Anycast Agility: Adaptive Routing to Manage DDoS

ASM Rizvi[1,2], Joao Ceron[3], Leandro Bertholdo[4], John Heidemann[1,2]

[1]USC/Information Sciences Institute, [2]USC/CS Dept., [3]SIDN Labs, [4]University of Twente

arXiv:2006.14058v1 [cs.NI] 24 Jun 2020

ABSTRACT

IP Anycast is used for services such as DNS and Content Delivery Networks to provide the capacity to handle Distributed Denial-of-Service (DDoS) attacks. During a DDoS attack service operators may wish to redistribute traffic between anycast sites to take advantage of sites with unused or greater capacity. Depending on site traffic and attack size, operators may instead choose to concentrate attackers in a few sites to preserve operation in others. Previously service operators have taken these actions during attacks, but how to do so has not been described publicly. This paper meets that need, describing methods to use BGP to shift traffic when under DDoS that can build a "response playbook". Operators can use this playbook, with our new method to estimate attack size, to respond to attacks. We also explore constraints on responses seen in an anycast deployment.

KEYWORDS

Anycast, DNS, DDoS, Traffic estimation

1 INTRODUCTION

Anycast routing is used by services like DNS or CDN where multiple sites announce the same prefix from geographically distributed locations. Defined in 1993 [43] anycast was first deployed by DNS roots in the early-2000s [4, 26, 56], and today it is widely used by DNS providers and Content Delivery Networks [15, 16, 21, 23, 65].

In IP anycast, BGP routes each network to a particular anycast site, dividing the world into *catchments*. BGP usually assigns user's networks to nearby anycast sites, providing generally good latency [54]. Anycast also helps during Distributed-Denial-of-Services (DDoS) attacks, since many sites provide a greater aggregate capacity than a single very large site. Each anycast catchment is independent, so if attackers overwhelm one site, others may still operate within capacity.

DDoS attacks are getting larger and more common. Different root servers and anycast services frequently report DDoS events [17, 36, 41, 42]. Different automated tools make it easier to make attacks [66], and by using DDoS-as-a-Service, unsophisticated users can now make attacks for only US$10 [33, 58]. DDoS intensity is still growing as we observe the recent memcached attack exceeds 1.35 Tb/s in size [40], and Internet-of-Things devices provide millions of vulnerable devices for botnets [30].

Operators depend on anycast during DDoS attacks, to provide capacity to handle the attack and to isolate attackers in catchments. Service operators would like to adapt to an ongoing attack, perhaps shifting load from overloaded sites to other sites with excess capacity. Prior studies of DDoS events have shown that operators take these actions but suggested that the best action to take depends on attack size and location compared to anycast site capacity [38]. While prior work suggested countermeasures and we know that operators use routing during attacks, there has been only limited evaluation of how routing choices change traffic [5, 24, 31, 46], and no specific, public guidance on how to use routing during an attack.

The goal of this paper to provide guidance to mitigate a DDoS attack using BGP routing to balance traffic across an anycast service.

Our first contribution is to define the mechanisms we need to balance traffic during DDoS. Defense requires two steps: first, one must *estimate traffic to each site* when it is under attack (§3.3). Sites may have different capacities and receive attacks at different intensity—the defender must have a rough estimate of these factors to make an informed response. Second, one must *employ BGP policy routing mechanisms* to shift traffic (§3.5). While general mechanisms for BGP policy are well known, their availability and details vary across Internet Exchange Points (IXPs) and ISPs. We document what approaches are typical and how they apply to DDoS defense.

The second contribution of our work is to *show how effective these tools are in practice* (§6). While the concept of specific BGP mechanisms is clear (for example, AS-prepending discourages traffic), understanding *how much traffic shifts* is a complex function. We evaluate how the AS policy, network topology, and the anycast deployment can affect the traffic shift (§7). We show that the *granularity of control* can be surprisingly limited, where sometimes large changes showing little or no effects, while some small changes can shift nearly all traffic to or from a site. These effects are strongly influenced by anycast deployment and ASes with whom they peer at each site. While we cannot show all possible deployments, we systematically explore different configurations to show typical ranges of response.

Our final contribution is to demonstrate that these approaches can defend against attacks successfully. We replay real-world attacks in a testbed and manipulate routing to defend (§8).

Our work uses publicly available datasets. Data from our experiments is available at no charge. We provide data from the experiments on the testbeds and attack traffic evaluated. Because our data concerns services and not individuals, we see no privacy concerns.

2 RELATED WORK

Anycast routing has been studied for a long time from the perspective of routing, performance, and DDoS-prevention.

BGP to steer traffic: Prior works showed BGP is effective to steer traffic to control link load [8, 24, 47]. But we need to consider anycast topology and upstream providers to make path prepending more effective [5]. Others proposed to manipulate BGP based on packet loss, latency and jitter [39, 46].

Chand *et al.* [13] suggested the concomitant use with BGP Communities [10, 12, 61] to manipulate traffic. More recently BGP communities have been used to create blackholes in IXPs and ISPs [19, 25].

Anycast performance: Most anycast research focused in efficient delivery and stability [11, 34, 35, 51, 64]. Later studies explicitly investigate the proximity of the clients [5, 11, 34]. Some studies try to improve anycast through topology changes [37, 54]. Anycasted services for DDoS is already used in commercial solutions *e.g.,* Amazon [55], Akamai[62] and AT&T [59]. However, none of them address how to use routing manipulations as a DDoS defense mechanism or make it clear.

Anycast catchment control as a DDoS mitigation tool: The idea of handling DDoS attacks by absorbing or shifting load across anycast sites has been long used by operators, but to our knowledge it was only formalized in 2016 by Moura *et al.* [38]. Kuipers *et al.* [31] extended refined Moura's work, defining the traffic shifting approaches that we review in §3.4 and explore through experiment.

To the best of our knowledge, our paper is the first to combine a full DDoS defense strategy. Our work goes from a new technique to estimate attack load to evaluate the probable routing decision, and their collateral effects over other sites. We prove our findings in a real attack and measure our results in two different anycast testbeds.

3 MECHANISMS TO DEFEND AGAINST DDOS

In this section, we describe an intervention against DDoS by estimating traffic sent to each site, and using BGP tools and defense strategies.

3.1 Overview and Decision Support

DDoS defense must include planning and response. Before the attack, the operator should map their anycast system (§3.2). Then during the attack, the operator detects the attack, they (a) deploy any filters they have, if they can, then (b) assesses the traffic sent to each site under attack (§3.3), (c) selects a strategy to shift traffic (§3.4), then (d) uses BGP to shift traffic (§3.5). The operator will then repeat this cycle to refine their response.

Our focus is on service operators that operate without using a third party scrubbing service, and that have a fixed amount of pre-deployed capacity. For operators who employ third-party scrubbing services, then the scrubbing service is the filter in step (a). Typically the scrubbing service takes over announcing the prefix under attack, thus the operators would use our approach. However, the *scrubbing service* must then handle the attack, and *they* may employ steps (b) through (d) to handle the attack on their sites.

We assume a network service operator monitors their service and is quickly alerted when a DDoS attack begins. The alert can be generated when a resource is nearly exhausted at a particular site. A site may consider several resources—ingress or egress network bandwidth, CPU or memory usage— since each of these resources may become a bottleneck under different attacks [50].

We expect the operators to continue to monitor the service during the attack and refine their response. Continuous motioning is necessary to guard against attacks that change over time (polymorphic attacks). It is also important because other parties may respond to the attack as well, for example, an upstream ISP may shift routing because of link oversubscription, resulting in changes to anycast catchments. Finally, monitoring allows the defender to refine their defense and recover from imperfect estimates of site traffic or capacity.

3.2 Mapping Anycast

During the attack we plan to manipulate routing to respond to traffic. But making *informed* choices about routing changes requires that we have some idea of what effect a change will have. We therefore *map* our anycast catchments and how they change in response to routing changes ahead of any attack.

When mapping catchments, we determine which networks are associated with which anycast sites. We use Verfploeter [18] to find the catchments of the anycast clients using *active probing*. Verfploeter uses a target hitlist of IPs, and probe them with ICMP ECHO request. ICMP REPLY from each target network ends up in its catchment site making the network to anycast site mapping. An alternative is to remember which customers are seen at each site over time

or to observe from distributed vantage points such as RIPE Atlas [3, 60].

Mapping should consider not only the current catchments but also *potential* shifts we might make during the attack. This full mapping is easy to do with Verfploeter, which can be continuously running in an adjacent BGP prefix to map the possible shifts. This mapping process is important to anticipate how traffic may be shifted. We will show later that BGP control is limited by the granularity of routing policy (§6) and by the deployment of the anycast sites (§7).

The actual attack traffic may distort anycast—if the attackers are concentrated in a few networks, routing changes may not spread them out. Even then, mapping helps anticipate how legitimate traffic will shift.

A challenge in pre-computed maps with routing alternatives is that routing is influenced by all ASes. Thus, the maps may shift over time due to changes in the routing policies of other ASes. Fortunately, prior work shows that anycast catchments are relatively slow to change [64].

