

Your Router is My Prober: Measuring IPv6 Networks via ICMP Rate Limiting Side Channels

Long Pan^{*†}, Jiahai Yang^{*†‡}, Lin He^{*†‡}, Zhiliang Wang^{*†‡}, Leyao Nie^{*†}, Guanglei Song^{*†}, Yaozhong Liu^{*†}

^{*} Institute for Network Sciences and Cyberspace, BNRist, Tsinghua University, Beijing, China

[†] Zhongguancun Laboratory, Beijing, China

[‡] Quan Cheng Laboratory, Jinan, Shandong, China

Abstract—Active Internet measurements face challenges when some measurements require many remote vantage points. In this paper, we propose a novel technique for measuring remote IPv6 networks via side channels in ICMP rate limiting, a required function for IPv6 nodes to limit the rate at which ICMP error messages are generated. This technique, iVANTAGE, can to some extent use 1.1M remote routers distributed in 9.5k autonomous systems and 182 countries as our “vantage points”. We apply iVANTAGE to two different, but both challenging measurement tasks: 1) measuring the deployment of inbound source address validation (ISAV) and 2) measuring reachability between arbitrary Internet nodes. We accomplish these two tasks from only one local vantage point without controlling the targets or relying on other services within the target networks. Our large-scale ISAV measurements cover $\sim 50\%$ of all IPv6 autonomous systems and find $\sim 79\%$ of them are vulnerable to spoofing, which is the most large-scale measurement study of IPv6 ISAV to date. Our method for reachability measurements achieves over 80% precision and recall in our evaluation. Finally, we perform an Internet-wide measurement of the ICMP rate limiting implementations, present a detailed discussion on ICMP rate limiting, particularly the potential security and privacy risks in the mechanism of ICMP rate limiting, and provide possible mitigation measures. We make our code available to the community.

I. INTRODUCTION

Recent years have witnessed the rapid growth of IPv6 adoption. Statistics [30], [2] show that $\sim 40\%$ of client systems have adopted IPv6 as of 2022. Therefore, developing understanding of IPv6 networks by effective measurements is of great importance. However, measuring IPv6 networks usually faces challenges. For one thing, a common challenge of Internet measurements (both IPv4 and IPv6) is the lack of vantage points. Many measurement tasks cannot be accomplished without owning one vantage point inside the target network. Therefore, a considerable cost of renting or deploying vantage points is usually necessary for an Internet-wide measurement. For another, IPv6 is considered to be more secure in many respects. For instance, the hugely expanded address space of IPv6 [34], [65] makes exhaustive scans no longer possible and poses challenges to the discovery of measurement targets such as active hosts or DNS resolvers. The removal of the long-

standing identification field (IPID) from the fixed IPv6 header [15] also makes many IPID side channel-based attacks and measurements [21], [3], [40], [57] inapplicable to IPv6.

Similarly, unlike Internet Control Message Protocol (ICMP) for IPv4 [60]¹, ICMP version 6 (ICMPv6) requires IPv6 nodes to limit the rate at which ICMP error messages are originated, to limit the bandwidth and forwarding costs, and reduce the risk of ICMP flooding attacks [12]. However, while the pervasiveness of ICMP rate limiting in IPv6 networks does reduce the risk of networks being flooded with ICMP packets and makes it more difficult to perform traceroute-based active topology discovery [7], [66], it opens up new opportunities for us to measure remote IPv6 networks via ICMP rate limiting side channels.

In this paper, we propose a novel ICMP rate limiting side channel-based measurement technique, iVANTAGE, to measure remote IPv6 networks. With the ability to “send” and “receive” packets on remote Internet nodes by exploiting ICMP rate limiting side channels, iVANTAGE can, to some extent, employ routers in remote networks as our “vantage points” without directly controlling them. We apply iVANTAGE to the following two tasks, both of which are difficult and even theoretically *impossible* from only one local vantage point.

Measuring Deployment of Inbound Source Address Validation. Source address validation reduces the risk of spoofing-based cyberattacks, such as distributed denial of service (DDoS) attacks (especially amplification attacks [50], [41]) and other kinds of infiltration [14], by verifying the legitimacy of the source addresses of packets. It can be divided into inbound source address validation (ISAV) and outbound source address validation (OSAV) [22], [14], [41]. While a majority of observed networks ($\sim 85\%$ prefixes and $\sim 70\%$ autonomous systems) have deployed OSAV according to the *Spoofers* project [47], [9], [6], the lack of ISAV is much more pervasive [14], [41]. Therefore, it is crucial to understand the Internet-wide ISAV deployment. However, ISAV deployment is difficult to measure because it is challenging to have vantage points in every network or autonomous system (AS). Without any vantage points inside the target network, we cannot know whether ISAV filters our probes with spoofed source addresses.

While previous work for ISAV deployment measurements mainly relied on in-network volunteers [47], [9], [6] or DNS resolvers [41], [14], all of our measurements can be performed

¹Though not required by the RFC [60], many IPv4 nodes implement ICMP rate limiting [74], [64], [31].

from one local vantage point with better coverage, and do not rely on in-network volunteers or in-network open DNS servers. Our measurements cover $\sim 9.7k$ ASes, finding that $\sim 79\%$ of them are (partially) vulnerable to spoofing, which is the most large-scale measurement study of IPv6 ISAV to date.

Measuring Reachability. As the word Internet implies, networks are interconnected to each other. We may take it for granted that every two nodes on the Internet can reach each other. However, many causes such as pervasive Internet censorship [1], [58], [68], [57], [52], [17], [62], [61], [63], [73], link failures [24], [38], and routing failures [76], [77], can lead to loss of reachability. Pinpointing network disruptions will help network administrators troubleshoot, and detecting censorship will help understand Internet-wide activities. As with measuring ISAV deployment, it is also challenging and even impossible to measure reachability between two remote nodes from one local vantage point without controlling any of them. Previous work primarily used DNS [58], [68], virtual private networks (VPNs) [52], echo servers [73], [63], or IPID side channels [17], [57] to perform measurements. However, their approaches relied on many particular services with limited coverage. They focused on various kinds of Internet censorship instead of network-level reachability itself. Many methods are also inapplicable to IPv6, for instance, there is no IPID in the fixed IPv6 header, discovering targets such as echo servers is also hard because of the large IPv6 address space.

Our iVANTAGE-based approach aims at measuring *network-level reachability* itself. It is applicable to IPv6 networks, does not rely on global IPID counters, DNS, echo servers or VPNs, and can be done from only one local vantage point. We evaluate our method for reachability measurements with ground truth, achieving an accuracy of over 90% and over 80% precision and recall, respectively.

Finally, we perform an Internet-wide measurement of the implementations of ICMP rate limiting, discussing the potential risks of ICMP rate limiting and providing possible mitigation measures.

We make our source code publicly available to the community at <https://github.com/iVantage-NDSS23/iVantage>, which includes tools to perform Internet-wide active measurements of IPv6 ISAV and reachability measurements between two arbitrary IPv6 nodes, using methods we proposed in this paper.

Contributions. We make the following contributions:

- We propose a novel technique based on ICMP rate limiting side channels, iVANTAGE, which can to some extent use 1.1M remote routers distributed in 9.5k ASes and 182 countries as our “vantage points” for active measurements without directly controlling them.
- We perform the most large-scale measurement study of IPv6 ISAV deployment to date by a novel method based on iVANTAGE, from only one local vantage point, covering $\sim 50\%$ of all IPv6 ASes.
- We propose a new method based on iVANTAGE for measuring the reachability (including network-level censorship) between arbitrary IPv6 Internet nodes from only one local vantage point without controlling any of them, and evaluate this method with ground truth (precision and recall of over 80%).

- We perform an Internet-wide measurement of the implementations of ICMP rate limiting, provide a detailed discussion on ICMP rate limiting, reveal the potential security and privacy risks in the mechanism of ICMP rate limiting, and provide possible mitigation measures.

II. BACKGROUND AND RELATED WORK

A. ICMP Rate Limiting

Internet Control Message Protocol (ICMP) [60], [12] is commonly used in troubleshooting [24], topology discovery [5], [7], [36], [74], and network management [10], [70]. Well-known tools like `ping`, `traceroute`, and their variations are all mainly based on ICMP. However, allowing nodes to originate ICMP messages without any restrictions can lead to waste of bandwidth and network resources, and potential risks of ICMP flooding attacks. Therefore, in the IPv6 version of ICMP (ICMPv6) [12], IPv6 nodes are required to limit the rate at which ICMP error messages are originated. Related work [74], [7] and our measurement results also demonstrate that ICMP rate limiting is more common in IPv6 networks than in IPv4.

It is worth noting that, though ICMPv6 specification recommends a token bucket-based implementation for ICMP rate limiting [12], in our measurements, it is not necessary to care about the mechanism or how the target nodes implement ICMP rate limiting. We just need to send packets, observe the rate limiting by receiving packets – and no more than that.

There has also been previous work on ICMP rate limiting. Ravaoli et al. [64] and Guo et al. [31] provided measurement studies of ICMP rate limiting. They mainly focused on the rate limiting of ICMP Echo Replies and ICMP Time Exceeded messages in IPv4 networks. However, as a mandatory function for all IPv6 nodes as specified in the RFC [12], our measurements and previous work [7], [74] have demonstrated that ICMP rate limiting is more common in IPv6 networks. We also mainly focus on the rate limiting of ICMP Destination Unreachable instead of other relatively common types.

Vermeulen et al. [74] exploited the shared ICMP rate limiting mechanism of different router interfaces for alias resolution. Similarly, Man et al. improved the success rate of port prediction in DNS cache poisoning attacks [48] relying on insecure implementations of ICMP rate limiting for Linux-based operating systems. Their work brings inspiration to us, revealing the insecurities in the mechanism of ICMP rate limiting.

B. SAV and Its Measurement

IP source address spoofing refers to the process of sending packets having arbitrary IP addresses as their source IP addresses. These spoofed packets are difficult to trace and usually result in reflection and amplification attacks [47], [50], [41], flooding the victim with traffic. For example, a spoofing-based reflection attack against Amazon Web Services in 2020 resulted in record-breaking traffic of 2.3 Tbps [69].

Source address validation (SAV) was introduced and formalized in 2000 [22], serving as an essential defense against various kinds of spoofing-based attacks. SAV mainly falls into

two categories [22], [14], [41]: 1) *Outbound SAV (OSAV)* is deployed on the network egress. OSAV checks whether the source addresses of outbound traffic are indeed IP addresses within this network and discards packets whose source address does not belong to this network. 2) *Inbound SAV (ISAV)* is deployed on the network ingress. Since the source address of inbound packets is unlikely to belong to the destination network, ISAV filters such packets to reduce the potential risk of spoofing-based attacks. Measurement studies already show that lack of ISAV is much more severe than OSAV [47], [9], [14], [41].

Many previous studies concentrated on measuring SAV deployment. Misconfiguration-based [42], [45] approaches relied on network misconfigurations, and passive traffic-based [44], [51] approaches relied on inter-domain traffic. All of them have limited coverage and cannot perform Internet-wide measurements of SAV deployment.

The *Spoofers* project [47], [9], [6] measures SAV deployment for over ten years by requiring in-network volunteers to run the *Spoofers* client. The main limitation of the *Spoofers* project is that it relies on volunteers. The coverage of the *Spoofers* project is also relatively small, only $\sim 1k$ IPv6 ASes have been measured.

Korczynski et al. [41] and Deccio et al. [14] measured ISAV deployment by sending spoofed and non-spoofed DNS queries and receiving queries on authoritative DNS servers. The main difference is that, Korczynski et al. scanned the whole IPv4 Internet to discover open and closed DNS resolvers, which is not feasible in IPv6 due to the large address space, while Deccio et al. utilized “Day in the Life” (DITL) data [53] sponsored by the DNS Operations, Analysis and Research Center (OARC) [54] to find potential DNS resolvers. The main limitation of their DNS-based approaches is that for networks where there is no open DNS resolvers deployed, their approaches cannot confirm the presence of ISAV. Their coverage is also limited. The former did not measure IPv6 networks, and the latter found 3,952 IPv6 ASes lacking ISAV. Compared to their work, our iVANTAGE-based approach has better coverage, and does not need one authoritative DNS server or DNS data.

Dai et al. [13] measured IPv4 ISAV via IPID and Path Maximum Transmission Unit Discovery (PMTUD) side channels. They also performed DNS-based measurements like Korczynski et al. [41] and Deccio et al. [14]. Their work presented a most comprehensive view of IPv4 ISAV to date. However, their IPID-based and IP fragmentation-based method are not applicable in IPv6. There is no IPID in fixed IPv6 headers. Their PMTUD-based approach requires sending probe packets with DF (don’t fragment) flags, which do not exist in IPv6 fixed headers or extension headers.

C. Measuring Reachability

In most cases, any two nodes on the Internet can reach each other. However, many causes including but not limited to link failures [24], [38], routing failures [76], [77] and pervasive Internet censorship [1], [58], [68], [57], [52], [17], [62], [61], [63], [73] may lead to loss of reachability. Recent research on the origin of scanning [32], [75] also shows that scanning the same targets from different origins will lead to 5% to

10% discrepancy in the response rates. Therefore, measuring reachability between remote Internet nodes is crucial. It will help pinpoint connectivity problems, improve the reliability of Internet paths, and understand the activities in cyberspace.

Measuring reachability between two remote nodes without directly controlling one of them can be hard and even impossible. Some previous work exploited DNS [58], [68]. However, DNS reachability cannot fully reflect the real reachability. Similarly, *Quack* [73] and *Hyperquack* [63] exploited echo servers, which can only detect application-level reachability, having limited coverage. It is also hard to discover echo servers in such a large IPv6 address space. *ICLAB* [52] mainly used VPNs as vantage points. However, though *ICLAB* has VPNs or in-country volunteers in 62 countries and 234 ASes, it’s still not enough compared with over 190 countries and over 20k IPv6 ASes. The cost of buying and deploying VPN services is also considerable. *Spoooky Scan* [17] and *Augur* [57] exploit IPID side channels to measure the connectivity disruptions. However, their methods are not applicable to IPv6 networks because there is no IPID field in the fixed IPv6 header [15].

