# Blackholing at IXPs:
# On the Effectiveness of DDoS Mitigation in the Wild

Christoph Dietzel[1,2], Anja Feldmann[1], and Thomas King[2]

[1] TU Berlin [2] DE-CIX

**Abstract.** DDoS attacks remain a serious threat not only to the edge of the Internet but also to the core peering links at Internet Exchange Points (IXPs). Currently, the main mitigation technique is to blackhole traffic to a specific IP prefix at upstream providers. Blackholing is an operational technique that allows a peer to announce a prefix via BGP to another peer, which then discards traffic destined for this prefix. However, as far as we know there is only anecdotal evidence of the success of blackholing.

Largely unnoticed by research communities, IXPs have deployed blackholing as a service for their members. In this first-of-its-kind study, we shed light on the extent to which blackholing is used by the IXP members and what effect it has on traffic.

Within a 12 week period we found that traffic to more than $7,864$ distinct IP prefixes was blackholed by $75$ ASes. The daily patterns emphasize that there are not only a highly variable number of new announcements every day but, surprisingly, there are a consistently high number of announcements ($> 1000$). Moreover, we highlight situations in which blackholing succeeds in reducing the DDoS attack traffic.

## 1 Introduction

Distributed Denial of Service (DDoS) attacks are and will continue to be a serious threat to the Internet. Indeed, the intensity and the dimension of such attacks is still rising, in particular due to amplification and reflection attacks [7,32,33]. DDoS attacks impact not only edge networks but can also overwhelm cloud services [36] or congest backbone peering links at Internet Exchange Points (IXP) [30]. Various DDoS detection and defense mechanisms strive to diminish the impact of attack traffic on the victim's infrastructure while minimizing the collateral damage to legitimate traffic. While there has been some progress towards limiting amplification [19], DDoS attacks remain a major security challenge as new protocol or implementation weaknesses are identified almost daily [38].

Various taxonomies [18,23,37] distinguish between proactive (preventive) and reactive techniques. Among the reactive defenses, we distinguish between *source-based*, *destination-based*, and *network-based* [39] mechanisms depending on where they are deployed. In this paper, we focus on how *blackholing* – a network-based reactive defense mechanism – is used at IXPs.

The term blackhole originates in physics and describes an object with such a strong gravitation that nothing can escape from it. In networking it refers to situations where

IP packets are silently discarded, often due to misconfiguration. Indeed, since the late-1980s, blackholing has been used – on a per device basis – to counter DDoS attacks [13]. In 2002, Greene [12] proposed to extend blackholing to routers within an Autonomous System (AS) via iBGP communities, see RFC 3882. In eBGP, an AS is able to communicate to another AS for which prefix the packets should be dropped via BGP communities [5]. In 2009, Kumari and McPherson extended the community ranges to include dropping by source addresses, see RFC 5635. Major Internet Service Providers (ISP), e.g., DT, NTT, and Hurricane Electric, use blackholing within their network and have been offering blackholing services since between 2005 and 2007 to their customers [9, 15, 27].

However, the use of eBGP blackholing services by a DDoS victim is not trivial as the victim has to contact its direct neighbors. The signaling has to be done on a per neighbor basis. IXPs simplify this by acting as a proxy. They offer a public peering infrastructure and the major IXPs have more than 500 member ASes. Due to this multiplication factor, IXPs are in principle convenient locations for blackholing. First ad hoc uses of blackholing occurred around 2010. The blackholing feature is now available at some major IXPs such as DE-CIX, MSK-IX, NETIX, NIX.CZ, and TPIX [8, 25, 26].

In this paper, we rely on three month's worth of routing and traffic measurements from one of the largest IXPs worldwide to examine the extent of blackholing usage and its effectiveness. We find a significant number of blackholes announced, mainly $/32$ but also less specific. Indeed, the usage considerably depends on the prefix length and the announcing member AS. Furthermore, we reveal that blackholing succeeds in reducing DDoS attack traffic.

## 2 Blackholing at IXPs

Blackholing is used as a DDoS mitigation strategy inside a single or between multiple ASes. Consequently, the victim AS announces the attacked destination IP prefix upstream network via BGP. Traffic towards these prefixes is discarded upstream, usually at the upstream AS ingress point. This reduces the amount of traffic not only for the destination network but also for all upstream ASes.

Historically, blackholing was implemented at the edge routers of an AS. However, over time it was moving from the edge (customer or provider networks) to the core of the Internet (ISPs and IXPs).

**IXPs:** IXPs are shared and settlement free peering platforms that operate a switching fabric to interconnect its members' networks. Among the member ASes that exchange traffic are typically a wide range of network types, e.g., Tier-1 ISPs, regional providers, hosters, content providers, CDNs, and even IXP resellers. Many IXPs offer route servers as a free value-added service [31]. It greatly simplifies the BGP session management for their connected members. Therefore, route servers collect routing information in a centralized manner and redistribute them to connected member routers.

