Detecting and Measuring In-The-Wild DRDoS Attacks at IXPs

Karthika Subramani[1(✉)], Roberto Perdisci[1,2], and Maria Konte[2]

[1] University of Georgia, Athens, USA
{ks54471,perdisci}@uga.edu
[2] Georgia Institute of Technology, Atlanta, USA
mkonte@gatech.edu

Abstract. Distributed reflective denial of service (DRDoS) attacks are a popular choice among adversaries. In fact, one of the largest DDoS attacks ever recorded, reaching a peak of 1.3 Tbps against GitHub, was a memcached-based DRDoS attack. More recently, a record-breaking 2.3 Tbps attack against Amazon AWS was due to a CLDAP-based DRDoS attack. Although reflective attacks have been known for years, DRDoS attacks are unfortunately still popular and largely unmitigated.

In this paper, we measure in-the-wild DRDoS attacks as observed from a large Internet exchange point (IXP) and provide a number of security-relevant insights. To enable our measurements, we first developed IXmon, an open-source DRDoS detection system specifically designed for deployment at large IXP-like network connectivity providers and peering hubs. We deployed IXmon at Southern Crossroads (SoX), an IXP-like hub that provides both peering and upstream Internet connectivity services to more than 20 research and education (R&E) networks in the South-East United States. In a period of about 21 months, IXmon detected more than 900 DRDoS attacks towards 31 different victim ASes. An analysis of the real-world DRDoS attacks detected by our system shows that most DRDoS attacks are short lived, lasting only a few minutes, but that large-volume, long-lasting, and highly-distributed attacks against R&E networks are not uncommon. We then use the results of our analysis to discuss possible attack mitigation approaches that can be deployed at the IXP level, before the attack traffic overwhelms the victim's network bandwidth.

Keywords: DDoS attack · DRDoS attack · IXP · Traffic analysis

1 Introduction

Large-scale distributed denial of service (DDoS) attacks pose an imminent threat to the availability of critical Internet-based operations [35], and have become part of sophisticated cyber-warfare arsenals [52]. DDoS attacks can take many different forms [43], and leverage weaknesses that span from the application-layer to the physical-layer. In particular, recent incidents have demonstrated that *bandwidth exhaustion* DDoS attacks are capable of bringing down even the most well-provisioned Internet services, such as highly popular websites (e.g., Twitter, Netflix, etc.) and cybersecurity services [27,39,49,67]. Among bandwidth exhaustion

L. Bilge et al. (Eds.): DIMVA 2021, LNCS 12756, pp. 42–67, 2021.
https://doi.org/10.1007/978-3-030-80825-9_3

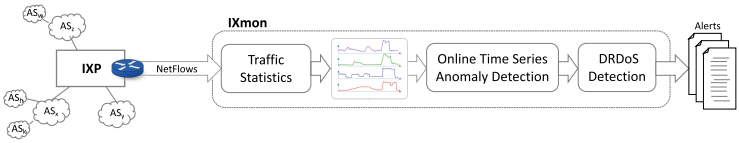

Fig. 1. IXmon system overview

attacks, distributed reflective denial of service (DRDoS) attacks are a popular choice among adversaries [44]. In fact, one of the largest DDoS attacks ever recorded, reaching a peak of 1.3 Tbps against GitHub, was a memcached-based DRDoS attack [20]. More recently, a record-breaking 2.3 Tbps attack against Amazon AWS was due to a CLDAP-based DRDoS attack [51] and attackers have started exploiting Microsoft's RDP for DDoS attacks [63].

Although reflective attacks have been known for years [54] and could be mitigated in part by filtering/throttling traffic to/from some UDP services (e.g., filtering memcached traffic at the edge of a network [23]), DRDoS attacks are unfortunately still popular [33] and largely unmitigated. At the same time, while some information about DRDoS attacks can be found in blog posts or white papers from security vendors (e.g., [22]), there is a lack of systematic studies that provide an in-depth measurement of the properties of *in-the-wild* DRDoS attacks, such as occurrence frequency, the distribution of their sources, duration, volume, targets, and what mitigation steps could be applied to combat them.

In this paper, we aim to partly fill this gap by measuring real-world DRDoS attacks as observed from a large Internet exchange point (IXP)[1]. IXPs are high-density peering and connectivity hubs that provide infrastructure used by autonomous systems (ASes) to interconnect with each other (e.g., public or private peering and other connectivity agreements). Because IXPs provide an increasingly large portion of the global Internet infrastructure used by ASes to exchange traffic, they can play a key role in detecting and mitigating DDoS attacks.

To enable our measurements, we first develop IXmon, an open-source DRDoS detection system specifically designed for deployment at large IXP-like network connectivity providers and peering hubs. While there exists several DDoS detection and mitigation solutions, such as *traffic scrubbing* services [12,21], these are typically expensive third-party commercial services. In addition, they are not designed for detecting DRDoS attacks at IXPs, and are instead more focused on inline DDoS traffic detection and traffic filtering. On the other hand, our IXmon system is fully open-source[2], can be deployed at large IXPs, and can also be used to enable IXP-based DDoS mitigations. IXmon's goal is not only to detect the occurrence of a DRDoS attack very early after its inception, but also to identify ASes that host the reflectors used in the attack. This capability could be used

[1] Whereas others may define IXPs purely as facilitating public peering, we refer to IXPs more broadly as hubs that facilitate both peering and commercial connectivity (e.g., transit) services.

[2] https://github.com/perdisci/IXmon.

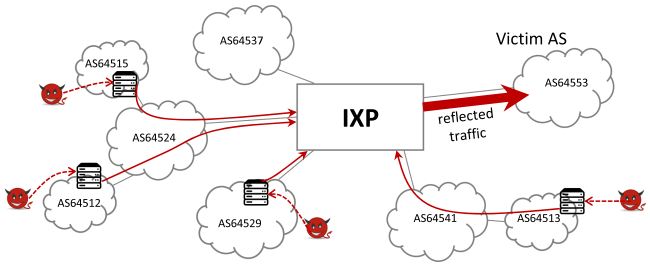

Fig. 2. Example of reflection attack traffic flowing through an IXP

to enable filtering of DRDoS attack traffic at IXP level before it is routed to the victim, thus preventing the victim's network bandwidth from being exhausted.

Figure 1 provides an overview of our IXmon system, whereas Fig. 2 shows an example of how reflected traffic belonging to a DRDoS attack may traverse an IXP's fabric to reach the victim network. To detect DRDoS attacks, IXmon takes in input network flow summaries (e.g., using Cisco's NetFlow v9 format [19]), which report flow statistics for all traffic from any source IP to a any destination IP that crosses the IXP. Because IXmon aims to detect DRDoS attacks, we focus on UDP flows whose source port is associated with services that can be abused for amplification attacks, such as *DNS*, *NTP*, *memcached*, *CLDAP*, etc. [54] (see Sect. 3 for a complete list). Given a specific service (e.g., *memcached*), we then aggregate all related UDP flows directed to each destination AS and compute the overall traffic volume of all flows belonging to the same (service, dstAS) pair. We update these aggregate flow statistics in an online fashion at regular (small) time intervals, and perform online time series anomaly detection to detect highly anomalous increases in traffic volume. Finally, every time an anomalous traffic volume increase is detected for a (service, dstAS) pair, we pass this information to the DRDoS detection module, which applies additional checks to filter out possible false positives and only issue an alert for events that are highly likely associated with actual DRDoS attacks. Additionally, the DRDoS detection module identifies the source ASes involved in an attack, and ranks them according to the attack traffic volume they contribute. By knowing the UDP source port number, the destination AS (i.e., the victim network), and the source ASes that contribute the highest amount of attack traffic, an IXP could then deploy traffic filtering rules to mitigate the attack in its very early stages. In fact, this filtering rule deployment process could be automated by automatically deriving BGPFlowSpec rules [1] from IXmon's alerts.

Notice that while time-series analysis has been previously used in other contexts to detect DDoS attacks and other network traffic anomalies [10,40,62], the contributions of our approach stems from adapting previous approaches to modeling IXP-level traffic and to measuring in-the-wild DDoS attacks at a real-world IXP.

We have deployed IXmon at Southern Crossroads (SoX) [60], an IXP-like hub that provides both peering and upstream Internet connectivity services to

more than 20 research and education (R&E) networks in the South-East United States. In a period of about 21 months, IXmon detected more than 900 DRDoS attacks towards 31 different victim ASes. In Sect. 4, we study the characteristics of these attacks and present a number of insights regarding their duration and intensity, what services are most abused, what networks are more often targeted, and whether the victim networks took action to mitigate the attacks.