While there exists rich literature with studies of measuring the censorship [58], [68], [57], [52], [17], [62], [61], [63], [73], in this paper, we will focus on the *network-level reachability* itself instead of Internet censorship because Internet censorship is only one potential cause of loss of reachability. Actually, some types of censorship will also not cause loss of network-level reachability [1].

Our iVANTAGE-based approach does not rely on DNS, global IPID counter, echo servers, or VPNs and can be used to measure reachability between arbitrary IPv6 nodes from only one local vantage point.

III. iVANTAGE

A common challenge with active network measurements is the lack of vantage points. Without the ability to *send* or *receive* packets on a vantage point within a remote network, many measurement tasks cannot be completed. However, the reality is that it is not possible to have vantage points in all networks. Therefore, we propose an ICMP rate limiting-based technique, iVANTAGE, to solve this challenge, which consists of the following steps: finding remote “*vantage points*” first and then having them “*send*” and “*receive*” packets.

“Vantage Point” Discovery. iVANTAGE’s first step is to find targets that may serve as our “vantage points”, even though they are not vantage points in the true sense. They are mainly routers in remote networks, which are not under our control. We refer to them as remote “vantage points” (RVPs). All the IPv6 nodes that will originate ICMP messages in response to particular probes can be used as RVPs. In §IV, we describe a simple but effective method for discovering RVPs.

“Send”. After discovering the available RVPs, we need to control the RVPs to send probes for active measurements, which is the most crucial part. However, since we have no control over the RVPs, we can only induce them to send probes. An intuitive way to do this is to use spoofed source addresses to send packets to the RVPs that require the RVPs to send responses, e.g., ICMP Echo Requests, TCP SYN. Upon receiving the probes, the RVPs will send replies to

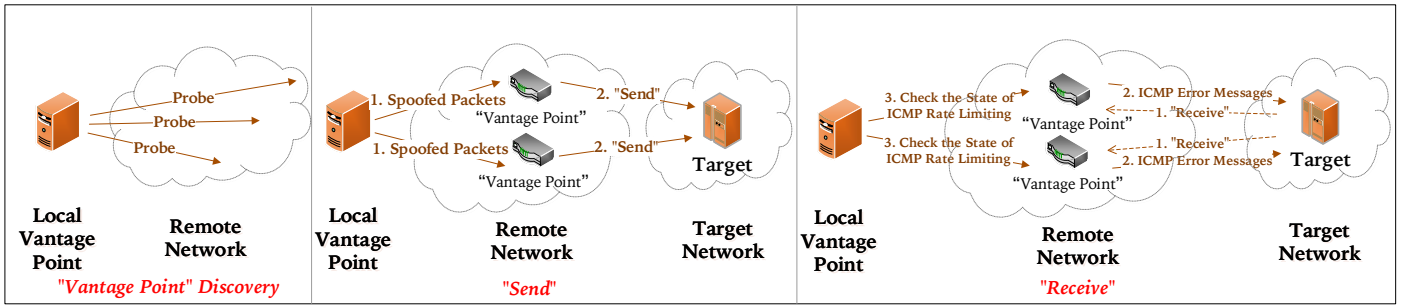


Fig. 1. An Overview of iVANTAGE. iVANTAGE first discovers potential “vantage points”, and then perform measurements from them with the ability to “send” and “receive” on the “vantage points”.

the spoofed source addresses, just as if we were asking the RVPs to send probes as we intended. Inducing RVPs to send replies is not necessarily helpful because, in general, the replies sent by the RVPs do not have any observable effect on the measurement targets. However, if the RVPs send replies that cause the measurement targets to send ICMP error messages, these replies will have an observable effect in the form of *different states of ICMP rate limiting*. Distinguishing between these different states can help us to measure targets.

“Receive”. In most measurements, there is also a need to capture packets from an RVP. However, without access to the RVP, capturing its received packets is difficult. Fortunately, for some measurements, it is generally not necessary to capture all packets on the RVP. We only care about whether (or even when) it receives a particular type of packets. The different states of ICMP rate limiting can help confirm whether (or even when) the packets are received. If the RVP receives a considerable number of packets that would cause it to send ICMP error messages, it can be observed that the RVP’s ICMP rate limiting is triggered. Observing this difference makes it possible to know if the RVP receives these packets, just like capturing packets on the RVP.

Figure 1 provides an overview of iVANTAGE. Suppose we want to perform measurements from a vantage point inside a remote network, but we do not have such a vantage point. iVANTAGE will first find RVPs in the remote network. Then, the key is how to let RVPs send or capture packets as we wish without directly controlling them. As mentioned earlier, if we want the RVPs to send packets to the target, we send packets to the RVPs using the target’s source address. Then the RVPs will send replies to the target, just as if we asked RVPs to send probes. Similarly, it is difficult to know whether RVPs can receive the packets sent by the target. However, suppose the packets that the target sends to RVPs will induce RVPs to originate ICMP error messages (e.g., these packets are sent to unreachable IP addresses). In that case, we can know whether RVPs have received these replies by checking the state of ICMP rate limiting on the RVPs.

iVANTAGE mainly aims at measuring filtering policy and network connectivity of remote networks with the ability to “send” and “receive” packets on the RVPs. Admittedly, RVPs can not perform all kinds of measurements as real vantage points. In the rest of this paper, we apply iVANTAGE to two different but both challenging measurement tasks. In the measurements of ISAV deployment (§V), we focus primarily

on how to “receive” packets on RVPs; in the measurements of reachability (§VI), we rely on RVPs to “send” and “receive” both. Both of these two tasks theoretically require in-network vantage points. See how iVANTAGE uses others’ routers as our “vantage points” and make the impossible possible.

IV. “VANTAGE POINT” DISCOVERY

To fully exploit ICMP rate limiting side channels, we need to collect a set of IP address pairs (hereafter referred to as *data pairs*). The data pairs consist of IP address pairs in the form of $\langle \text{target}, \text{periphery} \rangle$: by sending an ICMP Echo Request (i.e., ping) to the *target*, we can receive an ICMP error message from the *periphery*, which is usually different from the *target*. Since such ICMP error messages (mostly ICMP Destination Unreachable in our measurements) are usually originated by the last hop router on the path to the *target*, we call it *periphery*. We borrow the definition of *periphery* from [66], which refers to the last hop router connecting end hosts. These peripheries are actually the remote “vantage points” (RVPs) that will be used in the subsequent measurements.

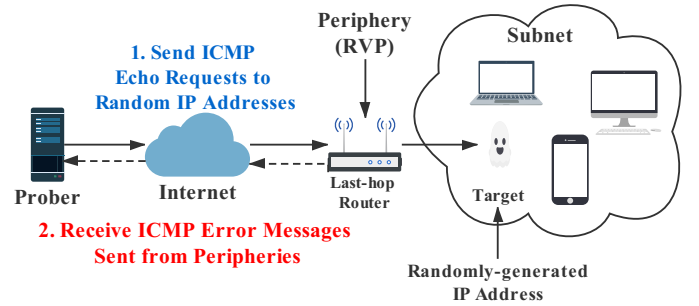


Fig. 2. Process of “Vantage Point” Discovery. The local prober sends lots of ICMP Echo Requests to the randomly-generated addresses and capture ICMP error messages.

To collect the data pairs, we implement a stateless scanner like many existing high-speed probes [16], [5], [36]. In consideration of the large address space of IPv6 networks, we cannot perform an exhaustive scan. Therefore, unlike other scanners that usually send probes to the IP addresses given, our scanner performs scans according to the $\sim 100k$ global IPv6 BGP prefixes present in the BGP routing table (maintained by RouteViews [55]), and generates targets by the following

²Some locations may be inaccurate.

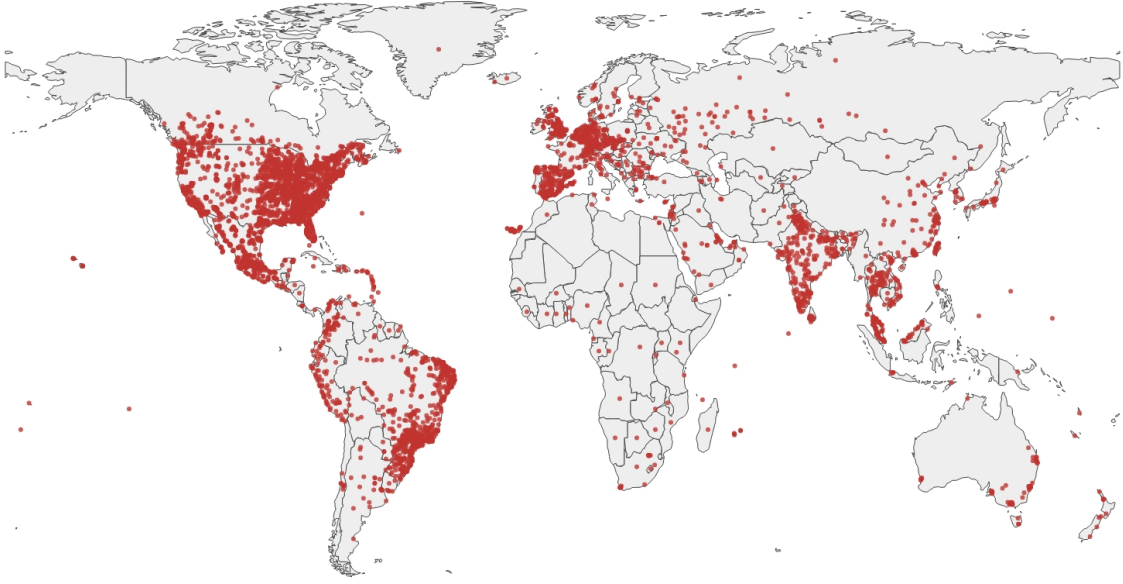


Fig. 3. Geographical Distribution of Our “Vantage Points”²

method [67], [43]: For every BGP prefix, we combine different /64 prefixes and random interface identifiers (the last 64 bits of IPv6 addresses). For instance, for BGP prefix 2000:1234::/40, we may have 2000:1234:0000:0000:adef:4983:19d2:3f12 first (the gray segments are generated randomly), and then 2000:1234:0000:0001:fed1:4f12:0894:349d is the next. The last target may be 2000:1234:00ff:ffff:1f3e:175a:4159:4de1. In practice, we randomize the probing sequence. When the scanner keeps probing each prefix, we capture *ICMP error messages* in response to our probes at the same time. The ICMPv6 specification [12] only requires IPv6 nodes to limit the rate of originating ICMP error messages³, and our measurements (see §VII) also reveal that the rate limiting of other types of ICMP packets (like ICMP Echo Reply) is less obvious and thereby more difficult to be exploited as side channels. Since the targets are randomly generated, they are very likely to trigger a great many ICMP error messages. We can easily extract both *target* and *periphery* from the ICMP error messages we received because it is required for ICMP error messages to quote the invoking packet (including, of course, its destination address) [12].

Figure 2 illustrates how the process works on one topology model. There is one periphery (usually CPE router or base station) before one customer subnet. Our probes are sent to the subnets, most of them are destined to unreachable IP addresses, triggering ICMP error messages sent from the peripheries.

³Though not required by the RFCs [12], [60], it’s common for Internet nodes to limit the rate of originating ICMP Echo Replies [74], [64], [31]. We also exploit the rate limiting of ICMP Echo Replies to perform supplemental ISAV measurements.

We limit the sending rate and use the multiplicative group of integers modulo n [27], [16] to randomize the probing sequence, refraining from sending probes to one network in succession. In addition, we do *not* send that many packets for each BGP prefix; instead, we stop scanning as soon as we find enough (e.g., 50) data pairs in a BGP prefix to limit the number of packets we send. We will also stop scanning if we are still unable to find any data pairs after sending a sufficient number (e.g., 1M) of packets to that network.

We repeated the scan several times in three months from a local vantage point on the university campus. Finally, we found **1,118,817** distinct data pairs that will be used as RVPs in subsequent measurements. Among all the ICMP error messages we received, 87.2% are ICMP Destination Unreachable, and the rest are ICMP Time Exceeded. It’s very surprising to receive many ICMP Time Exceeded messages because for all the packets we sent, we set the hop limit (a.k.a. Time to Live, TTL) to 64. It may be attributed to the routing loops or other kinds of misconfigurations [43].

Figure 3 shows the distribution of remote “vantage points” (using GeoLite2 [49] to locate IP addresses). Our “vantage points” are distributed in **182** different countries, **29,679** BGP prefixes and **9,498** ASes.