3.3 Estimating a Site's Offered Traffic

Estimating the offered load to each site is an important first step in DDoS defense to allow us to select our defense strategy (spread traffic or absorb, as described next in §3.4). (By "offered load", we mean all the traffic sent to the site, before any loss due to DDoS-driven congestion. Ideally the site would handle all this traffic.)

Here we describe *how* we estimate site traffic. We later evaluate the accuracy of this approach (§5) and use it to defend (§8).

Challenge and idea: Offered load is the combination of legitimate traffic and, during an attack, attack traffic that is sent to a site. The main consequence of a DDoS attack is the exhausted resources, and during an attack, the server and its access networks are overwhelmed. As a result, direct measurements at the server detect only *received* traffic. During an attack, received traffic is constrained by the access link and some offered traffic will be lost one or more hops upstream, before we can observe it. Our insight is that we can *directly* infer loss from *examination of end-to-end, known good traffic that is received*, and from loss we can estimate offered load to site.

Approach: We estimate site offered load by measuring the fraction of known good traffic that arrives at the service. We next describe each of these sub-problems.

We want to observe loss of legitimate traffic. Unfortunately, there is no general way to determine the current rate of legitimate traffic—traffic rates constantly change, sometimes unpredictably. Moreover, sophisticated attackers make attack traffic look just like good traffic, making the traffic rate impossible to measure when it is most needed.

We therefore use *subset of known good traffic* to represent all legitimate traffic. For DNS, RIPE Atlas provides a regular source of known-good traffic, sent from many places, with out-of-band reporting. We assume that most commercial services (in addition to DNS) have similar kinds of regular monitoring traffic.

We want offered load, or $T_{offered}$. We know the observed traffic rate $T_{observed}$—it is the access link bitrate, or it can be measured at the access link. We know that $\alpha \cdot T_{offered} = T_{observed}$, where α is the accept fraction (the traffic that is not dropped).

To determine α, we observe that known good traffic has the same loss on incoming links as does other good traffic and attack traffic. We know that RIPE Atlas sends measurement traffic at a known, constant rate T_{known}, so $\alpha \cdot T_{known,offered} = T_{known,observed}$. Solving for α and substituting back gives us: $\hat{T}_{offered} = \frac{T_{observed} \cdot T_{known,offered}}{T_{known,observed}}$.

3.4 Defense Strategies

With knowledge of the offered load the defender can select a defense strategy. A first question is: is the attack larger than the overall capacity or smaller?

For very large attacks, the defender's goal is to preserve successful service at some sites, while allowing other sites to operate in degraded mode as *absorbers* [38]. The defender may also choose to shift traffic away from some degraded sites to ease their pain. Unloading the overloaded sites is recognized as *breakwaters* [31].

For moderate-size attacks, the defender should try to serve all traffic, shifting traffic away from overloaded sites to less busy sites. In heterogeneous anycast networks, where some sites have more capacity than others, the defense approach can be different. In these cases, larger, "super"-sites can attract traffic from smaller sites. For moderate-size attacks, it may even be best for smaller sites to shut down if the super-sites can handle the traffic. We describe the BGP *mechanisms* to shift traffic in the next section.

3.5 BGP Tools to Shift Traffic

BGP is the standard protocol used to exchange routing information between Autonomous Systems (ASes) on the Internet [48]. While organizations with their own wide-area networks can use SDN or other techniques for traffic engineering (TE) on their internal WAN [27, 52], BGP is the only tool for TE between organizations. As a result, BGP has mechanisms and conventions to manage routing policy. We use two BGP mechanisms in the paper: AS-Path prepending and BGP communities.

AS-Path Prepending is a way to decrease the preference for a routing path. BGP's AS Path is the list of ASes back to

the route originator. The AS Paths' primary purposes it to prevent routing loops, and it also serves as a rough estimate for distance, with BGP routes with shorter AS Paths. By artificially inserting extra ASes into the AS Path, the route originator can dereference one site in favor of others. Path prepending is known to be a coarse routing technique for traffic engineering. We measure how fine the control AS path prepending provides to anycast in §6.1.

We define *Negative Prepending* as the use of AS-Path prepending to draw traffic towards a site. Prepending can only increase path lengths, but an anycast operator in control of all anycast sites can prepend at all sites except one, in effect giving that site a shorter AS path (relative to the other sites) than it had before. "Negative prepending by one at site S" is, therefore, shorthand for prepending by one at all sites other than S.

BGP Communities (or community strings) label specific BGP routes with 32 or 64 bits of information. How this information is interpreted is up to the ASes. While not officially standardized, a number of conventions exist where part of the information identifies an AS and the other part a policy such as blackholing, prepend, or set local-preference. Community strings are widely supported to allow ISPs to delegate some control over routing policy to their customers [1, 61].

4 EVALUATION METHODOLOGY

The next sections evaluate our defense. They use a common methodology for evaluation: we deploy an anycast testbed, modify routing to influence traffic, then map the catchments through active probing with Verfploeter. This evaluation method parallels how routing would be mapped in an operational system (§3.2), but there it would use a test prefix running in parallel with the operational service.

4.1 Anycast Testbeds

We evaluate our ideas on two existing anycast testbeds, since they support controlled experiments without hard to real-world services. We use two independently developed testbeds to show the generality of our results: `Peering` [53]and `Tangled` [18]. Table 1 summarizes information about each testing with its own set of geographically distributed sites, and lists the of site's location (`Peering` supports more sites but we used 8 sites). Some sites are connected to one or more transit providers, while others are connected to IXPs. `Peering` mainly has academic peers while `Tangled` has more commercial providers.

Both sites support BGP-based traffic engineering. Each testbed support prepending, although in `Peering`, we are limited to at most 3 prepends. Support for community strings varies by site and testbed and its detailed in §6.1.

Testbed	Used Sites	Total
`Peering`	Amsterdam*†(AMS), Boston* (BOS), Belo Horizonte*†(CNF), Seattle* (SEA) Athens* (ATH), Atlanta* (ATL), Salt Lake City* (SLC), Wisconsin* (MSN)	8
`Tangled`	Miami (MIA)*, London (LHR)*, Sydney (SYD)*, Paris (CDG)*, Los Angeles (LAX)*, Enschede (ENS)*, Washington (IAD)*, Porto Alegre (POA)*†	8

Table 1: Testbed and respective sites used in our experiments. Transit providers (*) and IXP (†).

4.2 Measuring Routing Changes

To measure a BGP policy effect we first change the route announcement at the site, give some time to propagate, confirm that it is accepted, and finally start the anycast measurement.

Methods to make a routing change differ by testbed. `Peering` provides an API so that testbed software can validate changes. `Tangled` allows us direct access to BGP at each site, so we used *exabgpcli* [20] to apply the desired configuration on sites.

Propagation time: After a change, we allow for *BGP route propagation*. We know that routes can be inconsistent (resulting in loops or blackholes) while routing is in flux [32, 57, 63]. Based on these prior studies we allow 15 minutes for routing to "settle".

Propagation of BGP policies: Policy constraints could limit the acceptance of announced routes, but in practice these limits do not affect our traffic engineering. Best practices for networks at the edge filter AS-paths longer than 10 hops, and ASes in the middle often accept up to 50, both more prepends than we need. In addition, even though we announce routes at IXPs, sometimes they are accepted at only some peers. We confirm the propagation of our routes using RIPE RIS [49]. It provides routing observations from multiple global locations. Using RIPE RIS, we confirm that configurations in our experiments are never blocked due to route filtering.

5 EVALUATING ESTIMATION OF OFFERED LOAD

We next evaluate our approach to estimate offered load. We first consider three real-world events, then build a mathematical model predicting estimation accuracy and validate it in controlled testbed experiments.

5.1 Case Studies

We first demonstrate offered load estimation on real-world DDoS events. We select a short-lived attack event in `B-root` from 2019-09-07 where multiple `B-root` sites were affected and had different levels of packet loss. We also examine two

Attack Date	Dur.	How	Site	Acc. Frac.	Obs. …(Mq/s)…	Est.	Rpt.
		UDP	LAX	36%	1.0	2.7	
2019-09-07	1m	vol.	MIA	80%	1.1	1.4	2.7
			ARI	12%	0.1	1.0	(tot.)
2016-06-25	3h	SYN-fl.	LAX	4%	0.1	2.5	10
2015-11-30	3h	UDP	LAX	2%	0.4	18.0	5.0

Table 2: Different attacks at B-root (with anonymized site codes), comparing observed, estimated, and externally reported traffic rates.

large DDoS events from 2015-11-30 and 2016-06-25 to look at the worst-case attacks, when B-root exhibited significant upstream loss.

In Table 2 we show each attack and its type (UDP volumetric attacks are DNS flooding), with observed traffic rates and known-good loss rate. The last two columns show our estimate and external reports. Each estimate uses §3.3, with a 5-minute window from RIPE (about 1200 observations).

