Challenges in Inferring Spoofed Traffic at IXPs

Lucas Müller
UFRGS / CAIDA
lfmuller@inf.ufrgs.br

Matthew Luckie
University of Waikato
mjl@wand.net.nz

Bradley Huffaker
CAIDA / UC San Diego
bradley@caida.org

kc claffy
CAIDA / UC San Diego
kc@caida.org

Marinho Barcellos
UFRGS and University of Waikato
marinho@inf.ufrgs.br

ABSTRACT

Ascertaining that a network will forward spoofed traffic usually requires an active probing vantage point in that network, effectively preventing a comprehensive view of this global Internet vulnerability. Recently, researchers have proposed using Internet Exchange Points (IXPs) as observatories to detect spoofed packets, by leveraging Autonomous System (AS) topology knowledge extracted from Border Gateway Protocol (BGP) data to infer which source addresses should legitimately appear across parts of the IXP switch fabric. We demonstrate that the existing literature does not capture several fundamental challenges to this approach, including noise in BGP data sources, heuristic AS relationship inference, and idiosyncrasies in IXP interconnectivity fabrics. We propose a novel method to navigate these challenges, leveraging *customer cone* semantics of AS relationships to guide precise classification of inter-domain traffic as in-cone, out-of-cone (*spoofed*), unverifiable, bogon, and unassigned. We apply our method to a mid-size IXP with approximately 200 members, and find an upper bound volume of out-of-cone traffic to be more than an order of magnitude less than the previous method inferred on the same data. Our work illustrates the subtleties of scientific assessments of operational Internet infrastructure, and the need for a community focus on reproducing and repeating previous methods.

CCS CONCEPTS

• **Networks** → **Network measurement**; **Network security**.

KEYWORDS

IP spoofing, Internet eXchange Point, Denial-of-service, Network filtering

ACM Reference Format:

Lucas Müller, Matthew Luckie, Bradley Huffaker, kc claffy, and Marinho Barcellos. 2019. Challenges in Inferring Spoofed Traffic at IXPs. In *The 15th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '19)*, December 9–12, 2019, Orlando, FL, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3359989.3365422>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CoNEXT '19, December 9–12, 2019, Orlando, FL, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6998-5/19/12...\$15.00

<https://doi.org/10.1145/3359989.3365422>

1 INTRODUCTION

Networks that allow spoofed source Internet Protocol (IP) addresses in packets are a cybersecurity risk on the global Internet, because they enable attacks such as spoofed denial-of-service (DoS) attacks that are operationally infeasible to trace back to the actual source. Recognizing that lack of *source address validation* (SAV) is fundamentally an architectural limitation [10, 60], the Internet Engineering Task Force (IETF) introduced best current practices recommending that networks block packets with spoofed source addresses [9, 29]. Compliance with these filtering practices has misaligned incentives i.e., it protects the *rest* of the Internet from attacks being sourced from the network that must pay a non-trivial cost for deploying and accurately maintaining the filters. Thus, despite many attempts to improve SAV deployment and mitigate the impact of DoS attacks, some of the most damaging DoS attacks in the Internet still leverage IP spoofing as a vector, setting new records each year for the volume of traffic launched at even highly provisioned networks, disrupting access to those networks [43, 44, 59, 71].

Identifying networks that do not filter spoofed packets is critical to global network infrastructure protection, because it provides a focus for remediation and policy interventions [53]. However, identification of these networks is challenging at Internet scale. The definitive method requires an active probing vantage point in each network being tested, to see if a spoofed packet successfully traverses the network [13, 15]. Since there are approximately 65K independently routed networks on the Internet in 2019 [6, 75], this method has limited feasibility for a comprehensive assessment of Internet spoofing.

Broader visibility into the spoofing problem may lie in the capability to infer lack of SAV compliance from large, heavily aggregated Internet traffic data, such as traffic observable at Internet Exchange Points (IXPs). Most Autonomous Systems (ASes) connect to an IXP to exchange traffic between their customers, i.e., via peering relationships where neither AS pays the other for transit. For these ASes, legitimate source addresses in packets will belong to direct or indirect customers of the AS sending the packets across the IXP fabric to their peers.

However, inferring SAV deployment at an IXP is remarkably challenging, more so than has been captured in the literature, due to a combination of operational complexities that characterize today's interconnection ecosystem. First, determining which source addresses are valid in packets arriving at a given port of an IXP switch fabric is challenging, because there is no registry of which addresses networks should forward; in practice, we must infer valid

source addresses. Second, while the original role of IXPs was to promote peering between ASes, networks now also use IXPs to obtain IP transit services from a provider [1], and we have found evidence of organizations joining their sibling network ASes across an IXP. For ASes offering transit across the IXP, and for sibling networks, it is infeasible to infer invalid source addresses from IXP traffic data – the set of valid addresses is potentially the entire address space. Third, while IXPs may be thought of as a single switching fabric, in practice IXPs and resellers offer complex services, including remote peering, layer-2 transport, and virtualized segmenting of traffic into multiple Virtual Local Area Networks (VLANs). These interconnection practices occur below and are thus not visible to the IP layer or in the Border Gateway Protocol (BGP).

Accurately inferring SAV deployment at an IXP requires navigating all of these aspects. In this paper, we describe a methodology that does so. One of our discoveries does not bode well for the ability to automate this method: identifying the myriad cases that explain patterns in traffic at a given IXP is largely manual in nature, and must be repeated at each IXP to accommodate IXP-specific architectural engineering and business decisions. However, we imagine its utility as part of an expert system suite of cybersecurity services or compliance practices of modern IXPs.

This paper makes the following contributions:

(1) We provide a detailed analysis of methodological challenges for inferring spoofed packets at IXPs. Based on IP routing, addressing, and IXP concepts, we analyze methodological challenges and their implications for building IP spoofing detection capabilities at IXPs (§2). We include a comprehensive analysis of previous work which also inferred spoofing at IXPs. We also analyze challenges specific to applying BGP-based SAV inference methods to modern IXP connectivity fabrics (§3).

(2) We develop a methodology to classify traffic flows for the purposes of accurately inferring spoofed traffic. We design and implement Spoofer-IX, a novel methodology to detect the transmission of spoofed traffic (which implies lack of source address validation) by AS members of IXPs (§4). Spoofer-IX addresses two fundamental issues not addressed in the existing literature [45]. First, Spoofer-IX considers the type of relationship between neighbors at an IXP when determining which source addresses are valid in IP packets crossing the IXP. Second, Spoofer-IX considers asymmetric routing and traffic engineering, by designing a prefix-level customer cone that includes addresses that may be valid source addresses for an AS to transit. The accuracy of this method depends on the quality of BGP data and AS relationship inferences, which we know to be imperfect [54]. However, our method is congruent with what network operators do when configuring static access control lists to deploy SAV [30, 37, 42].

(3) We use our methodology to classify packets at a IXP in Brazil with approximately 200 members. We apply our method to traffic and topology data (described in §5) from one of the largest IXPs in Brazil, with more than 200 member ASes using the IXP switching fabric. We report our analysis findings, and results of our interactions with IXP and network operators to validate the findings (§6). We investigate the impact of different filtering choices on inferred valid address space, and the likelihood of false negatives when classifying traffic according to different filtering choices. We also compare our method with a recently proposed method [45]

that did not consider AS relationships in its inference of spoofed traffic, reporting that the majority of members at the IXP sent spoofed packets, and demonstrate pitfalls of this approach. Indeed, at the medium-sized IXP we studied, with approximately 200 members, this previous method inferred spoofed traffic coming from 62.3% of addresses over a one-week period in May 2019, but our AS-relationship-aware method inferred spoofed traffic coming from less than 1 in 5 (18.7%) member ASes during our five-week observation period in May 2019.

