Contents

Software Systems for Surveying Spoofing Susceptibility

Responsive to TTA #1, we propose to develop, test, and deploy new tools to measure and report on the deployment of source address validation best practices. Our project includes applied research, software development, new data analytics, systems integration, operations and maintenance, and an interactive analysis and reporting service. First, to allow for testing and monitoring of individual networks, we will develop a polished open-source client-server system for Windows, MacOS, and UNIX-like systems that periodically tests a network's ability to both send and receive packets with forged source IP addresses (spoofed packets). Our system will address substantial deficiencies of the best available data on this global critical infrastructure vulnerability. Second, we will produce reports and visualizations to enable prioritization of source address validation (SAV) compliance attention where it will have the highest benefit, as well as real-time reports that focus operator attention based on observations of inadequate packet filtering as tests complete. Third, we will take advantage of data sources that magnify our view of SAV deployment on many networks without the need for a vantage point in each network, by developing a system to passively detect spoofed packets crossing an Internet exchange point (IXP) peering fabric. Some IXPs have hundreds of participating networks, which makes them an extremely large and entirely unexplored lens from which to measure and support expanded deployment of source address validation. Fourth, to promote testing of access networks, we will create an OpenWrt [2] package of our spoofer client and deploy it on the BISmark measurement platform [23]. Finally, to assist contractors in evaluating SAV compliance, we will prototype a portable touch-screen appliance using the Raspberry Pi platform featuring our client software.

We are uniquely qualified to pursue this work. First, we have extensive experience obtained from developing and operating the MIT spoofer project that informs the thorough approach we propose in this document. Second, we have access to unique sources of data that we will strategically utilize: the UCSD network telescope, which we will use to study the observable effects of SAV policy on spoofed DDoS attacks; and DNS-OARC traffic data for local-node (anycast) root server instances, which we will use to measure the deployment of best practices by ASes peering at public IXPs. Third, we have unparalleled expertise in developing Internet-scale active measurement software and AS topology relationship inferences, placing us in an ideal position to develop open source software for SAV assessment, as well as to develop and report SAV metrics and analysis.

1 Performance Goals

The Regents of the University of California; University of California, San Diego on the behalf of the San Diego Supercomputer Center's Center for Applied Internet Data Analysis (CAIDA) research program, offer this technical proposal which includes the following deliverables: (1) a production-quality client-server source address validation (SAV) testing system that builds on experiences we gained in building and operating the existing system first deployed by Robert Beverly at MIT; (2) a reporting and analysis system that optimizes compliance attention and assesses its impact; (3) a traffic-based SAV-analysis system that gauges SAV deployment using traffic data and peering matrices from Internet exchange points (IXPs) and customer prefix data; (4) a portable touchscreen system that provides a convenient form factor for independent contractors to test SAV compliance; (5) an open-source home-router testing system. The project will leverage the results of existing technologies and infrastructure funded by the Department of Homeland Security and the National Science Foundation.

The proposed work targets objectives outlined in TTA#1: Measurement and Analysis to Promote Best Current Practices. Specifically, we propose to build and operate multiple open-source software tools for anti-spoofing assessment that will allow a site to determine if it has successfully deployed source address validation, and provide on-going monitoring and testing to ensure SAV continues to operate correctly through network upgrades and reconfigurations. Our reporting and analysis system will promote the deployment of SAV by guiding compliance attention where it will have the most benefit, and provide independent measures of the effectiveness of the promoting SAV best-practices. To promote additional testing that will magnify our view of SAV deployment on many networks, we will pursue three additional goals: develop new analytics and software tools that detect spoofed packets crossing Internet exchange points; port our testing tools to the most popular open source home router platform; and prototype a portable appliance that government-approved agents could use in compliance testing. The resulting technologies and data will improve our ability to identify, monitor, and mitigate the infrastructure vulnerability that serves as the primary vector of massive DDoS attacks on the Internet.