In 2019-09-07 event we estimate each site separately. Each site showed a different loss rate, perhaps reflecting attacker location and site catchment size. Our estimation says S-1 saw more than half of all traffic which is roughly proportional to the catchment size at that time. There is no published ground truth, but M-root provided us their observed traffic rates, binned to 5 minutes, from which we projected a 2.7 Mq/s size. Because the event was shorter than M-root's reporting period and our estimation period, we suspect M-root binning provides an underestimate. This uncertainty makes it difficult to make strong statements about whether our estimate or M-root's report is more accurate. However, our direct observations and estimate bracket M-root's report, suggesting the attack was larger than we could directly observe.

In the two large events in 2016 and 2015, B-root had only one site and observed over 95% loss. For the 2016 event, we see only 0.1 Mq/s and estimate 2.5 Mq/s, much lower than the reported 10 Mq/s [42]. We believe this underestimate is due in part to a network bug at B-root where, in addition to attack traffic, the ingress network was flow-control-limited well below capacity. However, the estimate is much closer than direct observation.

Finally, for 2015, loss was at 98% (2% access), resulting in an overestimate of 18 Mq/s, three times reported values. At such high loss rates, estimation is very sensitive to noise.

Although traffic estimation is not perfect, in each case it provides much better information than direct observation. For 2019, it suggests traffic engineering would be warranted.

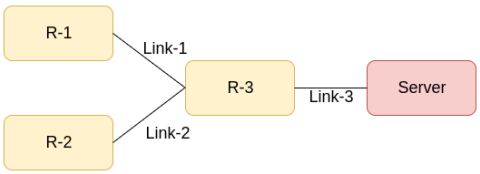

Figure 1: Topology with two upstream providers.

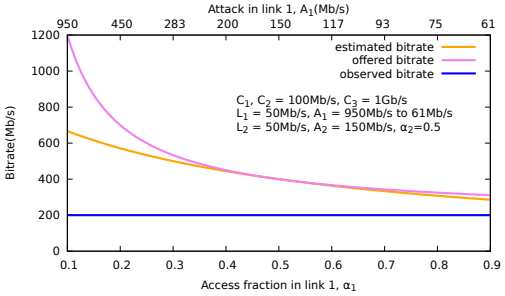

Figure 2: Modeling site traffic estimation.

5.2 Modeling Traffic Estimation

We next consider a simple mathematical model of loss rate to explore estimation accuracy when attacks are evenly distributed or concentrated on certain networks.

We consider a simple topology (Figure 1) where two access links (Link-1 and Link-2 with capacity C_1 and C_2) enter the service router (R3) to the server through Link-3 (at C_3). For simplicity, we assume $C_1 + C_2 < C_3$, so the internal network is never a bottleneck. Links-1 and -2 have legitimate traffic L_i. During an attack, attack traffic A_i causes total traffic $T_i = L_i + A_i$, resulting in loss on the links and a traffic accept fraction α_i ($\alpha_i = C_i/T_i$). We estimate offered load (T_i) based on loss of known-good traffic as §3.3.

We have explored many scenarios to understand when our estimator over- and under-estimates. In Figure 2, we consider a simple scenario to explain the case studies (§5.1). In this case, normal traffic on both links is the same, and when attacked, link 2 observes a fixed 50% loss ($\alpha_2 = 0.5$), and we vary attack strength on link 1 to see how the accuracy of estimation changes. We see that when both links have the same attack strength, our estimation is very accurate (the two curved lines are the same at $\alpha_1 = 0.5$). Our estimation is only slightly off when link 1 sees little or no attack (as α_1 approaches 1). However, as offered load on link 1 grows, more and more legitimate traffic is lost and so our underestimation grows (the estimated middle line does not track the true top line).

Examination of other scenarios (omitted for space) considered unequal access links and the potential for congestion on link-3. We find there are cases where we see slight over-

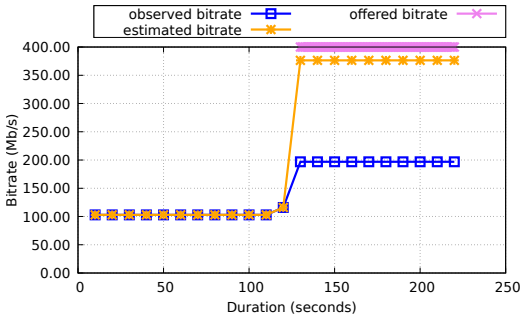

(a) Almost accurate estimation with overall 52% access fraction.

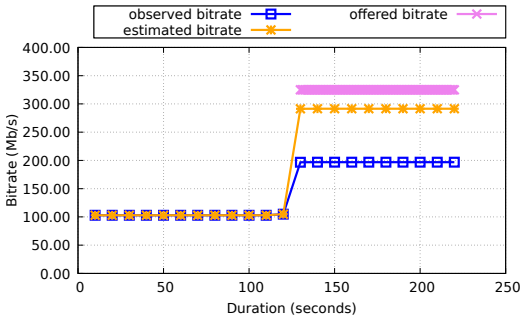

(b) Under-estimation with overall 68% access fraction.

Figure 3: DETERLab experiment to test the conditions for accurate and under estimation.

or under-estimation, but large estimation error occurs only as legitimate traffic drops to zero.

5.3 Model Validation

We next validate our model with experiments in a testbed. DETERLab is a configurable testbed that enable isolation in a controlled network [6].

We follow our modeled topology in Figure 1. Here we use unequal legitimate traffic, with L_1 of 80 Mb/s and L_2 of 20 Mb/s, so changes to attack traffic on L_1 have greater impact on our estimate. When loss is the same on both links (50% loss with A_1 = 120 Mb/s and A_2 = 180 Mb/s), estimation of site offered load should be accurate. Figure 3a shows this case in the testbed; we slightly underestimate.

We repeated this experiment with a smaller attack in link 1 (A_1 is 45 Mb/s, so loss is 20% and $\alpha_1 = 0.80$), so total offered load is 325 Mb/s. Here our model predicts we will underestimate, and the testbed experiments show 291 Mb/s.

Our model matches testbed experiments, providing confidence that we understand where it works best and where it mis-estimates. It provides some understanding of our evaluation of real-world cases.

6 FLEXIBILITY OF TRAFFIC ENGINEERING IN ANYCAST

Given an estimate of offered load, the operators want to shift traffic with BGP mechanisms. We next look at the two methods that are widely available: AS-path prepending and using community strings to control routing announcements. For each method we first consider where it can be used (not all methods are available at all sites). We then ask: how much traffic can it shift, and to what degree can this shift be controlled? A method with fine-granularity control is preferred over all-or-nothing.

Here we evaluate these questions across our two testbeds; in §7 we evaluate how choice of testbed sites matters.

6.1 Control With Path Prepending

First we consider AS-path prepending.

6.1.1 Coverage. Support for AS-Path prepending is quite complete—it requires no explicit support from the upstream provider, so we found prepending worked at all sites in both of our testbeds. In `Peering`, we are allowed to use a maximum of three prepends, and in `Tangled` we use up to five prepends. Previous study [13] shows a maximum of 5 prepends is sufficient because 90% of active ASes are located less than six AS hops away. We use RIPE RIS [49] to check the routing visibility when prepends are in place, and we do not observe changes in the routing propagation for both testbeds. Otherwise, this might reveal the existence of AS path length filters [28, 29].

6.1.2 Does Prepending Work? Since AS-Path prepending is widely supported, it seems like an attractive way to shift traffic. We next evaluate how much traffic it can shift.

To explore this question we examine a particular anycast configuration. We first establish the traffic baseline as how many /24 network blocks are in each anycast site's catchment. Weighting catchments by prior traffic loads gives result with different constants but similar shape.

Here we consider a representative scenario using `Peering` using three sites from three continents—Europe (Amsterdam-AMS), North America (Boston-BOS) and South America (Brazil-CNF). In §7 we explore how different configurations change these results.

Figure 4 shows the traffic from each site under different conditions. The middle bar in each graph is the baseline, the default condition with no prepending. We then add prepending at each site, with one, two or three prepends in each bar going to the right of center. We also consider negative prepending (§3.5) in one to three steps, with bars going left of center.

We first consider the baseline (the middle bar) of all three graphs in Figure 4. Amsterdam (AMS, the bottom, maroon

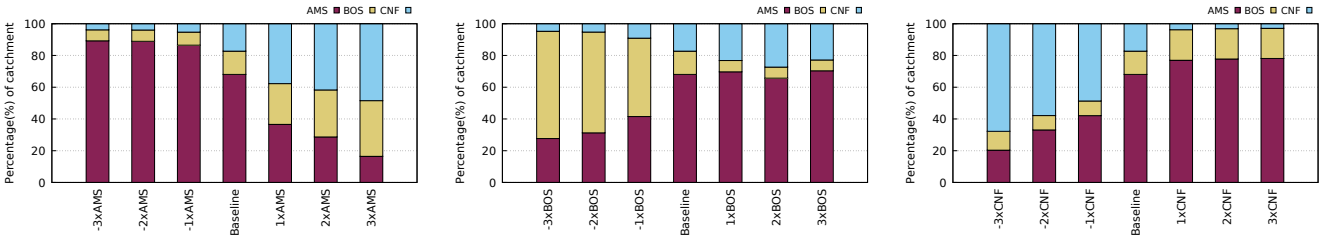

(a) AMS site. **(b) BOS site.** **(c) CNF site.**

Figure 4: `Peering`: Impact of path prepending in catchment distribution with AMS, BOS and CNF sites on 2020-02-24.

part of each bar) gets about 68% of the traffic. AMS receives more traffic than BOS and CNF because that site has two transit providers and several peers, and Amsterdam is very well connected with the rest of the world.