(4) We find evidence that epistemological and cross-validation challenges remain, and we publish our code to promote further work. When we compared our results with CAIDA’s crowdsourced measurements, we found that CAIDA received positive spoofing tests (lack of SAV) in 54% of the member ASes at this IXP. This is not necessarily inconsistent, since even at a heavily aggregating exchange point, one cannot detect lack of SAV without actually observing spoofed packets, which CAIDA’s crowdsourced approach explicitly injects. We conclude our paper with a discussion of lessons learned (§8), including that we believe further work is required to understand the degree to which IXPs can be used as a lens into SAV deployment, and why we think such work is important to future cybersecurity efforts. Our conclusions highlight the persistent tension between the need for reproducibility of methods and results [7, 8], and the opacity characteristic of commercial infrastructure. We publicly release our code [62] in hopes that other researchers and IXPs will use it to further improve our collective ability to measure and expand deployment of SAV filtering.

2 BACKGROUND AND RELATED WORK

2.1 Source Address Validation

The Internet architecture provides no explicit mechanism to prevent packets with forged headers from traversing the network. This vulnerability allows IP spoofing attacks, i.e., when hosts send IP packets using fake source addresses that cannot feasibly be traced back. To reduce the incidence of this type of attack, network operators can configure their routers to identify and block spoofed packets before these packets leave their networks. Such filtering is well-specified and a standardized IETF best current practice [29], frequently referred to as Source Address Validation (SAV) [38]. Network operators often implement SAV by using *ingress filters* in routers, which drop packets with source addresses outside the locally valid address space before they enter the global Internet.

2.2 Address Space Fundamentals

For the purposes of this study, we distinguish three main categories of IP address space: Bogon, Unassigned, and Routed. *Bogon* addresses are reserved by the IETF [22, 61] for specific uses such as private networks and loopback interfaces; they do not uniquely identify any host, and should not be routed on the Internet. *Unassigned* addresses [34, 35] have not been assigned by an Internet registry to an AS and should not be used or routed by anyone. *Routed* addresses have been assigned to some AS, and are thus potentially valid source addresses in inter-domain traffic.

Figure 1: Illustration of the architecture of modern IXPs. Modern IXPs typically construct a switching fabric using a core switch that interconnects other switches located in remote colocation facilities. ASes typically connect to a switch located in a colocation facility, and can form bilateral peering relationships with neighbors. These ASes may request a VLAN to isolate their traffic from other members at the IXP. Resellers can provide services such as remote peering and layer-2 transport.

2.3 IXPs as Observatories

IXPs are attractive vantage points to observe signals of SAV deployment, as hundreds of ASes may be present at a single logical location. The IXP operator assigns each member a unique IP address from a prefix controlled by the operator, which the member assigns to their router interface connected to the IXP, and uses to establish BGP routing with other members. When a member AS's router transmits a packet across the Ethernet switching fabric, the source and destination media access control (MAC) addresses in the Ethernet frame uniquely identify the AS pair exchanging the packet, and its direction.

Figure 1 illustrates the architecture of many modern IXPs [4, 23, 28, 30, 40, 41, 48]. The figure contains two separate IXPs and their switching fabrics #X and #Y, with a core switch for each IXP. While some IXPs may consist of a single core switch where participants interconnect, operators achieve the scale of modern large IXPs by placing switches at distinct physical colocation facilities, any of which can serve as an IXP attachment point. The figure shows that the switches are adjacent, but in practice colocation facilities are usually in different buildings. IXP operators often use sFlow [66] or NetFlow [19] to collect traffic flow statistics. A comprehensive view of all traffic from all services at the IXP would require flow data captured from all switches in the switching fabric, as traffic between participants at a single colocation facility will not travel to the core switch.

Participants can exchange traffic directly across the switching fabric in a bilateral session. In figure 1, ASes A and B exchange traffic directly. However, modern IXPs often use VLANs to provide logical isolation between different types of interconnection [18, 27]. For example, an IXP may provide a route server, but only offer

that route server on a specific VLAN. Similarly, traffic between two participants may be sufficiently sensitive or high volume that members request a VLAN from the IXP to isolate their communications [3, 24, 47]. In figure 1, ASes C and D exchange traffic in their own isolated VLAN.

To foster IXP growth and enable more networks to interconnect, IXPs have supported resellers, which provide value-added services at an IXP, such as remote peering and layer-2 transport [17, 39, 58, 64]. A reseller provides remote peering services so that an AS that is not physically present at a colocation facility can still reach other members at the IXP, without the AS incurring colocation facility fees or port charges from the IXP operator. These resellers require some cooperation with the IXP, e.g., [2, 46]. The IXP assigns the remote peers any VLAN tags they require to participate at the exchange as local members do.

An IXP may use different technical approaches to support remote peering providers [17, 41, 64]. A reseller can bridge Ethernet networks so that the MAC address of the customer router's interface will uniquely identify the origin of traffic in the peering fabric. A second approach is for a reseller to push a tag (reseller-tag) to uniquely identify their specific customer AS to the IXP, so that the MAC address of the Ethernet frame corresponds to the reseller's router. Figure 1 illustrates this second approach, where reseller J allows customer ASes F and G to reach other members. When the reseller transmits these packets into the IXP, the reseller also pushes a tag (reseller-tag) to uniquely identify their specific customer AS to the IXP. The IXP bridges traffic into the IXP switching fabric by removing the outer-most reseller-tag while keeping the IXP-tag. In figure 1, the sFlow tap sees the IXP-tag and the MAC address of the reseller, which uniquely identifies the AS that sent the packet.

Figure 2: The customer cone constrains the set of source addresses expected in valid inter-domain traffic transiting an AS behaving rationally in a c2p or p2p relationship. In the c2p relationship shown in (a), B transits traffic from its customers to A, but not its peers and providers. Similarly, in the p2p relationship shown in (b), C only transits traffic from its customers to D (likewise, from D to C). However, as shown in (c), the p2c relationship does not constrain the source addresses transited by E to F, and neither does the s2s relationship between G and H in (d).

A reseller can also provide remote peering to members colocated at one IXP that want to reach members in a different IXP. Figure 1 shows a more complicated example, where AS E bridges their network between metropolitan regions using the services of a reseller (K) present at both IXPs.

2.4 AS Relationships and Customer Cones

The three primary classes of AS relationships are customer-provider (c2p, p2c), peering (p2p) and sibling (s2s). In a c2p relationship (also known as transit), a customer buys access to achieve global reachability to all routed Internet address space. In a p2p relationship, two ASes agree to exchange traffic destined to prefixes they or their customers own, typically without either AS paying the other [31]. In a s2s relationship, a single organization operates both ASes, and may transit packets received from any source.

An AS's *customer cone* includes all ASes reachable through its customer ASes, i.e., direct and indirect customer ASes (in other words, ASes reachable only through p2c links) [54]. The customer cone constrains which source IP addresses one should see in valid inter-domain traffic transiting from a customer to its provider, or between peers. Figure 2 illustrates the subtleties: an AS in a c2p or p2p relationship with another AS should only send packets with a source address from within its customer cone – respectively, (a) and (b) in figure 2. In contrast, a link between a provider to its customer or between two siblings may forward packets with *any* routed source address – (c) and (d) in figure 2.