If an IXP-connected network (AS) is hit by a massive DDoS attack that causes large amounts of ingress traffic over the IXP link, either the network itself or the network interconnection link is at risk of congestion. As a last resort, either operators of the targeted AS can trigger blackholing for its own prefixes or blackholing is triggered on

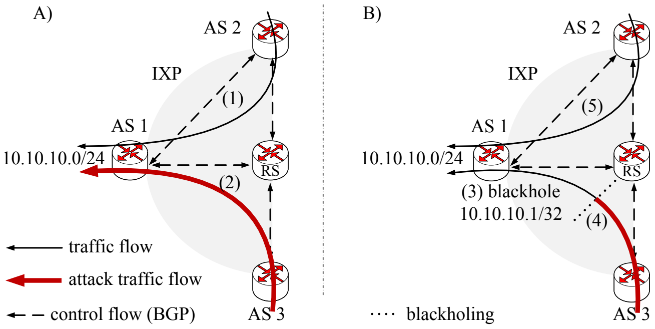

**Fig. 1.** DDoS attack at IXP member before/during blackholing.

the behalf of the prefix owner, e.g., through its upstream AS. Both scenarios render the attacked network unreachable for attackers and for everyone else.

**Explanatory Example Scenario:** Figure 1 depicts the traffic flow process at an IXP prior ($A$) and after ($B$) the activation of blackholing. The initial situation is that a member (AS1) receives traffic from its peers and while AS2 sends solely legitimate traffic (1), AS3 traffic contains significant amounts of DDoS traffic (2). Now AS1's IXP-connected router advertises the attacked prefix – usually a more specific – for blackholing towards the route server (3). This can be done either explicitly, i.e., using a BGP next hop with a predefined blackholing IP address, or implicitly, i.e., via a well-known BGP community. The community is then translated to the next hop blackholing IP address at the route server. All connected members receive the BGP update, learn the new BGP next hop address for the announced prefix, choose it as best path since it is more specific, and send their traffic to the blackholing IP.

The IXP handles this IP address and resolves it by means of the ARP into a predefined blackholing MAC address. All Ethernet frames with this destination MAC are discarded via ACL at the IXP layer-2 ingress switch interfaces (4). Note, this process is non-transparent for the traffic source, e.g., attacker. All other announced prefixes remain unaffected (5), but may do not suffer from congestions anymore. In cases where the DDoS traffic is mainly coming from a certain member's networks, the so-called *policy control* feature of route servers can be used to limit blackholing only to those ASes. In general, policy control allows the definition of white- and blacklists for BGP announcements by a well-defined set of BGP communities. These communities are interpreted by the route server.

**Blackholing Usage:** The implementation of blackholing at IXPs is beneficial because: i) route servers disentangle the configuration process for triggering blackholing. A single route update can address all members at once. ii) The large number of networks that meet at the IXP also increase the effectiveness. iii) Given the central position in the Internet, blackholing at IXPs allows the alleviation of the impact closer to the attack source. iv) It can protect the intermediate networks on the path through the Internet, but it is far enough from the source to be efficient.

However, while blackholing at IXPs shields member networks and the links from congestions, it cannot distinguish between legitimate and malicious traffic. All packets destined for the defined IP prefix are dropped and, thus, it is not reachable from all upstream networks on the data path.

Moreover, after detecting a massive DDoS attack, the operator must trigger blackholing. This is a manual process where the router configuration must be adjusted in order to announce via BGP an IP prefix under attack. Typically, a more specific IP prefix is announced to limit the impact on benign traffic to the minimum. The triggering AS is not necessarily the owner of the IP prefix. Thus, the announcing member must register this prefix in the IRR database to be accepted by the IXP.

## 3 Data Sources

In this paper, we rely on the following datasets from one of the largest European IXPs [6]. This IXP serves around 600 members and peaks to over 4 Tbit/s in 2015.

We used 5-minute interval snapshots from a publicly accessible looking glass at the IXP route servers to gather the BGP announcements for long-term control plane analysis. The announcements for blackholing can be discriminated by means of a well-defined next hop IP. Due to the sampling frequency, only announced prefixes that were active at these moments can be captured. Short-term new and withdrawn announcements are not caught. If a previously active prefix was absent in one measurement we considered it as a new announcement when it reappears. The data covers a 3-month period from December 2014 onward. From this dataset we identify 22,994 blackholing BGP announcements (after excluding measurement and looking glass outages, etc.)