In summary, we make the following contributions:

- To measure in-the-wild DDoS attacks, we develop IXmon, an open-source DRDoS detection system (available after publication) specifically designed to be deployed at large IXP-like peering and connectivity hubs.
- We deploy IXmon at a large IXP-like R&E peering and connectivity hub located in the South-East United Sates for a period of about 21 months, where we detected a large number and variety of real-world DRDoS attacks in near real time.
- We analyze the real-world DRDoS attacks detected by our system and report a number of security-relevant measurements and insights. For instance, we show that most DRDoS attacks are short lived, lasting only a few minutes, but that large-volume, long-lasting, and highly-distributed attacks against R&E networks are not uncommon.

2 Background on IXPs

IXPs have been traditionally established as infrastructures that primarily offer peering services. The primary role of an IXP is to serve as a physical exchange point to facilitate the exchange of Internet traffic between different autonomous systems (ASes). The minimum number of ASes that interconnect at an IXP should be at least three and there must be a clear and open policy for other ASes to join [4]. The ASes interconnect through a shared switching fabric that the IXPs offer. This interconnection infrastructure can vary widely in complexity. Some infrastructures can be very simple and minimal (as a single switch), or very complex (as a large scale distributed infrastructure that includes remote peering) [45].

Since their initial establishment, the role of IXPs has been evolving along with their offered services. Some services are offered as free value-added services and others are paid services. Many IXPs offer both public peering and private peering, multi-lateral and bi-lateral peering, data center services, multiple network management and other services including route servers, SDN-based network management, traffic engineering, and traffic blackholing.

IXPs have been recently further evolving towards becoming major peering and connectivity hubs, claiming a central role as part of the Internet's core infrastructures [11,16,53]. There are currently hundreds of IXPs worldwide, with more than 200 just in Europe [3]. IXP membership and traffic growth show their dynamic and evolving role in the Internet ecosystem. Some of the largest IXPs have several hundreds members, while they carry as much traffic as some of the largest global Tier-1 ISPs [3]. It should also be noted that IXPs may serve

different roles in different regions of the world. For example, there exist significant differences between traditional European IXP models and US-based IXPs [9]. In addition, non-profit, EDU-oriented IXPs such as SoX [60] exist with the purpose of helping EDU networks interconnect directly with each other (as in typical IXP peering) but also connect with upstream providers (i.e., providing an exchange point for access to upstream services). In this work, we refer to IXPs in this latter broader sense, as exchange points in which multiple ASes peer with each other and can also connect to upstream Internet connectivity services.

3 IXmon System

In this section, we describe how IXmon's components work, following the high-level overview shown in Fig. 1. It worth noting that IXmon relies on time-series analysis as a component of our detection pipeline. While time-series analysis been previously used in other contexts to detect DDoS attacks and other network traffic anomalies [10,40,62], the contributions of our approach stems from adapting previous approaches to modeling IXP-level traffic and to measuring in-the-wild DDoS attacks at a real-world IXP, as explained below.

Approach Overview: IXmon is designed to detect DRDoS attacks in near *real time* (e.g., with a delay of only one minute) in IXP-like network environments. Given the traffic towards a specific AS, A, to detect DRDoS attacks against A we look for the following factors:

1. Focus on traffic coming from a UDP source port typically associated with a service that can be abused for attack amplification.
2. For each of those source ports, has the traffic volume towards A increased in a highly anomalous way?
3. Is the anomalous traffic distributed across several contributing source ASes?

As an example, assume that a destination AS A usually receives very low amounts of traffic from source port UDP 123, which is typically associated with the NTP service. We monitor all traffic from port UDP 123 that flows towards A through IXP's fabric. All of a sudden, at time t we detect a spike in incoming NTP traffic, and notice that several different source AS numbers are contributing in a coordinated way to this traffic spike. This scenario meets the "recipe" for a DRDoS attack, which IXmon aims to detect automatically. Next, we explain how we translate the above high-level approach into a concrete DRDoS detection system.

3.1 Aggregate Traffic Statistics

IXmon is designed to monitor network traffic at large real-world IXP-like peering and connectivity hubs. Due to the sheer amount of traffic observed from such a vantage point, efficiency is a high priority goal. In particular, memory consumption is a main concern, given the large amount of network traffic statistics that

we need to track over all possible targets and sources of DRDoS attacks visible from an IXP. To this end, our first step is to condense detailed information about network flows crossing the IXP into *traffic sketches* containing aggregated traffic statistics.

IXmon receives network flow statistics as input. While our current implementation supports and has been tested only on Cisco NetFlow versions 9 and 10 [18], it is designed to also support other formats, including sFlow [48]. For simplicity, in the following we will simply use the term *flow* to refer to a network flow in NetFlow format. While NetFlow flows include many details about how the related network packets traversed the IXP (e.g., including the network interfaces involved in routing the flow), we will only refer to the properties that are used by our system. Let the tuple

$$f_i = (srcIP_i, srcPort_i, dstIP_i, dstPort_i, protocol_i, packets_i, bytes_i) \quad (1)$$

represent a network flow, where $packets_i$ and $bytes_i$ represent the number of packets and overall number of bytes sent from the source to the destination IPs/ports that have been "captured" by flow f_i.

The IXP collects all flows crossing its infrastructure by implementing a uniform packet sampling policy to reduce load on its routers and sends them to IXmon in a *stream* (flows are sent out when they are closed by a FIN packet, in case of TCP, or after a configurable timeout managed by the IXP operators). IXmon mines this stream of traffic flows to detect DRDoS attacks in near real time. Given our focus on DRDoS attacks, we keep only flows whose protocol is UDP and whose source port is related to a service that is known to be vulnerable to be used for attack amplification. The set of source port numbers and related UDP services we use in our current configuration of IXmon is inspired by previous work [54,66] and listed in Table 1.

Table 1. List of monitored UDP source ports

Service	Port	Bandwidth amplification factor
DNS	53	28 to 54
NTP	123	556.9
CLDAP	389	56 to 70
CharGen	19	358.8
Memcached	11211	10,000 to 51,000
SunRPC	111	7 to 28
SSDP	1900	30.8
SNMP	161	6.3
SRCDS	27005	–
Call of Duty	20800	–
NETBIOS	137	3.8
RIP	520	131.24
Quake	27960	63.9
Steam	29015	5.5
QOTD	17	140.3

To analyze the continuous large stream of UDP flows received by IXmon, we proceed as follows. First, IXmon partitions time into intervals of fixed length Δt (one minute, in our experiments). Given the set of all flows received during an interval Δt, we map $srcIP$ and $dstIP$ to their respective AS numbers, $srcAS$ and $dstAS$ (e.g., using RouteViews data [7]). This gives us flows:

$$F_i(t) = (srcAS_i, srcPort_i, dstAS_i, dstPort_i, packets_i(t), bytes_i(t)) \quad (2)$$

where t indicates the start of a time interval Δt, *protocol* is omitted since it is constant (always UDP), and the packets and bytes counts vary in time while the other flow parameters are fixed for a given subscript index. Then, given a time interval Δt, we aggregate all flows $F_i(t)$ that share the same source port

and destination AS numbers, and sum up all of their bytes. More formally, we obtain aggregate *sketch* flows of this form:

$$A_k(t) = (srcPort_k, dstAS_k, bytes_k(t)) \tag{3}$$

where $bytes_k(t)$ is the sum of the byte counts contributed by all flows aggregated into $A_k(t)$.

Notice that, given a fixed pair of source port, $srcPort_k$, and destination AS, $dstAS_k$, the AS-level flows $A_k(t)$ give us a time series of total traffic volume (i.e., $bytes_k(t)$) flowing through the IXP that originated from $srcPort_k$ (from any source IP) and destined towards $dstAS_k$ (to any destination IP belonging to that AS and any destination UDP port). Also, while not represented in the above *sketch*, for simplicity, we keep track of the contribution (in terms of total bytes) to flow $A_k(t)$ of each $srcAS_i$ whose traffic is aggregated into the sketch.

3.2 Online Time Series Anomaly Detection

Given a stream of flow sketches $A_k(t)$ related to a $(srcPort_k, dstAS_k)$ pair, we detect anomalous increases in traffic volume by performing an online analysis of the time series represented by $bytes_k(t)$. Specifically, we maintain a time series model consisting of an exponentially-weighted moving average and variance [29], as follows:

$$\mu(t) = \alpha \cdot \mu(t-1) + (1-\alpha) \cdot b(t) \tag{4}$$

$$\sigma^2(t) = (1-\alpha) \cdot (\sigma^2(t) + \alpha \cdot (b(t) - \mu(t-1))^2) \tag{5}$$

where α is a constant and where we omitted the subscript k and used $b(t)$ in place of $bytes_k(t)$, for brevity. Then, given the moving average, $\mu_k(t)$, and variance $\sigma_k^2(t)$ computed at time t for $A_k(t)$, we compute an anomaly (or *deviation*) score as:

$$\delta_k(t) = \max\left(0, \frac{b_k(t) - (\mu_k(t) + \theta \cdot \sigma_k(t))}{b_k(t) + \varepsilon}\right) \tag{6}$$

where θ is a tunable parameter (set to 3 in our experiments) and ε is a small constant (e.g., 10^{-6}) that is only needed to avoid division by zero. Essentially, $\delta_k(t)$ tells us how much $b_k(t)$ deviates (on the positive side) from the moving average plus a tolerance factor proportional to the standard deviation. Notice that $\delta_k(t) \in [0, 1]$, which we use as an anomaly score. The larger $\delta_k(t)$, the more strongly the current reading of A_k's traffic volume, $b_k(t)$, deviates from the expected value plus some tolerance that takes natural variations into account. If $\delta_k(t) > \tau$, where τ is a tunable detection threshold (set to 0.5 in our experiments), we say that the current reading of the traffic volume for the flows aggregated by A_k is anomalous.