2.5 Measuring Deployment of SAV

Many academic research efforts have described techniques to promote deployment of SAV [25, 49, 50, 77]. Fewer efforts have tried to empirically measure SAV compliance for networks attached to the global Internet. In 2005, Beverly, *et al.* developed a client-server technique to allow users to test networks to which they are currently attached [12], and operationalized a platform to track trends over time [13, 15]. The platform allows for inference of deployed SAV policy, but has limited coverage, because it relies on users downloading and running measurement software. To overcome this limitation, researchers have recently investigated techniques to infer lack of SAV using macroscopic Internet data sets. In 2017, Lone *et al.* reported a technique to infer spoofed traffic in massive traceroute archives, based on the assumption that an edge network should never appear to be providing transit in a traceroute path [51]. This method is limited by whatever appears in the traceroute archives, and can be hampered by traceroute artifacts caused by inconsistent Internet Control Message Protocol (ICMP) implementations in routers [57].

Most closely related to our study, in 2017 Lichtblau *et al.* used a large IXP as a vantage point for inferring which networks at the IXP had not deployed SAV [45]. For each member at the IXP, their method infers a set of IP prefixes containing addresses that may legitimately appear in the source field of IP packets crossing an IXP. They infer that a member AS that sends a packet into the IXP switching fabric with a source address outside of those prefixes has not deployed SAV. They argued against using AS relationships and

Figure 3: Example full cones (§3.1.1) for six ASes given these BGP paths. The full cone for an AS includes every prefix that contains that AS in the path for all routes observed by public route collectors, regardless of the underlying relationships.

AS customer cones which they claimed did not address asymmetric routing. However, their method did not consider ASes forming customer-provider or sibling relationships at the IXP, where all routed addresses may be legitimate source addresses in IP packets crossing an IXP – (c) and (d) in figure 2. In these cases, there is no way to infer SAV deployment across these links at the IXP.

3 TACKLING METHODOLOGICAL CHALLENGES

We describe the core of our methodology in the context of two complex challenges to inferring spoofed traffic in IXP traffic data. The first challenge (§3.1) is determining which addresses are valid source addresses in traffic transiting a given neighbor AS, i.e., packets with a source address that is *in-cone* for that AS. An incomplete set of valid addresses could yield false inferences of failure to deploy SAV, should a valid address appear in the observed packets but not be in the in-cone set, i.e., be *out-of-cone* for that AS. The second challenge (§3.2) is navigating the analytical implications of modern IXP interconnection practices that can impede the visibility of both topology and traffic. These practices complicate the analysis of which ASes exchanged traffic and their routing relationship. Once we address these challenges, the remainder of our method is IXP-specific but straightforward, and we describe it in §4.

3.1 Subtleties in Cone Construction

Inferring the set of valid source addresses for packets traveling from a specific AS to a specific adjacent AS at an IXP requires navigating a multidimensional parameter space. Precision in this process is crucial. Mistakenly excluding valid addresses could result in a misclassification of an AS as not performing source address validation (false positive). Similarly, including invalid source addresses could result in spoofed packets going undetected (false negatives). As mentioned in §1, there is no global registry that contains ground truth on which addresses are valid source addresses for packets transited by an AS; instead, we must infer them from BGP routing data sources [65, 68, 70], even though these sources may contain spurious announcements [52].

Figure 4: Example customer cones (§3.1.2) for six ASes using the same BGP paths from figure 3. In customer cone construction, we annotate each AS link with a c2p, p2c, or p2p relationship before inferring the prefix-level customer cone.

3.1.1 Full Cone. The full cone (used in [45]) is the more permissive of the two construction methods. Aiming to minimize false positives, Lichtblau *et al.* chose to “not distinguish between *peer-ing/sibling*, *customer-provider* and *provider-customer* links. Rather, whenever [the algorithm sees] two neighboring ASes on an AS path, [the algorithm] presumes a directed link between the two, where the left AS is considered upstream of the right AS.” The resulting cone for an AS, which they call its *full cone* (FC), includes every prefix that contains that AS in the BGP route’s AS path [45], for all routes observed by public route collectors in Routing Information Base (RIB) snapshots and updates during the measurement period.

They acknowledged that this method intentionally sacrifices specificity, i.e., inflating the address space considered legitimate for each AS pair, in the interest of avoiding false positives, i.e., avoiding mistakenly attributing a failure to deploy SAV. Using this method, a stub AS that provides a public BGP view containing all prefixes it received from its peers and providers will have *all* of these prefixes included in its full cone, i.e., the entire routed address space will be deemed valid. Figure 3 illustrates the full cones for six ASes; if A were a stub AS and a customer of B, all three prefixes would be included in A’s full cone even though no system in A should originate packets with those source addresses.

3.1.2 Customer Cone. The customer cone is the more restrictive of the two construction methods; it takes into account the semantics of AS relationships. As described in §2, the AS-level customer cone defines the set of ASes reachable using customer links from the AS, including the AS itself [54]. We use the *provider/peer-observed customer cone* (PPCC) algorithm defined in [54] to build an AS-level customer cone. Using the paths in figure 4, the PPCC method constructs the cone of AS C using routes observed from its providers and peers. The PPCC method accommodates hybrid relationships,

Figure 5: The cone construction approach significantly impacts the source addresses each method will consider valid. In (a) we show that 5.5% of all ASes had the equivalent of all routed addresses (175 /8 equivalents) in their full cone in April 2017. In (b) we show that while 90.5% of ASes had (full and customer) cones covering the same set of addresses, 58% of the IXP-EU members would have covered more addresses, with 42% of ASes having a full cone 100 times larger than their customer cone. Note, per discussion in §3.2, an AS announcing 0.01% /8 equivalents is announcing less than 0.006% of the routed address space.

where an AS may not propagate all of its customer routes to all of its peers and providers. Customer cone inference critically relies on accurate routing relationship inferences; a customer link incorrectly inferred to be a peer link will result in address space that the provider AS transits being incorrectly excluded from its customer cone. Figure 4 illustrates the AS-level customer cones for the same ASes and paths as figure 3, with link annotations to identify the inferred routing relationships between ASes. However, an AS-level customer cone does not define the set of valid source addresses in traffic transiting a given neighbor AS.

Once we have the AS-level customer cone for C, we transform it into its corresponding prefix-level cone by including all prefixes originated by ASes in the AS-level customer cone for C during the same observation window. This novel prefix-level cone construction accommodates traffic engineering practices, where an AS may announce different prefixes through different providers, but forward traffic from within these prefixes according to the best route to the destination. To illustrate, in figure 4, we include 203.0.113.0/24 in C's prefix-level customer cone, even though that prefix is not observed in any BGP paths involving C, because F is in C's customer cone. Importantly, we do not include these three prefixes in A's customer cone, because A has no customers. We also combine the prefix-level customer cones of siblings, because a sibling C may transit packets from the customer cone of any of C's siblings to C's peers or providers.

3.1.3 Impact of the Cone Construction Method. Figure 5 shows how the choice of cone construction method impacts inference of valid address space for all ASes (figure 5a) and for the ASes at the IXP-EU used in [45] and the IXP-BR in our study (figure 5b), in both cases using traffic and BGP data from April 2017 (see §5 for further detail on the datasets we used). In particular, 5.5% of all ASes in the Internet had a full cone that contained all routed address space. For 90.5% of ASes, the full cone and customer cone were congruent (included the same addresses), but 58% of IXP-EU member ASes had full cones covering more addresses than the customer cone,

Figure 6: The inferred out-of-cone traffic volume for the full cone is sensitive to changing BGP observation window sizes in the construction of the cone. While the 7 and 9 day lines are almost identical, the 5-day line contains an order of magnitude more traffic because the set of valid addresses for each AS is smaller.

and 42% of ASes had an FC 100 times larger than their CC. This disparity of cone sizes for all ASes compared to those at the IXP is because while over 80% of the Internet's ASes are stubs, i.e., do not provide transit, these are less likely to peer at an IXP. Further, IXPs are popular places to operate public route collectors because the collector can obtain BGP routing views from multiple ASes at a single place. Therefore, those ASes at an IXP that provide a routing view will have all of the prefixes they announce in routes to the collector, including those from their peers and providers, in their full cone. Figure 6 shows how the choice of BGP observation window impacts [20] the inference of out-of-cone traffic at our IXP in Brazil in April 2017 using the full cone. This effect is because of the FC's permissive nature, which exposes the cone inference to announcements across the whole Internet.

