

Flooding Attacks by Exploiting Persistent Forwarding Loops

Jianhong Xia, Lixin Gao, Teng Fei
University of Massachusetts at Amherst
{jxia, lgao, tfei}@ecs.umass.edu

ABSTRACT

In this paper, we present flooding attacks that exploit routing anomalies in the Internet. In particular, we focus on routing anomalies introduced by persistent forwarding loops. Persistent forwarding loops may share one or more links with forwarding paths to reachable addresses. An attacker can exploit persistent forwarding loops to overload the shared links to disrupt the Internet connectivity to those reachable addresses.

To understand the extent of this vulnerability, we perform extensive measurements to systematically study persistent forwarding loops and the number of network addresses that can be affected. We find that persistent forwarding loops do exist in the current Internet. About 0.2% of routable addresses experience persistent forwarding loops and 0.21% of routable addresses can be attacked by exploiting persistent forwarding loops. In addition, 85.16% of the persistent forwarding loops appear within destination domains and they can be observed from various locations, which makes it possible to launch attacks from many vantage points. We also find that most persistent forwarding loops are just two hops long, which enables an attacker to amplify traffic to persistent forwarding loops significantly. To the best of our knowledge, this is the first study of exploiting the vulnerability of persistent forwarding loops to launch DDoS attacks.

1 INTRODUCTION

Distributed denial of service (DDoS) attack is one of the most prevailing threats in the Internet. In general, DDoS attacks send traffic from a large number of compromised hosts to deplete network or host resources needed by the victim. In this paper, we present flooding attacks that exploit routing anomalies in the Internet. In particular, we focus on a critical vulnerability in network routing architecture that is caused by persistent forwarding loops. Forwarding loops have been observed in previous measure-




Figure 1: Flooding Attacks by Exploiting Persistent Forwarding Loops

ment studies [3, 8, 11, 12]. Although transient forwarding loops disappear after routing protocol convergence, forwarding loops caused by configuration errors can last for a long time. In addition to obvious issues that persistent forwarding loops can blackhole network addresses, they can also be exploited to overload links that appear in the persistent forwarding loops.

Fig. 1 shows an example of a flooding attack that exploits persistent forwarding loops. Traffic to host X traverses routers R_a , R_b , R_c and other network devices to reach host X . At the same time, traffic to host Y also traverses routers R_a , R_b and R_c . However, due to misconfigurations in router R_c , traffic to host Y will be forwarded back to R_b . Therefore, any packet destined to Y falls into the loop between R_b and R_c , and will be dropped until its time-to-live (TTL) expires. In this scenario, link L_{bc} can be flooded if malicious attackers deliberately send a large amount of traffic to host Y . Host X would experience denial of service. Since traffic traversing a persistent forwarding loop typically traverses the links in the loop several times before being dropped, attackers take much less effort to launch flooding attacks, and therefore making the attacks stealthy. Since network operators can see high congestion only on the shared link L_{bc} but not on other links such as L_{ab} , without packet or flow-level measurements on the shared link, this kind of attack is hard to detect.

We perform extensive measurements to systematically study persistent forwarding loops. We find that persistent forwarding loops do exist in the current Internet. About 0.2% of routable addresses experience persistent forward-

ing loops and 0.21% of routable addresses can be attacked by exploiting persistent forwarding loops. In addition, 85.16% of the persistent forwarding loops appear within destination domains and they can be observed from various locations, which makes it possible to launch attacks from many vantage points. We also find that most persistent forwarding loops are just two hops long, which enables an attacker to amplify traffic to persistent forwarding loops significantly.

The remainder of this paper is structured as follows. Section 2 introduces the concept of persistent forwarding loops. Section 3 describes measurement design and data collection. Section 4 characterizes persistent forwarding loops. Section 5 exploits flooding attacks using persistent forwarding loops. We summarize the paper in Section 6.

2 PERSISTENT FORWARDING LOOPS

2.1 Concept of Forwarding Loops

In general, a packet from source s traverses a sequence of routers to reach destination d . A packet experiences a *forwarding loop* if it traverses a set of routers more than once. One of the powerful tools for discovering the forwarding path from a source to a destination is traceroute, which returns a sequence of router interfaces on the forwarding path. We denote the sequence of router interfaces in a trace from s to d as (r_1, r_2, \dots, r_n) . If $r_i = r_j$ and $i \neq j$, then the trace contains a forwarding loop (r_i, \dots, r_j) . The *length of forwarding loop* (r_i, \dots, r_j) is $j - i$.