To understand the impact of blackholing on the traffic flow, we rely on IPFIX data from the IXPs switching fabric for the same period. IPFIX at the IXP is configured to randomly capture 1 out of 10,000 packets on every member link. The IPFIX data contains the MAC and IP addresses, IP protocol identifier, TCP/UDP port numbers, and length of the captured packets. For statements about traffic volumes we extrapolate from the sampled flows.

In addition we use route server and IPFIX data for policy control verification and a case study from July 2015.

## 4 Blackholing: A Usage Analysis

In this section, we elaborate on how blackholing is used in the wild from a control plane perspective. For the remainder of this paper the term ”*announcement*” refers to BGP announcements that trigger blackholing. Additionally, all notations about IP prefixes refer to blackholed IP prefixes if not otherwise stated.

### 4.1 A Prefix View of Blackholing

The IXP’s route server accepts BGP advertised blackholes with a prefix length $n$, with $/32 \leq n \leq /8$. We find that only prefixes $\geq /18$ are announced by the IXP members. Figure 2(a) shows the distribution of unique announcements (y-axis in log-scale) per prefix length. The mode on the far right indicates that mainly $/32$ prefixes are blackholed, indeed more than $97\%$ of all announcements. Another mode is between $/24$ and

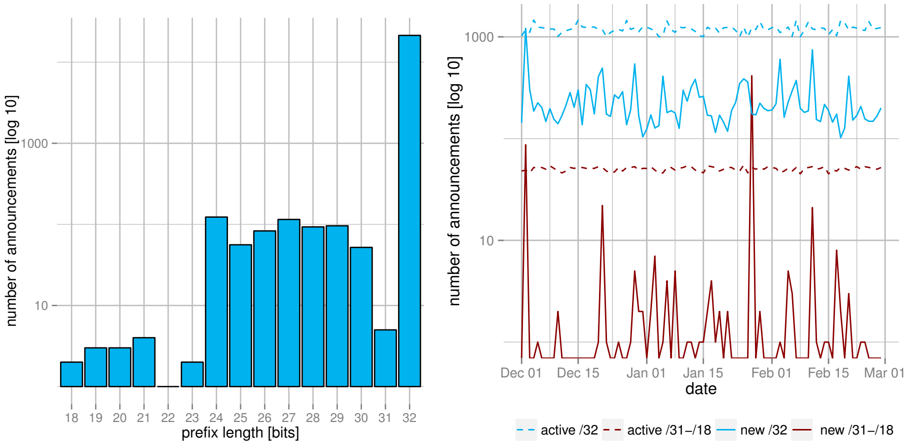

(a) Announcements by prefix length. (b) Avg. active and new announcements per day.

**Fig. 2.** Prefix views of blackholing.

$/30$, which accounts for $2.5\%$. Prefixes with the length of $\leq\ /23$ account for a very small fraction, namely $9$ announcements $(0.5\%)$. In summary, mostly host routes are used for blackholing.

Due to the employment of the policy control feature at the route server, prefixes are not necessarily announced to all peers. We randomly sampled the route server's RIB four times with a seven day interval. On average $25\%$ of all announcements carry a policy control community that limits its propagation.

To understand if the blackholing usage changed over time, Figure 2(b) shows the announcements per day, clustered by prefix length over a three-month period. We distinguish between new announcements per day and active on average per day. Unexpectedly, we find that the total number of active announcements is almost stable. In particular, the $/18 - /31$ prefix cluster contains eight $/24$ announcements that are active over the entire measured period. Unfortunatly, we did not get a response from the operator to fathom the intention for the long lasting announcements. For announcements between $/25 - /30$, we again see permanently announced prefixes that are active for a period of several weeks. Since the announcements with prefixes of the $/18 - /31$ are only short-lived, they do not impact the average number of active ones. In contrast, the most prevalent prefix class, the $/32$s, varies significantly. It ranges from an average minimum of $994$ to a maximum of $1,463$ and a mean of $1,195$ for all active announcements.

The number of new announcements per day differ notably compared to the averages. They show significant variations. From no activities over several days to large numbers of new announcements during one day for all but the $/32$. Indeed, focusing on the peak on January 27th, we see a total of $415$ newly announced unique prefixes. All prefixes are announced simultaneously by the same AS at 6 A.M. and last for about $10$ minutes. The number of new announcements for the $/32$ prefixes varies between $102$ and $1211$.