Notice that anomalies can be detected in real time, enabling a rapid detection (and a potential automated mitigation) of DRDoS attacks.

Additional Details: At every new time interval, we use Eqs. 4 and 5 to update our time series model. However, once an anomaly is detected, we stop updating the

model until the new traffic volume measurements go back to pre-anomaly levels. More formally, assume t_d is the first time in which an anomaly is detected, we do *not* use the new measurement at time t_d to compute $\mu(t_d+1)$ and $\sigma^2(t_d+1)$. Now, let

$$\delta_k(t+n,t) = \max\left(0, \frac{b_k(t+n) - (\mu_k(t) + \theta \cdot \sigma_k(t))}{b_k(t+n) + \varepsilon}\right) \tag{7}$$

and $t_d = t+1$. In other words, at the time when the anomaly is detected, $n = 1$. At the next time slot, $n = 2$, we compare the latest measurement of the traffic volume $b_k(t+n)$ to the time series model that was last updated at time t. If $\delta_k(t+n,t) > \tau$ this means that the anomalous traffic is still present at time $t+n$, and we continue to keep the same model computed at time t. Let us now assume that at $n = m$ the anomalous levels of traffic revert back to normal. Namely, $\delta_k(t+m,t) \leq \tau$. Then, we use $b_k(t+m)$ to update the values of μ_k and σ_k and keep updating the model at the following time intervals, until another anomaly is identified.

This approach of updating the average and standard deviation only during "normal times" allows us to more easily determine when a traffic volume anomaly, which may represent a DRDoS attack, starts and ends. Specifically, in the example above we can determine that the anomaly started at time $t+1$ and ended at time $t+m$.

3.3 Attack Detection

Let $A_k(t)$ be a traffic sketch time series, and assume that t_d is the time interval in which a time series anomaly has been detected using the approach described in Sect. 3.2. To detect DRDoS attacks in real time while filtering out possible traffic volume anomalies unrelated to reflection attacks, we introduce two additional conditions:

- *Minimum traffic volume*: Given the last aggregate traffic volume measurement, $b_k(t_d)$, we discard the detected anomaly if $b_k(t_d) < \nu$ (in our experiments we set ν to 5 Mbps). The reason is that if the aggregate traffic volume is very low, either the anomaly is not caused by an attack, or the effects of the attack on the target AS's bandwidth are negligible and can be ignored.
- *Source AS volume entropy*: Since we focus on DRDoS attacks, we expect the anomalous traffic volume increase to be distributed across multiple reflectors located in different source ASes.

To compute the source AS volume entropy, we first consider the set of source ASes whose traffic is aggregated into A_k, and take into account the overall number of bytes sent from each of this sources ASes to A_k's destination AS (i.e., the potential victim network). Let $S_k(t_d) = \{s_1, s_2, \ldots, s_n\}$ represents the set of traffic volume amounts contributed by each source AS at time t_d. We then normalize each element in the set as $s'_i = \frac{s_i}{\sum_{j=1}^{n} s_j}$. Finally, we treat s'_i as the

probability of "observing" the i-th source AS as contributor to A_k's aggregate traffic, and compute the entropy $\mathcal{H}(S'_k(t_d))$ of the set $S'_k(t_d) = \{s'_1, s'_2, \ldots, s'_n\}$. If $\mathcal{H}(S'_k(t_d)) = 1$, it means that the traffic from port $srcPort_k$ to $dstAS_k$ is evenly distributed across the contributing source ASes. On the other hand, low values of $\mathcal{H}(S'_k(t_d))$ mean that most of the traffic is contributed by only one (or very few) source ASes. Therefore, we set a threshold h so that traffic volume anomalies are labeled as DRDoS attacks only when $\mathcal{H}(S'_k(t_d)) > h$ (in our experiments, we use $h = 0.4$).

All time series anomalies detected based on the algorithm described in Sect. 3.2 that also meet the two above conditions are labeled as DRDoS attacks. Correspondingly, a DRDoS attack alert is issued, which contains all details of the attack as measured at time t_d, including the destination AS number, source port, current aggregate attack volume, and distribution of traffic amounts from the contributing source ASes. A new alert is issued for every new time interval $t_d + n$ for which the attack is sustained, allowing a network operator to identify whether the attack is still ongoing or has terminated (when no new alert is issued). On the other hand, time series anomalies that do not pass the checks discussed above are logged and can be sent to network operators but are not labeled as DRDoS attacks.

4 Analysis of In-the-Wild Attacks

In this section we provide some background information about SoX, describe how we setup and deployed IXmon at SoX, and present our measurements and analysis of the in-the-wild DRDoS attacks we detected during our deployment period. SoX's customer ASes rely on the IXP's infrastructure for both peering with each other and upstream connectivity. Therefore, SoX provided us with an important vantage point for measuring DRDoS attacks.

Notice that because sizable ground truth datasets of IXP traffic with labeled DRDoS attacks are very difficult to come by (we are not aware of any publicly available dataset of this kind), to tune IXmon's detection parameters we rely on domain knowledge and a manual analysis of IXmon's logs during the preliminary phases of our deployment. In addition, during our preliminary deployment phase we also contacted SoX and its participants to verify some of the attacks detected by IXmon, and we received positive confirmation from network operators that in fact the victim network identified by IXmon was under attack at the time when the alerts were issued. In practice, to tune our systems' detection parameters we take a conservative approach that favors minimizing possible false detections (see Sect. 4.1). While this may cause us to miss some smaller (i.e., lower volume and duration), more subtle DRDoS attacks, these attacks are unlikely to have a significant impact on their target networks.

One possible valuable alternative to enable gathering more ground truth could be to correlate our findings with traffic from DRDoS honeypots [38,64]. At the same time, concurrent work has found that the intersection between attacks observed at IXPs and attacks observed from DRDoS honeypots may be limited [37]. We plan to investigate the overlap between attacks detected by IXmon and DRDoS honeypots in followup work.

In the following analysis we anonymize all AS numbers related to autonomous systems involved in the detected DRDoS attacks, as some of this information may be sensitive (e.g., some of SoX's members may not want to publicly disclose how many attacks their network received and if/how they mitigated them). For instance, we replace AS 10490 with a consistent but randomly chosen identifier of the form "Anon.XXX" (where XXX is a positive integer).

4.1 IXmon Implementation and Setup

We implemented IXmon's flow parsing and traffic aggregation modules in C++, leveraging an open-source tools named FastNetMon [47]. FastNetMon is a DDoS detection system mainly geared towards enterprise networks or single ASes. Its detection approach is not designed to detect and track DRDoS attacks related to many possible large networks and involving large numbers of source and destination ASes, making it unusable for our purposes. For instance, we found that in FastNetMon one would need to explicitly specify all subnets that should be considered as DRDoS attack targets, and that attack detection is done per IP address. In an IXP environment in which many large ASes are the potential targets, in which there can be many sources of attack, and in which we are interested in tracking if the IXP customers are either victims or potentially contributors to DRDoS attacks, we found that FastNetMon would use an exceedingly large amount of resources. Therefore, while we leveraged and adapted the NetFlow parsing module of FastNetMon, we designed and implemented our own open-source IXP-focused online time series anomaly detection and DRDoS detection algorithms using Python. Our IXmon system code can be found on GitHub[3].

As explained in Sect. 3, the mining and aggregation of the NetFlow traffic, which are implemented in C++, allow IXmon to be scalable and process large volumes of traffic typically observed at IXPs (in the order of hundreds of Gbps). During our experiments, IXmon has had no issue keeping up with the large traffic volumes received from SoX, thanks to the use of efficient flow aggregation.

IXmon's online anomaly detection and DRDoS detection algorithms include a few tunable parameters (see Sect. 3). As mentioned earlier, to set our systems' parameters we take a conservative approach that favors minimizing possible false detections. We set the length of the time interval for traffic aggregation $\Delta t = 1$ minute. This interval is long enough to accumulate sufficient aggregate data from the stream of flows related to each $(srcPort_k, dstAS_k)$ pair and to compute meaningful traffic sketches, and at the same time it enables *near real-time* DRDoS detection. Specifically, after traffic sketches are computed they are immediately analyzed and an alert is triggered immediately as attacks are detected in the data stream.