2 Detailed Technical Approach

Despite source IP address spoofing being a known vulnerability for at least 25 years [4], and despite many efforts to shed light on the problem (e.g. [6, 8, 9]), spoofing remains a viable attack method for redirection, amplification, and anonymity, as evidenced most recently and publicly in February 2014 during a 400 Gbps DDoS attack against Cloudfare [20]. That particular attack used an amplification vector in some implementations of NTP [20]; a previous attack against Spamhaus [10] in March 2013 achieved 300+ Gbps using an amplification vector in DNS. While some application-layer patches can mitigate these attacks [24], attackers continuously search for new vectors. To defeat DDoS attacks requires operators to ensure their networks filter packets with spoofed source IP addresses [15], a *best current practice (BCP)* known as source address validation (SAV). However, a network that deploys source address validation primarily helps other networks, a classic tragedy of the commons in the Internet.

Testing a network's SAV compliance requires a measurement vantage point inside (or adjacent to) the network, because the origin network of arbitrary spoofed packets cannot be determined [3].

For the past nine years, our approach was to use a software client that volunteers across the Internet could download and run from their networks, testing their own network's ability to send various types of spoofed packets to our server, which collected and aggregated test results. Figure 1 illustrates a simplified view of our current system architecture, which includes: (1) a server instance that coordinates measurements and obtains results, (2) client software for Windows, MacOS, and UNIX-like systems, and (3) a set of distributed Ark nodes that receive spoofed packets and allow us to infer where along a path source address validation (SAV) may be taking place.

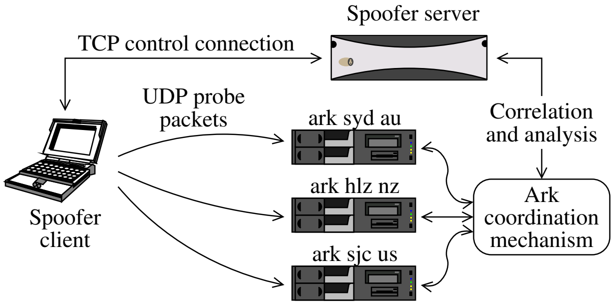

Figure 1: *Architecture of current client-server testing system. Spoofed and unspoofed UDP packets are sent to a distributed set of Ark nodes to test the ability to send spoofed packets, as well as where along a traversed path SAV might be taking place.*

We have used the resulting data to inform (but, due to sampling issues, not resolve) the continuing debate on which networks on the Internet permit spoofed packets to exit their networks, and have allowed network operators to retrieve outcomes of tests conducted from their network. Despite our initial success in obtaining and reporting data on SAV deployment, there are at least nine limitations to our current approach: (1) the software relies on volunteers running it from within the network being tested; (2) the software is only run on demand, and does not provide any continual or longitudinal data; (3) the software has a rudimentary user interface, discouraging some volunteers from using it; (4) the software uses a method of spoofing packets that allows operating system components (e.g., NAT) to rewrite the source address; (5) the software requires root privileges, restricting the class of vantage points we can use; (6) the software relies on operating system binaries whose upgrade path can break the software; (7) all results are sent to our server, and some networks (e.g., government) may be reluctant to involve others in evaluating the security hygiene of their network; (8) volunteers lack motivation, because validating traffic exiting a network primarily benefits others; and (9) reports are tailored toward network operators, and do not suggest where SAV compliance attention would have the most benefit.

We propose to develop, test, and deploy new systems and tools to measure and report on the deployment of source address validation (SAV). We will address all nine of the above limitations, informed by our experiences in operating the existing spoofer system. In addition to completely rewriting our client-server testing system and delivering reports that can be used to assess and promote the deployment of anti-spoofing best practices, we will integrate new software features that generate user incentives for persistent deployment of our measurement tools, and we will evaluate new measurement approaches that have a significantly larger and less biased lens through which to assess SAV deployment and its impact.