Next, we consider the durations and remove all announcements where we did not capture either the beginning or the end. We cluster these announcements $(100\%)$ by prefix length and show in Figure 3(a) the histogram of announcement durations in minutes using log-scale. The majority of long announced prefixes are $/32$s. Altogether,

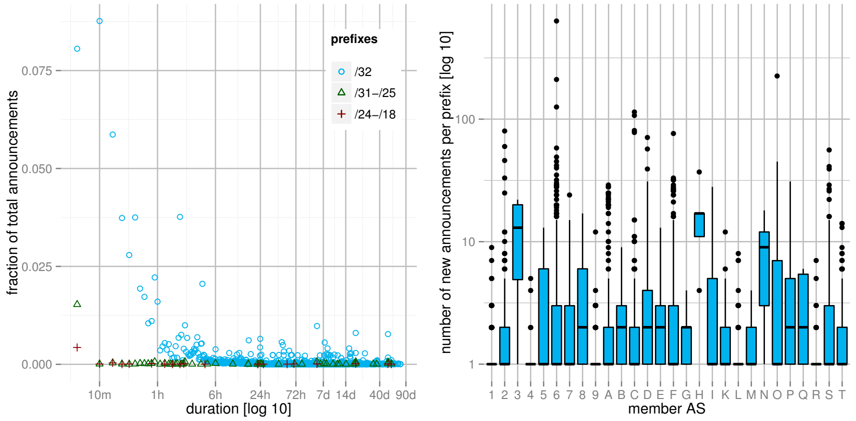

(a) Fraction of all announcements per minutes binned by prefix.

(b) Announcement frequency per prefix [log 10].

**Fig. 3.** Prefix and AS view of blackholing.

the largest fraction ($0.1$ of all prefixes) is announced for about five minutes. The two clusters with the less specific prefixes ($/18 - /24$ fraction of $0.004$ and $/25 - /31$ with $0.015$) have few announcements longer than five minutes. Their longest announcements last for $57.39$ days. The $/32$ prefixes are again more diverse. Notably, around $9\%$ of the announcements are active for about $10$ minutes while $38.5\%$ last longer than $240$ minutes. Interestingly, multiples of $1$ hour are more dominant. The longest duration we observe is $76.31$ days.

The operational background for such observation is that the members' monitoring capabilities for blackholing are limited. As soon as a prefix is announced for blackholing, the announcing AS is not aware of the amount of traffic that is dropped. Hence, some members turn blackholing off and on within a short period of time in order to check whether the DDoS attack is still on-going.

### 4.2 An AS View of Blackholing

To understand how the $75$ different ASes ($12\%$ of member ASes) use blackholing we take a closer look at the announcements from a member AS perspective. These ASes are categorized according to Peering DB: Network Service Providers (NSP) $50\%$, Cable/DSL/ISP $25\%$, and Content Providers $19\%$. The NSPs are overrepresented compared to IXP-wide $42\%$, while the latter two accord to those.

The ASes announced $7,864$ unique prefixes, of which $10\%$ were announced once and around $15\%$ between two and three times. $47$ ASes announced fewer than $50$ prefixes in total and were excluded to focus on the blackholing-heavy ASes. We then focus on the remaining $28$ ASes. The mean of all announcements for the same prefixes across all remaining ASes is $3.13$ with a median of only $1$. Figure 3(b) shows the number of announcements per unique prefix by AS in a boxplot. Overall, the median is almost always lower than $10$. Looking at the details we find that, despite the prevailing low announcing frequencies, there are also AS-wide high frequencies. Surprisingly, outliers spread from $10$ to $100$. The observed maximum number of a unique announced prefix is

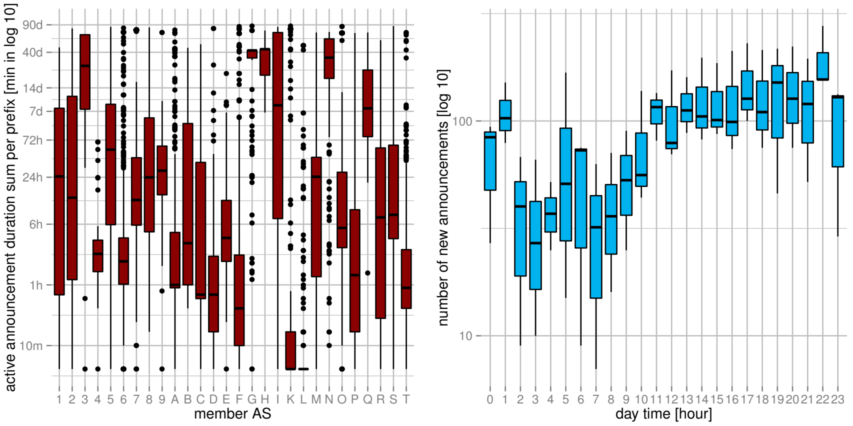

(a) # active duration per prefix [log 10]. (b) Selected ASes: announcements by time.

**Fig. 4.** AS views of blackholing.

623. This observation may provide further evidence for an operational procedure where blackholing is switched on and off many times for the very same IP prefix within a short time frame.

To check if the frequently announced prefixes have an impact over the total time span that they are active, we accumulate the duration for all unique prefixes of our selected ASes. Figure 4(a) shows these values at the y-axis with the same ASes as in Figure 3(b). We find that the majority of ASes announce prefixes for a duration longer than 1 hour but less than 7 days. Nevertheless, we observe some ASes that either announce their prefixes primarily over a short or long period of time.