We next consider prepending at each site (the bars to the right of center). In each case, *prepending succeeds at pushing traffic away from the site*, as expected. For AMS, each prepend shifts more traffic away, with the first prepend cutting traffic from 68% to 37%, then to 29%, then to about 16%. BOS and CNF start with less traffic and prepending has a stronger effect, with one prepend sending most traffic away (at BOS, from 15% to 7%) and additional prepends showing little further change. These non-linear changes are because changing BGP routing with prepending is based on path length, and the Internet's AS-graph is relatively flat [2, 14].

The bar graphs also show that when prepending pushed traffic away from a site, it all goes to some other site. Where it goes depends on routing and is not necessarily proportional to the split in other configurations. For example, after one prepend to AMS, more traffic goes to CNF (the top sky blue bar) than to BOS (the middle yellowish bar). These unexpected shifts are why we suggest pre-computing a "playbook" of routing options before an attack (§3.2) to guide decisions during an attack and anticipate the consequences of a change.

We also see that negative prepending succeeds at drawing traffic towards the site—in each case the bars to the left of center see more traffic in the site that is not prepending while the others prepend. AMS sees relatively little change (68% to 89%) since it already has most traffic, while BOS and CNF each gain up to 68% of traffic.

All three sites show some networks that are "stuck" on that site, regardless of prepending. One reason for this stickiness is when some networks are only routable through one site because they are downstream of that exchange. We confirm this by taking traceroute to two randomly chosen blocks that are stuck at BOS, 18.18.11.0/24 and 24.63.240.0/24. Traceroutes and geolocation (with Maxmind) confirm they are in

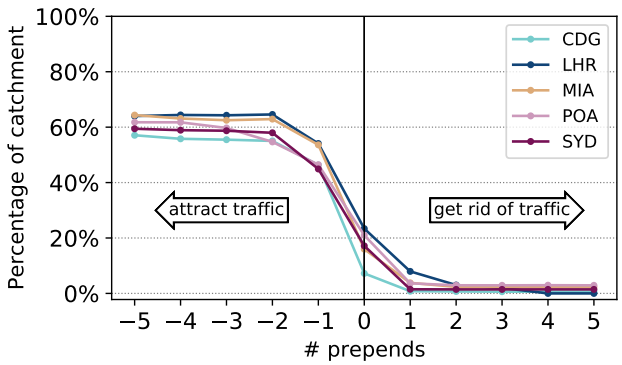

Figure 5: `Tangled`: catchment distribution in respective BGP prepend policy.

Boston, at MIT and a Comcast network (based on the penultimate traceroute hop). We have used the local-preference BGP attribute to move stuck blocks, but a systematic exploration of that option is future work.

In summary, the experiment shows that AS prepend does work and can shift traffic among sites, however, this traffic shift is not uniform.

6.1.3 What granularity does prepending provide? Having established that prepending can shift traffic, we next ask: how much control does it provide? This question has two facets: how much traffic can we push away from a site or attract to it, and how many different levels are there between minimum and maximum.

Limits: Figure 4 suggested that in `Peering`, with those three sites, there is a limit to the traffic that can shift. AMS, BOS, and CNF always get about 16%, 7% and 3% of blocks, regardless of prepending.

Figure 5 confirms this result with a 5-site deployment (two from Europe, one from North America, one from South America and one from Australia) in our other testbed (`Tangled`). X axis is presented with the number of prepends applied to each site. The number zero (0) represents the `baseline`, the positive numbers (1-5) are the number of prepending applied

Site	Transits	IXP peers	Options
AMS	2	122	20
BOS	1	0	1
CNF	1	5	1
ATH	1	0	1
SEA	1	0	1
SLC	1	0	1
ATL	1	0	1
MSN	1	0	1

Table 3: Selected sites from `Peering` along with the number of transits, IXP peers, and propagating ASes considering IPv4 BGP sessions on 2020-02-19.

and the negative numbers represent negative prepends. As depicted, each site can capture at most 55–65% of blocks, and can shed at most 95% of blocks, even with up to 5 prepends. We can also see that we do not get a granular control as only three points are between the minimum and maximum.

We conclude that while prepending can be a useful tool to shift traffic, it provides relatively limited control.

6.2 Control With BGP Communities

While AS-Path Prepending works everywhere, it provides limited control. We next show that BGP community strings have the opposite trade-off: what options they support vary from site to site, but when available, they usually provide more precise control over traffic.

6.2.1 Coverage. ASes must opt-in to community strings, as opposed to prepending's near-universal support (since AS paths are used for loop detection, prepending works unless it is explicitly filtered out), Explicit support is required because communities are only a tagging mechanism; the actions they trigger are at the discretion of peering AS. Prior work has studied the diverse options supported by community strings [25].

To evaluate coverage, we review support for BGP communities in the testbeds we use. The testbeds provide information about two dozen locations, a diverse, although potentially biased, sample.

`Peering` supports selective announcement to the transit providers and IXP peers that a site connects to, although each site provides different numbers of IXP peers and transits (see Table 3). Many sites have only one a transit and so provide no options. We experimented with all options at AMS, since with more than 100 peers over the IXP fabric it should provide flexibility. Unfortunately only a fraction of these respond to our announcements, 104 IXP peers seem to fail to propagate our routing announcement based on measurements from Verfploeter. Similarly, we found no IXP peers propagate our announcement when we use CNF. The *options* column of Table 3 summarizes these results, showing how many routing options we have using community strings and transits.

`Tangled` has different support over the sites. MIA supports communities for selective announcement and selective path pepending. They document support to well-known communities, but they do not suport those described in RFC8642 [7]. Four other sites (LHR, SYD, CDG, and PDA) provide all three communities (for selective announcement, well-known, and prepending). LAX and IAD do not support communities. The POA site is at an IXP with 208 IXP peers; it should provide very fine-grained control. (We plan to examine POA control to see its effects in practice.)

We conclude that, even at IXPs, selective announcement often requires direct negotiation with each peer, further limiting coverage.

6.2.2 Do Community Strings Work and At What Granularity? We next examine how well community strings work and what granularity of control they provide. We use community strings to specific BGP selective announcements, where we propagate our route only to specific transit providers or IXP peers.

For our experiment, we use `Peering`, varying announcements at AMS and observing traffic when anycast is provided from AMS, BOS and CNF (the same topology as §6.1.2). As described in §6.2.1 selective announcement community strings are provided only at AMS, and they affect only two transits and eighteen IXP peers, so community strings only work in this subset of the network.

To select the target ASes for selective announcement, we sort all the working peers of AMS site, based on the size of their customer cone using CAIDA's AS rank list [9]. We then choose the 6 largest IXP peers and the 12 largest, as the left two bars in Figure 6. We then examine the route server, announced separately (the next bar), and then all IXP peers including route servers. Finally, we see the coverage with each of the two transit providers, announced separately.

First, we see that selective announcement provides greater control than just prepending. AMS' catchment goes from 68% of blocks (the baseline), while other configurations provide 53 to 6% of blocks.

Second, we see that there is some overlap in some combinations. For example, the transits each reach more than half of all blocks reachable from AMS, so we know some blocks are reachable form both transit providers. Thus, while there is some control over how many blocks to route to AMS, some peers are very "strong" and will pick up many blocks if they are allow to announce our prefix.

Third, we see the important role of route servers. While direct coordination with 12 IXP peers brings only 7% blocks at AMS, a route server lets AMS reach more ASes and 14% of the blocks alone.

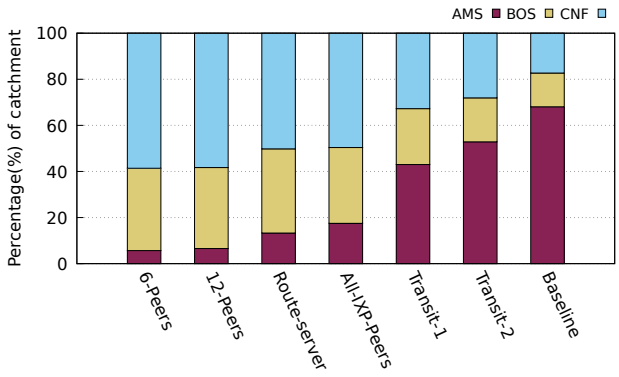

Figure 6: `Peering`: Impact of community string in catchment distribution with AMS, BOS and CNF sites (announced only from AMS site) on 2020-02-25.