V. MEASURING ISAV DEPLOYMENT

In this section, we apply iVANTAGE to the measurement of ISAV deployments. The goal is to identify whether the target network (or BGP prefix, AS) has deployed ISAV to filter spoofed inbound packets. The main challenge is how to

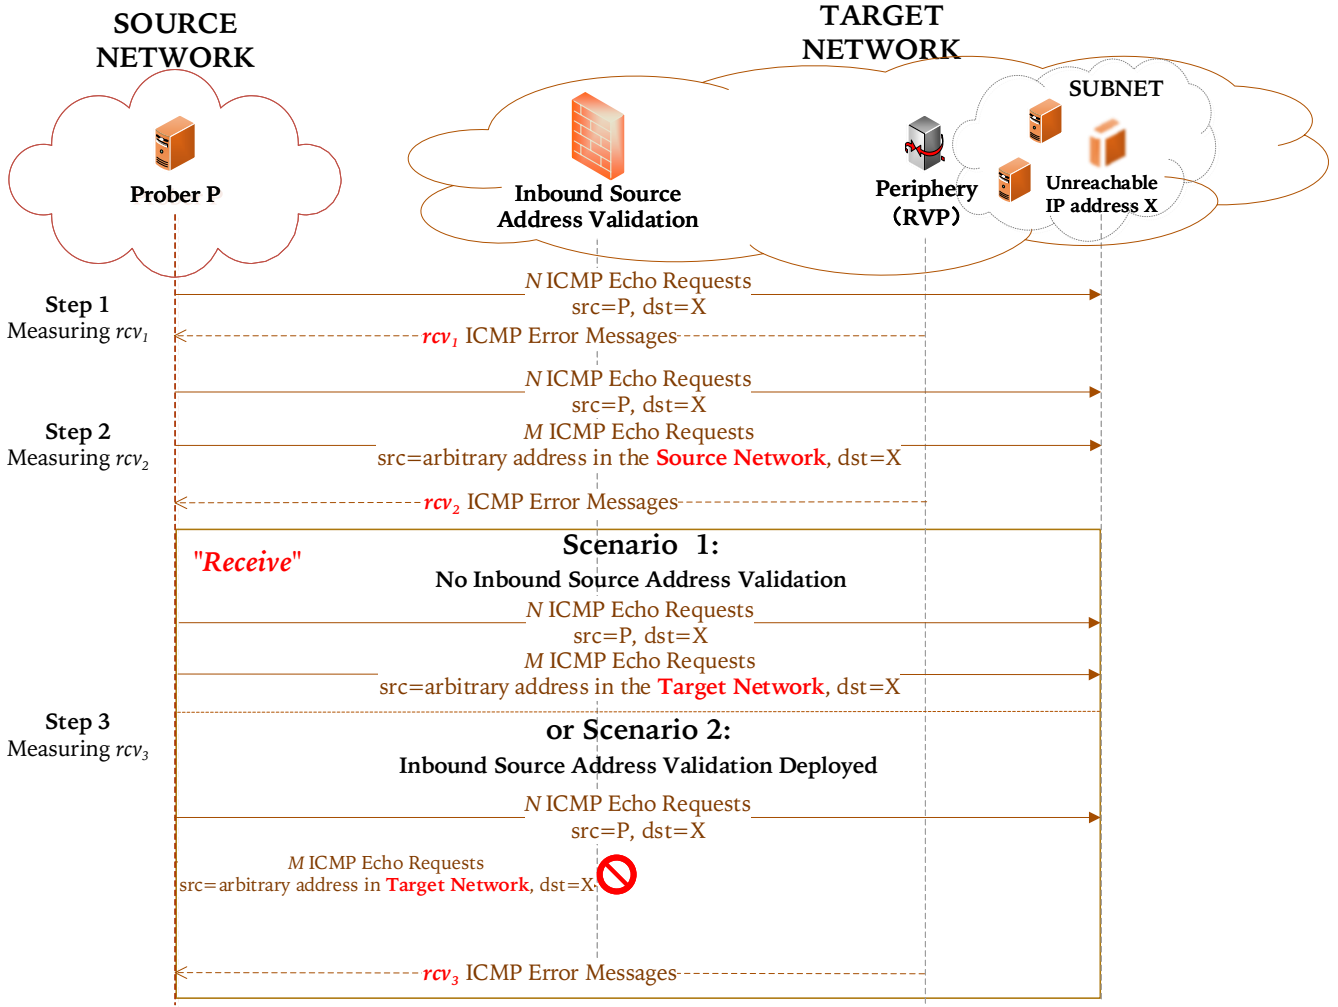


Fig. 4. Methodology of Measuring Inbound Source Address Validation Deployment

know whether the spoofed packets are filtered by ISAV, i.e., whether the RVPs can “receive” the spoofed packets, which, fortunately, IVANTAGE can easily do, as mentioned earlier.

Each target network will fall into one of the following three categories:

- **ISAV Deployed:** ISAV is deployed in the target network, which usually happens on the border of the target network.
- **Vulnerable to Spoofing:** ISAV is not deployed on the path toward the RVP in the target network. The target network is vulnerable to IP address spoofing attacks.
- **Uncertain:** The deployment of ISAV in the target network cannot be determined.

A. Methodology

Figure 4 illustrates the method for measuring ISAV deployment via ICMP rate limiting side channels. We have a local prober P in the source network. In contrast, we have a pair of IP addresses in the target network, the target (i.e., X in Figure 4) and the periphery (i.e., the RVP), which are collected

previously. We focus on measuring three values rcv_1 , rcv_2 , and rcv_3 :

- rcv_1 : The number of ICMP error messages (sent from RVP) P receives after sending N ICMP Echo Requests (i.e., ping) to the target. We refer to these packets as *probe packets*.
- rcv_2 : Similar to rcv_1 , but besides sending the same probe packets, P also sends M ICMP Echo Requests to the target simultaneously with a spoofed source address that belongs to the network of P⁴. We call such packets *noise packets*.
- rcv_3 : Similar to rcv_2 , but the spoofed source address of the noise packets belongs to the target network.

Our method is based on measuring and comparing the averages of the above three values.

rcv_1 vs. rcv_2 . Theoretically, as long as the RVP implements (global) ICMP rate limiting, we have $rcv_1 > rcv_2$. This is because noise packets exacerbate RVP’s ICMP rate limiting.

⁴It is not absolutely necessary to spoof the source address here. We spoof in order not to receive ICMP error messages caused by noise packets.

If there is $rcv_1 \approx rcv_2$, we can know that the RVP doesn't implement ICMP rate limiting or only implements very loose (or non-global) ICMP rate limiting, and therefore cannot be a good RVP in ISAV measurements. In practice, we will ignore these RVPs.

rcv_2 vs. rcv_3 . The process of measuring rcv_2 and rcv_3 is very similar. The only difference is the network to which the spoofed source address of noise packets belongs. If it belongs to the network of P, the spoofed source address is hard to be identified by the target network; if it belongs to the target network, it will be detected and filtered by ISAV (if deployed). Hence, if there is $rcv_2 < rcv_3$, we can infer that the target network has filtered the noise packet, and therefore, ISAV is deployed. In this case, we usually have $rcv_1 \approx rcv_3$. This is because there is no difference between the processes of measuring rcv_1 and rcv_3 if all the spoofed-source noise packets are filtered. Similarly, if there is $rcv_2 \approx rcv_3$, it makes no difference which network the spoofed source address of the noise packets belongs to, suggesting no ISAV on the path toward the target network. Actually, our measurements find that if ISAV is not deployed, we often find $rcv_2 > rcv_3$ instead of $rcv_2 \approx rcv_3$. This may be because remote Internet nodes prefer to reply to ICMP Echo Requests sent from their proximate Internet nodes [75], so noise packets with spoofed source addresses within the target network trigger a more obvious ICMP rate limiting.

rcv_1 vs. rcv_3 . As mentioned before, if there is $rcv_1 > rcv_3$, it means that noise packets with spoofed source addresses within the target network are not filtered, so ISAV does not exist; if $rcv_1 \approx rcv_3$, ISAV exists because the noise packets are filtered, so there is no essential difference between rcv_1 and rcv_3 .

B. Inferring ISAV Deployment

We infer the ISAV deployment by comparing the averages of these three values: rcv_1 , rcv_2 , and rcv_3 . As mentioned above:

- For networks with ISAV: $\overline{rcv_1} \approx \overline{rcv_3} > \overline{rcv_2}$.
- For networks without ISAV: $\overline{rcv_1} > \overline{rcv_2} \approx \overline{rcv_3}$ or $\overline{rcv_1} > \overline{rcv_2} > \overline{rcv_3}$.
- For networks without ICMP rate limiting (or too loose, too strict, non-global ICMP rate limiting): $\overline{rcv_1} \approx \overline{rcv_2} \approx \overline{rcv_3}$.

In practice, we infer the ISAV deployment based on simple but very effective rules. We introduce a factor $0 < \lambda < 1$ to avoid potential interference from fast-changing network environments, and we can be sure of $a < b$ only if $a < \lambda \times b$.

- If $\overline{rcv_3} < \lambda \times \overline{rcv_1}$, then ISAV is not deployed;
- Else if $\overline{rcv_2} < \lambda \times \overline{rcv_3}$, then ISAV is deployed;
- Else, the deployment of ISAV cannot be determined.

It's possible to have both $\overline{rcv_3} < \lambda \times \overline{rcv_1}$ and $\overline{rcv_2} < \lambda \times \overline{rcv_3}$. However, our measurements show that it rarely occurs (less than 5%). For these special cases, we compare the values of $\overline{rcv_3}/\overline{rcv_1}$ and $\overline{rcv_2}/\overline{rcv_3}$, and then rely on the smaller one. All three values will be measured multiple times to avoid the

interference from uncertain network states and possible packet loss. We only compare the averages of these three values.

We simply compare the averages of rcv 's instead of using other more statistically rigorous methods because we find the rcv 's (i.e., rcv_1 , rcv_2 , rcv_3) we measured on one same RVP in multiple measurements are relatively consistent in our early experiments. On average, 88.3% of them are same as their modes, i.e., the rcv value that appears most often in multiple measurements.

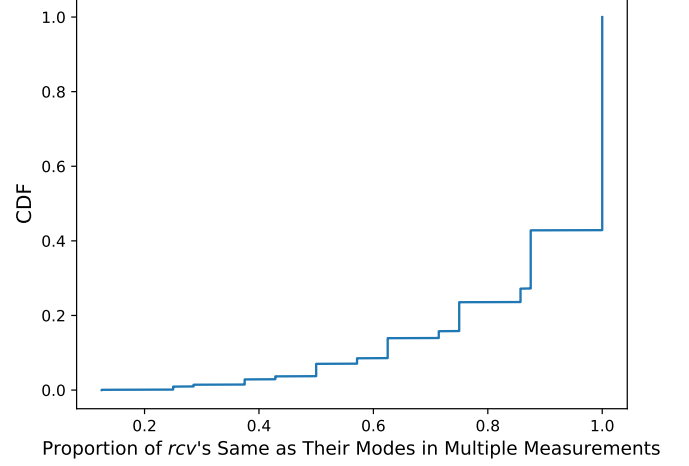


Fig. 5. Distribution of Proportions of rcv 's that are Same as Their Modes in Multiple Measurements. Results are from Early ISAV Measurements on $\sim 2k$ RVPs.

C. A Large-scale Measurement

1) Experiment Setup: We aim at measuring the deployment of ISAV in all the advertised BGP prefixes. We choose BGP prefix-level granularity because some (large) ASes may contain BGP prefixes with different ISAV policies. Note that our method can be used to measure networks of any size. We measure from a single vantage point on the campus network. Our measurements last for about 3 months.

Target Selection. After RVP discovery, we collected data pairs distributed in $\sim 30k$ BGP prefixes. For every BGP prefix, we will choose an appropriate RVP within it. It is common to have many RVPs within one target BGP prefix, so we prefer the RVPs implementing ICMP rate limiting that are neither too strict nor too loose. If the ICMP rate limiting of an RVP is too strict, e.g., replying with only one ICMP error message no matter how many ICMP Echo Requests it receives, it is difficult to observe the different states of its ICMP rate limiting. Similarly, suppose the ICMP rate limiting is too loose, e.g., replying with 50 ICMP error messages in response to 50 ICMP Echo Requests that we send out quickly. In that case, it is also difficult for us to observe the difference. Sending more ICMP Echo Requests, e.g., 100, may help us observe more obvious ICMP rate limiting of RVPs, but it may have a negative impact on the target network. Therefore, RVPs that implement moderate ICMP rate limiting are prioritized. In practice, we will ignore the RVPs with $rcv_1 = N$ or $rcv_1 = 1$ because the ICMP rate limiting is too strict or loose unless such RVPs are the only ones in the target network.

How Many Packets to Send? To determine the values of N and M , we first consult some documentations of well-known router manufacturers like Cisco, Juniper, Huawei and HPC [39], [20], [37], [19]. We found that both the implementations and default thresholds of ICMP rate limiting vary considerably with different devices, OS versions and manufacturers. Some of them implement a token bucket-based (e.g., a maximum of 10 tokens and new token placed at an interval of 100 milliseconds [19]) as recommended by the RFC [12], while there are also many devices that rate limit the packets per second (pps) of ICMP packets (ranging from 1 pps [39] to 100 pps [37]). Based on the findings, we believe it's most reasonable to choose a value lower than the maximum pps we found (i.e., 100 pps) as the value of N and choose a bigger value as M to make $N + M$ apparently higher than 100.

Actually, there is a trade-off when we choose the number of packets to send. More packets trigger more obvious ICMP rate limiting but may bring ethical issues or lead to waste of resource. Sending few packets is more friendly to target networks but may result in less observable rate limiting. We randomly choose ~ 1000 RVPs to perform early experiments.

First, we test how many packets (i.e., $N + M$) are sufficient enough to trigger observable rate limiting. If $N + M$ packets still fail to trigger obvious rate limiting, our method doesn't work. Therefore, we require $N + M$ packets to result in a relatively obvious rate limiting (e.g., 5%, 10% or 20% decline in the replies we received); otherwise, the packets are *insufficient*. Figure 6 shows the proportion of insufficient cases when sending different numbers of packets. Packets with a total number of 30-100 remain insufficient for about 15% RVPs, where we can only observe less than 5%, 10% or 20% decline. The proportion of insufficient cases decreases significantly when the packet number exceed 100 and 120. 130 packets are sufficient in more than 95% cases. The experimental results are consistent with our expectations according to preliminary findings. Therefore, a total of ~ 130 packets seem to be adequate for our measurements.

Secondly, we test how many probe packets and noise packets can make the impact of noise packets more observable. Since we have already total number of ~ 130 packets (i.e., $M + N$) are already sufficient in most cases, what we need to do is to determine the ratio of probe packets (N) to noise packets (M). Taking into account possible packet loss in more complicated network environments when it comes to Internet-wide measurements, we define $M + N = 150$ and then try different values of M/N . We define the *observability* as the decline in the proportion of the replies we can receive after we send noise packets (i.e., $1 - rcv_{after}/rcv_{before}$) and then measure the observability of different ratios of M to N (i.e., M/N). Table I implies that bigger M/N will result in better observability and the observability stabilizes at ~ 0.4 when $M/N \geq 2$. We finally choose $M/N = 2$ instead of 2.5 or larger because we think too many noise packets and too few probe packets cannot help us characterize the behavior of rate limiting well.