Figure 3(b) and 4(a) expose no clear correlation between frequency and duration of announcements for the same AS. On the one hand some ASes announce the same prefix frequently for short times, and on the other hand some ASes announce blackholing only once for a short duration. However, also the contrary can be observed: frequent and once for a longer duration. This indicates that there is no common operational procedure for triggering blackholing. This is not surprising as DDoS attacks often differ from attack to attack. Thus, operators respond to each attack individually.

Nevertheless, we want to see if there are operational patterns: one would expect either more announcements during daytime working hours, or during peak hours. Neither is the case and the behavior differs by AS. For some, we see clear patterns, e.g., see Figure 4(b). Here, the number of new announcements is substantially smaller but the variance is higher at night.

## 5 Blackholing Impact on Traffic

To study the effectiveness of blackholing on the data plane, we correlate the actual traffic with the BGP announcements on the control plane. Figure 5(a) depicts the bit rates of the blackholed traffic at the IXP during our three month observation period. It shows per day the hourly maximum and hourly average in Mbit/s. The per day maximum varies from 50 to $1,000$ Mbit/s with a peak at $2,100$ Mbit/s. At first glance, this may

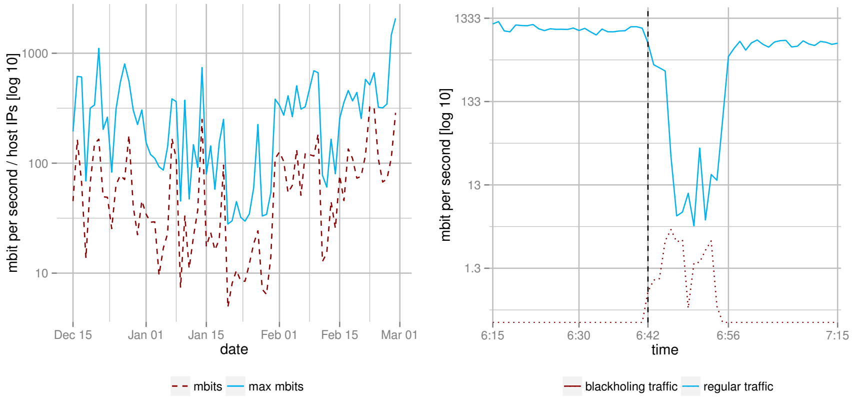

(a) Blackholed traffic for 3 month period.

(b) Traffic volumes for *Case Study I*.

**Fig. 5.** Traffic volumes over time.

seem small especially when compared to the average traffic rates at this IXP. However, keep in mind that this is the traffic that is discarded and should only effect short time-scale DDoS attacks. Moreover, as soon as the blackhole is in effect, this regulates traffic volume. For TCP, blackholing disables connections and for UDP the sender might notice the blackhole and therefore throttle the attack. Another reason is what we plot: the average across a full hour. Indeed, the daily averages (dashed line) are up to 10 times smaller than the daily maxima (solid line).

Given the large number of blackholing announcements, see Section 4, we next determine how many of these prefixes actually receive traffic. We find that in those hours with more than $1$ Mbit/s average blackhole traffic, a mean of 81 and a maximum of $871$ IPs receive traffic. Thus, we conclude that typically only a small number of IPs receive a substantial amount of blackholed traffic.

To assess the impact of blackholing we next examine the temporal correlation of blackhole announcements with traffic. We focus on two case studies: (Case Study I) an event that lasted a relatively short time period and (Case Study II) one lasting longer and involving a larger traffic volume.

**Case Study I:** Figure 5(b) is an example where the blackholed prefixes are in the range $/19 - /29$. This AS (AS $k$ in Figures 3(b) and 4(a)) announced $415$ prefixes for blackholing — all at the same time. Overall, the blackhole was active for roughly 10 minutes (dashed vertical line). Figure 5(b) shows the traffic volume as received by the AS for all these blackholed prefixes for 60 minutes, namely, $\sim$30 minutes before the first blackhole announcement, during the blackhole, and $\sim$20 minutes afterwards. In addition, we show the traffic for the same prefixes that is discarded, as well as the times when the blackhole announcements are made (vertical line)). Right after the blackhole is announced the traffic that the AS receives drops by roughly a factor of $100$. The blackholed traffic (dotted lines) is smaller than the "missing" traffic due to the reasons mentioned above. After the blackhole is deactivated the traffic volume rises to a level that is close to the previous one. The difference is roughly 300 Mbit/s. We also conclude that the objective of the blackhole was achieved as there were no further blackhole announcements for these prefixes by this AS.