In Eq. 6, we set the parameter $\theta = 3$. Essentially, θ controls how much the traffic volume can deviate from the mean, before an anomaly is detected. The value of $\theta = 3$ is quite conservative, and is inspired by the fact that for Gaussian distributions $Pr(\mu - 3\sigma \leq X \leq \mu + 3\sigma) \approx 99.73\%$. In addition, we set the anomaly

[3] https://github.com/perdisci/IXmon.

detection threshold $\tau = 0.5$. In other words, we tune the system to detect large anomalies, as compared to historic traffic volumes modeled by moving average and standard deviations. While this may cause us to miss small (i.e., low volume) attacks, it makes sure that the anomalies we detect are in fact highly likely related to attacks. This is further reinforced by additional constraints explained in Sect. 3.3.

As for the parameters defined in Sect. 3.3, we set $\nu = 5$ Mbps because DRDoS attacks whose peak traffic is lower are unlikely to cause much disruption to institutional networks (such networks typically have Internet connectivity bandwidth ranging from hundreds of Mbps to tens of Gbps). Finally we set the source AS entropy threshold $h = 0.4$. We tuned this threshold based on a data collected during a preliminary deployment of IXmon, and is meant to capture attacks whose traffic is fairly distributed across multiple sources, rather than all coming mostly from one single source AS.

An additional "operational" parameter is related to the packet sampling rate used by the network operator that provides the raw flows. In IXmon, we take the sampling rate into account, and adjust our traffic measurements accordingly (e.g., we adjust the average traffic volume measured per minute of observation).

4.2 Data Collection at SoX

As mentioned earlier, we deployed IXmon at a large IXP called SoX (AS 10490) that provides peering and Internet connectivity services to several research and education networks. Specifically, we deployed IXmon at one of two routers operated by SoX that enables peering among educational networks and upstream connectivity to Internet2 [5]. This provided us with visibility on most of the traffic crossing the SoX infrastructure (though not all).

Based on public data on AS-to-AS relationships provided by CAIDA [14,41], SoX has more than 20 direct customer networks (also called the IXP *members* or *participants*), peers with 9 other large ASes, and is connected to 5 upstream providers. Furthermore, SoX serves as upstream provider for a variety of smaller ASes that are reachable through it from the rest of the Internet. This study is based on data collected between *April,2018 - April,2020* (due to interruptions due to operational reasons, our traffic monitoring was only active during part of this time period). Overall, we collected traffic information for 634 days. During this period, the source/destination traffic crossing the IXP's fiber was related to a total of *5212* different autonomous systems.

4.3 Attack Measurements and Analysis

In this section, we present an analysis of the DRDoS attacks detected by IXmon to understand their behavior and gain insights that could prove useful for mitigating future attacks. As an example of the attacks that are included in the analysis provided below, Fig. 3 shows a snapshot of two different DRDoS attacks detected by IXmon. Notice that IXmon detected the represented attacks in near real-time (within about one minute from the attack inception). However, the plots in Fig. 3

are formed post-detection stage by combining consecutive attack alerts, and are shown here to visualize the intensity and duration of the attacks as a whole. The x axis shows the time window within which the attack occurred (including a duration of 30 min prior to and after the attack), whereas the y axis shows the volume of traffic contributed by each source AS involved in the attack (the graph is limited to the top 10 source ASes by volume). Each line in the graph represents the traffic sent to the victim AS from a single source AS. For instance, Fig. 3a (top) shows a DRDoS attack that leverages the CLDAP service (source port 389) directed towards AS Anon.2371. The aggregate traffic for the attack, which sums the contribution of all source ASes that sent traffic to AS Anon.2371 from UDP port 389 reached a peak of ≈210 Mbps. It is interesting to notice that before and after the attack there was little or no traffic sent by those source ASes to the destination AS from port 389. Then, all of a sudden all the source ASes start sending high volumes of traffic in a coordinated way, which is a telltale sign of an ongoing DRDoS attack. After all, inter-AS CLDAP use is rarely needed or justified, and it is therefore natural to have very low or no inter-AS CLDAP traffic outside of DRDoS attacks. In addition, having many source ASes sending CLDAP traffic to a common destination AS would be quite a big coincidence for this to be explained by normal activities.

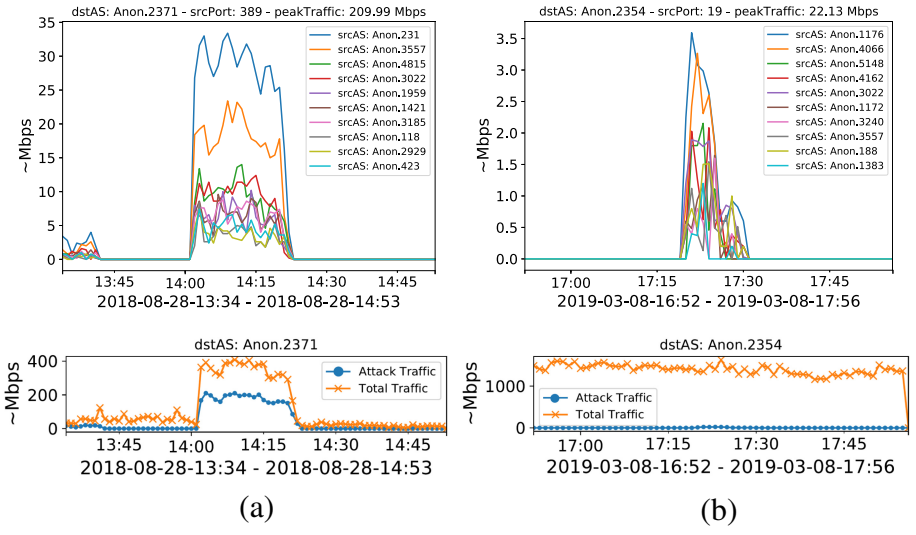

Fig. 3. Two examples of DRDoS attacks detected by IXmon. The top figures show the attack traffic contributed by the top 10 source ASes, while the bottom figures show the overall attack traffic volume compared to all traffic (TCP and UDP) flowing towards the destination AS. Notice also that while these figures span a large time window, IXmon detected the attacks in near real-time.

Volume and Duration. While large DDoS attacks have caught the attention of bloggers and news media, there is limited publicly available data on the overall distribution and characteristics of in-the-wild DRDoS attacks (some information can be found in a 2017 blog post by Cloudflare [22]).

To better understand in-the-wild DRDoS attacks, we analyze the characteristics of all attacks detected by IXmon. Specifically, during our deployment period IXmon detected 987 attacks. We use this large number of attacks to measure the distribution of the volume and duration of in-the-wild DRDoS attacks, which are reported in Figs. 4a and 4b. It can be seen that most of the observed attacks ($\approx 80\%$) have a duration of less than 10 min, whereas the median peak attack volume is less than 20 Mbps. Overall, only $\approx 8\%$ of the attacks reach a peak volume of more than 100 Mbps with a few attacks reaching peaks above 1 Gbps (the highest attack volume we observed was *1.5 Gbps*).

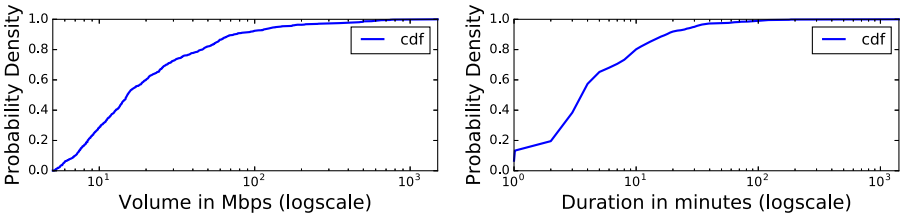

(a) Peak attack volume (b) Attack duration

Fig. 4. Distribution of peak attack volumes and durations

A number of factor may explain the relatively low volume of the attacks we observed, compared to measurements provided in other works [37]. First, we should note that low-volume DDoS attacks are not uncommon [2,8]. Also, tens of Mbps are often sufficient to overwhelm a single machine within a network, although the impact on the network overall may be low. For instance, such DRDoS attacks may be sufficient to knock a competing gamer offline [46,50]. In addition, as mentioned earlier, our system has access to only one of two SoX router and it is therefore possible that additional DRDoS attack traffic was not measurable by our IXmon deployment. In general, we should keep in mind that attacks towards different types of networks (e.g., educational vs. commercial) and measured from different vantage points (e.g., different types of IXPs), may present different characteristics.