We finally choose $N = 50$ and $M = 100$ in our measurements, which is in line with our previous expectations with N being a value smaller than 100 and $N + M$ being bigger than 100. The research community can also choose other

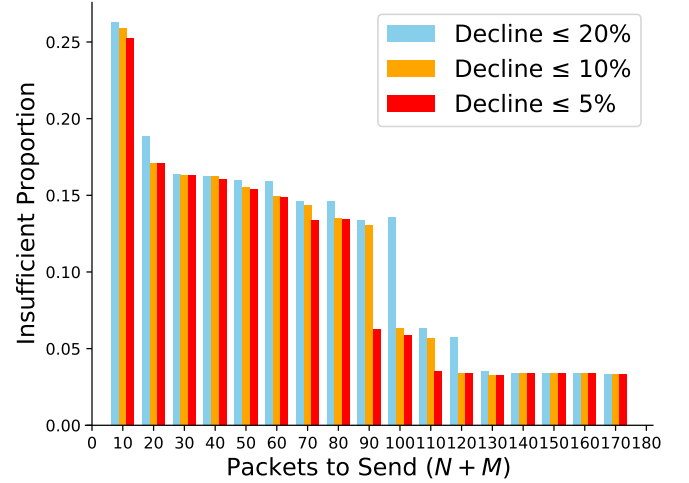


Fig. 6. Insufficient Proportion of RVPs When Sending Different Numbers of Packets

TABLE I. OBSERVABILITY OF DIFFERENT RATIOS OF NOISE PACKETS TO PROBE PACKETS

	M/N				
	0.5	1	1.5	2	2.5
M (Noise Packets)	50	75	90	100	107
N (Probe Packets)	100	75	60	50	43
Observability	0.3171	0.3650	0.4072	0.4192	0.4272

values to balance those trade-offs when performing similar measurements.

Choice of λ . The choice of λ can be very hard because of limited ground truth of ISAV deployment. The existing measurement results of ISAV are mostly private because of potential risks of those no-ISAV networks being attacked. Actually, the main problem is how obvious is obvious enough? For instance, $\overline{rcv} = 0$ obviously proves the presence of rate limiting, but how about $\overline{rcv} = 0.2N, 0.4N, \dots, 0.7N$? After comparing different λ in small-size ground truth by DNS-based measurements like [41], [14] and considering the evaluation result of the reachability measurement with ground truth (see Table III in §VI), we finally choose $\lambda = 0.6$ in the measurements. That is, when inferring the ISAV deployment by the rules introduced earlier, we confirm that $a < b$ only if $a < 0.6 \times b$. Experiments show that the value of λ in the range of 0.5 to 0.7 will not obviously affect the results. After all, in a general sense, a $\sim 40\%$ decline in the *average* number of ICMP messages is believed to be self-evident enough to reveal the triggering of ICMP rate limiting.

Spoofed Source Addresses. As for the spoofed source address, we choose an arbitrary address with the same /80 prefix as the local vantage point as the spoofed source address in the source network. Actually, we can choose any address in the source network, as long as the spoofed source packets can be sent to the Internet. Similarly, we choose an arbitrary IP address with the same /124 prefix as the RVP as the spoofed source address in the target network, which ensures that the spoofed source address is indeed in the same target network as the RVP.

Measurements. We continuously measure the values of rcv_1 , rcv_2 , and rcv_3 of each BGP prefix. However, we do not

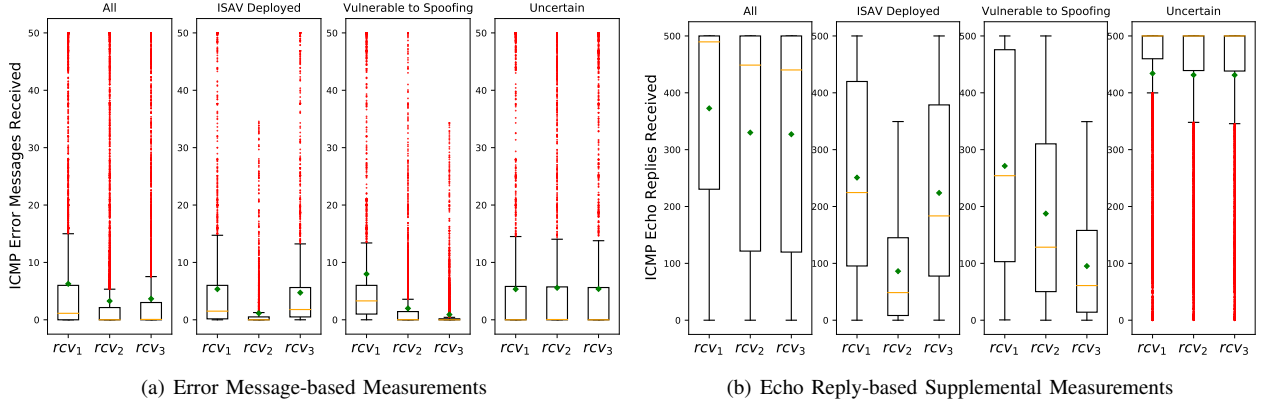


Fig. 7. Distributions of rcv_1 , rcv_2 and rcv_3 under Different Categories. (Orange Lines: Medians. Green Squares: Averages)

TABLE II. MEASUREMENT RESULTS OF ISAV DEPLOYMENT

BGP Prefixes		Autonomous Systems	
Vulnerable to Spoofing	20830 (75.66%)	Vulnerable to Spoofing	6520 (67.37%)
ISAV Deployed	6702 (24.34%)	ISAV Deployed	2020 (20.87%)
Uncertain	19319	Inconsistent	1137 (11.75%)
Total	46851	Total	9677

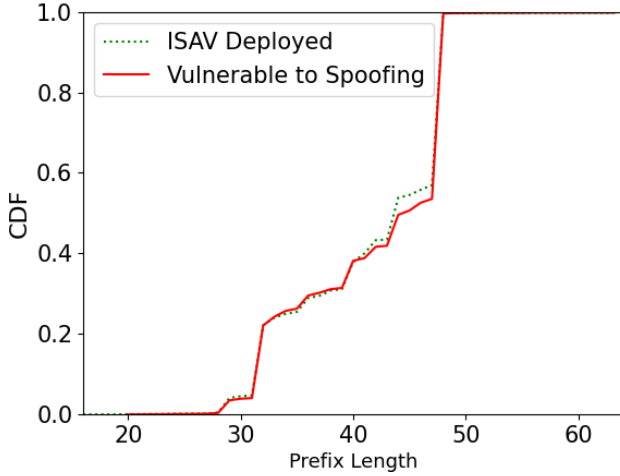


Fig. 8. Prefix Lengths vs. ISAV Deployment

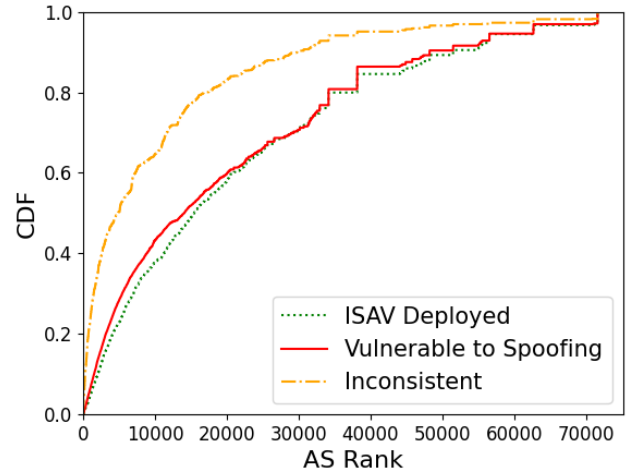


Fig. 9. AS Ranks vs. ISAV Deployment

measure these three values for a prefix consecutively. If we measure rcv_1 , rcv_2 , and rcv_3 for a prefix without interruption, the ICMP rate limiting for an RVP may be triggered repeatedly in a short time. Instead, we will first measure rcv_1 for the first prefix and then rcv_1 for the second prefix. After measuring all the values of rcv_1 , we will measure rcv_2 for the first prefix. For every prefix, we will measure ~ 10 times and calculate the averages of rcv_1 , rcv_2 , and rcv_3 to avoid potential random errors or statistical bias.

Supplemental Measurements. For those BGP prefixes that remain uncertain after measurement, we also perform additional measurements using the rate limiting of ICMP Echo Replies, though, the rate limiting of ICMP Echo Replies is much looser and challenging to observe (see §VII). That is to say, we directly send ICMP Echo Requests to RVPs instead of unreachable IP addresses and capture the ICMP Echo Replies

instead of ICMP error messages. To enrich the targets, we add active IP addresses within those uncertain BGP prefixes from the IPv6 hitlist (6.3M addresses) collected by Gasser et al. [25]. However, we prioritize the RVPs we previously discovered because many of the IP addresses from the hitlist are core routers instead of the peripheries. They implement extremely loose ICMP rate limiting, which is difficult for us to observe the difference. We let $N = M = 500$ in the ICMP Echo Reply-based supplemental measurements. Table IV shows that $N = 500$ is sufficient to trigger observable rate limiting of ICMP Echo Replies for $\sim 65\%$ of the targets. Larger N and M may help us observe more obvious rate limiting, but it may lead to ethical issues since numbers larger than 500 are too large.

2) *Results:* Figure 7 shows the distribution of rcv_1 , rcv_2 , and rcv_3 as measured in the ICMP error message-based

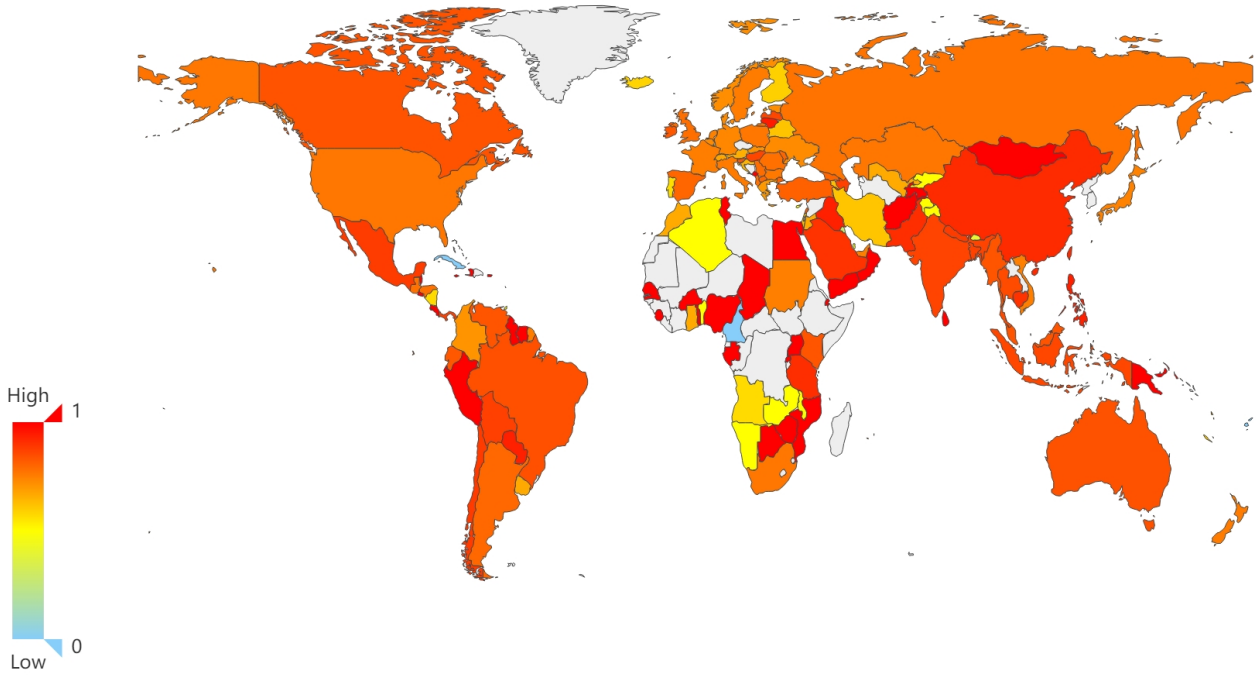


Fig. 10. Fraction of Vulnerable to Spoofing vs. All ASes Per Country

measurements and in the supplemental ICMP Echo Reply-based measurements. For all target networks, we find that rcv_1 is larger than rcv_2 and rcv_3 because the noise packets make a difference by triggering or exacerbating the ICMP rate limiting of the RVPs. We infer the ISAV deployment and then draw the box plots of the three values under three different categories. We find that the results are as expected. For networks where we infer that ISAV is deployed, we have $rcv_1 \approx rcv_3 > rcv_2$, while for networks lacking ISAV, $rcv_1 > rcv_2 \approx rcv_3$ or $rcv_1 > rcv_2 > rcv_3$, which also shows that our simple rules do work. In addition, there are networks for which we cannot infer their ISAV deployment because $rcv_1 \approx rcv_2 \approx rcv_3$.

Table II is the summary of the ISAV measurement results. In comparison with relatively well-deployed OSAV [47], [9], the lack of ISAV is alarming. Our measurements, the most large-scale study of IPv6 ISAV to date, cover 27,532 BGP prefixes, 9,677 ASes, and 186 countries, and identify $\sim 2\times$ more IPv6 ASes lacking ISAV than the state-of-the-art measurements [14].

The results show a pervasive absence of ISAV. Intuitively, one may think that some large BGP prefixes, high-rank (according to AS Rank [8]) ASes or developed countries are more likely to deploy ISAV, but the results are counterintuitive. Figure 8 and Figure 9 show no clear correlation between the ISAV deployment and the sizes (ranks) of the BGP prefixes (ASes). However, it is noteworthy that Figure 9 shows that high-rank ASes tend to have different ISAV policies, which is as we expected. It is not reasonable for some very large (or high-rank) AS to have only one ISAV policy. Figure 10 provides a global map of ISAV deployment.