2.2 Transient and Persistent Forwarding Loops

Forwarding loops can be transient or persistent. *Transient forwarding loops* are the forwarding loops that resolve themselves without human intervention or network topology changes. They may occur during routing protocol convergence [1]. Hengartner et al. [3] has demonstrated that forwarding loops exist in the Sprint network by analyzing packet traces. In general, forwarding loops are transient if they will disappear once the routing protocol converges. However, some forwarding loops will not disappear without human intervention or network topology changes. We refer to those forwarding loops as *persistent forwarding loops*.

2.3 Shadowed Address and Imperiled Address

If there is a persistent forwarding loop from source s to destination d , we refer to the IP address of d as a *shadowed address*. For example, the IP address of host Y in Fig. 1 is a shadowed address. Note that the links in persistent

Hop	Traceroute to shadowed address 81.181.31.127	Traceroute to imperiled address 80.96.192.10
1	128.119.91.254	128.119.91.254
2	128.119.2.238	128.119.2.238
...
18	166.49.147.134	166.49.147.134
19	195.39.208.82	195.39.208.66
20	193.226.179.18	193.226.179.18
21	193.226.130.226	193.226.130.226
22	194.176.189.42	194.176.189.42
23	193.226.130.226	194.105.11.178
24	194.176.189.42	80.96.192.10
25	193.226.130.226	
26	194.176.189.42	
27	193.226.130.226	
28	194.176.189.42	
29	193.226.130.226	
30	194.176.189.42	

Figure 2: A Shadowed Address and an Imperiled Address

forwarding loops may still be able to carry traffic to other reachable addresses. That is, a persistent forwarding loop may share one or more links with forwarding paths to the IP addresses other than shadowed addresses. If a destination d' is reachable and the forwarding path to d' shares one or more links with a persistent forwarding loop, we refer to the IP address of d' as an *imperiled address*. For example, the IP address of host X in Fig. 1 is an imperiled address. Imperiled address is named so because the address is in a dangerous situation and it may suffer from the potential threats posed by the persistent forwarding loops.

An example of a shadowed address and an imperiled address is shown in Fig. 2. In this example, traffic to the shadowed address *81.181.31.127* falls into a persistent forwarding loop between routers *193.226.130.226* and *194.176.189.42*. However, traffic to the imperiled address *80.96.192.10* relies on the same link that appears in the forwarding loop to reach the destination.

3 DATA COLLECTION

3.1 Measurement Design

We use traceroute to discover forwarding paths in our study. Our goal is to identify all possible persistent forwarding loops in the Internet. In order to reduce the measurement overhead, we select a set of representative IP addresses to trace. We intend to select as few IP addresses as we can, while trying to discover as many forwarding paths as possible. Most networks or their subnets are allocated a set of contiguous IP address blocks, and forwarding decision in a router is based on the destination IP address only. Therefore, tracing different IP addresses in the same contiguous IP block from a source may observe the same for-

Table 1: Summary on Measurement Design and Trace Data

Data Set	Fine-grained Prefix Selection	# of Selected IP Addresses/Prefix	# of Traces for Each Selected IP Address	# of Prefixes Traced	# of IP Addresses Traced	# of Traces Collected
D_A	all fine-grained prefixes	2	once	5,238,191	9,259,257	9,384,350
D_B	10% of candidate prefixes	2	5~7 times	12,983	21,212	139,739
D_C	35% of shadowed prefixes	50	once	3,705	171,218	187,980
D_{D1}	46% of shadowed prefixes	4	twice on average	4,894	19,762	41,556
D_{D2}	46% of shadowed prefixes	4	twice on average	4,894	20,691	44,969
D_{D3}	46% of shadowed prefixes	4	twice on average	4,894	20,657	44,936
D_{D4}	46% of shadowed prefixes	4	twice on average	4,894	17,825	34,873

warding path. Typically, a /24 block is used as the smallest unit, thus the forwarding path to any IP address in a /24 block should represent forwarding paths to all IP addresses in that block. Therefore, we design our measurement by selecting a couple of IP addresses in each /24 block to perform traceroute.

