairwise hosts (trace T_1), and from March 11, 2005 to March 21, 2005 between 9 pairwise hosts (trace T_2).

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.

2) BGP Routing Measurements: At each collection host, we collect BGP messages from the network's border router. The monitor receives BGP updates from the border router. Because of the configuration, the monitors do 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.

IV. Experimental Results

In this section, we study failures that are caused by routing dynamics, and understand the extent to which those failures are caused by routing dynamics. In particular, we analyze various characteristics of the end-to-end path failures that we observed on the data plane and characteristics of routing dynamics.

A. Path Failure Characteristics

We study data plane behavior of various types of path failures based on the following two categories: (1) those that we can reliably attribute to routing dynamics and (2) failures that we cannot reliably attribute to a control plane. It is probably reasonable to assume that failures in the second 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. Although we can reliably infer when a path failure is caused by routing dynamics, we unfortunately cannot attribute the second 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.

1) Overall Results: Table I and II summarize the number of each type of failure and the number of packets lost due to each type of failure, for trace T_1 and T_2, respectively. The two tables show the same statistics for the subset of failures longer than 30 seconds. That is, in most cases, failures involving routing dynamics were responsible for the majority of lost packets. In addition, Table I and II show that the number of routing loops and the number of packets lost due to routing loops are the smallest, which implies that most failures involving routing dynamics do not tend to result in routing loops.

2) Failure Duration: Figure 1 shows a cumulative distribution function (CDF) of the duration of each class of end-to-end path failures. (Figure 2 shows the CDF of the number of packets lost due to each class of end-to-end path failures. Because each host probed each path once every five seconds, the number of probes lost is simply another way to look at failure duration.) Furthermore, about 10% of the failures involving routing dynamics last longer than 15 minutes, though 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 [7], [14].

We observe that routing dynamics account for the majority of lost packets, and are also responsible for failures that last considerably longer than failures that are not caused by routing instability. 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.

B. Routing Dynamics

We use IP-level path changes to help us identify routing dynamics. Tables I and II summarize the number of failures caused by routing dynamics and packets lost due to routing dynamics for traces T_1, T_2, respectively. Recall that most routing dynamics do not involve loops. As discussed in Section II-B, IP-level paths with multiple failure points are caused by the propagation of BGP withdrawals. Among the failures caused by routing dynamics, we observe that for datasets T_1 and T_2, 10% (T_1), and 4% (T_2) of all failures caused by routing dynamics can be attributed to this case, respectively.

Further, to identify whether BGP instability was the cause of routing dynamics, we correlated IP-level path changes observed from 6 testbed hosts with BGP instability as observed from the networks where those hosts were located. Similar to the approach taken in previous work [7], 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 that about 48% for dataset T_1 and T_2 coincide with corresponding BGP updates. (Note that this observation is consistent with previous work, which observed that end-to-end path outages coincide with BGP instability roughly half of the time [7].) Note that, because IP-level path changes involving BGP do not always change the AS path, not all routing dynamics caused by BGP will be visible by observing IP-level path chances from end hosts.

	All Failures			**Failures $\geq$ 30 seconds**		
Failure Type	*Number*	*Lost Packets*	*Fraction*	*Number*	*Lost Packets*	*Fraction*
Routing Dynamics	800	12334	0.6904	380	10978	0.8113
- Routing Loops	12	281	0.0157	9	272	0.0201
- Loop-free Dynamics	788	12053	0.6747	371	10706	0.7912
Unknown	1307	5530	0.3096	131	2553	0.1887
Total	2107	17864	1.0000	511	13531	1.0000

TABLE I

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

	All Failures			**Failures $\geq$ 30 seconds**		
Failure Type	*Number*	*Lost Packets*	*Fraction*	*Number*	*Lost Packets*	*Fraction*
Routing Dynamics	175	5920	0.8475	165	5852	0.8652
- Routing Loops	5	160	0.0229	3	147	0.0217
- Loop-free Dynamics	170	5760	0.8246	162	5705	0.8434
Unknown	76	1065	0.1525	53	912	0.1348
Total	251	6985	1.0000	218	6764	1.0000

TABLE II

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

V. Related Work

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 [13]; 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 [5], [6], [8]. Other work has characterized failures that are correlated with IS-IS routing updates [4]. 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 [10]. 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 [11], [15].

Conversely, other studies have examined the correlation between packet delay and packet loss and model congestion-induced packet loss [16], [22], 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 [17] 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 [7]. 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.

VI. 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 failures 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 technique 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 the IP-level path information 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. 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

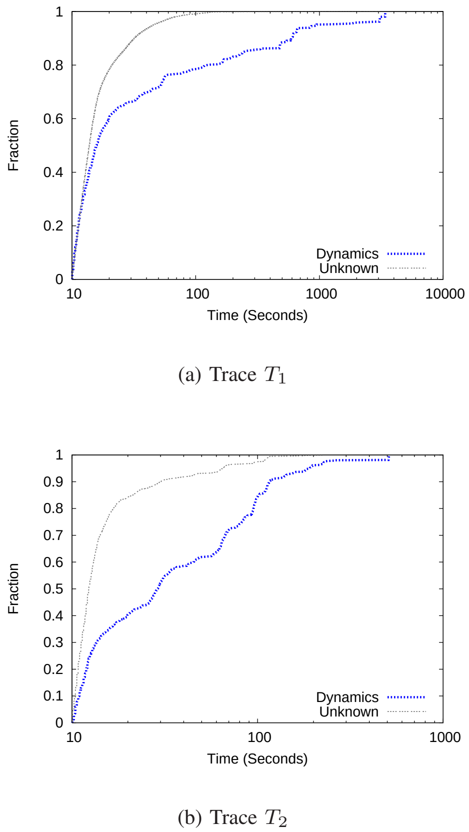

(a) Trace T_1

(b) Trace T_2

Fig. 1. CDF of failure duration for failures of each type.