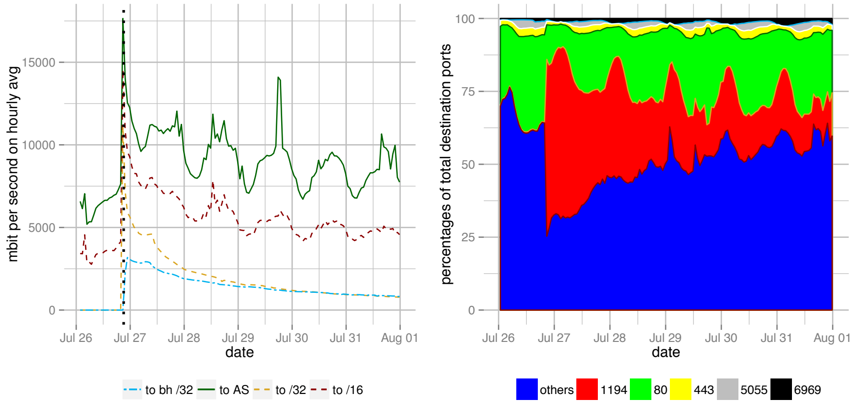

(a) Traffic over time for different subsets. (b) Distribution of dst ports over time.

**Fig. 6.** Traffic and port mix for *Case Study II*.

One could expect that the amount of traffic that is blackholed is the same as the reduction of the regular traffic. This is not the case and explanations are: i) Depending on the BGP router configuration, it cannot be guaranteed that a particular peer accepts more specific prefixes than $/24$. ii) The AS under attack may take other corrective actions besides blackholing at IXPs. For instance, blackholing at upstreams, moving traffic from peering to transit, or activating DDoS traffic filtering services (e.g., CloudFlare or Prolexic). iii) If the blackholed traffic contains a large fraction of TCP traffic and these TCP connections are broken by the blackhole, this can reduce the data traffic drastically.

**Case Study II:** Figure 6(a) shows the same data but for another AS for a six day period from 26th of July 2015 onwards. We picked this case study as it involved substantial amounts of traffic, an interesting application (port) mix, and a single IP — a $/32$ prefix ($to\ bh/32$ and $to\ /32$ legend in Fig. 6(a)). In addition, the plot contains the traffic to the covering $/16$ prefix ($to\ /16$) as well as the overall traffic for this AS ($to\ AS$).

Notice the radical increase in traffic shortly before the blackhole announcement. The traffic to the AS spikes from roughly 6 Gbit/s to 17.6 for an AS with an aggregated port capacity of 20 Gbit/s. The plot highlights that the root cause is in the non-blackholed $/16$ and in particular in the $/32$ that is then blackholed. The blackhole announcement for this $/32$ is highly effective as the regular traffic for the IP as well as for the $/16$ and the AS drops significantly immediately after the announcement. Note, that we can still see traffic to the $/32$ because not necessarily all peers accept more specific announcements than $/24$. We also notice the peak in the blackholed traffic for the $/32$, which increases to 3.2 Gbit/s. Thus, the blackhole reduces the traffic to the AS and the prefixes by about one third. Over the next days the traffic to the IP gradually decreases while the non-blackholed traffic to the $/16$ and $/32$ shows clear daily patterns. We captured several updates for the blackholed prefix. The blackhole is not revoked, but just updated with different communities which are not honored by the IXP's route server.

To understand why the blackhole is effective, we plot in Figure 6(b) the relative transport TCP/UDP port distribution for the traffic to the AS. Over the whole period port 80 (http), 1194 (OpenVPN), 443 (https), 5055, 6969 (BitTorrent) are the most prominent ports. Accordingly, the plot is a stacked barplot with these ports and other ports at the bottom.

Initially, the traffic share of http is $\sim 30\%$. But with the blackhole trigger event the OpenVPN traffic drastically increases. Indeed, it constitutes about 50% of all traffic to this AS. As time passes and the blackhole takes effect the port mix slowly converges to the initial distribution. The dominance of OpenVPN is also reflected in the blackholed traffic for the blackholed IP. $99.9\%$ of the traffic is UDP and involves port 1194. Thus, this change of ports is also reflected in the distribution of transport protocols.

We find that blackholing is effective in numerous situations. However, the observed volumes of traffic depend on numerous factors, e.g., prefix length, announcement duration, general traffic utilization, attack pattern, and/or policy control settings. We also highlight that the traffic mix can vary significantly between non-blackholed and blackholed traffic.

## 6 Related Work

While this work focuses on blackholing, a network-based reactive measure to diminish massive DDoS attacks in the core of the Internet, this section summarizes other reactive DDoS defense mechanisms and highlights other measurement studies.

Source-based defense techniques are deployed near the source of an attack and aim to impede service of intermediate and destination networks. Common mechanisms are IP source address filtering and heuristics on ingress/egress traffic flows [1, 10, 22].