Two factors that particularly challenge SAV deployment as well as SAV measurement are *complex traffic engineering requirements* and *incentive misalignment*. We briefly review these two sets of complications. Traffic engineering practices are a complication because they render unusable the most obvious and simplest engineering approach to automate a router's inference of valid source addresses. This simple approach is known as *unicast reverse-path forwarding (uRPF)*, which restricts a router to forwarding only packets that come from the router's best route to the source of the packet [3]. The problem with this simple approach is that routers might inadvertently filter legitimate packets because traffic engineering requirements might prevent announcement of all prefixes to all transit providers, resulting in intentionally asymmetric paths.

However, an under-appreciated consequence of the exhaustion of the IPv4 address space and emergence of IPv6 is that it may minimize the need for automated validation of source addresses, and improve the practicality of static ingress access lists in the future. Using empirical data from Route Views and RIPE's Routing Information Service (RIS), Figure 2 shows a remarkable reduction in the rate at which ASes announce new prefixes: only 6% and 3% of ASes announce reachability to different prefixes month-to-month in IPv4 and IPv6 respectively. In IPv4, this is 1/5th the monthly rate of change it was in 1998. In IPv6, the monthly rate of change has barely changed over the last decade, even though IPv6 deployment has grown exponentially since 2008 [14]. We suspect that today it is considerably more feasible for stub networks to provide their transit providers with a stable set of IP prefixes assigned to devices in their stub network for use in an ingress access list [3]. In Section 2.2 we propose to leverage this insight to automatically produce and report ingress access lists [3] for stub ASes; a transit network operator could validate such an access list with their customer ASes and then deploy it. We focus on stub ASes because existing data [5] indicates that smaller ASes, as measured by the number of connections (degree) they have with other ASes, are disproportionately able to send spoofed packets, consistent with their having fewer resources to deploy, or perhaps understand how to deploy, defenses. Providing enough knowledge to reduce the cost of deploying SAV for these smaller stub ASes allow a substantial increase in the scale of SAV deployment.

Incentive misalignment also inhibits measurement of BCP38 (SAV) deployment. The canonical measurement method is basically crowd-sourcing spoofing attempts from as many distinct networks as possible to infer the global scope of SAV deployment. The fundamental challenge in this case is that some users lack motivation to even measure their own network's SAV deployment, since such deployment primarily benefits others. To improve user incentives to run the measurement, we will augment the tool with capabilities that assess best practices that directly affect them, e.g., measuring whether appropriate *inbound filtering* of potentially malicious (spoofed) traffic exists [3, 21], and reporting whether the network appears on various security reputation-based blacklists of networks observed engaging in malicious activity.

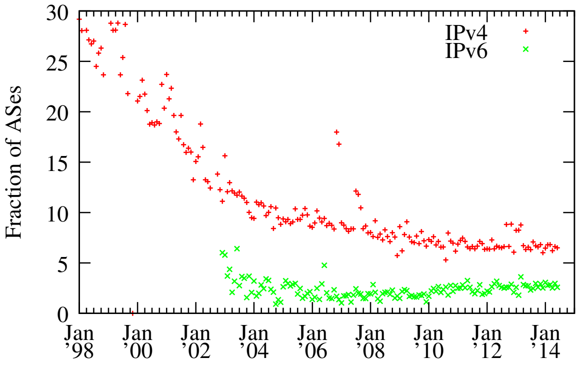

Figure 2: *Fraction of ASes whose address space announcements change month-to-month. During 2014, $\approx$ 6% and $\approx$ 3% of ASes announce different IPv4 and IPv6 addresses month-to-month, respectively. Perhaps a consequence of IPv4 exhaustion, this trend toward stable announcement patterns may diminish the need for automated validation of source addresses, and improve the practicality of static ingress access lists.*

2.1 Production-quality source address validation testing system

We will build a production-quality client-server testing system that will periodically test the ability of a vantage point to send and receive packets with forged source addresses. This system will include client and server software implementations required to accomplish testing; we will build both from scratch to overcome limitations in the current system.