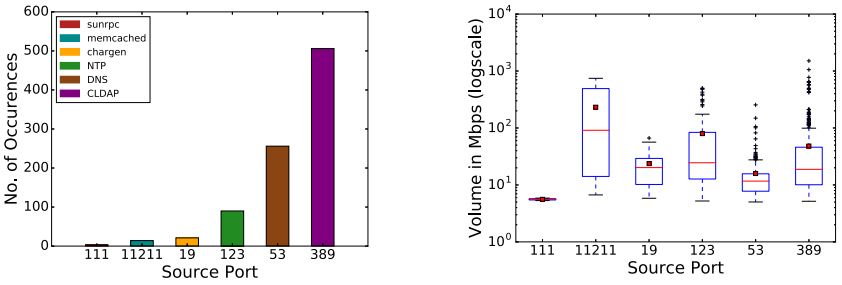

Fig. 5. (a) Number of attack instances per (reflection) source port (b) Distribution of attack volume per (reflection) source port

Services Abused for Attack Amplification. IXmon monitors traffic from the UDP ports listed in Table 1. However, only some of these ports were used in DRDoS attacks visible from SoX. Figure 5a shows the distribution of source ports (ab)used for reflecting traffic against DRDoS victims, with the y axis showing the number of attacks in which a given port was used. As can be seen, CLDAP (port 389) appears to be the most abused service for attack amplification, followed by DNS (port 53) and NTP (port 123). Figure 5b reports a boxplot showing the distribution of peak attack volume per port (the red line represents the median, while the red square shows the average value). This shows that some CLDAP-based attacks reached peak volumes above 1 Gbps.

Multi-vector Attacks. DRDoS attacks can be launched by abusing more than one UDP service at a time. Currently, IXmon separately tracks traffic from a given source port and detects DRDoS attacks independently for each abused service. However, attackers can abuse multiple services at the same to increase the number of reflectors to be aimed against the victim and thus further amplify the attack bandwidth. To analyze these attacks in our alerts dataset, we can retrieve DRDoS attacks related to individual source ports that have a common destination AS and that overlap in time. By doing so, we found 36 multi-vector attack instances (out of 987) involving up to 4 different source ports simultaneously.

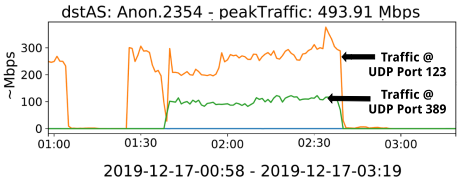

Fig. 6. Example of a multi-vector attack

Figure 6 shows an example of attack detected by IXmon that simultaneously leverages NTP (port 123) and CLDAP (port 389) to reflect the attack traffic towards AS Anon.2354. A coordinated surge in traffic volume can be seen from both source ports, clearly indicating a multi-vector attack.

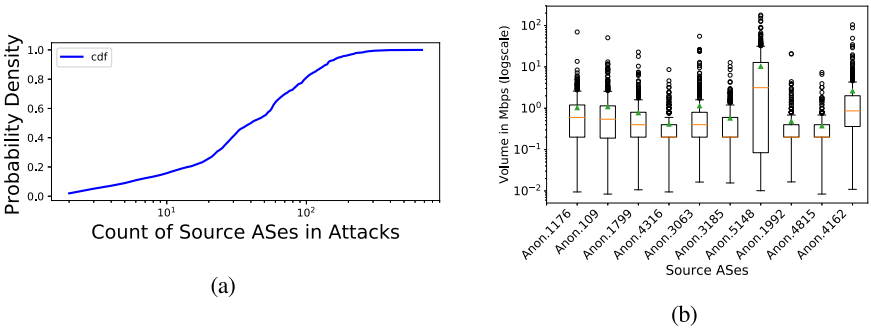

Fig. 7. (a) Distribution of number of source ASes involved in attacks (b) Distribution of peak traffic volume contributed by the top 10 Source ASes in Attacks

Distribution of Reflectors. DRDoS attacks are executed by exploiting a (at times large) number of publicly reachable reflection servers. In this section, we analyze where reflected attack traffic originates from. Figure 7a shows the distribution of the number of different source ASes that contribute to each attack (notice that, due to packet sampling, reflectors that only contribute very low amounts of traffic may not be visible in our data). The median is 40, indicating that at least half of all attacks are highly distributed across many different source networks that are themselves abused to reflect and amplify attack traffic. In Fig. 7b we show the distribution of peak traffic contributed to different attacks by the top 10 source ASes (ranked based on the number of DRDoS attacks each source AS participates to). As can be seen, the median (red line) peak volume for reflected traffic from each AS is relatively limited, typically around ≈ 1 Mbps, though there are also significant outliers with high peak traffic volumes. Either way, when combining together all contributing source ASes the attacks these ASes facilitate can easily reach hundreds of Mbps.

To analyze the geographical distribution of the networks where reflection servers reside, we plot the location of the source ASes that contributed to the DRDoS attacks detected by IXmon. To map the geolocation of a given AS we first obtain the prefixes owned by the AS, based on BGP traffic from the day before the AS participated in an attack. Next, we select a random IP address belonging to one of the prefixes and map the IP address to its geolocation via a IP geolocation API [6]. While this is only an approximate method for determining the geolocation of an AS (some AS numbers span multiple regions), it gives an idea of how geographically distributed the reflectors typically are.

As an example, Fig. 8 shows the geolocation of both destination ASes (i.e., the victims) and source ASes (i.e., the networks that host the servers abused for reflection and amplification) related to NTP-based attacks detected by IXmon. Naturally, given the fact that SoX serves as a peering hub for research and education networks in the South-East USA, the destination ASes are clustered in that region. It is easy to see that the sources of NTP traffic are distributed widely across the world. This is evidently anomalous, in that in normal (i.e., non-attack) cases the vast majority of NTP responses would be coming from NTP servers that are geographically closer to the requesting IP address. Combined with the fact that no NTP requests are sent from a victim AS to those reflection servers, this lack of "locality" could be used as a way to develop an attack mitigation strategy.

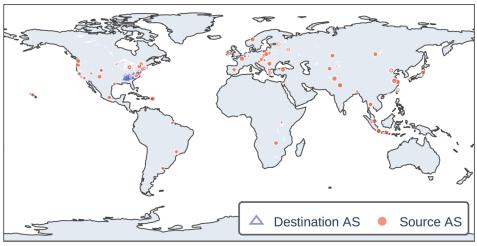

Fig. 8. Geo-locations of source and destination ASes for NTP-based DRDoS attacks

4.4 Attack Mitigation

We now analyze whether the operators of the victim networks attempted to mitigate the attacks detected by IXmon. Specifically, we focus on mitigations that require BGP actions. Afterwards, we discuss how IXmon could help mitigate future attacks by (a) detecting DRDoS attacks in near real time (with a delay of about $\Delta t = 1$ min); (b) determining the AS being targeted by the attack and what service (i.e., source UDP port) is being abused to reflect/amplify attack traffic; and (b) identifying the source ASes that contribute the most to the attack, so that attack traffic originating from those ASes can be filtered out.

Mitigation Strategies. Multiple ways exist to respond to DDoS attacks [55]. However, as we focus on bandwidth exhaustion DRDoS attacks, we ignore mitigations implemented locally at the victim network. Instead, we focus on mitigations that are implemented upstream, with the help of third-party networks such as traffic providers or scrubbing centers, and that make use of BGP to drop or redirect traffic before it reaches the victim network:

– *Blackholing*: BGP-based blackholing redirects all traffic towards a victim AS (both legitimate and malicious traffic) into a null interface, or "blackhole." Although multiple variations of blackholing exist, they are primarily achieved by adjusting the next-hop attribute and BGP communities in BGP announcements [36]. The next-hop method involves the trigger source sending a BGP update to the edge routers with the next-hop attribute set to an IP address that is pre-configured to a *null* interface. The most commonly used next-hop IP for blackholing is 192.0.2.1, which is reserved by IANA for test networks [15].

– *Traffic re-routing*: In this method, all traffic towards the victim network is redirected to third-party services, such as a traffic scrubbing center that is capable of detecting and dropping DDoS attack traffic. Then, legitimate traffic is forwarded back to the original destination (i.e., the victim AS). To re-route traffic, a BGP announcement can be issued by the scrubbing center AS taking ownership of the victim's targeted IP prefixes, essentially performing an *authorized* BGP hijacking. After these BGP announcements propagate, all traffic destined to the victim AS will instead reach the scrubbing center. After the attack has ended, another BGP messages can be issued to reinstate the original IP prefix ownership and again route all traffic to the true destination.

Detecting BGP-Based DRDoS Mitigations. To detect whether mitigations were put into place to counter the attacks detected by IXmon, we perform an analysis of BGP announcements related to the victim ASes before, during and after a DRDoS attack occurrence. To this end, we leverage routing information from RouteViews [7], as explained below:

– *Blackholing*: To detect the use of blackholing mitigations, we monitor the BGP updates involving all IP prefixes owned by a victim AS, and check if any of these updates announce the next-hop to be 192.0.2.1. In addition, we look for BGP updates with a community value set to 666, which is commonly used to implement balckholing [36].
– *Traffic re-routing*: To detect cases in which traffic is re-routed to a third-party AS (e.g., to a scrubbing center), we gather all BGP updates made around a DRDoS attack time window and consider all updates related to IPs that fall within the victim's network ownership. Then, for each such BGP update, we check if the origin AS (extracted from RIB records) has changed, compared to before the attack (e.g., compared to the previous day). If the origin AS in the BGP updates observed during the attack does not match the previously seen origin AS, we mark this as a temporary change in ownership, and check whether future BGP messages also show another change of AS ownership from the third-party AS back to the previous origin AS. We implement this approach using `PyBGPStream` library and Routeview data.