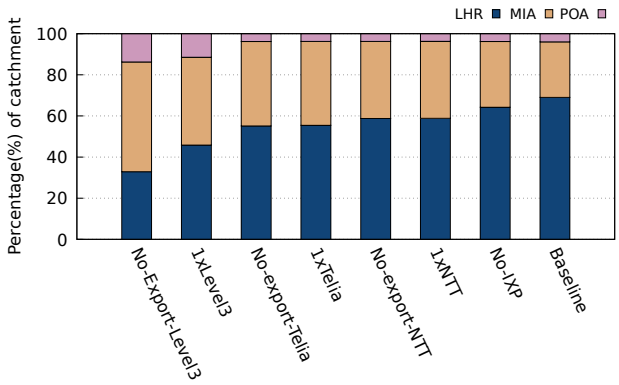

Figure 7: `Tangled`: using different communities to shift traffic on site LHR on 2020-04-05.

Finally, we see that transit providers play in important role. AMS site has two transit providers—BIT BV (AS-12859) and Netwerkvereniging Coloclue (AS-8283). Announcing to AS-8283 attracts more traffic to AMS than announcing to AS-12859. Different AS relationship of these two transits with their upstream provides us a different traffic distribution.

When compared to AS path prepending, BGP communities provide way more options. As shown in our experiments, the different options indeed can provide better control over the traffic distribution.

To investigate if the results found on `Peering` can be generalized, we made a set of experiments on `Tangled`. Like `Peering`, we select 3 sites from three continents—London(LHR), Miami (MIA) and Porto Alegre (POA), and use communities for selective prepending and selective announcement. In Figure 7, we show the catchment distribution after using the community strings from LHR. In the baseline, when no communities are used, LHR handles 69% of traffic. From right to left, we see a gradual decrease in the catchment distribution from 69% to 33%. Stop announcing to IXP peers reduces traffic from 69% to 64%. But using prepending and no export communities in AS2914 (NTT America), AS1299 (Telia Company) and AS3356 (Level 3), we can get 30-60% of the catchments in LHR.

Both testbeds show that community strings are not widely available in all sites, and that even well-known communities are not fully adopted. When, community strings can provide finer-grained control. Selective announcement mostly provides more "flexibility" depending on how many IXP peers and transits are connected. We also find that some sites do not provide the support that they claim which means community strings require an extra step like contacting the transit provider for an explicit agreement.

6.3 Playbook Flexibility

Based on our understanding of prepending and communities, we can now build a playbook of possible traffic configurations for an anycast network.

Table 4 shows a sample playbook with selected configurations. Operators will be aware of the baseline, and when the site is under attack, if they wish to shift a site to a different balance they can read it off from the table. Of course, if attackers are concentrated on certain locations, traffic may not shift exactly as predicted, but this table is a starting point. Finally, this table also suggests *where traffic ends up* after a reconfiguration. The consequences of operator's actions on other sites are as important as reducing the load on one site.

Finally, the table helps us to quantify the "flexibility" that traffic engineering allows us in this anycast deployment. If we divide the traffic mix into 10% bins, we see that AMS has 9 options, while CNF has 7, and BOS has only 5. Because AMS and CNF mostly exchange traffic within them after a BGP change, and because BOS is less well connected, no configuration with three sites allows BOS to take traffic within 40-60% range.

Playbook stability:

A flexible playbook is of limited use if routing immediately changes. We know routing changes when links fail, or when ISPs begin new peering or purchase new transit. How long will a playbook be stable?

To answer this question, we observe the catchment of the baseline case with the sites of Table 4. Table 5 shows the percent of /24 blocks that go to each site over time. We see that the fraction of blocks is generally quite stable, with only about 5% of blocks shifting in or out of a site. In addition, prior work has shown very strong anycast stability over hours to days [64]. While catchments are relatively stable, we encourage operators to refresh playbooks periodically (perhaps weekly or monthly).

Bins(%) / Site	AMS	BOS	CNF
0-10	6-peers, 12-peers	{1, 2, 3}×BOS, 1×CNF, {-1, -2, -3}×AMS, {-1, -2}×CNF	{1, 2, 3}×CNF, {-1, -2, -3}×AMS, {-1, -2, -3}×BOS
10-20	Route-server, All-IXP-peers, 3×AMS	Baseline, {2, 3}×CNF, Transit-2, -3×CNF	Baseline
20-30	-3×CNF, 2×AMS, -3×BOS	Transit-1, 1×AMS,	Transit-2, {1, 2, 3}×BOS
30-40	1×AMS, -2×BOS, -2×CNF	-1×BOS, 3×AMS, 6-peers, 12-peers, Route-server, All-IXP-peers, 2×AMS	1×AMS, Transit-1
40-50	-1 ×CNF, -1×BOS, Transit-1	–	3×AMS, -1×CNF, All-IXP-peers, 2×AMS
50-60	Transit-2	–	-2×CNF, 6-peers, 12-peers, Route-server
60-70	Baseline, {1, 2}×BOS	{-2,-3}×BOS	-3×CNF
70-80	3×BOS	–	–
80-90	{1, 2, 3}×CNF, {-1, -2, -3}×AMS	–	–
90-100	–	–	–
Number of options	9	5	7

Table 4: A playbook for `Peering` with AMS, BOS, and CNF.

Months	AMS(%)	BOS(%)	CNF(%)
2020-02	68.1	14.6	17.3
2020-04	70.4	14.2	15.4
2020-06	65.3	14.1	20.6

Table 5: Percent blocks in each catchment over time.

7 CONSTRAINTS OF ANYCAST DEPLOYMENT

While §6 showed that BGP can shift traffic, it steers traffic between the sites of a given deployment. Very large anycast operators may have multiple anycast deployments on subsets of their points-of-presence, and so are able to activate new capacity. Here we instead assume a smaller operator where all possible sites are active in the same anycast topology of an attack. We examine how site choice and number of sites alter the options for traffic engineering during an attack.

7.1 Effects of Choice of Anycast Sites

First, we consider how site location affects our BGP playbook. Site location and peering at that location influences its catchment, so use of an Australian or Europen might may result in different catchments than a sites in North America.

In §6.1, we studied catchments with three specific `Peering` sites on three continents: AMS, at a large, commercial IXP in Europe; CNF with an academic backbone transit in Brazil; and BOS, an academic site in the U.S. We now switch to three educational sites all in the United States: SEA, at University of Washington on the west coast; SLC, at the University of Utah in the Rockies; and BOS, at Northeastern University in Boston on the east coast.

More important than just geographical location, site connectivity is the most important factor in choosing sites. Multiple transit providers increase the chance of having more BGP options to effect traffic control and granularity. While a poorly connected site inside a university network tends to provide less traffic control options.

Prepending baseline: Figure 8 shows catchment sizes for the three sites with positive and negative prepending. For this network the baseline distribution is unbalanced, but less so than before, with SEA capturing 50% of blocks. We contacted `Peering` people to know why SEA gets an unbalanced traffic. They suspect University of Washington/Northwestern GigaPoP is connected to the Seattle IXP, which may lead to its routes propagating faster (being shorter) than the others. Which catchment has the greatest visibility depends on the peers each site's organization has.

Prepending coverage and granularity: As with our prior experiments, we can adjust prepending to see how traffic shifts. With these three sites, traffic shifts very quickly for BOS and SEA after one positive or negative prepend. SLC has more flexibility, perhaps because it has the smallest catchment at the baseline, and gains more coverage with each step of negative prepending, to 42%, 63%, and 91% of blocks. Often (but not always), we see that academic sites exhibit less granularity because either they have few peers, or their peers are academic networks with similar connectivity. As a result, minor changes in AS-path length place one site further from the others. In addition, this less granular control shows the

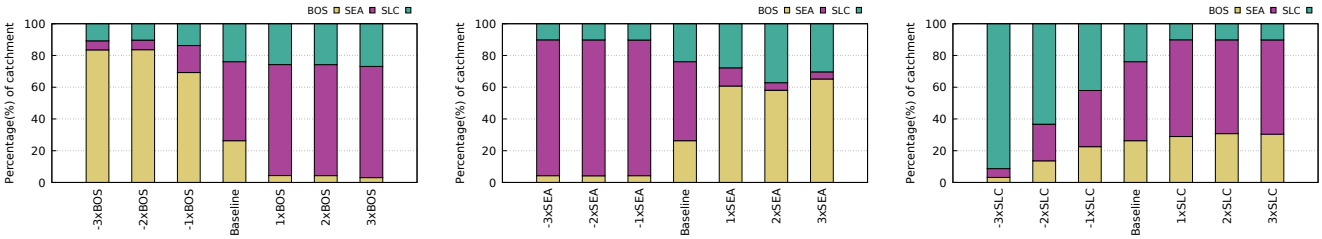

(a) BOS site. **(b) SEA site.** **(c) SLC site.**

Figure 8: `Peering`: Impact of choosing BOS, SEA and SLC sites on 2020-02-28

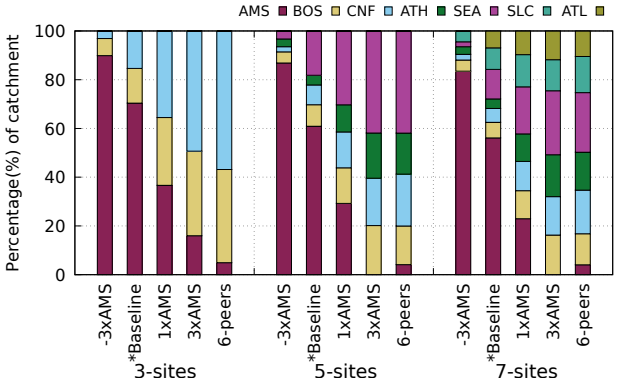

Figure 9: `Peering`: Impacts of changing the number of anycast sites from 2020-04-07 to 2020-04-10.

importance of building a playbook that is specific to a given deployment, or when the anycast topology changes.