It should be noted that not all IP addresses have been allocated. Even we can obtain a list of all allocated IP addresses [5], not all allocated IP addresses have been used in the current Internet. To reduce the overhead of our measurement, we select the set of IP addresses based on the information from BGP routing tables in the RouteViews Project [10]. The RouteViews server peers with BGP routers in many large ISPs such as AT&T and Sprint. We refer to the prefixes in the BGP routing tables in the RouteViews server as *routable prefixes*, and refer to the IP addresses covered by routable prefixes as *routable addresses*.

We divide all routable prefixes whose lengths are shorter than 24 into multiple /24 prefixes, and keep routable prefixes whose lengths are no shorter than 24 unchanged. For example, prefix 12.0.0.0/8 is divided into 65,536 prefixes represented by 12.x.x.0/24. All prefixes in our measurement have a length of at least 24. We refer to these prefixes as *fine-grained prefixes*.

Since the forwarding paths to only a limited number of IP addresses are used to represent forwarding paths to all IP addresses in a fine-grained prefix, we extend the concept of persistent forwarding loops to fine-grained prefixes. We say that there is a persistent forwarding loop to a fine-grained prefix p from source s if we find a destination d_p in p that experiences persistent forwarding loops from s . Similarly, if a fine-grained prefix p contains a shadowed address, we refer to prefix p as a *shadowed prefix*. If a fine-grained prefix q contains an imperiled address, we refer to prefix q as an *imperiled prefix*.

3.2 Data Sets

We have collected four sets of trace data in this study. Most traces are collected from one location. We also collect additional traces on different hosts from various locations to

support our findings. Unless otherwise specified, all traces are collected from the campus network of University of Massachusetts at Amherst.

The first data set, D_A , is for detecting forwarding loops in all fine-grained prefixes. From a total of 0.17 million routable prefixes, we obtain about 5.36 million fine-grained prefixes. Due to security and privacy concerns posed by networks owned by governmental and military agencies, we filter out their prefixes according to WHOIS [4]. After filtering, 5.24 million fine-grained prefixes are traced. To reduce the overhead of our measurement, we perform traceroute to two IP addresses in each prefix, the first one and a random one. We collect 9.38 million traces in D_A by tracing to 9.26 million IP addresses. We refer to those fine-grained prefixes with forwarding loops in D_A as *candidate prefixes*. Since some selected IP addresses are traced twice during the experiment, the number of traces is slightly greater than the number of IP addresses in D_A .

The second data set, D_B , is for detecting persistent forwarding loops. To identify persistent forwarding loops, we trace to candidate prefixes and perform traceroute multiple times. Although we can observe forwarding loops from a single trace, it is impractical to monitor the network forever to identify persistent forwarding loops. Thus, we adopt an approximate criterion with respect to the general time scale of routing convergence. We trace an IP address d multiple times within four days. If there is a forwarding loop for all traces to d , we classify the forwarding loop as a persistent forwarding loop. In this case, d is a shadowed address and the fine-grained prefix that contains d is a shadowed prefix. To collect D_B , we trace to 10% of candidate prefixes and select two IP addresses in each prefix, a first one and a random one. Each selected IP address is traced at least 5 up to 7 times. We collect 139,739 traces by tracing to 21,212 IP addresses. Since in some prefixes only one IP address is traced during the experiment, the number of selected IP addresses is less than twice of number of prefixes in D_B .

After we identify persistent forwarding loops and shadowed prefixes from D_B , we further examine the forwarding consistency to multiple IP addresses in shadowed prefixes, and the observability of persistent forwarding loops from different locations. For these purposes, we collect two ad-

ditional data sets, D_C and D_D . In D_C , we trace to about 35% of the shadowed prefixes and select 50 random IP addresses in each prefix. We collect data D_D from four hosts in PlanetLab [9] that are located in Asia, Europe, US east coast and US west coast. We denote data sets from four hosts as D_{D1} , D_{D2} , D_{D3} , and D_{D4} respectively. For each of them, we trace to about 46% of shadowed prefixes and select 4 random IP addresses in each prefix. Table 1 summarizes our measurement design and data sets.

4 CHARACTERIZING PERSISTENT FORWARDING LOOPS

A trace of traceroute normally contains a sequence of router interface addresses. However, some traces may contain “*” or “!” when routers do not send back ICMP packets, replies get lost or filtered, or destinations cannot be reached. To reduce ambiguity, we filter out the traces that contain “*” or “!” between two appearances of a same address. We also filter out the traces where the same address appears continuously because forwarding loops could not be constructed by a single router interface.