Measuring In-the-Wild Mitigations. Using the BGP-based analysis explained earlier, we measure whether a mitigation effort was deployed for the DRDoS attacks detected by IXmon. With respect to mitigating attacks via traffic re-routing traffic, we found 56 BGP relevant announcements that occurred during 3 different DRDoS attacks. These BGP announcements effectively changed the origin AS of IP prefixes owned by the victim network and redirected traffic to a known traffic scrubbing provider. All of these mitigation efforts were related to attacks directed towards AS Anon.2354, with traffic being re-routed to the AS Anon.1890. All 3 attacks for which mitigation was deployed had a duration greater than 30 min. With respect to mitigation via BGP-based blackholing, we did not find any evidence that blackholing was used for remediating any of the DRDoS attacks we detected.

Figure 9 shows an example of DRDoS attack for which a traffic rerouting mitigation was implemented. As can be seen, in this case the attack had been ongoing for around 45 min, before traffic was re-routed to a scrubbing center. Traffic rerouting is identified by the BGP announcements to change origin (as seen in Fig. 9) of the victim prefix to scrubbing center's AS Anon.1890. Considering this specific AS Anon.2354 that had employed scrubbing services, we performed an experiment to test if our system had missed any attacks for which similar mitigation by rerouting traffic was deployed. To this end, we collected BGP updates for a period of 6 months related to prefixes belonging to our victim AS Anon.2354 whose origin was changed to the scrubbing AS Anon.1890. However, we did not find any evidence in BGP updates to denote an attack that our system missed.

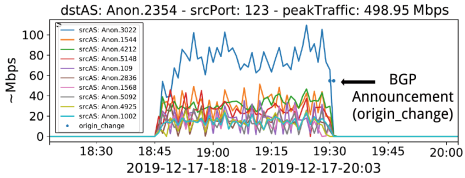

Fig. 9. Example of DRDoS attack and traffic re-routing mitigation

While it was a bit surprising that only 3 attacks and only one network operator used attack mitigation, personal communications with the SoX operators confirmed that only that one member network made use of a DDoS mitigation plan available to all of SoX's customers/members. Another surprising observation is that only long-lived attacks are considered for mitigation. It is possible that one of the main issue is that currently DDoS detection happens "manually," once the attack has started to cause noticeable disruption and perhaps network users start complaining to the operators. Our IXmon system can reduce such detection delay significantly, by performing DRDoS attack detection in near real time with an inexpensive open-source solution.

Improving Attack Mitigation at IXPs. While re-routing traffic to third-party scrubbing services is a commonly used strategy for mitigating DDoS attacks, it can become a quite expensive depending on the size and duration of the attack that a victim is trying to defend against. Another possibility for mitigation is to rely on IXPs and upstream ISPs to implement traffic blackholing. However, as explained earlier, currently blackholing is either an "all or nothing" or very coarsely selective strategy that can cause significant collateral damage [30,59], because it filters out both legitimate and attack traffic. In this section, we discuss how IXmon could enable IXPs to help their customer/members who fall victim of DRDoS attacks, by making traffic blackholing more "surgical" so that only traffic associated with specific services and with specific attack-contributing source ASes is blocked. This has the potential of significantly reducing collateral damage.

The strategy we propose is the following. Let V be the victim AS of a DRDoS attack detected by IXmon, p be the source UDP port abused for reflecting attack traffic towards V, and $S = \{s_1, s_2, \ldots, s_n\}$ be the set of source ASes that send traffic from port p to V during the attack. The IXP could implement a filtering rule that only blocks all traffic from each source AS s_i and port p directed towards V. Because all information necessary to create these filtering rules is contained in IXmon's DRDoS alerts, it would be possible to simply automatically translate each alert into a BGP Flowspec rule that can be propagated to the IXP's routers thus greatly reducing the mitigation time compared to manual intervention.

To understand what is the potential impact to the above strategy, we investigate the extent of the "collateral damage" (i.e., blocked non-DRDoS traffic) a target network may incur. To this end, let us consider the measurements shown in Fig. 10. Each heatmap corresponds to one of the UDP source ports reported in Fig. 5a, from which we observed at least one DRDoS attack. All four heatmaps are related to one single destination AS, which we select as the AS number for which we observed the largest number of distinct DRDoS attacks, during IXmon's deployment period. The x axis reports a period of 30 consecutive days of traffic monitoring, whereas the y axis reports a randomly selected set of 20 source ASes. These source ASes were selected among all source ASes that during the 30 days period in the x axis sent at least some traffic from any of the six source UDP ports. The color of each heatmap cells indicates the total number of MBytes sent by a source AS to the destination AS during each day. Gray cells indicate zero bytes, whereas other cell colors indicated the "intensity" of the daily traffic. From all these graphs we exclude attack traffic detected by IXmon. The reason is that we want to highlight the volume of normal (i.e., non-DRDoS attack) traffic typically sent by any source AS to a destination AS, as seen from the vantage point of an IXP.

Let us consider first Fig. 10a, which is related to CLDAP traffic (port 389). As we can see, it is rare to observe any inter-AS traffic for this service. This makes sense, in that CLDAP is primarily meant as an authentication protocol to be used within a local network. Similarly, ports 19, 111, and 11211 are unlikely to be used for legitimate inter-AS communication purposes. Therefore, blocking inter-AS traffic from these ports at the IXP level is unlikely to cause much collateral damage at all. Services such as NTP (port 123) and DNS (port 53) have a different traffic profile. Inter-AS traffic in these cases is not uncommon, though the overall volume can be quite low, and therefore traffic filtering can produce some observable collateral damage. For example, filtering all source port 53 traffic towards a destination AS may impact DNS resolutions for domains whose authoritative name servers are located within the destination AS. However, let us assume IXmon detects an attack related to one of these ports/services. A BGP Flowspec rule automatically derived from IXmon's DRDoS alert would suggest that the IXP filter all traffic coming from the identified attack source port directed to the victim network. In addition the filtering rule would specify what source ASes are contributing to the attack, so that the IXP could block only traffic from a specific source port and a specific subset of source ASes, thus

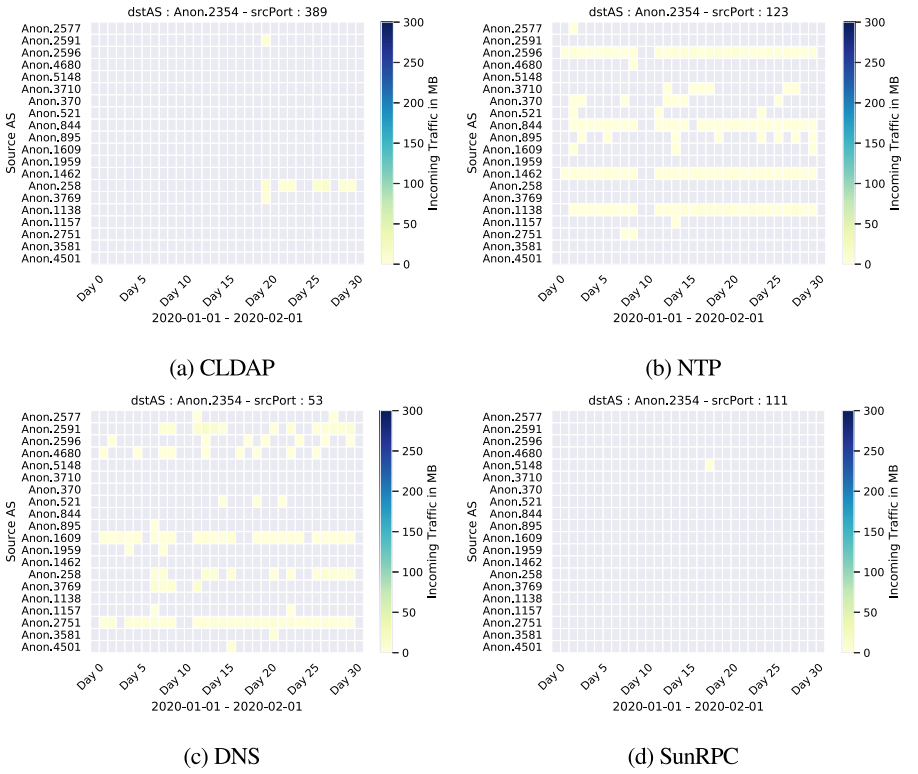

Fig. 10. Daily traffic (in MBytes per day) to destination AS Anon.2354 from a set of 20 legitimate ASes not involved in DRDoS attacks.

further limiting possible collateral damage. Furthermore, filtering could be limited to the duration of the attack. As soon as IXmon detects that the DRDoS attack is over, a new BGP Flowspec rule could be issued so that the IXP would stop filtering any traffic towards the target AS. This approach could help IXPs protect their downstream customer/member networks from bandwidth exhaustion DRDoS attacks with minimal collateral damage.