Community coverage: While communities are common at IXPs and transit providers, academic networks (NRENs) have a more simple set of communities. However, some of our sites are located behind university networks and do not provide BGP communities.

None of these academic sites on `Peering` provide community strings. This observation confirms our prior coverage observation: community string support is not uniformly available. We also looked at two other combinations of sites in `Peering` and found similar results (details in Appendix A).

7.2 Effects of Number of Anycast Sites

Next, we vary the *number* of sites and see how that changes traffic control. We select 3, 5 and 7 sites from each testbeds. We announce our prefix from each of the sites to establish a baseline, then select different traffic engineering configurations to shift traffic. Figure 9 shows selected configurations, grouped by number of sites.

Baseline: With more sites, overall capacity increases and *baseline traffic at each site is reduced* as traffic spreads out. For example, in Figure 9, baselines are marked with *, and the largest site (AMS) shifts from 70% of blocks with three sites to 61% and 56% with 5 and 7 sites. Smaller sites shift less (BOS goes from 14% to 6% and 6%, and CNS from 15% to 8% and 6%). Greater capacity and distribution requires a larger and distributed attacker to fully exhaust the overall service.

This spread of traffic happens when sites are physically distributed and using different peers. At first, we consider 3 sites in Europe and North and South America (AMS, BOS and CNF). The AMS catchment shrinks 9% with new sites in Europe and North America (ATH in Greece and SEA in the west). Most of this shift is blocks that move from AMS to SEA (208k blocks move, and 101k, about half are in the U.S.). More North American sites (SLC and ATL) reduce AMS by only 5% (68k blocks). Traffic shifts from the European catchment diminish in part once all North American blocks already stay on the continent.

In the same way, we performed experiments on `Tangled` to investigate the effects of the number of anycast sites. We could confirm the findings on `Tangled` and more details are presented on Appendix A. An interesting finding was the site shadowing. At `Tangled` we have two sites (IAD and LAX) that share the same transit provider. When both sites are running only LAX received traffic. This happens because the transit provider of both sites selects one to deliver traffic based on its own policy.

Traffic flexibility: Although more sites make it harder to overwhelm the whole service, if attackers are concentrated in a particular catchment it is easier to exhaust a single site. We need to explore the BGP options even when there are more number of sites.

With more sites, *the largest site usually shows the largest changes*, usually with a few possible catchment sizes. Comparing the baseline with one prepending in Figure 9 we noticed AMS shifts from 70% to 37% with three sites, from 61% to 29% with five, and from 56% to 23% with seven, always being cut in half.

Even with more number of sites, some blocks are often "stuck" at a particular site. With three negative prepends,

AMS gets most of the traffic, but it tops out at 90% with three sites, and only 87% and 84% with five and seven. We conclude that each site has its own set of "stuck blocks" that are captive to it and will not move with traffic engineering.

With more sites, the fine control of BGP communities becomes more important because path-prepending becomes less sensitive. For example, when we have 5 or 7 sites, prepending three times from AMS sheds all its traffic. However, selective announcement only to 6 peers can keep ~5% catchment at AMS (bars with 6-peers in Figure 9). Other selective announcements can provide us even more control which is not possible with path prepending.

New sites: Our examination of site location also suggests what happens when new sites are brought online. Traffic shifts to the new site, but exactly which traffic and how much is difficult to predict, and sometimes a site with similar peers or locations as another. A playbook to anticipate these changes is important if new sites are a possible response to DDoS attack.

8 DEFENSES AT WORK

We now evaluate our system (§3.1), showing how offered load estimation (§5) lets us select traffic engineering choices from our playbook (§6) to support defending against an attack. We then evaluate defense given smart adversaries.

8.1 Defense for One Real-Life Scenario

For an example defense we reproduce a small but real-world attack from `B-root`, evaluating it against a three-site anycast network. We evaluate how we could manage the attack using measured routing catchments.

We take normal and attack traffic from a real event on 2017-03-06 in `B-root` [45]. This event was a volumetric attack where the attack queries have common formats like RANDOM.qycl520.com\032 or RANDOM.cailing168.com\032\032, and the server mostly replies with NXDomain against these queries. In this event we do not believe source IP addresses were spoofed. The event was small enough that `B-root` was able to fully capture it across all active anycast sites at the time. The event lasted about 5 hours, but we show only the first 1000 s.

We evaluate this event by simulating traffic rates against two three-site anycast networks, both using `Peering`: first our AMS, BOS, CNF configuration from §6, then BOS, SEA, SLC from §7.1. We assign traffic to each anycast site based on measurements of catchments from our experiments. Rather than run a DNS service, we examine traffic levels of each site and compare it to a target capacity. The overall event had peak rates of 120k queries/s (55 Mb/s), and our target capacity is 55k queries/s (25 Mb/s) at each site. (Typical services today and difficult attacks will both be much larger; rather than scaling the attack up, we select a relatively low capacity.)

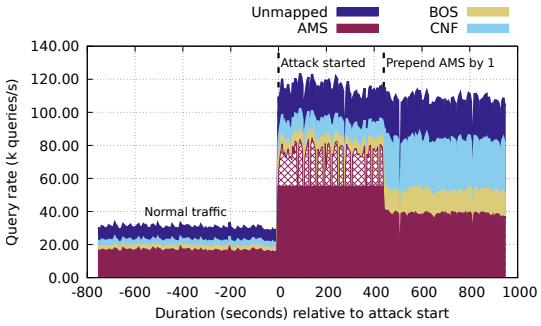

(a) Breakwaters: steering traffic using positive prepending.

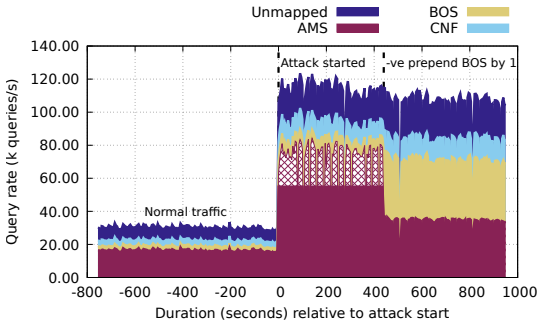

(b) Breakwaters: steering traffic using negative prepending.

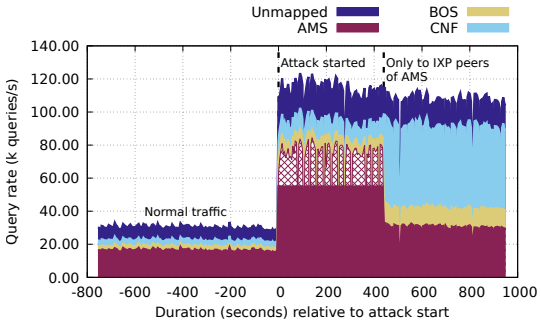

(c) Breakwaters: steering traffic using community string.

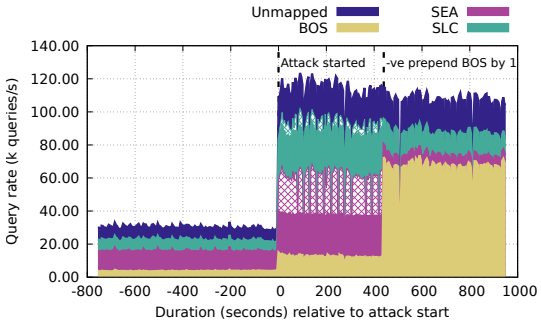

(d) Super-site: pushing traffic towards the super-site using negative prepending.

Figure 10: Redistributing attack traffic using different strategies.

In our replay we treat the event as starting at 0 s and make an anycast change at 400 s. The operator will notice the attack from query rate (or bitrate) alerts set at AMS. The operator will then carry out offered load estimation, which will report that the attack will be over capacity at AMS (shown as the maroon striped area). Consuliting the pre-computed playbook (Table 4), the operator can then select a response.

Figure 10 presents four experiments, first showing three different responses and then an alternate topology. In each graph we show absolute number of queries to each site. As we cannot know the catchment of every block using Verfploeter, we also show the number of unmapped queries. For the first experiment, Figure 10a, the operator selects path prepending at AMS. The playbook suggests this will reduce AMS traffic to about two thirds its prior rate (from 80% to 50% in the playbook), and most of the shifted traffic will go to CNF. After the route change, we see all three sites are under capacity, as shown by the elimination of the AMS striped region.

The operator could take other responses. For example, suppose there were concerns about potential overload at CNF. The playbook suggests that negative prepending at BOS will shift traffic to BOS without increasing traffic at CNF. Figure 10b shows traffic after negative prepending once. AMS is no longer over capacity, and CNF traffic only changes slightly as BOS picks up the load.