4.1 Prevalence of Shadowed Prefixes and Imperiled Prefixes

4.1.1 Shadowed Prefixes

In our measurement, we identify the candidate prefixes from D_A and perform traceroute to these candidate prefixes to collect D_B . We then analyze D_B to identify persistent forwarding loops and shadowed prefixes.

Among 5.24 million prefixes traced in D_A , 139,278 of them are identified as candidate prefixes. If we convert them into IP addresses, they cover about 2.66% of routable IP addresses.

From data D_B that traces to 10% of candidate prefixes, we obtain 9,630 persistent forwarding loops, and identify 10,569 prefixes as shadowed prefixes. If we convert shadowed prefixes into IP addresses, we find that 81.39% of IP addresses in our sampled space have persistent forwarding loops. This number constitutes 0.2% of all routable IP addresses. Shadowed prefixes are located in 2,950 ASes, which suggests that IP addresses experiencing forwarding loops are originated from a large number of ASes. We believe that shadowed addresses in the Internet could be much more than what we have found in D_B because we trace to only 10% of candidate prefixes.

Note that we trace a limited number of IP addresses in each fine-grained prefix to collect forwarding paths. In order to confirm that, not only the selected IP addresses in shadowed prefixes experience forwarding loops, but other IP addresses in shadowed prefixes also experience forwarding loops, we use D_C for this study. 67.96% of shadowed

prefixes in D_C confirm that all additional sampled hosts have forwarding loops. We further investigate the reason that not all additional sampled hosts have forwarding loops in shadowed prefixes. We find that 73.41% of them are caused by the fact that infrastructure addresses (deployed for router interfaces) are sampled. For example in Fig. 1, although there is a forwarding loop when tracing to host Y , there is no forwarding loop if we trace to the interface address of R_c . It suggests that multiple IP addresses in the shadowed prefixes experience forwarding loops.

4.1.2 Imperiled Prefixes

As mentioned in Section 2, the vulnerability of persistent forwarding loops does not come from the shadowed addresses themselves. Rather, it comes from the shared links between persistent forwarding loops and the forwarding paths to imperiled addresses. To understand the extent of this vulnerability, we estimate the prevalence of imperiled addresses in the Internet.

The basic idea on identifying imperiled addresses is to find those IP addresses that are reachable and their forwarding paths share one or more links with persistent forwarding loops. It is not easy to fully identify the imperiled addresses in the Internet without a global view of forwarding paths from a source to a destination. In our experiment, we estimate the number of imperiled addresses/prefixes from D_A . Any reachable address in D_A that uses one or more links in a persistent forwarding loop is marked as an imperiled address. The fine-grained prefixes containing any imperiled address are marked as imperiled prefixes. Based on the persistent forwarding loops found in Section 4.1.1 and the traces in D_A , 10,828 of fine-grained prefixes are identified as imperiled prefixes. If we convert them into IP addresses, about 0.21% of all routable IP addresses are imperiled addresses. These imperiled addresses could be the potential victims when the vulnerability on persistent forwarding loops is exploited. These imperiled prefixes are originated from 1,516 ASes, so the potential victims widely spread in various domains. We show an imperiled address discovered in our measurement in Fig. 2.

Note that not all persistent forwarding loops share their links with forwarding paths to imperiled prefixes. Among 9,630 persistent forwarding loops, only 6.33% of them share links with forwarding paths to imperiled addresses. We call those shadowed addresses (prefixes) that can be exploited for attacking imperiled addresses *dark addresses (prefixes)*. Among 10,569 shadowed prefixes, only 5.64% of them are dark prefixes. With the growth and evolution of the Internet, some shadowed prefixes may become dark prefixes. Generally, a persistent forwarding loop shares one or more links with forwarding paths to only up to two imperiled prefixes. However, some persistent forwarding loops may share one or more links with forwarding paths to

as many as 1,000 imperiled prefixes. Flooding such shared links can result in denial of service to a large number of imperiled addresses.

4.2 Properties of Persistent Forwarding Loops

4.2.1 Location of Persistent Forwarding Loops

Identifying the location of persistent forwarding loops is helpful for us to understand where they occur. Persistent forwarding loops may occur within the destination domains, or across one or more other domains. It is difficult to accurately map infrastructure IP addresses to AS numbers [7]. However, because the most serious inaccuracies occur at AS boundaries, the accuracy may not be a problem if we only identify persistent forwarding loops that occur within destination domains. We consider that a persistent forwarding loop occurs within the destination domain if all interface addresses that appear in the loop are originated from the same AS as the shadowed address.