Our client software will contain a GUI for Windows, MacOS, and UNIX-like systems that allows a user to initiate a test and receive feedback on the outcome of the test: i.e. can the local network forward packets with forged source IP addresses, and if so, are the forged addresses limited to local subnet addresses, or a larger prefix? We will give the user an option, enabled by default, to have the software client run spoofing tests periodically in the background, initiating tests on any attached networks at most once per week. To prevent operating system interference when using raw sockets, the client will construct all packets as layer-2 frames regardless of operating system. Our client and server software will support IPv4 and IPv6, allowing the user to check if security policy is being applied consistently for both protocols. Where we find SAV deployment, we will utilize a built-in implementation of *tracefilter* [7] to infer where SAV filters are deployed, by sending spoofed TTL-limited packets and observing where in the path ICMP *time exceeded* messages are no longer generated.

We will write new server software that is easily deployable by others, such as independent government agencies and transit network operators. We will use Transport Layer Security (TLS) to prevent tampering with our spoofing tests. To support user communities, e.g., government networks, who are not comfortable testing to our servers, we will enable the client to configure its own selected server address. To support a more flexible approach to private testing, we will explore a redirection capability whereby other server operators instruct our spoofer server to redirect certain clients (based on the IP address they use to connect to our server instance) to their instance.

To improve user incentives to deploy the tool, we will add support to our client to test if the network has appropriate ingress filtering in place; specifically, we will test if the client is able to receive traffic from our server with spoofed source IP addresses in the same /24 subnet as the client. The tested network should filter such addresses at the edge of their network, per IETF Best Current Practice 84 [3]. This feature will also help operators detect weaknesses in their network that could be exploited by attacks such as triangular spamming [21]. To further incentivize deployment we will provide the user with additional visibility into their network's hygiene, such as the reputation of their network in the security community based on known IP-reputation blacklists.

2.2 Use measurement results to inform compliance efforts

Directly responsive to objective 2 of TTA #1, we propose a further development task that will utilize our expanded view of the spoofing landscape to focus anti-spoofing compliance attention where it will have the highest benefit. We propose to build a new web-based reporting system that will focus efforts of stakeholders and policy makers by correlating our expanded coverage of SAV tests (section 2.1) with characteristics of the tested networks such as their type (e.g., access, transit), country of operation, IP reputation, and their country's transparency of governance. Many small enterprises at the edge will never deploy SAV best practices, but we can help their upstream transit providers to deploy SAV on behalf of these edge networks by generating ingress access lists for ASes that are customers of a given transit provider, who could validate these access lists with these customers and deploy them to discard packets with forged source IP addresses. Using AS rank data [18], we will identify the transit providers in each country whose filtering practices could have the greatest impact in reducing spoofed traffic based on the number and types of customer networks they provide transit for. To motivate transit providers to deploy ingress access lists, we will annotate each transit provider with information on the observed ability of networks beneath them to spoof traffic. The reports will also serve to focus efforts of network operators in deploying mitigation strategies to reduce harms from inadequate filtering deployments. Our reporting and analysis system will have an option to restrict reports to an individual country, to assist countries in planning and implementing their own policies.

In order to provide an independent view of the impact of mitigation strategies against spoofed DDoS attacks, we will leverage the UCSD network telescope [12] to track evidence of DDoS attacks over time. Network telescopes are able to indirectly observe randomly-spoofed denial of service attacks worldwide by capturing a portion of the responses sent back from the victim to the spoofed IP addresses (*backscatter* traffic) [19]. Based on the technique originally presented in [19], we will develop software to automatically detect randomly-spoofed DDoS attacks worldwide, and we will extract data in order to observe trends in the targets (e.g., by country and AS) and magnitude of attacks (duration and volume). CAIDA has collected and stored backscatter traffic since 2004, enabling historic longitudinal view in such trends. While the telescope cannot be used to measure the deployment progress of BCP38, it does allow an independent view of anti-spoofing efforts on the measured impact on spoofed DDoS attacks.