Alternatively, destination-based DDoS mitigation attempts to combat attacks near the victim-end of the Internet. Common places for their deployment are edge routers or access routers of the destined AS. Proposed mechanisms include adaptive rate limiting [16, 21], network reconfiguration [3, 4, 35], and traceback [2, 34].Additionally, a multitude of filtering techniques such as time-based [14], history-based [29], and hop count-based [17] have been introduced.

Whereas source-based DDoS defense often suffers from its limited scope and the lack of a representative fraction of the attack traffic to be efficient, destination-based approaches come in too late on the path through the Internet. Thus, they jeopardize the destination AS or even intermediate networks. Network-based approaches seek to overcome these drawbacks and are deployed inside intermediate networks. They mainly incorporate distributed or trust-based detection and already presented reconfiguration or filtering mechanisms, e.g., [11, 24, 28].

Despite the large body of available approaches, effective reactive techniques that are deployed in practice are rare. Thus, there is a demand for defense techniques which are efficient, easy and quick to apply, and which ensure the continuing availability of the services, system, or network. However, none of the mentioned taxonomies for DDoS defense techniques [18, 23, 37] takes blackholing into consideration.

Although blackholing has not been examined to date, other recent measurement studies focus on attack amplification potential [7, 20, 32, 33, 38] and on progress towards diminishing their impact [19].

## 7 Summary and Future Work

In this paper, we perform a first study on the usage of blackholing at an IXP in the wild. We find that not only is blackholing frequently used with about $23,000$ announced blackholes over our measured period of 3 months, but also that they have a considerable prefix size — up to $/18$. While short-lived blackholes are prevalent, we also spot others that lasted for months. Moreover, we observe an apparently stable number of active blackholing announcements (about $1200$).

The frequent usage of blackholing on the control plane correlates with significant amounts of blackholed traffic on the data plane. Using two case studies we show that blackholing successfully reduces the amount of traffic. This emphasises that blackholing at IXPs can be a very useful tool to diminish massive DDoS attacks. Indeed, our analysis of the application (port) mix of one of the blackhole incidents indicates that blackholing is successful in reducing unusual OpenVPN traffic, likely a DDoS attack.

In general, IXPs are great locations for countering DDoS attacks via blackholing, as the IXP infrastructure is a multiplication factor. Still, blackholing is a relatively new feature and there is room for increased efficacy, e.g., effective monitoring and reporting, partially retracting blackholing, as well as common operation practices at the ASes (acceptance of more specific than $/24$ prefixes), and transitive blackholing. Moreover, the blackholed data can be used to better mitigate attacks in the Internet.

**Acknowledgments:** We thank all our colleagues for their feedback, and the reviewers for their suggestions. This work is supported by European Unions Horizon 2020 research and innovation programme under the ENDEAVOUR project (grant agreement 644960) and by the German Federal Ministry of Education and Research (BMBF Grant 01IS14009D BDSec).

## References

1. S. Abdelsayed, D. Glimsholt, C. Leckie, S. Ryan, and S. Shami. An Efficient Filter for Denial-of-Service Bandwidth Attacks. In *GOLBECOM*, 2003.
2. M. Adler. Trade-Offs in Probabilistic Packet Marking for IP Traceback. *JACM*, 2005.
3. S. Agarwal, T. Dawson, and C. Tryfonas. DDoS Mitigation via Regional Cleaning Centers. Technical report, Sprint ATL Research Report, 2003.
4. D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris. *Resilient Overlay Networks*. ACM SOSP, 2001.
5. T. Battles, D. McPherson, and C. Morrow. Customer-Triggered Real-Time Blackholes. NANOG 30, 2004.
6. N. Chatzis, G. Smaragdakis, J. Böttger, T. Krenc, and A. Feldmann. On the Benefits of Using a Large IXP as an Internet Vantage Point. In *ACM IMC*, 2013.
7. J. Czyz, M. Kallitsis, M. Gharaibeh, C. Papadopoulos, M. Bailey, and M. Karir. Taming the 800 Pound Gorilla: The Rise and Decline of NTP DDoS Attacks. In *ACM IMC*, 2014.
8. DE-CIX. DE-CIX Blackholing Support. www.de-cix.net/products-services/de-cix-frankfurt/blackholing/.
9. Deutsche Telekom. AS3320 BGP Communities, August 2005. www.onesc.net/communities/as3320/AS3320_BGP_Communities_v1.1.pdf.
10. T. M. Gil and M. Poletto. MULTOPS: A Data-Structure for Bandwidth Attack Detection. In *USENIX Security Symposium*, 2001.
11. J. M. Gonzalez, M. Anwar, and J. Joshi. A Trust-based Approach Against IP-Spoofing Attacks. In *IEEE PST*, 2011.