3) *Validation*: As there is no ground truth of Internet-wide ISAV deployment, we directly or indirectly validate the measurement results from the following aspects:

Validation by DNS Resolvers. We replicate the previous work [41], [14] on measuring ISAV based on DNS resolvers by registering a dedicated domain name and deploying an authoritative DNS server. We scanned the entire IPv4 address space to discover all open DNS resolvers by adding an extensive probe module to ZMap [16] and found 1,950,278 potential open DNS resolvers, of which 1,212,406 could forward queries to our authoritative DNS server. For each potential IPv4 open DNS resolver, we obtained its IPv6 address using DNS based on the method introduced by Hendriks et al. [33] and finally found 4,692 potential IPv6 addresses of open DNS resolvers. We repeat the measurements in [41], [14], and we find that for BGP prefixes identified by our measurements as vulnerable to spoofing, 87.94% of them were confirmed by ISAV measurements based on DNS resolvers. The conflicting portion of the results could be attributed to different ISAV policies in one BGP prefix. The open DNS resolvers we find are mainly in surprisingly large BGP prefixes, with 4,692 IPv6 addresses belonging to BGP prefixes with an average length of 23.91. Thus, different ISAV policies may be deployed in different subnets of such a large network prefix.

Comparison with Previous Work. Luckie et al. [47], [9] found that 68.6%-74.2% of tested IPv6 ASes lacking or partially lacking ISAV deployment. Their results are compelling because *Spoofers* relies on in-network volunteers. Our measurements also found that 67% -79% of tested IPv6 ASes lacking or partially lacking ISAV deployment. Deccio et al.

[14] found that $\sim 50\%$ of tested IPv6 ASes lacking ISAV based on DNS resolvers. Their number is relatively small because there are no open DNS resolvers deployed in some ASes. Similar to our results, Dai et al. [13] found 69.8% ASes lacking ISAV.

D. Limitations

The main limitation of our approach is that it will not work well in networks where we cannot find one RVP or the RVPs (or the active IP addresses from the hitlist [25]) we find do not implement (global) ICMP rate limiting, or just implement very loose or strict ICMP rate limiting. However, all of these obstacles will not result in misjudgments because those networks will remain uncertain.

VI. MEASURING REACHABILITY

In this section, we apply iVANTAGE to the measurements of reachability between two remote nodes on the Internet, neither of which is under our control. Measuring the reachability between two remote nodes requires us to “send” and “receive” packets without controlling them, which can be done by iVANTAGE as mentioned before.

A. Methodology

To measure the reachability between two Internet nodes, namely A and B, we first need to find an RVP. It will be easier if either A or B is already an RVP that we have discovered in the process of “vantage point” discovery. However, suppose that neither A nor B is an RVP. In that case, we refer to the previously collected RVPs to find an appropriate RVP as close to A or B as possible (i.e., a *proxy RVP*). Usually, we can find at least one appropriate proxy RVP for A or B since the 1.1M RVPs are widely distributed across $\sim 30k$ BGP prefixes, $\sim 9.5k$ ASes, and 182 countries. A proxy RVP can be a compelling

proxy for A (or B) because they are usually in the same BGP prefix (or even a smaller prefix). The assumption here is that the loss of reachability between two very close Internet nodes is unlikely to occur. Assuming that A is an RVP, our method consists of the following steps, as shown in Figure 11.

- 1) Local prober P sends N ICMP Echo Requests (very rapidly) to an unreachable IP address X (i.e., the *target* in the data pairs collected previously), which will result in ICMP error messages sent from A.
- 2) P receives rcv_1 ICMP error messages from A. In most cases, $rcv_1 < N$ because of the ICMP rate limiting.
- 3) P sends M ICMP Echo Requests to B with a spoofed source address, which is the same unreachable IP address we use in step 1.
- 4) B replies with M ICMP Echo Replies sent to that unreachable IP address. Note that ICMP Echo Replies are much more difficult to trigger the ICMP rate limiting, so B replies with M ICMP Echo Replies.
- 5) Since these replies are sent to an unreachable IP address, the last hop router A replies with several ICMP error messages. *If A is unreachable from B, this step does not happen.*
- 6) At approximately the same time, P again sends N ICMP Echo Requests to that unreachable IP address without spoofing the source address.
- 7) P receives rcv_2 ICMP error messages.

B. Inferring Reachability

We infer reachability by comparing rcv_1 and rcv_2 . If B can reach A, A will receive ICMP Echo Replies sent from B to an unreachable IP address, which will trigger ICMP rate limiting on A. So in step 7, P cannot receive as many ICMP error messages as before. If B cannot reach A, A will not receive the messages sent by B. Then step 5 will not happen. Therefore, its rate limiting will not be triggered, and we can receive as many ICMP error messages as before. We also infer the reachability based on the average of repeated experiments to avoid possible packet loss or other interference in a single experiment.

C. Measurements & Evaluation

1) *Ground Truth*: Almost any two nodes on the Internet can reach each other, so it is difficult to find some unreachable pairs of IP addresses in such a large IPv6 address space. We first obtain 10M active IPv6 addresses by the method proposed by Song et al. [71]. Then, we deploy two vantage points A and B located in two different continents, approximately 8,000 km from each other. By scanning those 10M addresses using ZMapv6 [26] from both vantage points, we finally find 149 consistently abnormal IP addresses as ground truth (“needles in a haystack!”). These addresses are reachable from A, but packets sent from them cannot reach B. This could be attributed to link failures or inter-domain routing failures because most of these addresses are in relatively small ASes (with an average AS rank [8] of 12059.7) and are also geographically distant (geolocated by GeoLite2 [49]) from the vantage point B (only 6/149 are in the same continent as the vantage point B). In addition, we select additional 851 normal IP addresses reachable from both A and B as a control group, adding up to 1,000 IP addresses. After all, abnormal

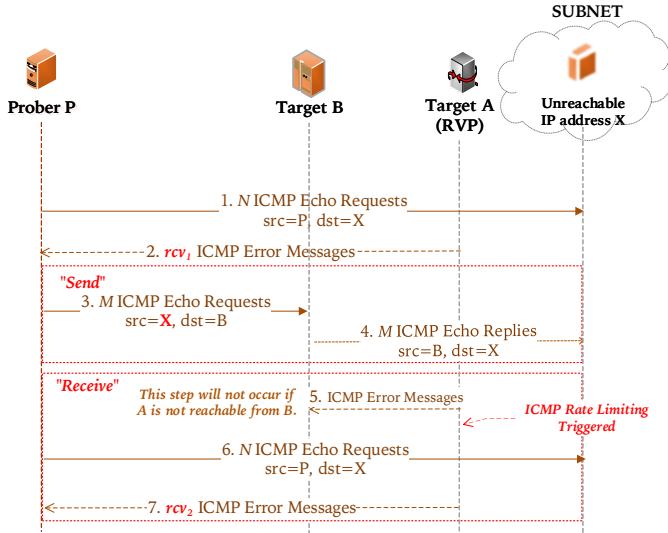


Fig. 11. Methodology for Measuring Reachability between Internet Node A and B. We “send” probes using Target B, and then “receive” packets on Target A. Though only one RVP is shown in the figure, actually, Target B is also used as a “vantage point” to “send” our probes.

IP addresses with reachability problems are much fewer than those normal IP addresses if we perform real measurement on the Internet.

2) *Measurements*: We perform measurements from vantage point A, aiming at distinguishing these 149 abnormal IP addresses that are unconnected with vantage point B from other 851 normal IP addresses. From the data pairs we previously collected, we find 3 RVPs in the same /56 prefix of B as the proxy RVPs. To prevent continuous ICMP rate limiting on the same router, these 3 RVPs are used in rotation, and also with an interval of five minutes.

With the method introduced earlier, we continuously record how many ICMP error messages the proxy RVP can receive in step 7 (i.e., rcv_2) because it is not necessary to repeat step 1 and 2 in a multi-target measurement. We send $N = 50$ probe packets and $M = 100$ noise packets. We measure every rcv_2 6 times and then calculate the average.

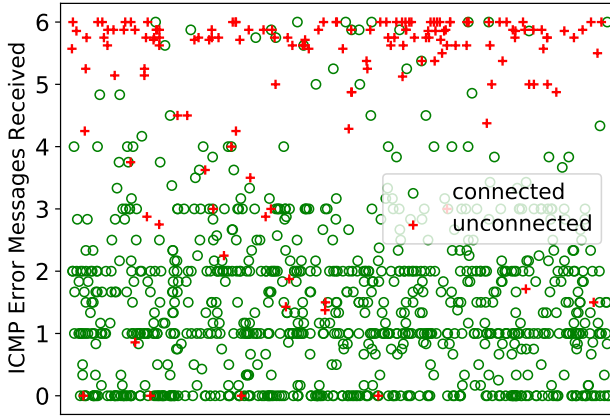


Fig. 12. Average Numbers of ICMP Error Messages Received ($\overline{rcv_2}$) for Connected and Unconnected IP Addresses

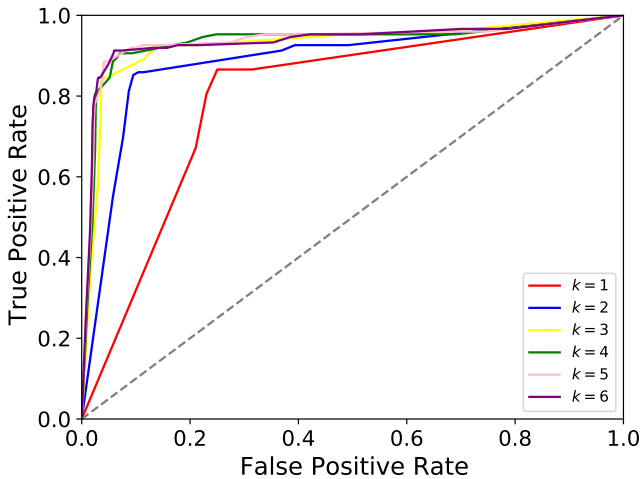


Fig. 13. ROC Curve of Our Method for Measuring Reachability⁵

3) *Evaluation*: Figure 12 shows the average values of rcv_2 when measuring the reachability between vantage point B and connected or unconnected IP addresses. Note that for unconnected IP addresses, we obviously receive more ICMP

error messages (specifically 5.131 vs. 1.472 on average) as expected.

TABLE III. EXPERIMENTAL RESULTS OF OUR REACHABILITY MEASUREMENTS ($k = 6$)

λ	Precision	Recall	Accuracy	F-Score
0.5	0.532	0.899	0.867	0.668
0.6	0.710	0.872	0.928	0.783
0.7	0.814	0.852	0.949	0.833
0.8	0.829	0.812	0.947	0.820
0.9	0.853	0.691	0.937	0.766

We also introduce a threshold λ , as we did in the ISAV measurements. If there is $\overline{rcv_2}/\overline{rcv_1} \geq \lambda$, we will infer that this IP address is unconnected with the vantage point B. Figure 13 is the receiver operating characteristic (ROC) curve of our method when we apply different λ for k -time measurements. When $k \geq 4$, the area under the curve (AUC) stabilizes at 0.92, which demonstrates the capability of our method and suggests that four measurements have been sufficient to make the results stable and convincing. Table III lists the experimental results of our measurements with different λ and $k = 6$, revealing a relatively stable accuracy of over 90%, precision and recall of over 80%, which we believe is already very good for a theoretically impossible measurement task.

D. Limitations & Challenges

Estimation of RTTs. The main challenge is how we control the interval between step 3 and step 6 so that the reflected packets sent by B (i.e., the ICMP Echo Replies) reach A at the same time as the probe packets sent in step 6 arrive. In practice, we estimate the RTT between A and B first (very roughly!) based on their geolocation information [11] (using GeoLite2 [49]) and the triangle principle [59], [35], where d represents the distance between A and B, and c is the speed of light:

$$\begin{cases} \frac{d}{2c/3} \leq \widehat{RTT}_{AB} \leq \frac{d}{c/3} \\ |RTT_A - RTT_B| < \widehat{RTT}_{AB} < RTT_A + RTT_B \end{cases}$$

We randomly choose different value within this range as the estimated RTT for each measurement. We then set the time interval (Δt) between step 3 and step 6 to $\Delta t = (RTT_B - \widehat{RTT}_{AB} + RTT_A)/2$ so that the reflected packets and our probe packets arrive simultaneously. Note that $\Delta t < 0$ is possible because we may do step 6 first instead of step 3.

The RTT estimation does not need to be very accurate for two main reasons. First, in practice, a slightly longer Δt is still fine because it is acceptable to let the reflected packets arrive a little earlier than the probe packets since the ICMP rate limiting lasts for a short period of time. Second, we randomly try different estimates in each measurement, so it would be sufficient if *some* estimates were relatively accurate, which would result in a significant decrease in $\overline{rcv_2}$. However, it is unacceptable if our probe packets arrive earlier. Further analysis shows that there is usually an unexpectedly small RTT between the IP address that is misclassified in the evaluation

⁵Positive: Unconnected, Negative: Connected

and the vantage point B. An incorrect estimate of RTT is likely to trigger a premature or late ICMP rate limiting, posing an obstacle to our measurements. However, this challenge also reveals a novel approach for estimating the latency between two arbitrary nodes, which we leave for future work (§VIII).

Coverage of RVPs. Our RVPs cover $\sim 30\%$ of BGP prefixes, $\sim 50\%$ of ASes, and almost all countries. Usually, our main concern is the reachability between nodes in different prefixes, ASes, or countries. Therefore, we have theoretically $\sim 51\%$, $\sim 75\%$, or almost 100% probability of finding at least one appropriate proxy RVP in the same prefix, AS or country as either of two targets, respectively. Considering that networks where we can't discover any RVPs may have very few active IPv6 hosts, the probabilities can be higher in our practical measurements. Actually, in terms of the numbers of ASes, the coverage of our RVPs is already better than all the existing censorship monitoring platforms [52], [58], [68], [73], [63], [23].

Note that generally, this method will not be affected by the ISAV deployed in the target networks because the spoofed source addresses of the packets we send in step 3 usually do not belong to the target networks, and thus will not be filtered by ISAV.