Our reporting system will publicly report summary data for all networks for which it has SAV tests. We will issue a weekly report for operators which summarizes SAV results from tested networks, similar in purpose to the CIDR report [16] which highlights unnecessary deaggregation of IP prefix announcements. We anticipate peer pressure will drive deployment of SAV filters where measurements have revealed the ability of a client in the network to send packets with forged source IP addresses. We will use our SpooferProject twitter account to strategically report test results, similar to how Comcast tweets when DNSSEC validation fails [1].

2.3 Traffic SAV analysis system

Most operators peer their networks at IXPs to exchange traffic between their *customer cones* – i.e. their customers, their customers' customers, and so on, to avoid paying a transit provider to carry traffic between those networks. The essential insight underlying our method is that ASes peering at IXPs (*participating ASes*) should only be transmitting traffic from their customers. If an IXP switch receives packets directly from a participating AS with source addresses outside of that AS's customer cone, these packets are likely spoofed from inside the AS's network. Some IXPs have hundreds of participants [25], so they are an extremely large

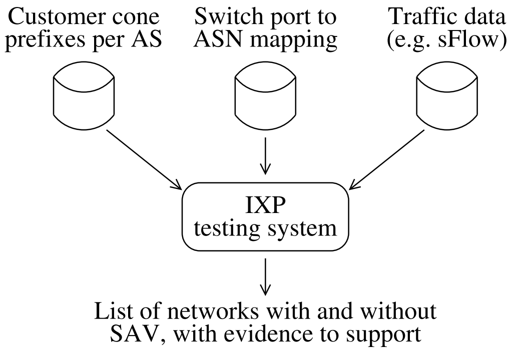
Customer cone prefixes per AS
Switch port to ASN mapping
Traffic data (e.g. sFlow)
IXP testing system
List of networks with and without SAV, with evidence to support

Figure 3: *Architecture of IXP-based testing system. Using CAIDA's customer cone inferences, a mapping of switch ports to ASes, and traffic data, we can infer which networks have deployed SAV.*

and entirely unexplored lens from which to measure and support expanded deployment of SAV.

We propose to build an open-source traffic analysis system to infer evidence that ASes participating at an IXP have themselves not correctly deployed SAV best practices. Figure 3 illustrates the architecture of our IXP-based testing scheme. Our tool will infer an IXP participant AS as likely allowing traffic with a spoofed source address to egress their network into the IXP if we observe an IP packet from a participant's switch port with a source address that falls outside of the participant's inferred customer cone. This inference process requires three sources of data. First, we need a list prefixes found in a each AS's customer cone, which we can readily infer using our existing system to infer routing relationships between ASes and AS customer cones [18, 11]. Second, we need a sample of packets captured from the switching fabric, such as available using the sFlow protocol used by many high-end IXP switches [17]. Third, we need to know which participating AS originated the packet, revealed by either its source Ethernet address, or the port number in the sFlow record.

We will partially validate this approach with known characteristics of spoofed packets [13]. As part of operating the UCSD network telescope, we have developed a set of filters to identify and remove spoofed traffic. For example, non-ICMP packets with a IP Time-to-Live (TTL) $>$ 200 are likely spoofed [22], as well as IP packets with the same source and destination IP addresses, and IP packets possessing a source address where the least significant byte is zero [13].

The probability that an IXP will receive spoofed traffic depends on the prefixes that participating ASes announce at the IXP and the destination addresses in the spoofed packets. We hypothesize that the probability an IXP switch will observe spoofed packets is correlated with the amount of address space advertised by participating ASes at the IXP. NTP and DNS servers [20, 10] are also often deployed at IXPs to maximize reachability, which also makes them easily reachable for use in amplification-based DDoS attacks, providing another opportunity to observe spoofed DDoS attack traffic and correlate it with a participating AS at the IXP.