12. B. R. Greene. Remote Triggering Black Hole Filtering. *Cisco Systems*, 2002.
13. B. R. Greene and P. Smith. *Cisco ISP Essentials*. Cisco Press, 2002.
14. Y. Hu, H. Choi, and H-A Choi. Packet Filtering to Defend Flooding-Based DDoS Attacks. In *Advances in Wired and Wireless Communication*, 2004.
15. Hurricane Electric. Customer Blackhole Community, 2006. `www.he.net/adm/blackhole.html`.
16. J. Ioannidis and S. M. Bellovin. Implementing Pushback: Router-Based Defense Against DDoS Attacks. *Columbia University Academic Commons*, 2002.
17. C. Jin, H. Wang, and K. G. Shin. Hop-Count Filtering: An Effective Defense Against Spoofed DDoS Traffic. In *ACM CCS*, 2003.
18. A. Keshariya and N. Foukia. DDoS Defense Mechanisms: A New Taxonomy. In *DPM*. '10.
19. M. Kührer, T. Hupperich, C. Rossow, and T. Holz. Exit from Hell? Reducing the Impact of Amplification DDoS Attacks. In *USENIX Security Symposium*, 2014.
20. D. C. MacFarland, C. A. Shue, and A. J. Kalafut. Characterizing Optimal DNS Amplification Attacks and Effective Mitigation. In *PAM*, 2015.
21. R. Mahajan, S. M. Bellovin, S. Floyd, J. Ioannidis, V. Paxson, and S. Shenker. Controlling High Bandwidth Aggregates in the Network. *ACM SIGCOMM CCR*, 2002.
22. J. Mirkovic, G. Prier, and P. Reiher. Source-end DDoS Defense. In *IEEE NCA*, 2003.
23. J. Mirkovic and P. Reiher. A Taxonomy of DDoS Attack and DDoS Defense Mechanisms. *ACM SIGCOMM CCR*, 2004.
24. A. T. Mizrak, S. Savage, and K. Marzullo. Detecting Compromised Routers via Packet Forwarding Behavior. *IEEE Network*, 2008.
25. MSK-IX. Protection against DDoS-attacks by blackholing. `www.msk-ix.ru/eng/routeserver.html\#blackhole`.
26. NETIX. Blackholing. `www.netix.net/services/14/NetIX-Blackholing`.
27. NTT Communications. Terms and conditions for use of global IP network services, August 2007. `http://www.ntt.net/english/library/pdf/terms.pdf`.
28. K. Park and H. Lee. On the Effectiveness of Route-Based Packet Filtering for Distributed DoS Attack Prevention in Power-Law Internets. In *ACM SIGCOMM CCR*, 2001.
29. T. Peng, C. Leckie, and K. Ramamohanarao. Protection from Distributed Denial of Service Attacks Using History-Based IP Filtering. In *IEEE ICC*, 2003.
30. M. Prince. The DDoS That Almost Broke the Internet, March 2013. `www.blog.cloudflare.com/the-ddos-that-almost-broke-the-internet/`.
31. P. Richter, G. Smaragdakis, A. Feldmann, N. Chatzis, J. Boettger, and W. Willinger. Peering at Peerings: On the Role of IXP Route Servers. In *ACM IMC*, 2014.
32. C. Rossow. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In *NDSS*, 2014.
33. F. Ryba, M. Orlinski, M. Wählisch, C. Rossow, and T. Schmidt. Amplification and DRDoS Attack Defense - A Survey and New Perspectives. *arXiv preprint arXiv:1505.07892*, 2015.
34. S. Savage, D. Wetherall, A. Karlin, and T. Anderson. Network Support for IP Traceback. *Networking, IEEE/ACM Transactions on*, 9(3):226–237, 2001.
35. E. Shi, I. Stoica, D. G. Andersen, and A. Perrig. OverDoSe: A Generic DDoS Protection Service Using an Overlay Network. *Computer Science Department*, 2006.
36. Sipgate. The Sipgate DDoS Story, October 2014. `https://medium.com/@sipgate/ddos-attacke-auf-sipgate-a7d18bf08c03`.
37. S. M. Specht and R. B. Lee. Distributed Denial of Service: Taxonomies of Attacks, Tools, and Countermeasures. In *ISCA PDCS*, pages 543–550, 2004.
38. R. van Rijswijk-Deij, A. Sperotto, and A. Pras. DNSSEC and Its Potential for DDoS Attacks: A Comprehensive Measurement Study. In *ACM IMC*, 2014.
39. S. T. Zargar, J. Joshi, and D. Tipper. A Survey of Defense Mechanisms Against Distributed Denial of Service (DDoS) Flooding Attacks. *Com. Surveys & Tutorials, IEEE*, 2013.