VII. DISCUSSION: ICMP RATE LIMITING

This section provides a detailed discussion of ICMP rate limiting with respect to Internet-wide implementation, security and privacy risks, and possible mitigation measures, respectively.

A. How do Internet Nodes Implement ICMP Rate Limiting?

We perform a large-scale measurement of the implementation of ICMP rate limiting. We select 25,741 RVPs among 25,741 different longest-match BGP prefixes belonging to 8,834 ASs, focusing on *breadth* rather than quantity, to provide a comprehensive view of ICMP rate limiting implementation.

We aim at measuring the rate limiting of three types of ICMP messages: *Time Exceeded* (by adjusting the hop limits), *Destination Unreachable* (by sending ICMP Echo Requests to unreachable addresses, which are collected previously), and *Echo Reply* (by sending ICMP Echo Requests to the RVPs directly). We also test whether the ICMP rate limiting is global by sending additional noise packets. Global, in this case, means limiting the rate of generation of ICMP error messages sent to all IP addresses, even if triggered by only one IP address. As we did in our ISAV measurements, we continuously measure the values of rcv_1 and rcv_2 for these three types of ICMP error (or informational) messages. As in the ISAV measurements, we let $N = 50$, $M = 100$ for ICMP error messages, and $N = M = 500$ for ICMP Echo Replies.

Based on the measurement results (Figure 14), we further calculate the percentage of global, strict, and loose ICMP rate limiting implementations (Table IV):

- **Global:** $\overline{rcv_2} < \lambda \times \overline{rcv_1}$ (we set $\lambda = 0.6$ in practice). This kind of ICMP rate limiting can be well exploited by iVANTAGE.

- **Strict:** $0.95 \leq \overline{rcv_1} \leq 1.05$. This kind is more secure because it is difficult (but still possible) to observe the difference before and after the rate limiting is triggered.
- **Loose:** $\overline{rcv_2} \geq 0.95N$, i.e., very loose (or even no) ICMP rate limiting is implemented, which can not be exploited as side channels but is vulnerable to ICMP flooding attacks.

TABLE IV. PERCENTAGES OF GLOBAL, STRICT, AND LOOSE ICMP RATE LIMITING IMPLEMENTATIONS

Type	Global	Strict	Loose
ICMP Destination Unreachable	72.16%	15.46%	2.41%
ICMP Time Exceeded	38.84%	1.94%	21.03%
ICMP Echo Reply	40.11%	0.88%	35.63%

Findings. We make the following findings:

- ICMP rate limiting is prevalent, with at least 65%-98% of the tested targets implementing significant ICMP rate limiting⁶. Of these, the rate limiting of ICMP Destination Unreachable is more stringent and has been observable in more than 97% of cases with $N = 50$. In contrast, the rate limiting of other two types is looser and more difficult to observe. Therefore, iVANTAGE can make good use of ICMP Destination Unreachable without sending a large number of packets.
- Global ICMP rate limiting is indeed common, especially in the case of ICMP Destination Unreachable. We estimate that more than 50% of Internet nodes enforce global rate limiting of all ICMP messages.
- The majority of tested Internet nodes ($> 70\%$) implement global rate limiting of ICMP Destination Unreachable. Moreover, the rate limiting for ICMP Destination Unreachable is easy to trigger and observe, with only $\sim 18\%$ implementing strict or loose rate limiting, indicating that iVANTAGE can be widely used for different RVPs distributed across the Internet.

B. Potential Risks

Internet standards keep striving to remove global things to protect from potential side channel-based attacks and measurements [28], [29]. Researchers exploit global IPID counters to perform alias resolution [40], [46], stealthy scans [3], and TCP hijacking [21]. In case that IPID counters are not global, global SYN caches are used as substitutes [18], [78]. Global ICMP rate limiting, though less harmful, can also be dangerous. By taking advantage of the large IPv6 address space, it is easy to induce IPv6 nodes to originate ICMP Destination Unreachable messages. Our measurements also confirm that the rate limiting of ICMP Destination Unreachable messages is easily observable and mostly global and thus can be well exploited as a side channel. Therefore, global ICMP rate limiting can be a good substitute for the well-known global IPID counter for side channel-based measurements. After all, the IPID field has been removed from the IPv6 fixed header. The only difference is that we observe the state of ICMP rate

⁶As we limit the amount of packets we send because of ethical concerns, some loose implementations may no longer be loose if M and N increase.

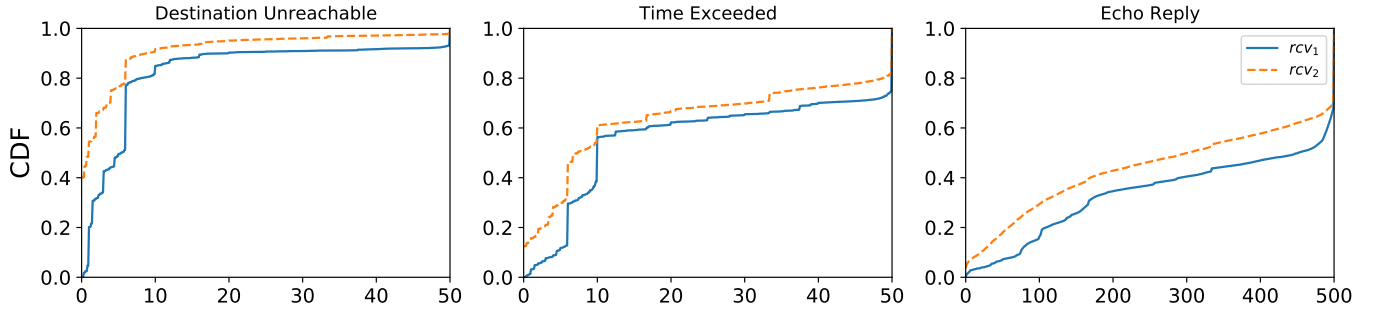


Fig. 14. Distributions of rcv_1 and rcv_2 of Different Types of ICMP Messages

limiting, not the increment of the IPID. It is often easier to observe the state of ICMP rate limiting in the presence of noise. In addition to our work, the side channels of ICMP rate limiting also reveal router configurations [74] and open ports [48]. All these works demonstrate that ICMP rate-limiting side channels may lead to security and privacy issues.

C. Mitigation Measures

We further discuss some possible mitigation measures to prevent ICMP rate limiting from being exploited as side channels.

Strict or Non-global ICMP Rate Limiting is Recommended. Non-global rate limiting is an intuitive and effective solution. For example, rate limiting for ICMP port 444 unreachable will not interfere with the rate limiting state that generates ICMP port 445 unreachable, and rate limiting for ICMP messages sent to source A will not interfere with the rate of ICMP messages sent to source B. However, non-global rate limiting is difficult to implement and deploy because too many rate limiting counters need to be maintained (e.g., token buckets as recommended by the RFC [12]). An easy-to-deploy alternative solution that we recommend is to implement strict ICMP rate limiting, i.e., to send only one ICMP message regardless of how many packets are received in a short period of time. iVANTAGE and other efforts to exploit ICMP rate limiting [74], [48] rely on receiving a different number of ICMP messages before and after triggering ICMP rate limiting. Even though still global, strict rate limiting makes the differences much less observable. However, strict ICMP rate limiting cannot cope with bursty traffic and is therefore not recommended by the RFC [12]. Thus, there may be a trade-off, and it is still difficult to find a perfect solution. Side channels of ICMP rate limiting may be exploitable for a long time to come.

ICMP Destination Unreachable Should be Restricted.

It is easy to find an unreachable IP address in such a large IPv6 address space, so it is also easy to induce IPv6 nodes to initiate ICMP Destination Unreachable messages. The node initiating an ICMP Destination Unreachable message exposes itself, which can then be exploited. Just like the process of RVP discovery, actually we do not have any active IPv6 addresses at first. However, by sending ICMP Echo Requests to those fabricated destination IP addresses, we can discover a great many active IPv6 addresses by receiving ICMP error messages. Then, they can even be used as our “vantage

points”! Therefore, allowing IPv6 nodes to generate ICMP Destination Unreachable messages without any restrictions will be dangerous. For example, when a router receives a series of packets destined for very strange destination addresses within its subnet (especially if these packets are sent from a remote network), it may be a safer choice to ignore them than to initiate ICMP destination unreachable messages for each packet. Especially, we also recommend not initiating ICMP error messages in response to ICMP Echo Replies⁷.

VIII. LIMITATIONS AND FUTURE WORK

iVANTAGE makes use of remote routers as “vantage points” to perform active measurements via ICMP rate limiting side channels. Admittedly, there are still some limitations of iVANTAGE and our work. We look forward to improving our work in the future from following aspects:

A. More iVANTAGE-based Measurement Applications

We present two typical applications of iVANTAGE in this paper, which seems to be kind of limited, but we believe that our work is just the beginning. There are still many other measurements or attacks that can be performed based on iVANTAGE to be discovered. Note that it is also possible to send other packets (e.g., TCP-SYN, UDP, DNS queries, NTP, etc.) instead of ping as probe packets. **All packets that the destination needs to reply to can be used as probe packets, and we can measure the reachability of these packets based on iVANTAGE.**

For instance, Figure 15 shows a novel idea we have just come up with based on iVANTAGE to discover hidden machines in remote network. Just like the famous TCP idle scan [3] and its variations [18], [78], [21], this new idea exploits a zombie machine (i.e., RVP) to discover hidden machines that only respond to specific devices in its own networks (e.g., only reply to the ping sent from the network it belongs to). This is common, especially for some local servers, such as local DNS resolvers. We will deeply dive into this new idea in the future, and there may also be other novel and also interesting ideas based on iVANTAGE like that to be found.

⁷Though IPv6 nodes are not allowed to originate ICMP error messages as a result of receiving ICMP error messages according to the RFC [12], ICMP Echo Replies are not ICMP error messages.

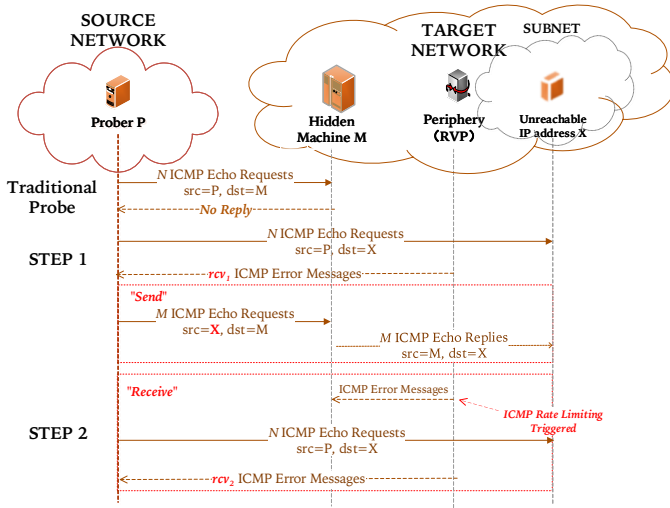


Fig. 15. Novel Idea based on iVANTAGE to Discover Hidden Machines. The existence of the hidden machine can be identified by comparing rcv_1 and rcv_2 .

B. Demystifying ICMP Rate Limiting Behaviors

This paper presents an Internet-wide measurement on ICMP rate limiting implementations, but we think there is more to take a deep dive into. For example, we just identify ICMP rate limiting in a very intuitive (but we believe effective) way: by comparing the number of replies we received. We are going to focus on the characterization of behaviors and the implementations of ICMP rate limiting, which may help us identify rate limiting more accurately and efficiently with fewer packets to be sent.

C. Accurate Latency Estimation between Remote IPv6 Nodes

As mentioned earlier in §VI, by checking the state of ICMP rate limiting, we can know whether the packet from step 4 or step 6 arrives first at A. Thus, by controlling Δt between step 3 and step 6 and checking the state of ICMP rate limiting on A, we can estimate the RTT between A and B. This can also be considered to be another measurement application of iVANTAGE.

D. IPv4 Applications

iVANTAGE faces challenges in IPv4 networks for two reasons: 1) ICMP rate limiting is not required [60] (though many may implement [74], [64], [31], not that pervasive [74], [7]) in IPv4 networks; 2) It's also more difficult to find unreachable addresses to induce possible RVPs to initiate ICMP Destination Unreachable, which is the type of ICMP error message that we mainly exploit in iVANTAGE. How to make iVANTAGE well applicable to IPv4 networks is what we may consider in the future.

IX. CONCLUSION

In this paper, we focus on ICMP rate limiting side channels. We propose a novel technique, iVANTAGE, and apply it to two difficult and theoretically impossible measurement tasks: measuring ISAV deployment and reachability between remote

Internet nodes from only one local vantage point. In addition, we measure the implementation of ICMP rate limiting, reveal the security and privacy risks of existing ICMP rate limiting implementations, and provide possible mitigation measures. We will further study on ICMP rate limiting and its side channels in the future.

X. ETHICAL CONSIDERATIONS

Before performing our measurements, we looked through some key guidelines for Internet measurements [56], [72], [4]. Since our department does not have an Institutional Review Board, we consulted the academic board of our department. Feedback from the academic board mainly concerned the possible effect on the target devices when triggering ICMP rate limiting. We replied that: 1) ICMP rate limiting is a required and basic function of IPv6 nodes, and triggering ICMP rate limiting is also common on the real Internet, generally not considered harmful; 2) our measurements will never trigger continuous ICMP rate limiting on the same target (see below); 3) there already existed several published papers exploiting the ICMP rate limiting mechanism [48], [74], particularly, lab experiments in [74] show that ICMP rate limiting does not have a discernible negative effect on the devices. The academic board approved our study after serious consideration. In practice, we have taken into account following aspects for ethical considerations in our measurements.

A. Anonymity

As many previous work on ISAV, we totally ensure the anonymity of prefixes and ASes we measured to prevent those vulnerable-to-spoofing networks from being attacked by spoofing-based cyberattacks. We also do not make the IP addresses of RVPs we discovered public to the community.