5 Related Work

In this section we are presenting prior work in the area of DDoS detection and mitigation both in the context of IXPs and large Tier-1 ISPs.

Detection: Concurrently to our work, Kopp et al. [37] also studied amplification attacks from an IXP. Many of our findings agree with their results [37]. However, our work differs from [37] in the following ways: 1) IXmon can detect low volume attacks, whereas [37] only focuses on attacks with volume ≥ 1 Gbps; 2) IXmon provides insights into traffic from research and education networks in the

US, whereas [37] focuses mostly on commercial networks; 3) IXmon is an open-source system that can be used for near real-time detection of DRDoS attacks, whereas [37] appears to present offline traffic analysis results.

Sekar et al. [56] proposes LADS, a multi-stage flow collection and monitoring infrastructure for DDoS detection at Tier-1 ISPs that relies on SNMP and NetFlow feeds from routers. While LADS's detection approach also relies on detecting traffic volume anomalies, IXmon uses a more lightweight approach based on time series anomaly detection that is entirely focused towards an IXP-based deployment. Rossow et al. [54] provide a detailed study of how different protocols can be abused for amplification attacks, and analyze DRDoS traffic at a large ISP. The authors set up multiple bait services and monitored their abuse by attackers and also propose ways to identify DRDoS victims and legitimate reflectors that are abused to amplify the attacks. For instance, traffic asymmetries between the victim and reflectors are used as a telltale sign that the victim never requested traffic from the reflector, and that the incoming traffic is instead due to spoofing. We also explored using a similar feature in our system. However, as we attempted to measure such traffic imbalances to detect spoofed traffic, we observed that the heavy traffic sampling applied by SoX did not allow us to detect spoofed traffic with high confidence, and we therefore chose not to use this feature in IXmon.

Hsieh and Chan [34] propose a neural networks approach to detect DDoS attacks. They rely on network features such as number of packets, number of bytes, time interval variance, packet rate and bit rate. Similarly, [68] proposes to detect DDoS attacks based on Naive Bayes and Random Forest trees. The drawback of these approaches is that they are not designed for real-time traffic analysis and deployment at large IXPs. Furthermore, they require large volumes of historical labeled data for reliable model training, which is often difficult to collect.

BGP-Based Mitigation: Past research [26,30] has developed BGP-based techniques that an infrastructure operator can use to mitigate DDoS attacks. These techniques work in the premise that a network operator has already deployed a tool to detect DDoS attacks. Once a DDoS attack is detected then the network operator can inform an upstream provider, for example, a higher-tier ISP or an IXP, to enforce BGP-based rules and redirect the attack traffic away from the victim network. The techniques are primarily based on: a) BGP Blackholing, and b) BGP Flowspec rules. [25,30,44] offer a detailed description and measurements of the BGP blackholing technique that has become popular and is offered as a service at many IXPs. [1] offers an example application of BGP Flowspec rules. Another study [57,58,65] proposes an additional BGP-based technique, called BGP poisoning, to filter out attack traffic. Our work differs from these approaches because we focus on designing a detection system that can be deployed at IXPs to enable the measurement of DRDoS attack characteristics and that could also be used to enable faster and more selective attack mitigation.

SDN-Based Mitigation: Previous works propose systems that leverage the capabilities of Software Defined Networking (SDN) technologies and Network Functions Virtualization (NFV) to detect and mitigate DDoS attacks. To overcome BGP-based mitigation techniques [13,61], Fayaz et al. [28] propose an OpenDayLight [42] controller and a network of Virtual Machines (VMs) for increased scalability. The controller is designed to route the traffic through the VMs to scrub the traffic. Gupta et al. [17,31,32] and Dietzel et al. [24] have proposed SDN enabled applications as a network management solution for IXPs. Our approach is not based on the SDN and NVF paradigms. Instead, our system can complement these approaches because it can be deployed on infrastructures that do not have SDN-based capabilities, and could be adapted to work with SDN-based traffic routing infrastructure at IXPs to mitigate DRDoS attacks in a very selective way with low collateral damage.

6 Conclusion

In this paper, we studied in-the-wild DRDoS attacks as seen from a large Internet exchange point (IXP). To enable this study, we first developed IXmon, an open-source DRDoS detection system specifically designed for deployment at large IXP-like network connectivity providers and peering hubs. We then deployed IXmon at Southern Crossroads (SoX), an IXP-like hub that provides both peering and upstream Internet connectivity services to more than 20 research and education (R&E) networks in the South-East United States. In a period of about 21 months, IXmon detected more than 900 DRDoS attacks towards 31 different victim ASes. An analysis of the real-world DRDoS attacks detected by our system shows that most DRDoS attacks are short lived, lasting only a few minutes, but that large-volume, long-lasting, and highly-distributed attacks against R&E networks are not uncommon. We then used the results of our analysis to discuss possible attack mitigation approaches that can be deployed at the IXP level, before the attack traffic overwhelms the victim's network bandwidth.

Acknowledgments. We would like to thank the anonymous reviewers for their constructive comments and suggestions on how to improve this paper, and Prof. Christian Rossow for serving as our shepherd. Also, many thanks to the SoX network operators for their help with IXmon's deployment. This material is based in part upon work supported by the National Science Foundation (NSF) under grants No. 1741607 and 1741608. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

References

1. BGP flowspec. https://archive.nanog.org/sites/default/files/wed.general.trafficdiversion.serodio.10.pdf
2. DDoS attack frequency grows 40%, low volume attacks dominate. https://www.helpnetsecurity.com/2018/09/13/ddos-attack-frequency-grows/

3. Euro IX- internet exchange points. https://www.euro-ix.net/media/filer_public/d5/84/d584495f-b8ae-4f24-b589-7b9efed3594b/ixp_report_2018-2019-final.pdf
4. European internet exchange association 2012 report on European IXPs. https://www.euro-ix.net/documents/1117-Euro-IX-IXP-Report-2012-pdf
5. Internet2: Regional research and education networks. https://internet2.edu/network/state-and-regional-r-e-networks/
6. IP geolocation mappingk. https://ipgeolocation.io/ip-location-api.html
7. Routeviews project. http://www.routeviews.org/routeviews/
8. Threat actors target remote learning during COVID-19. https://www.netscout.com/blog/threat-actors-target-remote-learning-during-covid-19
9. The U.S. vs. European internet exchange point models. http://drpeering.net/HTML_IPP/chapters/ch12-9-US-vs-European-Internet-Exchange-Point/ch12-9-US-vs-European-Internet-Exchange-Point.html
10. Fouladi, R.F., Ermiş, O., Anarim, E.: A DDoS attack detection and defense scheme using time-series analysis for SDN. J. Inf. Secur. Appl. **54**, 102587 (2020). https://doi.org/10.1016/j.jisa.2020.102587
11. Ager, B., Chatzis, N., Feldmann, A., Sarrar, N., Uhlig, S., Willinger, W.: Anatomy of a large European IXP. In: Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication, pp. 163–174 (2012)
12. Akamai: Why Akamai cloud security for DDoS protection? https://www.akamai.com/us/en/solutions/products/cloud-security/ddos-protection-service.jsp
13. Butler, K., Farley, T.R., McDaniel, P., Rexford, J.: A survey of BGP security issues and solutions. Proc. IEEE **98**(1), 100–122 (2010)
14. CAIDA: As relationship. http://data.caida.org/datasets/as-relationships/
15. Network Startup Resource Center: Remote blackhole filtering lab. https://nsrc.org/workshops/2019/mnnog1/riso/networking/routing-security/en/labs/RTBH-local.html
16. Chatzis, N., Smaragdakis, G., Feldmann, A., Willinger, W.: There is more to IXPs than meets the eye. ACM SIGCOMM Comput. Commun. Rev. **43**(5), 19–28 (2013)
17. Chiesa, M., et al.: Inter-domain networking innovation on steroids: empowering IXPs with SDN capabilities. IEEE Commun. Mag. **54**(10), 102–108 (2016)
18. CISCO: Netflow layer 2 and security monitoring exports. http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/netflow/configuration/12-4/nf-12-4-book/nf-lay2-sec-mon-exp.html
19. CISCO: Netflow v9. https://www.cisco.com/en/US/technologies/tk648/tk362/technologies_white_paper09186a00800a3db9.html
20. CloudFlare: Famous DDoS attacks learning objectives. https://www.cloudflare.com/learning/ddos/famous-ddos-attacks/
21. Cloudflare: How cloudflare's architecture allows us to scale to stop the largest attacks. https://blog.cloudflare.com/how-cloudflares-architecture-allows-us-to-scale-to-stop-the-largest-attacks/
22. CloudFlare: Reflections on reflection (attacks) (2017). https://blog.cloudflare.com/reflections-on-reflections/
23. CloudFlare: Memcrashed - major amplification attacks from UDP port 11211 (2018). https://blog.cloudflare.com/memcrashed-major-amplification-attacks-from-port-11211/
24. Dietzel, C., Antichi, G., Castro, I., Fernandes, E.L., Chiesa, M., Kopp, D.: SDN traffic engineering and advanced blackholing at IXPs. In: Proceedings of the Symposium on SDN Research (2017)