Among 91,090 traces with persistent forwarding loops, 85.16% of them occur in destination domains. It suggests that most persistent forwarding loops are close to the shadowed addresses rather than in the core of Internet. When persistent forwarding loops occur in destination domains, we conjecture that traffic to the shadowed addresses from different locations will most likely fall into these loops although they may traverse different paths in the core of Internet.

To confirm our conjecture, we collect additional traces, D_{D1} , D_{D2} , D_{D3} and D_{D4} on various hosts to verify the observability of persistent forwarding loops from different locations. We find that, persistent forwarding loops to about 90% of shadowed prefixes can still be observed from all four locations. Given that most persistent forwarding loops happen in destination domains, it is not surprising that they can be observed from various locations. However, comparing with the result that 85.16% of traces with persistent forwarding loops occur in destination domains, we conclude that although some persistent forwarding loops may not occur in destination domains, they can still be observed from different locations. It suggests that attackers are able to exploit persistent forwarding loops from different locations, which make this vulnerability more critical.

4.2.2 Length of Persistent Forwarding Loops

The length of persistent forwarding loops is important for us to understand traffic amplification factor in the links that appear in the persistent forwarding loops. When a packet enters a persistent forwarding loop, it may traverse the links in the loop multiple times before its TTL expires. The shorter the length of a persistent forwarding loop is, the

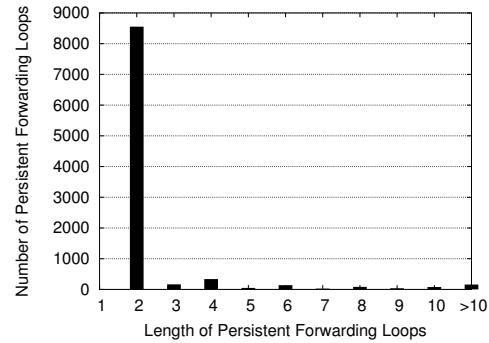


Figure 3: Distribution of Length of Persistent Forwarding Loops

more times a packet traverses the links in the loop. We find that, among 9,630 persistent forwarding loops, over 88.82% of them have a length of 2, which can significantly amplify the amount of traffic to shadowed addresses in the links that appear in the loops. About 8.71% of them have a length of 3 to 10. The rest of them have a length of 11 or longer. Several persistent forwarding loops have a length as long as 20. The distribution of length of persistent forwarding loops is shown in Fig. 3.

4.3 Possible Causes of Persistent Forwarding Loops

It is hard to identify the root causes of persistent forwarding loops without information about configurations on the involved routers. We conjecture that, the most possible cause of persistent forwarding loops is misconfiguration of the common usages of default routes and static routes. Several examples in [2] have shown that forwarding loops can happen if BGP or static routes are incorrectly configured. BGP misconfigurations are common today in the Internet [6].

To understand how misconfigurations can easily lead to persistent forwarding loops, we show an example that a network administrator neglects to configure a “pull-up route” at a border route to his upstream provider. Provider P owns 18.0.0/8 and delegates 18.1.0/16 to its customer C . The provider’s border router might have a static route directing traffic for 18.1.0/16 to the customer’s border router. The customer’s border router, in turn, might have routes for some subnets of 18.1.0/16, such as 18.1.1.0/24 and 18.1.2.0/24, but not for others. The customer’s border router may also have a default route (e.g., 0.0.0/0) pointing to the link back to the provider’s router, for access to the Internet. That would cause a persistent forwarding loop for all traffic destined to addresses in the range of 18.1.3.0 to 18.1.255.255.

In the above case, the forwarding loop is two hops long and near the destination domain. However, when the customer C is multi-homed, the same misconfiguration may

also lead to a persistent forwarding loop that occurs across multiple domains. For example, if the customer C has another provider B and prefers to use its link to provider B for outbound traffic. Therefore, the customer C 's border router has a default route 0.0.0.0/0 to provider B . The customer C also prefers to use the link from provider P for inbound traffic. In this case, when the customer C receives any traffic destined to addresses in the range of 18.1.3.0 to 18.1.255.255, it will forward the traffic to its provider B . The provider B will forward the traffic to its own provider or neighbors. Eventually the traffic will return to the provider P and reach the customer C again. That would cause a persistent forwarding loop occurring across multiple domains and has a loop length longer than 2.

To prevent this, the customer needs to configure a "null route" for 18.1.0.0/16 to discard packets to any destinations in 18.1.0.0/16 that do not have a more specific route.