Figure 10c shows a third strategy where we use BGP communities at AMS to shift traffic away. In this case we retain routes to the all IXP peers at AMS but drop both transits. According to the playbook, this approach can balance the traffic among three sites. Any one of these three approaches was effective to get a traffic distribution where all the query rates are below than the assumed capacity.

Finally, we consider a different set of sites (BOS, SEA, SLC from §7.1) and we assume SEA and SLC have capacity for 25k queries/s while BOS supports 75k queries/s. Figure 10d shows that at attack time, both smaller sites are overwhelmed (as shown by the striped regions). We use negative prepending at BOS to bring all sites under capacity, with BOS shouldering most of the load.

8.2 Adversarial Traffic

DDoS attacks are inherently adversarial, so we must assume the attacker will become aware of our defenses and seek to counter them. We consider some responses and counters below.

Attackers often spoof source IP addresses to prevent easily filtering. Although best practices suggest filtering [22], and the quantity of IoT devices perhaps obviate the need, spoofing is still common. While spoofed source addresses will result in incorrect predictions about where load will go, our primary decision is based on the actual load that arrives at a site and that load is unaffected by spoofing.

Attackers may try to concentrate attacks on one site, or spread attacks evenly across all sites. In either case, the defender can select their desired response (absorb and sacrifice a site or spread traffic). In fact, load estimation helps the defender's choice.

Polymorphic attackers changes attack types or intensity during an attack. Although we increase defense agility, an agile attacker will be a more difficult adversary.

Finally, our current proposal assumes human-in-the-loop decisions. If they were integrated into an automatic defense, one may need to moderate the rate of routing changes to avoid interactions with route-flap damping [44].

9 CONCLUSIONS

This paper provides the first complete description of using anycast to defend against DDoS attacks. Our approach starts with pre-computing a playbook of responses, then when the attack arrives, the operator estimates offered load, selects a strategy, and uses BGP to shift traffic to best defend. We answer key questions like how available and flexible this approach is for different configurations across two testbeds. We show that prepending is widely available but offers limited control, while BGP communities are the opposite We evaluate defenses with modeling and against a real-world event. While operators have always used their best judgment when under fire, our work provides guidance for responses.

ACKNOWLEDGMENTS

ASM Rizvi and John Heidemann's work on this paper is supported, in part, by the DHS HSARPA Cyber Security Division via contract number HSHQDC-17-R-B0004-TTA.02-0006-I.

Joao Ceron and Leandro Bertholdo's work on this paper is supported by European Unionfis Horizon 2020 research and innovation programme under grant agreement No 830927.

REFERENCES

[1] AMPATH. Ampath resources. https://ampath.net/AMPATH_BGP_Policies.php, 2020. [Online; accessed 4-April-2020].

[2] APNIC. BGP-stats routing table report—Japan view. https://mailman.apnic.net/mailing-lists/bgp-stats/archive/2020/05/msg00001.html, May 1 2020.

[3] Vaibhav Bajpai, Steffie Jacob Eravuchira, and Jürgen Schönwälder. Lessons learned from using the ripe atlas platform for measurement research. *ACM SIGCOMM Computer Communication Review*, 45(3):35–42, 2015.

[4] Hitesh Ballani and Paul Francis. Towards a deployable IP anycast service. In *WORLDS*, 2004.

[5] Hitesh Ballani, Paul Francis, and Sylvia Ratnasamy. A measurement-based deployment proposal for IP anycast. In *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*, pages 231–244, 2006.

[6] Terry Benzel, Robert Braden, Dongho Kim, Cliford Neuman, Anthony Joseph, Keith Sklower, Ron Ostrenga, and Stephen Schwab. Experience with deter: a testbed for security research. In *2nd International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities, 2006. TRIDENTCOM 2006.*, pages 10–pp. IEEE, 2006.

[7] J. Borkenhagen, R. Bush, R. Bonica, and S. Bayraktar. Policy behavior for well-known bgp communities. Technical Report 1546, RFC Editor, 2019.

[8] Matthew Caesar and Jennifer Rexford. BGP routing policies in ISP networks. *IEEE network*, 19(6):5–11, 2005.

[9] CAIDA. AS rank. https://asrank.caida.org/, 2020. [Online; accessed 9-March-2020].

[10] CAIDA. CAIDA UCSD BGP community dictionary. https://www.caida.org/data/bgp-communities/, 2020. [Online; accessed 4-April-2020].

[11] Matt Calder, Ashley Flavel, Ethan Katz-Bassett, Ratul Mahajan, and Jitendra Padhye. Analyzing the performance of an anycast CDN. In *Proceedings of the 2015 Internet Measurement Conference*, pages 531–537, 2015.

[12] R. Chandra, P. Traina, and T. Li. BGP communities attribute. Technical Report 1997, RFC Editor, 1996.

[13] Rocky KC Chang and Michael Lo. Inbound traffic engineering for multihomed ASs using AS path prepending. *IEEE network*, 19(2):18–25, 2005.

[14] Yi-Ching Chiu, Brandon Schlinker, Abhishek Balaji Radhakrishnan, Ethan Katz-Bassett, and Ramesh Govindan. Are we one hop away from a better Internet? In *Proceedings of the ACM Internet Measurement Conference*, pages 523–529, Tokyo, Japan, October 2015. ACM.

[15] Danilo Cicalese, Jordan Augé, Diana Joumblatt, Timur Friedman, and Dario Rossi. Characterizing IPv4 anycast adoption and deployment. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, pages 1–13, 2015.

[16] Danilo Cicalese and Dario Rossi. A longitudinal study of IP anycast. *ACM SIGCOMM Computer Communication Review*, 48(1):10–18, 2018.

[17] Cloudflare. Famous DDoS attacks — the largest DDoS attacks of all time. https://www.cloudflare.com/learning/ddos/famous-ddos-attacks/, 2020. [Online; accessed 4-May-2020].

[18] Wouter B. de Vries, Ricardo de O. Schmidt, Wes Hardaker, John Heidemann, Pieter-Tjerk de Boer, and Aiko Pras. Verfploeter: Broad and load-aware anycast mapping. In *Proceedings of the ACM Internet Measurement Conference*, London, UK, 2017.

[19] Christoph Dietzel, Anja Feldmann, and Thomas King. Blackholing at IXPs: On the effectiveness of DDoS mitigation in the wild. In *International Conference on Passive and Active Network Measurement*, pages 319–332. Springer, 2016.

[20] ExaBGP. The BGP swiss army knife of networking. https://github.com/Exa-Networks/exabgp, 2020. [Online; accessed 4-April-2020].

[21] Xun Fan, John Heidemann, and Ramesh Govindan. Evaluating anycast in the domain name system. In *2013 Proceedings IEEE INFOCOM*, pages 1681–1689. IEEE, 2013.

[22] P. Ferguson and D. Senie. Network ingress filtering: Defeating denial of service attacks which employ IP source address spoofing. RFC 2827, May 2000. also BCP-38.

[23] Ashley Flavel, Pradeepkumar Mani, David Maltz, Nick Holt, Jie Liu, Yingying Chen, and Oleg Surmachev. Fastroute: A scalable load-aware anycast routing architecture for modern CDNs. In *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*, pages 381–394, 2015.

[24] Ruomei Gao, Constantinos Dovrolis, and Ellen W Zegura. Interdomain ingress traffic engineering through optimized AS-path prepending. In *International Conference on Research in Networking*, pages 647–658. Springer, 2005.

[25] Vasileios Giotsas, Georgios Smaragdakis, Christoph Dietzel, Philipp Richter, Anja Feldmann, and Arthur Berger. Inferring BGP blackholing activity in the internet. In *Proceedings of the 2017 Internet Measurement Conference*, pages 1–14. ACM, 2017.

[26] T. Hardie. Distributing authoritative name servers via shared unicast addresses. Technical Report 3258, RFC Editor, 2002.

[27] Chi-Yao Hong, Subhasree Mandal, Mohammad Al-Fares, Min Zhu, Richard Alimi, Kondapa Naidu B., Chandan Bhagat, Sourabh Jain, Jay Kaimal, Shiyu Liang, Kirill Mendelev, Steve Padgett, Faro Rabe, Saikat Ray, Malveeka Tewari, Matt Tierney, Monika Zahn, Jonathan Zolla, Joon Ong, and Amin Vahdat. B4 and after: Managing hierarchy, partitioning, and asymmetry for availability and scale in Google's software-defined WAN. In *Proceedings of the ACM SIGCOMM Conference*, Budapest, Hungary, August 2018. ACM.

[28] Geoff Huston. BGP in 2017. https://labs.apnic.net/?p=1102, Jan 8 2018. [Online; accessed 11-June-2020].

[29] Team Cymru Inc. Secure Cisco IOS BGP template. https://www.team-cymru.com/secure-bgp-template.html. [Online; accessed 11-June-2020].

[30] Brian Krebs. Krebsonsecurity hit with record DDoS. *KrebsOnSecurity, Sept*, 21, 2016.

[31] Jan Harm Kuipers. Anycast for DDoS. https://essay.utwente.nl/73795/1/Kuipers_MA_EWI.pdf, 2017. [Online; accessed 5-May-2020].