Neither the full cone nor the customer cone handle the complexities that sibling ASes (ASes under the same ownership) bring. Because siblings may provide mutual transit to each other, the set of valid addresses that can transit between each AS is the entire

routed address space. To observe this behavior in public BGP data, which both the FC and CC use, would require a view from each sibling AS. Current sibling relationship inference methods [14, 32] use WHOIS data, which is not only inconsistently formatted across regions, but also becomes stale if not updated as mergers occur, leading to false and missing inferences [32].

3.2 Topology and Traffic Visibility

While the original role of IXPs was to promote peering between ASes physically present and connected to a switching fabric, in practice IXP services have become more complicated. For example, many networks now obtain transit services from a provider at the IXP [1]. Or, an organization can connect its sibling networks using the IXP switching fabric. IXPs may also offer services such as remote peering and layer-2 transport, as well as virtualized segmenting of traffic into multiple VLANs. These services present three challenges to accurate inference of SAV deployment.

First, the BGP routing relationship between two IXP members impacts whether the customer cone can constrain inference of valid source address space. As discussed in §2.4, a provider AS may forward packets with a source address from any routed prefix in the Internet to their customer, and a sibling may forward packets from the provider of one sibling to the customer of another sibling. In these cases, we cannot apply a cone of valid addresses to infer the SAV policy of the transmitting member. We can only make this inference when that member has a peering or transit relationship with another member. In contrast to prior work [45], we consider the routing relationship between the two IXP member ASes exchanging traffic when evaluating the source address of a packet crossing the IXP.

Second, there are traffic visibility impediments. As discussed in §2.3, traffic between members connected to the same switch will stay within the switch. In a distributed switching fabric, observing all member traffic requires traffic capture from all switches. Similarly, ASes may establish private interconnections with other ASes at the same colocation facility; their traffic exchange does not use the core IXP switching fabric. Further, to infer SAV policy of an IXP member, we require hosts in the cone of the IXP member to attempt to send spoofed packets to hosts they would reach across the IXP. Because most ASes peer at an IXP, only destinations in the customer cone of the receiving AS would receive that packet, i.e., the victim or the amplifier must be reached via the IXP. Because most customer cones are small (figure 5a, where only 5% of ASes have more than 0.006% of the routed address space in their customer cone) the chance of a victim or amplifier also being reached via a peering relationship at the IXP is small; a victim or amplifier is more likely to be reached via a transit relationship at the IXP.

Third, shared use of IXP ports creates attribution challenges. While the IXP can supply the AS number of record for a given port, with the associated Ethernet MAC address, that port does not necessarily uniquely identify the sending AS when a reseller uses the port to provide layer-2 transport, in cases of remote peering and port resale (§2.3), or when the port connects to another exchange. Prior work has illustrated measurement challenges of inferring remote peering [17, 64]. In this work, the IXP provided us the

Figure 7: Spoofed-IX Inference Method Overview.

reseller and IXP tags they used to bridge remote peers. This IXP-specific knowledge exemplifies why we believe a customer-cone-based approach to SAV inference will ultimately be integrated into expert system capabilities rather than be amenable to complete layer-3 automation.

4 IMPLEMENTING CLASSIFICATION PIPELINE

The customer cone construction method described in §3 underpins our traffic classification method - how we infer invalid source addresses (presumably spoofed) in packets crossing an IXP, and the ASes responsible for transmitting them. We describe how these pieces fit together in our system implementation, which relies on IXP traffic measurements and topological information, i.e., BGP data and IXP switching fabric forwarding databases. The implementation, illustrated in figure 7, has two stages: (1) build an accurate *prefix-level customer cone* from BGP data, and (2) verify that the customer cone can serve to constrain our inference, and if so classify traffic as *in* or *out* of the transmitting AS's customer cone.

4.1 Stage 1: Build the Customer Cone

The first stage has three phases, as follows.

Phase 1: Filter and Sanitize AS Paths. To avoid incorrectly identifying non-existent links between ASes, we use the method from [54] to discard paths with artifacts, such as loops, non-adjacent Tier-1 ASes, and reserved/unassigned ASes [33]. We also discard paths to prefixes longer than /24 or shorter than /8.

Phase 2: Infer AS Relationships. We use the sanitized AS Paths from phase 1 to derive AS relationships on a weekly basis, also according to the algorithm in [54]. This algorithm applies heuristics

Figure 8: Flowchart showing our traffic classification pipeline (stage 2 of methodology, described in §4.2).

to annotate each AS link with either a transit (*C2P*, *P2C*) or peering (*P2P*) relationship.

Phase 3: Construct the Prefix-Level Customer Cone. An AS's *prefix-level customer cone* is the set of prefixes covering source addresses from the AS and its customers, for which the AS will transit traffic. Conceptually, constructing this cone is the most complicated part of our method, and where mistakes can impact its accuracy. We construct a prefix-level customer cone using the method we described in §3.1.2.

4.2 Stage 2: Classify IXP Traffic

The second stage has three phases, illustrated in figure 8.

Phase 1: Filter Bogon and Unassigned Addresses. We first classify traffic with *bogon* and *unassigned* source IP addresses, according to Team Cymru [73], as described in §5. Networks sending packets with unassigned source IP addresses are unlikely to have implemented SAV correctly, since the most obvious implementation blocks traffic from such addresses because they are not routed, therefore have no feasible return path. This phase is independent of any routing semantics, unlike the subsequent two phases, which consider the sending and receiving ASes for the monitored link, the routing relationship between them, and the prefix-level customer cone of the sending AS.

Phase 2: Filter Unverifiable Packets. This phase classifies traffic flows as suitable to inference of spoofing using the customer cone, marking unsuitable traffic as *Unverifiable*. Verifiable traffic must satisfy all of the following:

- (1) It must have a valid MAC-to-AS mapping for both the sending and receiving MAC addresses.
- (2) It must not have a known router IP address in the source IP address of the packet. Such a source IP address could be from any interface on the router, which might be assigned by an AS whose address space is not in the customer cone of the router's owner.
- (3) It must not have a known IP address of the IXP LAN prefix. These prefixes are assigned to the IXP operator and should not be publicly announced, but sometimes member ASes mistakenly announce them.

- (4) It must not have a source MAC address from a remote peer or layer-2 transport provider.
- (5) It must not have a source MAC address from a known provider or sibling of the receiving AS.

Phase 3: Classify Packets with Customer Cone. The remaining traffic has a MAC-to-AS mapping, and is either transmitted by a customer of a transit provider at the IXP, or by a peer of another AS at the IXP. If a relationship was not visible in BGP, then we assume the traffic between these members was p2p and use the cones to classify the traffic exchanged. For these transmitting ASes, we classify traffic as *in-cone* or *out-of-cone* using the prefix-level customer cone (henceforth *customer cone* or *CC*) created in the previous stage. We classify a packet whose source IP belongs to the sending AS's customer cone address space as *in-cone*. Otherwise, we classify the packet as *out-of-cone*.