B. Relatively Harmless Probes

We only send ICMP Echo Requests (i.e., ping) in our measurements. Compared with other types of scans like port scans and sending DNS queries, sending ICMP Echo Requests is relatively harmless. We prevent probing one network in succession in all parts of our measurements. For instance, in the process of RVP discovery, we randomized the probing sequence by the multiplicative group of integers modulo n [27], [16] like many existing high-speed probers [16], [7], [36]. Network administrators can easily contact us by the e-mail address in the WHOIS database or the reverse DNS record. We received 5 complaints during our measurements, and their networks (the whole ASes) are excluded from our subsequent measurements.

C. Limiting the Amount of Packets

It is inevitable to send many packets to trigger ICMP rate limiting. However, we strictly limit the amount of packets we send. We only send 50 ICMP Echo Requests to check the state of ICMP rate limiting (500 for the rate limiting of ICMP Echo Replies). If the ICMP rate limiting is relatively loose (e.g., receive 50/50 ICMP error messages), we will not further increase the amount of our packets, even though we know this will help trigger more obvious ICMP rate limiting. Instead, we try another RVP. Similarly, our noise packets will

be no more than 100 (500 for ICMP Echo Replies), even though more noise packets may help us observe more obvious difference. While other studies exploiting ICMP rate limiting like [74] usually send thousands of packets per second, our measurements send much fewer packets.

D. Preventing Continuous ICMP Rate Limiting

Our measurements will **never** trigger ICMP rate limiting on one Internet node in succession. In RVP discovery, probes in random order prevent ICMP error messages from being continuously sent from one periphery; in our ISAV measurements, as introduced before, we will not measure the values of rcv_1 , rcv_2 and rcv_3 of one network without break, instead, we measure rcv_1 of the first network, and then measure rcv_1 of the second network; in our measurements of reachability, we use different RVPs in rotation, and with a relatively long interval of five minutes; when measuring ICMP rate limiting implementations, we also prevent continuous ICMP rate limiting as we did in the ISAV measurements. In our measurements, due to our randomness, after rate limiting is triggered on an RVP (which we find usually lasts for a very short time), there is a long time before rate limiting is triggered again. Thus, the duration of rate limiting is a really tiny fraction compared to their normal operation time and does not interfere with their normal function.

E. Previous Laboratory Experiments

We refer to previous work on rate limiting, e.g., [74], where they conducted real laboratory experiments on routers that sent more packets and triggered more severe rate limiting than our work. They found that rate limiting does not lead to destructive problems or high CPU usage, so our approach of much fewer probe packets and larger intervals does not severely affect networks. As a basic and required function of IPv6 nodes, we believe ICMP rate limiting will not lead to a disruptive impact on either the data plane or the control plane of the target device.

F. Real Internet Experiments

We request access to several edge routers in our campus network from the network administrators and monitor the changes in CPU usages, memory usages when we use them as RVPs in our ISAV and reachability measurements. However, we cannot observe *any* observable changes in CPU usages or memory usages on them, even after we increase the number of packets we send by a factor of 2 or 3. We also cannot observe any abnormal behaviors and we believe our measurements do not interfere with routers' basic functions. We think the reasons may include but is not limited to: 1) For every router, our measurement last only a short time over a long period of time (usually much less than 1 second), and routers usually cannot provide real-time performance monitoring with a granularity of less than 1 second to capture the changes. 2) The number of packets we send is also very small compared with thousands or even millions of packets the router forwards per second. 3) Common token-bucket implementations of ICMP rate limiting, just as recommended by the RFC [12], are really resource-saving and do not result in obvious increase of CPU and memory usages.

ACKNOWLEDGEMENTS

The authors would like to thank Erik Rye et al. [66], [67] and Xiang Li et al. [43] for their previous contributions to the discovery of IPv6 peripheries, which are crucial preliminary to our work. Previous research on ICMP rate limiting side channels by Man et al. [48] and Vermeulen et al. [74] are also great inspirations for us. This work is supported by the National Key Research and Development Program of China under Grant No. 2018YFB1800200 and Beijing Natural Science Foundation under Grant No.4222026. Lin He is the corresponding author of this paper: he-lin@tsinghua.edu.cn.

REFERENCES

- [1] G. Aceto and A. Pescapè, "Internet censorship detection: A survey," *Comput. Networks*, vol. 83, pp. 381–421, 2015. [Online]. Available: <https://doi.org/10.1016/j.comnet.2015.03.008>
- [2] Akamai. (2022) Ipv6 adoption visualization. <https://www.akamai.com/visualizations/state-of-the-internet-report/ipv6-adoption-visualization>.
- [3] Antirez, "New tcp scan method," *Posted to Bugtraq Mailing List*, 1998.
- [4] M. Bailey, D. Dittrich, E. Kenneally, and D. Maughan, "The menlo report," *IEEE Secur. Priv.*, vol. 10, no. 2, pp. 71–75, 2012. [Online]. Available: <https://doi.org/10.1109/MSP.2012.52>
- [5] R. Beverly, "Yarp'ing the internet: Randomized high-speed active topology discovery," in *Proceedings of the 2016 ACM on Internet Measurement Conference, IMC 2016, Santa Monica, CA, USA, November 14-16, 2016*. ACM, 2016, pp. 413–420. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2987479>
- [6] R. Beverly, A. W. Berger, Y. Hyun, and kc claffy, "Understanding the efficacy of deployed internet source address validation filtering," in *Proceedings of the 9th ACM SIGCOMM Internet Measurement Conference, IMC 2009, Chicago, Illinois, USA, November 4-6, 2009*. ACM, 2009, pp. 356–369. [Online]. Available: <https://doi.org/10.1145/1644893.1644936>
- [7] R. Beverly, R. Durairajan, D. Plonka, and J. P. Rohrer, "In the IP of the beholder: Strategies for active ipv6 topology discovery," in *Proceedings of the Internet Measurement Conference 2018, IMC 2018, Boston, MA, USA, October 31 - November 02, 2018*. ACM, 2018, pp. 308–321. [Online]. Available: <https://dl.acm.org/citation.cfm?id=3278559>
- [8] CAIDA. (2022) Caida as rank. <http://as-rank.caida.org/>.
- [9] —. (2022) Spoofer project. <https://www.caida.org/projects/spoofers/>.
- [10] J. Chen, C. Yang, W. Chen, Y. Huang, and H. Chu, "An icmp-based mobility management approach suitable for protocol deployment limitation," *EURASIP J. Wirel. Commun. Netw.*, vol. 2009, 2009. [Online]. Available: <https://doi.org/10.1155/2009/983594>
- [11] D. Cicalese, D. Z. Joumblatt, D. Rossi, M. Buob, J. Augé, and T. Friedman, "Latency-based anycast geolocation: Algorithms, software, and data sets," *IEEE J. Sel. Areas Commun.*, vol. 34, no. 6, pp. 1889–1903, 2016.
- [12] A. Conta, S. E. Deering, and M. Gupta, "Internet control message protocol (icmpv6) for the internet protocol version 6 (ipv6) specification," *RFC*, vol. 4443, pp. 1–24, 2006. [Online]. Available: <https://doi.org/10.17487/RFC4443>
- [13] T. Dai and H. Shulman, "Smap: Internet-wide scanning for spoofing," in *ACSAC*. ACM, 2021, pp. 1039–1050.
- [14] C. Deccio, A. Hilton, M. Briggs, T. Avery, and R. Richardson, "Behind closed doors: A network tale of spoofing, intrusion, and false DNS security," in *IMC '20: ACM Internet Measurement Conference, Virtual Event, USA, October 27-29, 2020*. ACM, 2020, pp. 65–77. [Online]. Available: <https://doi.org/10.1145/3419394.3423649>
- [15] S. E. Deering and R. M. Hinden, "Internet protocol, version 6 (ipv6) specification," *RFC*, vol. 8200, pp. 1–42, 2017. [Online]. Available: <https://doi.org/10.17487/RFC8200>
- [16] Z. Durumeric, E. Wustrow, and J. A. Halderman, "Zmap: Fast internet-wide scanning and its security applications," in *Proceedings of the 22th USENIX Security Symposium, Washington, DC, USA, August 14-16, 2013*. USENIX Association, 2013,

- pp. 605–620. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity13/technical-sessions/paper/durumeric>
- [17] R. Ensafi, J. Knockel, G. Alexander, and J. R. Crandall, “Detecting intentional packet drops on the internet via TCP/IP side channels,” in *Passive and Active Measurement - 15th International Conference, PAM 2014, Los Angeles, CA, USA, March 10-11, 2014, Proceedings*, ser. Lecture Notes in Computer Science, vol. 8362. Springer, 2014, pp. 109–118. [Online]. Available: https://doi.org/10.1007/978-3-319-04918-2_11
 - [18] R. Ensafi, J. C. Park, D. Kapur, and J. R. Crandall, “Idle port scanning and non-interference analysis of network protocol stacks using model checking,” in *19th USENIX Security Symposium, Washington, DC, USA, August 11-13, 2010, Proceedings*. USENIX Association, 2010, pp. 257–272. [Online]. Available: http://www.usenix.org/events/sec10/tech/full_papers/Ensafi.pdf
 - [19] H. P. Enterprise. (2022) Configuring rate limit for icmp error messages. https://techhub.hpe.com/eginfolib/networking/docs/switches/5130ei/5200-3942_13-ip-svcs_cg/content/483572479.htm.
 - [20] C. C. Expert. (2022) Icmp unreachable rate limiting. <https://www.cccexpert.us/isp-essentials/icmp-unreachable-rate-limiting.html>.
 - [21] X. Feng, C. Fu, Q. Li, K. Sun, and K. Xu, “Off-path TCP exploits of the mixed IPID assignment,” in *CCS ’20: 2020 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, USA, November 9-13, 2020*. ACM, 2020, pp. 1323–1335. [Online]. Available: <https://doi.org/10.1145/3372297.3417884>
 - [22] P. Ferguson and D. Senie, “Network ingress filtering: Defeating denial of service attacks which employ IP source address spoofing,” *RFC*, vol. 2827, pp. 1–10, 2000. [Online]. Available: <https://doi.org/10.17487/RFC2827>
 - [23] A. Filastò and J. Appelbaum, “OONI: open observatory of network interference,” in *2nd USENIX Workshop on Free and Open Communications on the Internet, FOCI ’12, Bellevue, WA, USA, August 6, 2012*. USENIX Association, 2012. [Online]. Available: <https://www.usenix.org/conference/foci12/workshop-program/presentation/filast%C3%B2>
 - [24] R. Fontugne, C. Pelsser, E. Aben, and R. Bush, “Pinpointing delay and forwarding anomalies using large-scale traceroute measurements,” in *Proceedings of the 2017 Internet Measurement Conference, IMC 2017, London, United Kingdom, November 1-3, 2017*. ACM, 2017, pp. 15–28. [Online]. Available: <https://doi.org/10.1145/3131365.3131384>
 - [25] O. Gasser, Q. Scheitle, P. Foremski, Q. Lone, M. Korczynski, S. D. Strowes, L. Hendriks, and G. Carle. (2022) Ipv6 hitlist service. <https://ipv6hitlist.github.io/>.
 - [26] O. Gasser, Q. Scheitle, P. Foremski, Q. Lone, M. Korczyński, S. D. Strowes, L. Hendriks, and G. Carle. (2022) Zmapv6. <https://github.com/tumi8/zmap>.
 - [27] C. F. Gauss and W. C. Waterhouse, *Disquisitiones Arithmeticae*. Yale University Press, 1966.
 - [28] F. Gont, “Security assessment of the internet protocol version 4,” *RFC*, vol. 6274, pp. 1–75, 2011. [Online]. Available: <https://doi.org/10.17487/RFC6274>
 - [29] —, “Security implications of predictable fragment identification values,” *RFC*, vol. 7739, pp. 1–20, 2016. [Online]. Available: <https://doi.org/10.17487/RFC7739>
 - [30] Google. (2022) Google ipv6 adoption. <https://www.google.com/intl/en/ipv6/statistics.html>.
 - [31] H. Guo and J. S. Heidemann, “Detecting ICMP rate limiting in the internet,” in *Passive and Active Measurement - 19th International Conference, PAM 2018, Berlin, Germany, March 26-27, 2018, Proceedings*, ser. Lecture Notes in Computer Science, vol. 10771. Springer, 2018, pp. 3–17. [Online]. Available: https://doi.org/10.1007/978-3-319-76481-8_1
 - [32] J. S. Heidemann, Y. Pryadkin, R. Govindan, C. Papadopoulos, G. Bartlett, and J. A. Bannister, “Census and survey of the visible internet,” in *Proceedings of the 8th ACM SIGCOMM Internet Measurement Conference, IMC 2008, Vouliagmeni, Greece, October 20-22, 2008*. ACM, 2008, pp. 169–182. [Online]. Available: <https://doi.org/10.1145/1452520.1452542>
 - [33] L. Hendriks, R. de Oliveira Schmidt, R. van Rijswijk-Deij, and A. Pras, “On the potential of ipv6 open resolvers for ddos attacks,” in *Passive and Active Measurement - 18th International Conference, PAM 2017, Sydney, NSW, Australia, March 30-31, 2017, Proceedings*, ser. Lecture Notes in Computer Science, vol. 10176. Springer, 2017, pp. 17–29. [Online]. Available: https://doi.org/10.1007/978-3-319-54328-4_2
 - [34] R. M. Hinden and S. E. Deering, “IP version 6 addressing architecture,” *RFC*, vol. 4291, pp. 1–25, 2006. [Online]. Available: <https://doi.org/10.17487/RFC4291>
 - [35] S. M. Hotz, “Routing information organization to support scalable interdomain routing with heterogeneous path requirements,” Ph.D. dissertation, University of Southern California, 1994.
 - [36] Y. Huang, M. Rabinovich, and R. Al-Dalky, “Flashroute: Efficient traceroute on a massive scale,” in *IMC ’20: ACM Internet Measurement Conference, Virtual Event, USA, October 27-29, 2020*. ACM, 2020, pp. 443–455. [Online]. Available: <https://doi.org/10.1145/3419394.3423619>
 - [37] Huawei. (2022) Limiting the rate of icmp packets - ar500, ar510, ar531, ar550, ar1500, and ar2500 v200r009 cli-based configuration guide - security - huawei. <https://support.huawei.com/enterprise/en/doc/EDOC1000177804/81605d06/limiting-the-rate-of-icmp-packets>.
 - [38] G. Iannaccone, C. Chuah, R. Mortier, S. Bhattacharyya, and C. Diot, “Analysis of link failures in an IP backbone,” in *Proceedings of the 2nd ACM SIGCOMM Internet Measurement Workshop, IMW 2002, Marseille, France, November 6-8, 2002*. ACM, 2002, pp. 237–242. [Online]. Available: <https://doi.org/10.1145/637201.637238>
 - [39] Juniper. (2022) icmp6 (error message rate limit) — junos os — juniper networks. <https://www.juniper.net/documentation/us/en/software/junos/transport-ip/topics/ref/statement/icmp6-edit-chassis.html>.
 - [40] K. Keys, Y. Hyun, M. J. Luckie, and K. C. Claffy, “Internet-scale ipv4 alias resolution with MIDAR,” *IEEE/ACM Trans. Netw.*, vol. 21, no. 2, pp. 383–399, 2013. [Online]. Available: <https://doi.org/10.1109/TNET.2012.2198887>
 - [41] M. Korczynski, Y. Nosyk, Q. Lone, M. Skwarek, B. Jonglez, and A. Duda, “Don’t forget to lock the front door! inferring the deployment of source address validation of inbound traffic,” in *Passive and Active Measurement - 21st International Conference, PAM 2020, Eugene, Oregon, USA, March 30-31, 2020, Proceedings*, ser. Lecture Notes in Computer Science, vol. 12048. Springer, 2020, pp. 107–121. [Online]. Available: https://doi.org/10.1007/978-3-030-44081-7_7
 - [42] M. Kührer, T. Hupperich, C. Rossow, and T. Holz, “Exit from hell? reducing the impact of amplification ddos attacks,” in *Proceedings of the 23rd USENIX Security Symposium, San Diego, CA, USA, August 20-22, 2014*. USENIX Association, 2014, pp. 111–125. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/kuhrer>
 - [43] X. Li, B. Liu, X. Zheng, H. Duan, Q. Li, and Y. Huang, “Fast ipv6 network periphery discovery and security implications,” in *Proceedings of the 2021 IEEE/IFIP International Conference on Dependable Systems and Networks*, ser. DSN ’21, 2021. [Online]. Available: <https://idealeer.github.io/publication/dsn21/dsn21-ipv6-paper.pdf>
 - [44] F. Lichtblau, F. Streibelt, T. Krüger, P. Richter, and A. Feldmann, “Detection, classification, and analysis of inter-domain traffic with spoofed source IP addresses,” in *Proceedings of the 2017 Internet Measurement Conference, IMC 2017, London, United Kingdom, November 1-3, 2017*. ACM, 2017, pp. 86–99. [Online]. Available: <https://doi.org/10.1145/3131365.3131367>
 - [45] Q. Lone, M. J. Luckie, M. Korczynski, and M. van Eeten, “Using loops observed in traceroute to infer the ability to spoof,” in *Passive and Active Measurement - 18th International Conference, PAM 2017, Sydney, NSW, Australia, March 30-31, 2017, Proceedings*, ser. Lecture Notes in Computer Science, vol. 10176. Springer, 2017, pp. 229–241. [Online]. Available: https://doi.org/10.1007/978-3-319-54328-4_17
 - [46] M. J. Luckie, R. Beverly, W. Brinkmeyer, and kc claffy, “Speedtrap: internet-scale ipv6 alias resolution,” in *Proceedings of the 2013 Internet Measurement Conference, IMC 2013, Barcelona, Spain, October 23-25, 2013*. ACM, 2013, pp. 119–126. [Online]. Available: <https://doi.org/10.1145/2504730.2504759>
 - [47] M. J. Luckie, R. Beverly, R. Koga, K. Keys, J. A. Kroll, and kc claffy, “Network hygiene, incentives, and regulation: Deployment of source address validation in the internet,” in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS 2019, London, UK, November 11-15, 2019*. ACM, 2019, pp. 465–480. [Online]. Available: <https://doi.org/10.1145/3319535.3354232>