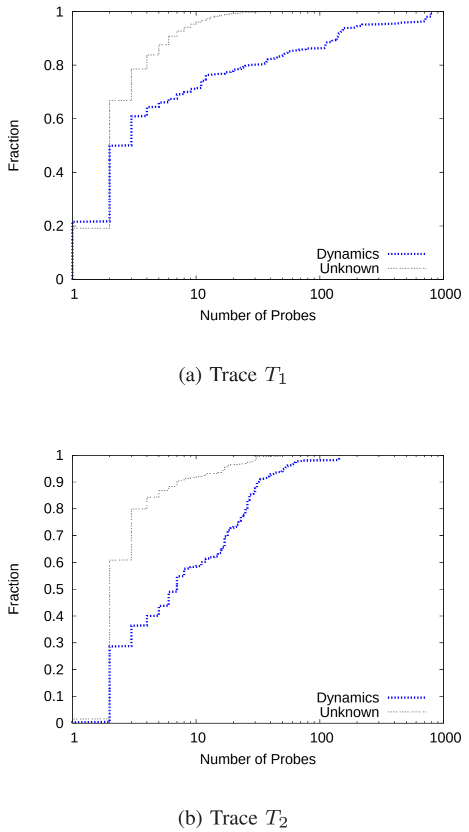

(a) Trace T_1

(b) Trace T_2

Fig. 2. CDF of number of probes lost for failure of each type.

eliminate the vast majority of end-to-end path failures caused by routing dynamics.

Acknowledgments

We would like to thank anonymous reviewers for their constructive comments. The work is partially supported by NSF grants CNS-0325868 and CNS-0208116.

References

[1] Alaettinoglu, C., and Casner, S. Detailed Analysis of ISIS Routing Protocol on the Qwest Backbone: A recipe for subsecond ISIS convergence. NANOG, Feb. 2002.

[2] Alaettinoglu, C., Jacobson, V., and Yu, H. Towards Millisecond IGP congergence. Internet draft, draft-alaettinoglu-ISIS-convergence-00.txt, Nov. 2000.

[3] Basu, A., and Riecke, J. Stability Issues in OSPF routing. In *SIGCOMM* (San Diego, CA, 2001), pp. 225–236.

[4] Boutremans, C., Iannaccone, G., Bhattacharyya, S., Chuah, C., and Diot, C. Characterization of Failures in an IP Backbone. In *SIGCOMM Internet Measurement Workshop* (Nov. 2002).

[5] Caesar, M., Subramanian, L., and Katz, R. H. Root Cause Analysis of Internet Dynamics. In *NANOG* (Feb. 2004).

[6] Chang, D. F., Govindan, R., and Heidemann, J. The Temporal and Topological Characteristics of BGP Path Changes. In *IEEE ICNP* (Nov. 2003).

[7] Feamster, N., Andersen, D., Balakrishnan, H., and Kaashoek, M. F. Measuring the Effects of Internet Path Faults on Reactive Routing. In *SIGMETRICS* (San Diego, CA, June 2003).

[8] Feldmann, A., Maennel, O., Mao, Z. M., Berger, A., and Maggs, B. Locating Internet Routing Instabilities. In *SIGCOMM* (2004).

[9] Francois, P., Filsfils, C., Evans, J., and Bonaventure, O. Achieving sub-second IGP Convergence in Large IP networks. *SIGCOMM Computer Communications Review 35*, 3 (2005), 35–44.

[10] Griffin, R. T. A. S. T., and Rexford, J. Dynamics of Hot-Potato Routing IP Networks. In *SIGMETRICS* (June 2004).

[11] Griffin, T. G., and Premore, B. J. An Experimental analysis of BGP Convergence Time. In *IEEE International Conference on Network Protocols (ICNP)* (Nov. 2001).

[12] Labovitz, C., and Ahuja, A. The Impact of Internet Policy and Topology on Delayed Routing Convergence. In *IEEE INFOCOM* (Anchorage, Alasks, Apr. 2001).

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

[14] Labovitz, C., Ahuja, A., and Jahanian, F. Experimental Study of Internet Stability and Backbone Failures. In *FTCS* (1999), pp. 278–285.

[15] Mao, Z. M., Govindan, R., Varghese, G., and Katz, R. Route Flap Damping Exacerbates Internet Routing Convergence. In *SIGCOMM* (Aug. 2002).

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

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

[18] Rekhter, Y., Li, T., and Hares, S. A Border Gateway Protocol 4 (BGP-4). *RFC 4271* (2006).

[19] RON Testbed Hosts. `http://www.datapository.net/tb/`.

[20] Wang, F., Gao, L., Wang, J., and Qiu, J. On Understanding of Transient Interdomain Routing Failures. In *IEEE ICNP* (2005).

[21] Wang, F., Mao, Z. M., Wang, J., Gao, L., and Bush, R. A Measurement Study on the Impact of Routing Events on End-to-End Internet Path Performance. In *SIGCOMM* (Pisa, Italy, Sept. 2006).

[22] Yajnik, M., Moon, S., Kurose, J., and Towsley, D. Measurement and modelling of the temporal dependence in packet loss. In *INFOCOM* (1999).