5 DATASETS

IXP-BR: traffic and routing data. We used sFlow [66] traffic data from a Brazilian IXP [40]. This IXP transports up to 200 Gbps of traffic among 200+ members. The IXP operators configured a sample rate of 1:4096 packets, and we used two datasets from 1 April to 6 May 2017, and 1 May to 5 June 2019, to evaluate our method.

Topology data over connectivity fabric. To identify the pair of adjacent ASes sending and receiving each flow across the IXP fabric, we used layer-2 information (i.e., MAC addresses) since the source and destination IP addresses in the IP headers of the observed packets contain the communication endpoints. To map MAC addresses to sending and receiving ASes of each flow (the MAC-to-AS mapping), we relied on information from the forwarding database of each switch that is part of the IXP switching fabric.

Router IP addresses. For comparability with previous work [45], we used CAIDA's Internet Topology Data Kit (ITDK) [16] to identify router interface IP addresses. We used the ITDK snapshot closest in time to the IXP traffic capture window. We consider traffic from ITDK-inferred router interfaces to be *unverifiable* (§4.2) because the source IP address could be from any of the interfaces of the router, which might be assigned by an AS whose address space is not in the Customer Cone of the router's owner (§4.2).

Figure 9: Five weeks of traffic for 2017 and 2019 classified with our method. We omit the unassigned class, which is negligible. For all ten weeks, we inferred almost no out-of-cone traffic – a maximum of 40 Mbps for an IXP with a peak of 200 Gbps.

Bogons and unassigned addresses. We used Team Cymru’s Full-bogons feed [72, 73] to filter out traffic with source IP addresses that are bogons (e.g., private, special use, reserved) [22, 61, 76] or unassigned. Unassigned prefixes are allocated by IANA to an RIR [34, 35], but not subsequently assigned by the RIR to an end-user (e.g., an ISP) [75]. We used the lists compiled by Team Cymru in each 4h interval per day for the same time windows as our IXP traffic data collection.

Public BGP Data. Our traffic filters rely on Customer Cones inferred from public BGP routing table snapshots collected by Route Views (RV) and RIPE’s Routing Information Service (RIS) [65, 70]. We downloaded one BGP RIB table per day from all available (18 and 16 in 2017, 19 and 18 in 2019 from RIS and RV, respectively) collectors for the same time windows as our traffic data. We extracted all AS paths in these tables that announced reachability to IPv4 prefixes, repeating this process for each week.

AS Siblings. We used CAIDA’s AS to Organization classification of ASes into sets that likely belong to the same organizations [32]. CAIDA’s method parses the Regional Internet Registries’ WHOIS dumps and delegation files to create a unified mapping between ASes and organization names, then uses hints in the name strings, delegation files, identifiers, and email addresses to infer AS sets with common ownership. For each measurement period, we used the AS-to-Organization mapping that CAIDA constructed using WHOIS data collected closest to the traffic capture window.

6 RESULTS

Figure 9 shows the volumes of traffic we classified into each category for two different five-week periods in 2017 and 2019. We present these two five-week periods to show our results are consistent at least for these time periods. In 2017, the peak rate across

the core switch during the period was 120 Gbps; in 2019 the peak had grown to 200 Gbps, and as expected the majority of the traffic across the exchange is classified as in-cone.

In 2017, the peak out-of-cone traffic we inferred was 3.7 Mbps, and in 2019, 40 Mbps. We believe these values are upper-bounds for out-of-cone traffic at the IXP core switch, and we derived these volumes after investigating the underlying properties of traffic between pairs of members, in rank order of contribution to the out-of-cone traffic volume at the IXP. For packets that had a signal they were not spoofed – e.g., a Transmission Control Protocol (TCP) packet with payload, or packets towards a known transport provider, we manually investigated the relationships between the parties. We found 27 sibling ASes in 11 distinct organizations that were exchanging traffic across the IXP, but missing from CAIDA’s public AS-to-Org dataset (§5). To determine which ASes were siblings, we consulted the official website of those ASes to find information on their ownership, contacted the ASes directly to enquire, or contacted the IXP operators to understand the relationship between two ASes at the IXP. Further, through the IXP operators, we approached 36 members of the IXP, and 34 of those members responded with explanations of the behavior we saw.

Although the number of members was similar between 2017 and 2019 (208 and 203, respectively), 28 new members were present in the 2019 analysis. Because we focused our manual investigations on the 2017 data, we believe that there are additional sibling relationships and routing behaviors in the 2019 data that we have not discovered yet. We hypothesize that these missing sibling inferences are the likely cause of the increase in out-of-cone traffic between 2017 and 2019. Table 1 summarizes the number of unique AS pairs we observed to exchange traffic for the five week periods beginning 1 April 2017 and 1 May 2019. While we inferred more

Figure 10: Classification of unverifiable traffic. 61.8% of the unverifiable traffic was sent by a provider to a customer across the exchange. Because a provider can transit packets from any source address in the Internet, there are no invalid addresses which would allow detection of spoofed packets. For completeness, we further classify traffic from each provider as being in or out of their customer cone.

Relationship	April 2017	May 2019
p2p	19,161 (98.7%)	12,057 (98.4%)
p2c	222 (1.1%)	183 (1.5%)
s2s	21 (0.1%)	10 (0.1%)
total	19,404	12,250

Table 1: Unique AS pairs observed exchanging traffic at the IXP in each 5-week period. Approximately 1.4% of AS pairs had a non-p2p relationship. (This IXP was rearchitected in 2019, which may explain the drop in observed peers.)

than 98% of the AS pairs had a p2p relationship, approximately 1.4% of AS pairs had a different class of relationship that impacts our ability to infer SAV policy of the transmitting AS.

Figure 9 also shows the volume of traffic with bogus source addresses, with a peak of approximately 100 Mbps across the exchange for the Wednesday at the end of week 3 (9b-iv). We found these networks deliberately used RFC1918 private addresses as source addresses of packets used to tunnel traffic between members – Generic Routing Encapsulation (GRE) and IP-in-IP represented 61.1% of the traffic, while the other 38.9% were ICMP, TCP, and User Datagram Protocol (UDP).

For both the 2017 and 2019 observation periods, there was a peak of approximately 25 Gbps of unverifiable traffic across the exchange, representing 15.3% of total traffic exchanged at the IXP (figures 9a-ii and 9b-ii). Figure 10 provides a classification of the traffic involved for the first week of May 2019. 61.9% of the unverifiable traffic was sent from a provider to a customer across the exchange, where no cone of valid addresses applies (§2.4). If we had applied the customer cone approach to this p2c traffic, we would have inferred 52% of it was from within the provider’s customer cone, with the remaining 48% of traffic being from outside of the provider’s customer cone. Because a provider can transit packets from any source address in the Internet (§2.4), there are no invalid addresses that would allow detection of spoofed packets. This potential for erroneous inference

Spoofers-CAIDA	Spoofers-IX		Sum
	In-cone	Out-of-cone	
Spoof-received	17	2	19 (54.3%)
Spoof-blocked	14	2	16 (45.7%)
Sum	31 (88.6%)	4 (11.4%)	35

Table 2: Congruity between CAIDA’s public spoofer dataset and inferences using the IXP. Of the 35 overlapping ASes, CAIDA’s spoofer dataset inferred 54% of them had not deployed SAV, because CAIDA received a packet with a spoofed source address. Only 4 of these 35 (11%) were observed to forward an out-of-cone packet into the IXP; 2 of these 4 were in CAIDA’s spoofer dataset as not deploying SAV.

is why we must classify all packets from a transit provider to a customer as unverifiable. Another 21.4% of the unverifiable traffic was because we did not have an AS mapping for either the source or destination MAC addresses (the IXP lacked historical data for this mapping), and for 14.1% of traffic we could not determine the origin AS because the source MAC address and VLAN tag indicated the traffic was from a remote peering provider. Finally, all of the other categories summed to only 2.6% of the traffic, so we do not discuss these categories further.