25. Dietzel, C., Feldmann, A., King, T.: Blackholing at IXPs: on the effectiveness of DDoS mitigation in the wild. In: Proceedings of Passive and Active Measurement: 17th International Conference, PAM 2016, Heraklion, Greece, 31 March–1 April 2016 (2016)
26. Dietzel, C., Wichtlhuber, M., Smaragdakis, G., Feldmann, A.: Stellar: network attack mitigation using advanced blackholing. In: Proceedings of the 14th International Conference on emerging Networking EXperiments and Technologies, pp. 152–164 (2018)
27. Digital Attack Map: DDoS attacks worldwide. http://www.digitalattackmap.com
28. Fayaz, S.K., Tobioka, Y., Sekar, V., Bailey, M.: Bohatei: flexible and elastic DDoS defense. In: 24th USENIX Conference on Security Symposium. USENIX Association, USA (2015)
29. Finch, T.: Incremental calculation of weighted mean and variance (2009). https://fanf2.user.srcf.net/hermes/doc/antiforgery/stats.pdf
30. Giotsas, V., Smaragdakis, G., Dietzel, C., Richter, P., Feldmann, A., Berger, A.: Inferring BGP blackholing activity in the internet. In: 2017 Internet Measurement Conference (2017)
31. Gupta, A., et al.: An industrial-scale software defined internet exchange point. In: 13th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 2016), pp. 1–14 (2016)
32. Gupta, A., et al.: SDX: a software defined internet exchange. ACM SIGCOMM Comput. Commun. Rev. **44**(4), 551–562 (2014)
33. Hao, M.: DDoS attack landscape (2020). https://nsfocusglobal.com/ddos-attack-landscape-3/
34. Hsieh, C., Chan, T.: Detection DDoS attacks based on neural-network using apache spark. In: 2016 International Conference on Applied System Innovation (ICASI), pp. 1–4 (2016)
35. Kang, M.S., Lee, S.B., Gligor, V.D.: The crossfire attack. In: Proceedings of the 2013 IEEE Symposium on Security and Privacy (2013). http://dx.doi.org/10.1109/SP.2013.19
36. King, T., Dietzel, C., Snijders, J., Doering, G., Hankins, G.: Blackhole BGP community for blackholing. https://tools.ietf.org/html/draft-ietf-grow-blackholing-00
37. Kopp, D., Dietzel, C., Hohlfeld, O.: DDoS never dies? An IXP perspective on DDoS amplification attacks. In: Hohlfeld, O., Lutu, A., Levin, D. (eds.) PAM 2021. LNCS, vol. 12671, pp. 284–301. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-72582-2_17
38. Krämer, L., et al.: AmpPot: monitoring and defending against amplification DDoS attacks. In: Bos, H., Monrose, F., Blanc, G. (eds.) RAID 2015. LNCS, vol. 9404, pp. 615–636. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-26362-5_28
39. Krebs, B.: Krebsonsecurity hit with record DDoS. https://krebsonsecurity.com/2016/09/krebsonsecurity-hit-with-record-ddos/
40. Li, D., Chen, D., Goh, J., kiong Ng, S.: Anomaly detection with generative adversarial networks for multivariate time series (2019)
41. Luckie, M., Huffaker, B., Dhamdhere, A., Giotsas, V., Claffy, K.: As relationships, customer cones, and validation. In: 2013 Conference on Internet Measurement Conference (2013). https://doi.org/10.1145/2504730.2504735
42. Medved, J., Varga, R., Tkacik, A., Gray, K.: Opendaylight: towards a model-driven SDN controller architecture. In: Proceeding of IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks 2014, pp. 1–6. IEEE (2014)
43. Mirkovic, J., Reiher, P.: A taxonomy of DDoS attack and DDoS defense mechanisms. SIGCOMM Comput. Commun. Rev. **34**(2), 39–53 (2004)

44. Nawrocki, M., Blendin, J., Dietzel, C., Schmidt, T.C., Wählisch, M.: Down the black hole: dismantling operational practices of BGP blackholing at IXPs. In: Proceedings of the Internet Measurement Conference, pp. 435–448. Association for Computing Machinery, New York (2019)
45. Norton, W.: The Internet Peering Playbook: Connecting to the Core of the Internet. DrPeering Press (2011). https://books.google.com/books?id=rkDz6fvX_XkC
46. NSFOCUS: Have rich game customers who suffered DDoS attacks turned to you? https://nsfocusglobal.com/have-rich-game-customers-who-suffered-ddos-attacks-turned-to-you/
47. Odintsov, P.: Fastnetmon - very fast DDoS analyzer with sflow/netflow/mirror support. https://github.com/pavel-odintsov/fastnetmon
48. Phaal, P., Lavine, M.: sFlow version 5. http://sflow.org/sflow_version_5.txt
49. Prince, M.: The DDoS that knocked spamhaus offline. https://blog.cloudflare.com/the-ddos-that-knocked-spamhaus-offline-and-ho/
50. Radware: 3 attack surfaces that can take your game offline. https://blog.radware.com/security/ddosattacks/2020/10/3-attack-surfaces-that-can-take-your-game-offline/
51. Computer Business Review: AWS hit with a record 2.3 Tbps DDoS attack. https://www.cbronline.com/news/record-ddos-attack-aws
52. Richards, J.: Denial-of-service: the Estonian cyberwar and its implications for U.S. National Security. http://www.iar-gwu.org/node/65
53. Richter, P., Smaragdakis, G., Feldmann, A., Chatzis, N., Boettger, J., Willinger, W.: Peering at peerings: on the role of IXP route servers. In: Proceedings of the 2014 Conference on Internet Measurement Conference, pp. 31–44 (2014)
54. Rossow, C.: Amplification hell: revisiting network protocols for DDoS abuse. In: Proceedings of the 2014 Network and Distributed System Security (NDSS) Symposium, February 2014
55. Ryba, F.J., Orlinski, M., Wählisch, M., Rossow, C., Schmidt, T.C.: Amplification and DRDoS attack defense - a survey and new perspectives (2015)
56. Sekar, V., Duffield, N.G., Spatscheck, O., van der Merwe, J.E., Zhang, H.: Lads: large-scale automated DDoS detection system. In: USENIX Annual Technical Conference (2006)
57. Smith, J.M., Schuchard, M.: Routing around congestion: defeating DDoS attacks and adverse network conditions via reactive BGP routing. In: 2018 IEEE Symposium on Security and Privacy (2018)
58. Smith, J., Birkeland, K., McDaniel, T., Schuchard, M.: Withdrawing the BGP re-routing curtain: understanding the security impact of BGP poisoning through real-world measurements (2020)
59. Snijders, J.: DDoS damage control, cheap and effective. https://ripe68.ripe.net/presentations/176-RIPE68_JSnijders_DDoS_Damage_Control.pdf
60. SoX: Southern crossroads. https://www.sox.net/
61. Streibelt, F., et al.: BGP communities: even more worms in the routing can. In: Proceedings of the Internet Measurement Conference 2018, pp. 279–292 (2018)
62. Tabatabaie Nezhad, S.M., Nazari, M., Gharavol, E.A.: A novel DoS and DDoS attacks detection algorithm using Arima time series model and chaotic system in computer networks. IEEE Commun. Lett. **20**(4), 700–703 (2016). https://doi.org/10.1109/LCOMM.2016.2517622
63. arsTECHNICA: DDoSers are abusing Microsoft RDP to make attacks more powerful. https://arstechnica.com/information-technology/2021/01/ddosers-are-abusing-microsoft-rdp-to-make-attacks-more-powerful/

64. Thomas, D.R., Clayton, R., Beresford, A.R.: 1000 days of UDP amplification DDoS attacks. In: 2017 APWG Symposium on Electronic Crime Research (eCrime), pp. 79–84 (2017)
65. Tran, M., Kang, M.S., Hsiao, H., Chiang, W., Tung, S., Wang, Y.: On the feasibility of rerouting-based DDoS defenses. In: 2019 IEEE Symposium on Security and Privacy (SP) (2019)
66. US CERT: UDP-based amplification attacks. https://www.us-cert.gov/ncas/alerts/TA14-017A
67. York, K.: DYN statement on 10/21/2016 DDoS attack. http://dyn.com/blog/dyn-statement-on-10212016-ddos-attack/
68. Zhang, B., Zhang, T., Yu, Z.: DDoS detection and prevention based on artificial intelligence techniques. In: 2017 3rd IEEE International Conference on Computer and Communications (ICCC), pp. 1276–1280 (2017)