We will rely on a given IXP operator voluntarily deploying and running our open source software, and notifying participants from whom the IXP received spoofed traffic. We will provide technical assistance to IXPs interested in running the software, and demonstrate its importance and effectiveness by using available DNS root server traffic data collected from anycast root server instances located at IXPs. Such traffic data is available to CAIDA as part of our membership of the DNS Operations Analysis and Research Center (OARC) project. A local-node (anycast) DNS root-server instance should only receive packets with source addresses from customer networks attached to the same IXP as the DNS server. If a DNS root-server instance receives packets from outside these prefixes, we can infer that a participant at the IXP has not deployed SAV best practices, and then encourage the IXP operator to run our software to identify participating ASes who should verify their SAV filtering configurations. The Internet Society has committed to collaborate with CAIDA, leveraging their relationships with IXP operators to facilitate cooperation in support of SAV analysis. DHS and its international partners can also encourage IXP operators to monitor SAV compliance of their participants and use our measurement and analysis results to promote deployment of SAV best practices.

2.4 Enabling SAV testing in home networks

Our client-server testing system (Section 2.1) requires the client to have a globally routable ("public") IP address. However, customers of broadband access networks obtain a single public IP address, and then use Network Address Translation (NAT) technology on the home router so that multiple devices can share a single public address. If a customer installs our client-server system on a device behind the NAT, the NAT software will re-write the source IP address on its way out of the home network, rendering an IPv4 spoofing test useless. (IPv6 tests will still work.) However, some home routers run the OpenWrt [2] open-source operating system, a platform we can modify to integrate, deploy, and evaluate our SAV testing functionality.

We propose to build software specifically to operate on the OpenWrt [2] platform and will ask the OpenWrt team to include a weekly spoofer test in the default OpenWrt build configuration, in order to improve user visibility into SAV best practice deployment on access networks. As a first step, we are cooperating with Georgia Tech's BISmark project [23] to include our software in BISmark home routers. BISmark home routers are based on the OpenWrt platform, and will help to prove the utility of our system to the OpenWrt developers. This task includes implementation of periodic tests with spoofed packets, and updates to the BISmark web interface to show test results to the home user.

2.5 Develop appliance to enable cost-effective SAV testing

We propose to build a Raspberry Pi-based appliance with a 2.8" touchscreen and plastic case that is designed to run a specialized GUI application and provide immediate feedback. Our target user is a contractor tasked with ensuring certain (e.g., government-serving) networks have deployed SAV. Our appliance has four desirable features: (1) the device can be locked down before being issued to a contractor, (2) the device has a wired Ethernet interface, (3) the device has a convenient form factor for travel with personnel, and (4) the device is small and inexpensive enough ($92) to be shipped to different government departments without the expectation it be returned. This system will be easily configured to interact with an instance of our server software, should the government wish to operate its own SAV-compliance project, and could be configured to probe weekly to verify persistence of best practice deployment.

We will use a $39.95 Raspberry Pi Model-B, a $34.95 compact touchscreen designed for the Pi, and a $21.95 case designed to secure the Pi and touchscreen. In total, the hardware for a complete unit currently costs $91.85. Our goal with this unit would be to demonstrate the feasibility of the approach, and we would release detailed documentation on building and operating the touchscreen units so that others can build and operate their own infrastructures.

References

[1] Comcast DNS Twitter Account. `https://twitter.com/ComcastDNS`.

[2] OpenWrt: Wireless freedom. `http://openwrt.org/`.

[3] F. Baker and P. Savola. Ingress filtering for multihomed networks, March 2004. IETF BCP84, RFC 3704.

[4] S.M. Bellovin. Security problems in the TCP/IP protocol suite. *ACM/SIGCOMM Computer Communication Review (CCR)*, 19(2):32–48, April 1989.