We inferred out-of-cone traffic for 38 of the 203 members (18.7%) at the IXP between 1 May and 5 July 2019. Of the 203 members, 35 (17.2%) were also in CAIDA’s public spoofer dataset [15], which requires a volunteer to have been present in the network to run an active measurement test that explicitly sends packets with spoofed source addresses to CAIDA’s servers to test SAV deployment of the volunteer’s network (§2.5). Table 2 summarizes the (in)congruity between the two datasets. Of the 35 ASes that overlapped, CAIDA’s spoofer dataset indicated 54% of them had *not* deployed SAV. Only 4 of these 35 ASes (11%) were inferred by Spoofers-IX to forward an out-of-cone packet into the IXP, implying that this IXP may not provide effective visibility into SAV deployment, because participants were not forwarding spoofed packets, at least during our five-week observation window.

Figure 11 shows the volume of out-of-cone traffic inferred by both the Spoofers-IX and full cone methods for traffic data captured during the first week of May 2019. The Spoofers-IX method inferred a peak of 40 Mbps of out-of-cone traffic (best seen in figure 9b), whereas the full cone method inferred a peak of 2.5 Gbps. The diurnal pattern of the inferred out-of-cone traffic matches user-demand for content, with no observable peaks suggesting a volumetric spoofed-source attack launched from within member ASes of the IXP. The second row of figure 11 shows churn in source IP addresses [11, 69] seen in each five minute window. For the full cone method, the absolute volume of source addresses observed follows the traffic volume profile as a whole, and is concentrated in 20-40 ASes per five minute window, which is not a typical pattern of attacks that utilize randomly-spoofed source addresses.

The discrepancy between the size of the traffic classified as out-of-cone by the full cone and Spoofers-IX methods is because the full cone classified some provider-to-customer traffic as being out-of-cone (§2.5), whereas Spoofers-IX classified provider-to-customer traffic as unverifiable. Figure 10 shows Spoofers-IX classified 1 – 5

Figure 11: Comparison of metrics for out-of-cone traffic inferred by the Spoofer-IX and full cone methods for the first week of May 2019. We compute each metric per 5-minute window of traffic data, and use the same range on Y axes between methods to allow for comparison. For the IXP we studied, the full cone method inferred an average of 1.5 Gbps of out-of-cone traffic, whereas our method inferred a maximum of 40 Mbps (best seen in figure 9b-iii).

Figure 12: Spoofer-IX classification of traffic classified as out-of-cone by the full cone method. Spoofer-IX infers that 92.6% of this out-of-cone traffic was from a provider to customer across the IXP, and therefore unverifiable, because a provider can transit traffic from any source IP address to their customer, and it is therefore not feasible to identify spoofed packets by their IP address alone.

Gbps of out-of-cone traffic from providers to customers as part of the unverifiable traffic that Spoofer-IX classified. When we classified the full cone's out-of-cone traffic using the Spoofer-IX method, 92.6% of the traffic was from a provider to a customer across the exchange, carrying 0.5 – 2 Gbps of traffic (figure 12).

Figure 13: Classification of in-cone traffic for the full cone that Spoofer-IX classified as unverifiable. The traffic profile is similar to that in figure 10, with some unverifiable provider-to-customer traffic classified as out-of-cone by the full cone method (figure 12).

Finally, the traffic volume classified as in-cone by the full cone method is larger than that by the Spoofer-IX method. 85.5% of the traffic that the full cone method classified as in-cone was also classified as in-cone by the Spoofer-IX method, with the remaining 14.5% classified as unverifiable by Spoofer-IX. Figure 13 shows how the Spoofer-IX method classified 59.9% of this unverifiable traffic as from a provider to a customer across the IXP, and 26.4% of the unverifiable traffic as out-of-cone for the provider. We hypothesize

that this traffic is classified as in-cone for the full-cone method because some provider ASes (or their customers) provided a BGP view, so the full cone included these addresses as in-cone for these provider ASes (§3.1.3). Note that the traffic profiles in figure 10 and figure 13 are similar: the discrepancy is mostly due to the full cone method classifying some of Spoofer-IX's unverifiable provider-to-customer traffic as out-of-cone (figure 12). However, all routed addresses may be legitimate source addresses in IP packets crossing an IXP from a provider to customer, and no cone of valid addresses can infer the SAV policy of the provider for these packets.

7 DISCUSSION AND INSIGHTS

Challenges of Validation. We could not acquire ground truth data to validate our results, in part due to the negligible amount of out-of-cone traffic we observed, and the challenge of asking any network to validate a small volume of packets. Due to lack of accessible ground truth, we instead verified that our prefix-level customer cone inferences (§3.1.2) were consistent with BGP data extracted from the IXP's route servers. The only inconsistencies we found were due to ASes that had been returned to their RIR and still appeared in public BGP announcements, but did not appear in routes from the IXP route servers.

Generality of the methodology. Assessing the generality of our approach requires applying our method to traffic from other IXPs, which is challenging because it requires the cooperation of other IXP operators. However, we believe our method is generalizable, as we designed and developed Spoofer-IX to accommodate the Best Current Operational Practices (BCOPs) defined by a group of IXPs [28, 37] that describe how IXP operators should configure IXPs. These documents describe how IXP operators should securely configure VLANs and route servers. As such we believe our methodology can be applied to other IXPs; more generally, any other method to infer spoofed traffic in IXP traffic data will have to address the same challenges we encountered.

Applying our method requires two data sets: the traffic data sets themselves, and the metadata that maps IXP infrastructure – VLAN tags on each packet, and MAC addresses to ASes. Our method is automated except for inference of the siblings (§6), which requires some manual effort. However, there are a wide variety of IXP architectures that affect traffic visibility (§3.2), and new IXP architecture innovations to support advanced services will require careful consideration of their impact on our method. Our use of traffic characterization was limited to the packet headers available to us; full payload would enable improvements in traffic analysis, and additional cross-checks.

Emerging IXP trends and their impact on the inference of SAV policy. New IXP services allow networks to self-provision private, on-demand bandwidth in seconds between data center locations (a.k.a, colocation facilities) or cloud service providers, [21, 26, 56, 58, 67]. In 2019, AMS-IX, DE-CIX and LINX joined to develop an API to provision and configure interconnection services at multiple IXPs [55]. The resulting IX-API [5] will allow users to manage their interconnection services, from ordering new ports, to configuring, changing, and canceling services at multiple IXPs. These proposals share a common goal: enable a more dynamic interconnection environment, where networks and IXPs can adapt

to changing conditions. They do not propose to change methods to implement the configurations tackled in this paper, but rather create abstractions to facilitate configuration changes.

8 LESSONS LEARNED

The use of IXPs as a focal point for SAV deployment has received recent attention by both the research [45] and policy communities [36, 63, 74]. However, inferring SAV deployment at an IXP is remarkably challenging, more so than has been captured in the literature, due to a combination of operational complexities that characterize today's interconnection ecosystem, and the inherently heuristic nature of topology and traffic inferences on persistently opaque network infrastructure. Many of our discoveries were eye-opening, although not cause for optimism for those interested in infrastructure protection.

First, although we approached this project aware of several methodological challenges for inferring spoofed packets at IXPs, the reality was even more daunting. We recognized the importance of using the semantics of AS relationships, which is conceptually straightforward but even more painstakingly complicated in practice than we expected. We designed, implemented, and applied a method that accounts for both epistemological and operational challenges, and showed how this method reveals inaccuracies in methods that are agnostic to AS relationship semantics.