- [48] K. Man, Z. Qian, Z. Wang, X. Zheng, Y. Huang, and H. Duan, "DNS cache poisoning attack reloaded: Revolutions with side channels," in *CCS '20: 2020 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, USA, November 9-13, 2020*. ACM, 2020, pp. 1337–1350. [Online]. Available: <https://doi.org/10.1145/3372297.3417280>
- [49] MaxMind. (2022) Geolite2. <https://dev.maxmind.com/geoip/geoip2/geolite2/>.
- [50] S.-J. Moon, Y. Yin, R. A. Sharma, Y. Yuan, J. M. Spring, and V. Sekar, "Accurately measuring global risk of amplification attacks using ampmap," in *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, Aug. 2021. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity21/presentation/moon>
- [51] L. F. Müller, M. J. Luckie, B. Huffaker, K. Claffy, and M. P. Barcellos, "Challenges in inferring spoofed traffic at ixps," in *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies, CoNEXT 2019, Orlando, FL, USA, December 09-12, 2019*. ACM, 2019, pp. 96–109. [Online]. Available: <https://doi.org/10.1145/3359989.3365422>
- [52] A. A. Niaki, S. Cho, Z. Weinberg, N. P. Hoang, A. Razaghpanah, N. Christin, and P. Gill, "Iclab: A global, longitudinal internet censorship measurement platform," in *2020 IEEE Symposium on Security and Privacy, SP 2020, San Francisco, CA, USA, May 18-21, 2020*. IEEE, 2020, pp. 135–151. [Online]. Available: <https://doi.org/10.1109/SP40000.2020.00014>
- [53] OARC. (2022) Dittl data. <https://www.dns-oarc.net/oarc/data/dittl>.
- [54] —. (2022) Domain name system operation, analysis, and research center (oarc). <https://www.dns-oarc.net/>.
- [55] U. of Oregon. (2022) Routeviews. <http://www.routeviews.org/routeviews/>.
- [56] C. Partridge and M. Allman, "Ethical considerations in network measurement papers," *Commun. ACM*, vol. 59, no. 10, pp. 58–64, 2016. [Online]. Available: <https://doi.org/10.1145/2896816>
- [57] P. Pearce, R. Ensafi, F. Li, N. Feamster, and V. Paxson, "Augur: Internet-wide detection of connectivity disruptions," in *2017 IEEE Symposium on Security and Privacy, SP 2017, San Jose, CA, USA, May 22-26, 2017*. IEEE Computer Society, 2017, pp. 427–443. [Online]. Available: <https://doi.org/10.1109/SP.2017.55>
- [58] P. Pearce, B. Jones, F. Li, R. Ensafi, N. Feamster, N. Weaver, and V. Paxson, "Global measurement of DNS manipulation," in *26th USENIX Security Symposium, USENIX Security 2017, Vancouver, BC, Canada, August 16-18, 2017*. USENIX Association, 2017, pp. 307–323. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/pearce>
- [59] P. Popovski, J. J. Nielsen, C. Stefanovic, E. De Carvalho, E. Strom, K. F. Trillinggaard, A.-S. Bana, D. M. Kim, R. Kotaba, J. Park et al., "Wireless access for ultra-reliable low-latency communication: Principles and building blocks," *IEEE Network*, vol. 32, no. 2, pp. 16–23, 2018.
- [60] J. Postel, "Internet control message protocol," *RFC*, vol. 792, pp. 1–21, 1981. [Online]. Available: <https://doi.org/10.17487/RFC0792>
- [61] R. S. Raman, L. Evdokimov, E. Wustrow, J. A. Halderman, and R. Ensafi, "Investigating large scale HTTPS interception in kazakhstan," in *IMC '20: ACM Internet Measurement Conference, Virtual Event, USA, October 27-29, 2020*. ACM, 2020, pp. 125–132. [Online]. Available: <https://doi.org/10.1145/3419394.3423665>
- [62] R. S. Raman, P. Shenoy, K. Kohls, and R. Ensafi, "Censored planet: An internet-wide, longitudinal censorship observatory," in *CCS '20: 2020 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, USA, November 9-13, 2020*. ACM, 2020, pp. 49–66. [Online]. Available: <https://doi.org/10.1145/3372297.3417883>
- [63] R. S. Raman, A. Stoll, J. Dalek, R. Ramesh, W. Scott, and R. Ensafi, "Measuring the deployment of network censorship filters at global scale," in *27th Annual Network and Distributed System Security Symposium, NDSS 2020, San Diego, California, USA, February 23-26, 2020*. The Internet Society, 2020. [Online]. Available: <https://www.ndss-symposium.org/ndss-paper/measuring-the-deployment-of-network-censorship-filters-at-global-scale/>
- [64] R. Ravaioli, G. Urvoy-Keller, and C. Barakat, "Characterizing ICMP rate limitation on routers," in *2015 IEEE International Conference on Communications, ICC 2015, London, United Kingdom, June 8-12, 2015*. IEEE, 2015, pp. 6043–6049. [Online]. Available: <https://doi.org/10.1109/ICC.2015.7249285>
- [65] J. P. Rohrer, B. LaFever, and R. Beverly, "Empirical study of router ipv6 interface address distributions," *IEEE Internet Comput.*, vol. 20, no. 4, pp. 36–45, 2016. [Online]. Available: <https://doi.org/10.1109/MIC.2016.52>
- [66] E. C. Rye and R. Beverly, "Discovering the ipv6 network periphery," in *Passive and Active Measurement - 21st International Conference, PAM 2020, Eugene, Oregon, USA, March 30-31, 2020, Proceedings*, ser. Lecture Notes in Computer Science, vol. 12048. Springer, 2020, pp. 3–18. [Online]. Available: https://doi.org/10.1007/978-3-030-44081-7_1
- [67] E. C. Rye, R. Beverly, and K. C. Claffy, "Follow the scent: defeating ipv6 prefix rotation privacy," pp. 739–752, 2021.
- [68] W. Scott, T. E. Anderson, T. Kohno, and A. Krishnamurthy, "Satellite: Joint analysis of cdns and network-level interference," in *2016 USENIX Annual Technical Conference, USENIX ATC 2016, Denver, CO, USA, June 22-24, 2016*. USENIX Association, 2016, pp. 195–208. [Online]. Available: <https://www.usenix.org/conference/atc16/technical-sessions/presentation/scott>
- [69] A. W. Services. (2020) Aws shield: Threat landscape report - q1 2020. https://aws-shield-tlr.s3.amazonaws.com/2020-Q1_AWS_Shield_TLR.pdf.
- [70] A. Sinha, L. Sejwal, N. Kumar, and A. Yadav, "Implementation of icmp based network management system for heterogeneous networks," in *2015 2nd International Conference on Computing for Sustainable Global Development (INDIACom)*. IEEE, 2015, pp. 382–387.
- [71] G. Song, L. He, Z. Wang, J. Yang, T. Jin, J. Liu, and G. Li, "Towards the construction of global ipv6 hitlist and efficient probing of ipv6 address space," in *28th IEEE/ACM International Symposium on Quality of Service, IWQoS 2020, Hangzhou, China, June 15-17, 2020*. IEEE, 2020, pp. 1–10. [Online]. Available: <https://doi.org/10.1109/IWQoS49365.2020.9212980>
- [72] J. van der Ham, "Ethics and internet measurements," in *2017 IEEE Security and Privacy Workshops, SP Workshops 2017, San Jose, CA, USA, May 25, 2017*. IEEE Computer Society, 2017, pp. 247–251. [Online]. Available: <https://doi.org/10.1109/SPW.2017.17>
- [73] B. VanderSloot, A. McDonald, W. Scott, J. A. Halderman, and R. Ensafi, "Quack: Scalable remote measurement of application-layer censorship," in *27th USENIX Security Symposium, USENIX Security 2018, Baltimore, MD, USA, August 15-17, 2018*. USENIX Association, 2018, pp. 187–202. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity18/presentation/vandersloot>
- [74] K. Vermeulen, B. Ljuma, V. Addanki, M. Gouel, O. Fourmaux, T. Friedman, and R. Rejaie, "Alias resolution based on ICMP rate limiting," in *Passive and Active Measurement - 21st International Conference, PAM 2020, Eugene, Oregon, USA, March 30-31, 2020, Proceedings*, ser. Lecture Notes in Computer Science, vol. 12048. Springer, 2020, pp. 231–248. [Online]. Available: https://doi.org/10.1007/978-3-030-44081-7_14
- [75] G. Wan, L. Izhikevich, D. Adrian, K. Yoshioka, R. Holz, C. Rossow, and Z. Durumeric, "On the origin of scanning: The impact of location on internet-wide scans," in *IMC '20: ACM Internet Measurement Conference, Virtual Event, USA, October 27-29, 2020*. ACM, 2020, pp. 662–679. [Online]. Available: <https://doi.org/10.1145/3419394.3424214>
- [76] F. Wang and L. Gao, "On inferring and characterizing internet routing policies," in *Proceedings of the 3rd ACM SIGCOMM Internet Measurement Conference, IMC 2003, Miami Beach, FL, USA, October 27-29, 2003*. ACM, 2003, pp. 15–26. [Online]. Available: <https://doi.org/10.1145/948205.948208>
- [77] F. Wang, J. Qiu, L. Gao, and J. Wang, "On understanding transient interdomain routing failures," *IEEE/ACM Transactions on Networking*, vol. 17, no. 3, pp. 740–751, 2009.
- [78] X. Zhang, J. Knockel, and J. R. Crandall, "Original SYN: finding machines hidden behind firewalls," in *2015 IEEE Conference on Computer Communications, INFOCOM 2015, Kowloon, Hong Kong, April 26 - May 1, 2015*. IEEE, 2015, pp. 720–728. [Online]. Available: <https://doi.org/10.1109/INFOCOM.2015.7218441>