5 FLOODING ATTACKS USING PERSISTENT FORWARDING LOOPS

In this section, we analyze the impact on the bandwidth consumption of the links in persistent forwarding loops and the effort that an attacker requires in order to launch such flooding attacks.

When a packet is sent to a shadowed address, it will fall into the persistent forwarding loop. It traverses the links in the loop and will be dropped only when its TTL expires. Therefore such a packet may traverse the links in the loop multiple times before being dropped and will consume more bandwidth. We define *traffic amplification factor* as the average number of times that a packet traverses a link in a persistent forwarding loop. Typically, a packet has a TTL value with 64 when created at its origin. From our measurement, we find that persistent forwarding loops occur on average 14 \sim 15 hops away from the source. Without losing generality, we consider that a packet traverses about 14 routers to fall into persistent forwarding loops. The persistent forwarding loops typically have a length of 2 as shown in Section 4.2.2. With this statistics, the traffic amplification factor can be estimated to be $\frac{64-14}{2} = 25$. It means that a packet will traverse the links in forwarding loops 25 times of what is expected. So the persistent forwarding loops can induce much more traffic than expected.

Due to the amplification on the traffic by persistent forwarding loops, attackers are expected to take much less effort to launch flooding attacks on imperiled addresses. For example in Fig. 1, if the available bandwidth for the link L_{bc} is 50Mbps, and traffic amplification factor is 25, then an attacker needs to send traffic at the rate of 2Mbps to flood L_{bc} . If an attacker has compromised 100 computers in the Internet and launches such an attack, the average traffic rate on each machine is only 20Kbps. Such a rate can be

easily performed by most users and be hard to detect from the source.

6 SUMMARY

In this paper we investigate the vulnerability on flooding attacks by exploiting persistent forwarding loops. We emphasize that such vulnerability can be exploited from various locations, and can severely affect the Internet connectivity to a significant number of network addresses. These findings suggest that this vulnerability could be a critical threat to the Internet security. In our future work, we plan to study the causes of persistent forwarding loops and how to eliminate these hidden troubles.

7 ACKNOWLEDGMENTS

The authors would like to thank our shepherd Jennifer Rexford for helpful suggestions and anonymous reviewers for valuable comments on the paper. This work was supported in part by NSF grant ANI-0208116, ANI-0085848, and the Alfred P. Sloan fellowship.

References

- [1] FRANCOIS, P., AND BONAVENTURE, O. Avoiding Transient Loops during IGP Convergence in IP Networks. In *Proc. IEEE INFOCOM* (March 2005).
- [2] HALABI, B. *Internet Routing Architectures*. Cisco Press, 1997.
- [3] HENGARTNER, U., MOON, S., MORTIER, R., AND DIOT, C. Detection and Analysis of Routing Loops in Packet Traces. In *ACM Sigcomm Internet Measurement Workshop* (November 2002).
- [4] [HTTP://WWW.ARIN.NET/WHOIS/ARINWHOIS.HTML](http://www.arin.net/whois/arinwhois.html).
- [5] INTERNET PROTOCOL V4 ADDRESS SPACE. <http://www.iana.org/assignments/ipv4-address-space>.
- [6] MAHAJAN, R., WETHERALL, D., AND ANDERSON, T. Understanding BGP Misconfiguration. In *Proc. ACM SIGCOMM* (August 2002).
- [7] MAO, Z. M., REXFORD, J., WANG, J., AND KATZ, R. Towards an Accurate AS-Level Traceroute Tool. In *Proc. ACM SIGCOMM* (August 2003).
- [8] PAXSON, V. End-to-End Routing Behavior in the Internet. In *IEEE/ACM Trans. Networking* (October 1997).
- [9] PLANETLAB. <http://www.planet-lab.org/>.
- [10] ROUTE VIEWS PROJECT. <http://www.antic.uoregon.edu/route-views/>.
- [11] SRIDHARAN, A., MOON, S., AND DIOT, C. On The Correlation Between Route Dynamics and Routing Loops. In *Proc. ACM Sigcomm Internet Measurement Conference* (October 2003).
- [12] ZHANG, M., ZHANG, C., PAI, V., PETERSON, L., AND WANG, R. PlanetSeer: Internet Path Failure Monitoring and Characterization in Wide-Area Services. In *6th Symposium on Operating Systems Design and Implementation (OSDI'04)* (December 2004).