But we also found epistemological challenges remain. While we infer out-of-cone traffic with our method at our IXP, there are still edge cases we have not yet explained, as some of the traffic appears to have signatures of legitimate traffic. More importantly, we believe further effort is required to understand the degree to which any IXP could be used as a SAV deployment lens. We publicly release our code [62] in hopes that other researchers and IXPs will use it to further improve our collective ability to measure and expand deployment of SAV filtering. Finally, this work illustrates the deep subtleties of scientific assessments of operational Internet infrastructure, which exemplifies the persistent tension between the need for reproducibility of methods and results [7, 8], and the opacity of commercial infrastructure.

ACKNOWLEDGMENTS

We thank the anonymous reviewers and our shepherd, Sergey Gorinsky, for their valuable feedback on our paper. We are also thankful to Leandro Bertholdo, Bruno Lorensi, Cesar Loureiro, Julio Sirota, Milton Kashiwakura, Demi Getschko – all from IX.br – for their support, feedback, and discussions that allowed this work to be possible. We are also thankful to Anja Feldmann and Franziska Lichtblau who helped to improve our work. This material is based in part on research sponsored by the Department of Homeland Security (DHS) Science and Technology Directorate, Homeland Security Advanced Research Projects Agency, Cyber Security Division via contracts D15PC00188 and 140D7018C0010, the National Science Foundation (NSF) via awards OAC-1724853 and OIA-1937165, Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq) grant 310408/2017-2, and by CAPES/Brazil via Finance Code 001. The published material represents the position of the authors and not necessarily that of the sponsors.

REFERENCES

- [1] Bernhard Ager, Nikolaos Chatzis, Anja Feldmann, Nadi Sarrar, Steve Uhlig, and Walter Willinger. 2012. Anatomy of a Large European IXP. In *ACM SIGCOMM*. 163–174.
- [2] AMS-IX. 2019. AMS-IX Partner Program. <https://www.ams-ix.net/ams/partners>.
- [3] AMS-IX. 2019. AMS-IX Private Interconnect Service. <https://www.ams-ix.net/ams/service/private-interconnect>.
- [4] AMS-IX. 2019. Amsterdam Internet Exchange (AMS-IX). <https://www.ams-ix.net/>.
- [5] AMS-IX and DE-CIX and LINX. 2019. IX-API Simplify your IX services. <https://ix-api.net/>.
- [6] APNIC. 2019. Weekly Routing Table Report. <http://thyme.apnic.net/current/data-summary>
- [7] V. Bajpai, O. Bonaventure, k. claffy, and D. Karrenberg. 2019. Encouraging Reproducibility in Scientific Research of the Internet. *Dagstuhl Reports* 8, 10 (Jan 2019), 41–62.
- [8] Vaibhav Bajpai, Anna Brunstrom, Anja Feldmann, Wolfgang Kellerer, Aiko Pras, Henning Schulzrinne, Georgios Smaragdakis, Matthias Wählisch, and Klaus Wehrle. 2019. The Dagstuhl Beginners Guide to Reproducibility for Experimental Networking Research. *ACM SIGCOMM Computer Communication Review (CCR)* 49, 1 (Feb. 2019), 24–30.
- [9] F. Baker and P. Savola. 2004. Ingress Filtering for Multihomed Networks. RFC 3704 (BCP 84).
- [10] S. M. Bellovin. 1989. Security Problems in the TCP/IP Protocol Suite. *ACM SIGCOMM Computer Communication Review (CCR)* 19, 2 (April 1989), 32–48.
- [11] Karyn Benson, Alberto Dainotti, kc claffy, Alex C. Snoeren, and Michael Kallitsis. 2015. Leveraging Internet Background Radiation for Opportunistic Network Analysis. In *ACM Internet Measurement Conference (IMC)*. 423–436.
- [12] R. Beverly and S. Bauer. 2005. The Spoofer Project: Inferring the Extent of Source Address Filtering on the Internet. In *USENIX Steps to Reducing Unwanted Traffic on the Internet (SRUTI)*.
- [13] Robert Beverly, Arthur Berger, Young Hyun, and k claffy. 2009. Understanding the Efficacy of Deployed Internet Source Address Validation Filtering. In *ACM Internet Measurement Conference (IMC)*. 356–369.
- [14] Xue Cai, John Heidemann, Balachander Krishnamurthy, and Walter Willinger. 2010. Towards an AS-to-organization Map. In *ACM Internet Measurement Conference (IMC)*. 199–205.
- [15] CAIDA. 2019. CAIDA Spoofer Project. <https://www.caida.org/projects/spoofers/>.
- [16] CAIDA. 2019. The CAIDA Internet Topology Data Kit. <http://www.caida.org/data/internet-topology-data-kit>.
- [17] Ignacio Castro, Juan Camilo Cardona, Sergey Gorinsky, and Pierre Francois. 2014. Remote Peering: More Peering Without Internet Flattening. In *ACM Conference on emerging Networking EXperiments and Technologies (CoNEXT)*. 185–198.
- [18] Nikolaos Chatzis, Georgios Smaragdakis, Anja Feldmann, and Walter Willinger. 2013. There is More to IXPs than Meets the Eye. *ACM SIGCOMM Computer Communication Review (CCR)* 43, 5 (Oct. 2013), 19–28.
- [19] B. Claise. 2004. Cisco Systems NetFlow Services Export Version 9. RFC 3954.
- [20] Giovanni Comarela, Gonca Gürsun, and Mark Crovella. 2013. Studying Inter-domain Routing over Long Timescales. In *ACM Internet Measurement Conference (IMC)*. 227–234.
- [21] Console. 2019. Console - The Cloud Connection Company. <https://www.consoleconnect.com/>.
- [22] M. Cotton, L. Vegoda, Ed. R. Bonica, and B. Haberman. 2013. Special-Purpose IP Address Registries. RFC 6890 (BCP 153). Updated by RFC 8190.
- [23] DE-CIX. 2019. DE-CIX Internet Exchange. <https://www.de-cix.net/en/>.
- [24] DE-CIX. 2019. DE-CIX MetroVLAN. <https://www.de-cix.net/en/de-cix-service-world/metrovlan>.
- [25] Z. Duan, X. Yuan, and J. Chandrashekar. 2006. Constructing Inter-Domain Packet Filters to Control IP Spoofing Based on BGP Updates. In *IEEE INFOCOM*. 1–12.
- [26] Epsilon. 2019. Epsilon Telecommunications Limited – Connectivity made simple. www.epsilontel.com/.
- [27] Euro-IX. 2019. IXP BCOPs (Best Current Operational Practices). <https://www.euro-ix.net/en/forixps/set-ixp/ixp-bcops/>.
- [28] Euro-IX. 2019. IXP BCOPs (Best Current Operational Practices), Technical Recommendations. <https://www.euro-ix.net/en/forixps/set-ixp/ixp-bcops/technical-recommendations/>.
- [29] P. Ferguson and D. Senie. 2000. Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing. RFC 2827 (BCP 38). Updated by RFC 3704.
- [30] David Freedman, Brian Foust, Barry Greene, Ben Maddison, Andrei Robachevsky, Job Snijders, and Sander Steffann. 2019. Mutually Agreed Norms for Routing Security (MANRS) Implementation Guide. <https://www.ripe.net/publications/docs/ripe-706>.
- [31] Lixin Gao. 2001. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Transactions on Networking* 9, 6 (Dec. 2001).
- [32] B. Huffaker, K. Keys, R. Koga, M. Luckie, and kc claffy. 2019. CAIDA inferred AS to organization mapping dataset. <https://www.caida.org/data/as-organizations/>.
- [33] IANA. 2019. Autonomous System (AS) Numbers. <https://www.iana.org/assignments/as-numbers/as-numbers.xml>.
- [34] IANA. 2019. IANA IPv4 Address Space Registry. <https://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xhtml>.
- [35] IANA. 2019. Internet Protocol Version 6 Address Space. <https://www.iana.org/assignments/ipv6-address-space/ipv6-address-space.xhtml>.
- [36] Internet Society. 2019. IXP Participants. <https://www.manrs.org/participants/ixp/>.
- [37] Internet Society. 2019. MANRS IXP Programme. <https://www.manrs.org/ixps/>.
- [38] Internet Society. 2019. Mutually Agreed Norms for Routing Security (MANRS). <http://www.manrs.org/manrs>.
- [39] IX Reach. 2019. IX Reach Remove Peering Services. <https://www.ixreach.com/services/remote-peering/>.
- [40] IX.br. 2019. IX.br – Internet Exchange Brazil. <http://ix.br>.
- [41] IX.br Forum 12. 2019. Remote Peering Panel with DEC-IX, AMS-IX, LINX and IX.br. <https://www.youtube.com/watch?v=K283b3AKZ94>.
- [42] Job Snijders. 2016. Practical everyday BGP filtering: Peer Locking (NANOG67). <https://www.youtube.com/watch?v=CSLpWBrHy10>.
- [43] O. Klabi. 2019. 1.3Tbps DDoS mitigated by our VAC. <https://twitter.com/olesovhcom/status/969328679410110466>.
- [44] S. Kottler. 2019. February 28th DDoS Incident Report. <https://githubengineering.com/ddos-incident-report/>.
- [45] Franziska Lichtblau, Florian Streibelt, Thorben Krüger, Philipp Richter, and Anja Feldmann. 2017. Detection, Classification, and Analysis of Inter-domain Traffic with Spoofed Source IP Addresses. In *ACM Internet Measurement Conference (IMC)*. 86–99.
- [46] LINX. 2019. ConneXions at London Internet Exchange Point. <https://www.linx.net/join-linx/connexions/>.
- [47] LINX. 2019. LINX Private VLAN. <https://www.linx.net/products-services/private-vlan/>.
- [48] LINX. 2019. London Internet Exchange (LINX). <https://www.linx.net/>.
- [49] B. Liu, J. Bi, and A. V. Vasilakos. 2014. Toward Incentivizing Anti-Spoofing Deployment. *IEEE TofS* (2014), 436–450.
- [50] Xin Liu, Ang Li, Xiaowei Yang, and David Wetherall. 2008. Passport: Secure and Adoptable Source Authentication. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. 365–378.
- [51] Qasim Lone, Matthew Luckie, Maciej Korczyński, and Michel van Eeten. 2017. Using Loops Observed in Traceroute to Infer the Ability to Spoof. In *Passive and Active Measurement (PAM)*. 229–241.
- [52] Matthew Luckie. 2014. Spurious Routes in Public BGP Data. *ACM SIGCOMM Computer Communication Review (CCR)* 44, 3 (July 2014), 14–21.
- [53] Matthew Luckie, Robert Beverly, Ryan Koga, Ken Keys, Joshua A. Kroll, and k claffy. 2019. Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*.
- [54] Matthew Luckie, Bradley Huffaker, Amogh Dhamdhere, Vasileios Giotsas, and kc claffy. 2013. AS Relationships, Customer Cones, and Validation. In *ACM Internet Measurement Conference (IMC)*. 243–256.
- [55] Lynsey Buckingham. 2019. IX-API For the Good of the Internet. <https://www.linx.net/ix-api-for-the-good-of-the-internet/>.
- [56] Pedro Marcos, Marco Chiesa, Lucas Muller, Pradeeban Kathiravelu, Christoph Dietzel, Marco Canini, and Marinho Barcellos. 2018. Dynam-IX: a Dynamic Interconnection eXchange. In *ACM Conference on emerging Networking EXperiments and Technologies (CoNEXT)*.
- [57] Alex Marder, Matthew Luckie, Amogh Dhamdhere, Bradley Huffaker, Jonathan Smith, and k claffy. 2018. Pushing the Boundaries with bdrmapIT: Mapping Router Ownership at Internet Scale. In *ACM Internet Measurement Conference (IMC)*.
- [58] Megaport. 2019. Megaport - A Better way to connect. <https://www.megaport.com>.
- [59] C. Morales. 2019. NETSCOUT Arbor Confirms 1.7 Tbps DDoS Attack; The Terabit Attack Era Is Upon Us. <https://www.arbortnetworks.com/blog/aserit/netscout-arbor-confirms-1-7-tbps-ddos-attack-terabit-attack-era-upon-us/>.
- [60] R. Morris. 1985. A Weakness in the 4.2BSD Unix TCP/IP Software Technical Report 117, AT&T Bell Laboratories. (1985).
- [61] R. Moskowitz, D. Karrenberg, Y. Rekhter, E. Lear, and G. de Groot. 1996. Address Allocation for Private Internets. RFC 1918 (BCP 5).
- [62] L. Muller, M. Luckie, B. Huffaker, kc claffy, and M. Barcellos. 2019. Spoofer-IX sourcecode. <https://github.com/spoofers-ix/spoofers-ix>.
- [63] NIC.br. 2019. Programa por uma Internet mais segura. <https://bcp.nic.br/i+seg/>.
- [64] George Nomikos, Vasileios Kotronis, Pavlos Sermpezis, Petros Gigis, Lefteris Manassakis, Christoph Dietzel, Stavros Konstantaras, Xenofontas Dimitropoulos, and Vasileios Giotsas. 2018. O Peer, Where Art Thou?: Uncovering Remote Peering Interconnections at IXPs. In *ACM Internet Measurement Conference (IMC)*. 265–278.
- [65] University of Oregon. 2019. Route Views Project. <http://www.routeviews.org/>.
- [66] P. Phaal and S. Panchen, and N. McKee. 2001. InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks. RFC 3176.
- [67] PacketFabric. 2019. PacketFabric. <https://www.packetfabric.com/>.