[5] Robert Beverly. Spoofer project: State of ip spoofing. `http://spoofer.cmand.org/summary.php`.

[6] Robert Beverly and Steven Bauer. The spoofer project: Inferring the extent of source address filtering on the Internet. In *Proceedings of USENIX SRUTI*, July 2005.

[7] Robert Beverly and Steven Bauer. Tracefilter: A tool for locating network source address validation filters. In *USENIX Security Poster*, August 2007.

[8] Robert Beverly, Arthur Berger, Young Hyun, and k claffy. Understanding the efficacy of deployed Internet source address validation filtering. In *Proceedings of the 9th ACM SIGCOMM Internet Measurement Conference*, November 2009.

[9] Robert Beverly, Ryan Koga, and kc claffy. Initial longitudinal analysis of IP source spoofing capability on the Internet, July 2013. `http://www.internetsociety.org/`.

[10] Peter Bright. Spamhaus DDoS grows to Internet-threatening size, March 2013.

[11] CAIDA. Cartographic capabilities for critical cyberinfrastructure. DHS S&T contract N66001-12-C-0130.

[12] Center for Applied Internet Data Analysis. UCSD Network Telescope, 2010. `http://www.caida.org/data/passive/network_telescope.xml`.

[13] Alberto Dainotti, Karyn Benson, Alistair King, kc claffy, Michalis Kallitsis, Eduard Glatz, and Xenofontas Dimitropoulos. Estimating Internet address space usage through passive measurements. *ACM SIGCOMM Computer Communications Review*, 44, Jan. 2014.

[14] Amogh Dhamdhere, Matthew Luckie, Bradley Huffaker, k claffy, Ahmed Elmokashfi, and Emile Aben. Measuring the deployment of IPv6: Topology, routing and performance. In *ACM SIGCOMM Internet measurement Conference*, pages 537–559, November 2012.

[15] P. Ferguson and D. Senie. Network ingress filtering: Defeating denial of service attacks which employ IP source address spoofing, May 2000. IETF BCP38, RFC 2827.

[16] Geoff Huston. Cidr report, November 2014. `http://www.cidr-report.org/`.

[17] Elisa Jasinska. sFlow - I can feel your traffic, December 2006. 23rd Chaos Communication Congress.

[18] Matthew Luckie, Bradley Huffaker, Amogh Dhamdhere, Vasileios Giotsas, and k claffy. AS relationships, customer cones, and validation. In *ACM SIGCOMM Internet measurement conference*, pages 243–256, October 2013.

[19] David Moore, Colleen Shannon, Douglas J. Brown, Geoffrey M. Voelker, and Stefan Savage. Inferring Internet denial-of-service activity. *ACM Trans. Comput. Syst.*, 24, May.

[20] Matthew Prince. Technical details behind a 400Gbps NTP amplification DDoS attack, February 2014. `http://blog.cloudflare.com/technical-details-behind-a-400gbps-ntp-amplification-ddos-attack`.

[21] Zhiyun Qian, Z. Morley Mao, Yinglian Xie, and Fang Yu. Investigation of triangular spamming: A stealthy and efficient spamming technique. In *IEEE Security and Privacy*, May 2010.

[22] Albin Sebastian. Default time to live (TTL) values, December 2009. `http://www.binbert.com/blog/2009/12/default-time-to-live-ttl-values/`.

[23] Srikanth Sundaresan, Sam Burnett, Nick Feamster, and Walter de Donato. Bismark: A testbed for deploying measurements and applications in broadband access networks. In *Proceedings of the USENIX Annual Technical Conference (ATC)*, June 2014.

[24] Paul Vixie. Rate-limiting state: The edge of the Internet is an unruly place. *ACM Queue*, 12(2):1–5, February 2014.

[25] Wikipedia. List of Internet exchange points by size.