- [68] PCH. 2019. Raw Routing Data. https://www.pch.net/resources/Raw_Routing_Data/.
- [69] Philipp Richter, Georgios Smaragdakis, David Plonka, and Arthur Berger. 2016. Beyond Counting: New Perspectives on the Active IPv4 Address Space. In *ACM Internet Measurement Conference (IMC)*. 135–149.
- [70] RIPE. 2019. Routing Information Service (RIS). <http://www.ripe.net/ris/>.
- [71] T. Scheid. 2016. Defending the Olympics from DDoS. <https://blog.apnic.net/2016/10/17/defending-olympics-ddos/>.
- [72] Team CYMRU. 2019. IPv4 Fullbogons. <https://www.team-cymru.org/Services/Bogons/fullbogons-ipv4.txt>.
- [73] Team CYMRU. 2019. The Bogon Reference. <http://www.team-cymru.com/bogon-reference.html>.
- [74] Tech Accord. 2019. Cybersecurity Tech Accord. <https://cybertechaccord.org/>.
- [75] The Number Resource Organization. 2019. NRO Extended Allocation and Assignment Reports. <https://www.nro.net/statistics/>.
- [76] J. Weil, V. Kuarsingh, C. Donley, C. Liljenstolpe, and M. Azinger. 2013. IANA-Reserved IPv4 Prefix for Shared Address Space. RFC 6598 (BCP 153).
- [77] A. Yaar, A. Perrig, and D. Song. 2006. StackPi: New Packet Marking and Filtering Mechanisms for DDoS and IP Spoofing Defense. *IEEE Journal on Selected Areas in Communications (JSAC)* (2006), 1853–1863.