[32] Craig Labovitz, Abha Ahuja, Abhijit Bose, and Farnam Jahanian. Delayed Internet routing convergence. *ACM SIGCOMM Computer Communication Review*, 30(4):175–187, 2000.

[33] Jason Lackery. A new twist on denial of service: DDoS as a service. https://blogs.cisco.com/security/a_new_twist_on_denial_of_service_ddos_as_a_service, 2010. [Online; accessed 6-June-2019].

[34] Zhihao Li, Dave Levin, Neil Spring, and Bobby Bhattacharjee. Internet anycast: Performance, problems, & potential. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, pages 59–73, 2018.

[35] Ziqian Liu, Bradley Huffaker, Marina Fomenkov, Nevil Brownlee, et al. Two days in the life of the DNS anycast root servers. In *International Conference on Passive and Active Network Measurement*, pages 125–134. Springer, 2007.

[36] Marek Majkowski. Memcrashed - major amplification attacks from UDP port 11211. https://blog.cloudflare.com/memcrashed-major-amplification-attacks-from-port-11211/, 2018. [Online; accessed 4-May-2020].

[37] Stephen McQuistin, Sree Priyanka Uppu, and Marcel Flores. Taming anycast in the wild Internet. In *Proceedings of the Internet Measurement Conference*, pages 165–178, 2019.

[38] Giovane C. M. Moura, Ricardo de O. Schmidt, John Heidemann, Wouter B. de Vries, Moritz Müller, Lan Wei, and Christian Hesselman. Anycast vs DDoS: Evaluating the November 2015 root DNS event. In *Proceedings of the ACM Internet Measurement Conference*, November 2016.

[39] Priyadarsi Nanda and AJ Simmonds. A scalable architecture supporting QoS guarantees using traffic engineering and policy based routing in the Internet. *International Journal of Communications, Network and System Sciences*, 2009.

[40] Lily Hay Newman. Github servived the biggest DDoS attack ever recorded. https://www.wired.com/story/github-ddos-memcached/, 2018. [Online; accessed 19-March-2018].

[41] Root Server Operators. Events of 2015-11-30. https://root-servers.org/news/events-of-20151130.txt, 2015. [Online; accessed 9-July-2019].

[42] Root Server Operators. Events of 2016-06-25. https://root-servers.org/news/events-of-20160625.txt, 2016. [Online; accessed 9-July-2019].

[43] Craig Partridge, Trevor Mendez, and Walter Milliken. Host anycasting service. Technical Report 1546, RFC Editor, 1993.

[44] Cristel Pelsser, Olaf Maennel, Pradosh Mohapatra, Randy Bush, and Keyur Patel. Route flap damping made usable. In *Proceedings of the*, pages 143–152, Atlanta, GA, USA, March 2011.
[45] LANDER project. Lander:b root anomaly-20170306. https://ant.isi.edu/datasets/readmes/B_Root_Anomaly-20170306.README.txt, 2019. [Online; accessed 27-April-2020].
[46] Bruno Quoitin, Cristel Pelsser, Olivier Bonaventure, and Steve Uhlig. A performance evaluation of BGP-based traffic engineering. *International journal of network management*, 15(3):177–191, 2005.
[47] Bruno Quoitin, Cristel Pelsser, Louis Swinnen, Olivier Bonaventure, and Steve Uhlig. Interdomain traffic engineering with BGP. *IEEE Communications magazine*, 41(5):122–128, 2003.
[48] Y. Rekhter, T. Li, and S. Hares. A border gateway protocol 4 (bgp-4). Technical Report 4271, RFC Editor, 2006.
[49] RIPE Network Coordination Centre. RIPE - Routing Information Service (RIS). https://https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris, 2020.
[50] ASM Rizvi, John Heidemann, and Jelena Mirkovic. Dynamically selecting defenses to DDoS for DNS (extended). Technical Report ISI-TR-736, USC/Information Sciences Institute, May 2019.
[51] Sandeep Sarat, Vasileios Pappas, and Andreas Terzis. On the use of anycast in DNS. In *Proceedings of 15th International Conference on Computer Communications and Networks*, pages 71–78. IEEE, 2006.
[52] Brandon Schlinker, Hyojeong Kim, Timothy Cui, Ethan Katz-Bassett, Harsha V. Madhyastha, Italo Cunha, James Quinn, Saif Hasan, Petr Lapukhov, and Hongyi Zeng. Engineering egress with Edge Fabric: Steering oceans of content to the world. In *Proceedings of the ACM SIGCOMM Conference*, pages 418–431, Los Angeles, CA, USA, August 2017. ACM.
[53] Brandon Schlinker, Kyriakos Zarifis, Italo Cunha, Nick Feamster, and Ethan Katz-Bassett. Peering: An AS for us. In *Proceedings of the 13th ACM Workshop on Hot Topics in Networks*, pages 1–7, 2014.
[54] Ricardo de O. Schmidt, John Heidemann, and Jan Harm Kuipers. Anycast latency: How many sites are enough? In *International Conference on Passive and Active Network Measurement*, pages 188–200, Sydney, Australia, March 2017.
[55] Thomas Bradley Scholl. Methods and apparatus for distributed backbone internet DDoS mitigation via transit providers, February 3 2015. US Patent 8,949,459.
[56] A. Shaikh, R. Tewari, and M. Agrawal. On the effectiveness of DNS-based server selection. In *Proceedings IEEE INFOCOM 2001. Conference on Computer Communications. Twentieth Annual Joint Conference of the IEEE Computer and Communications Society (Cat. No.01CH37213)*, volume 3, pages 1801–1810 vol.3, 2001.
[57] R. B. da Silva and E. Souza Mota. A survey on approaches to reduce BGP interdomain routing convergence delay on the Internet. *IEEE Communications Surveys & Tutorials*, 19(4):2949–2984, 2017.
[58] Daniel Smith. Antithe growth of DDoS-as-a-service: Stresser services. https://blog.radware.com/security/2017/09/growth-of-ddos-as-a-service-stresser-services/, 2017. [Online; accessed 18-June-2018].
[59] Oliver Spatscheck, Zakaria Al-Qudah, Seunjoon Lee, Michael Rabinovich, and Jacobus Van Der Merwe. Multi-autonomous system anycast content delivery network, December 10 2013. US Patent 8,607,014.
[60] RIPE NCC Staff. Ripe atlas: A global internet measurement network. *Internet Protocol Journal*, 18(3), 2015.
[61] One Step. One step BGP community guides. https://onestep.net/communities/, 2020. [Online; accessed 4-April-2020].
[62] Eric Sven-Johan Swildens, Zaide Liu, and Richard David Day. Global traffic management system using IP anycast routing and dynamic load-balancing, March 8 2011. US Patent 7,904,541.
[63] Renata Teixeira, Steve Uhlig, and Christophe Diot. BGP route propagation between neighboring domains. In *International Conference on Passive and Active Network Measurement*, pages 11–21. Springer, 2007.
[64] Lan Wei and John Heidemann. Does anycast hang up on you? In *2017 Network Traffic Measurement and Analysis Conference (TMA)*, pages 1–9, Dublin, Ireland, July 2017. IEEE.
[65] Fernanda Weiden and Peter Frost. Anycast as a load balancing feature. In *Proceedings of the 24th international conference on Large installation system administration*, pages 1–6. USENIX Association, 2010.
[66] Curt Wilson. Attack of the Shuriken: Many hands, many weapons. https://www.arbornetworks.com/blog/asert/ddos-tools/, 2012. [Online; accessed 31-Jan-2018].

APPENDIX A BGP IN OTHER ANYCAST SET-UP

In this section, we present the complementary results from the experiments performed in both testbeds. Here we present data used to generalize our findings.

A.1 A Small Site in Europe

AMS in `Peering` is well-connected with two transits, and several IXP peers. Next, instead of AMS, we take ATH in Europe which is connected through a research network in Greece. Our goal is to see whether the findings from §6 are still valid in this anycast setup. Like the previous setup, we also take BOS and CNF.

Figure 11 shows the catchment distribution. In the baseline case, as ATH is not connected like AMS, it gets only 20% traffic. CNF serves almost 50% traffic in this anycast setup. So, the traffic distribution is still skewed in this setup. Even if both AMS and ATH are in Europe, we see a different catchment control which indicates the importance of connectivity of the anycast sites.

Prepending works similarly in this setup. BOS and CNF can cut most of their traffic after first prepend. However, ATH can only shift 7% traffic after first prepend, and for more prepends it does not show any effect. In all these sites, we can also see that some blocks are always "stuck" to a particular site. Using negative prepending, we can push most of the traffic to BOS and CNF. However, we can only push 40% traffic to ATH site.

A.2 Sites in Nearby Location

Next, we take three sites that are in nearby location and have similar connectivity. We select sites from Boston (BOS), Atlanta (ATL) and Wisconsin (MSN) in `Peering`. All these sites are located within the eastern half of the U.S., and they are connected through education network—BOS with Northeastern University, ATL with Georgia Institute of Technology, and MSN with University of Wisconsin - Madison.

When the sites are located in nearby geo-location, and connected by similar network, path prepending can result in an "all or none" outcome. When we prepend from ATL

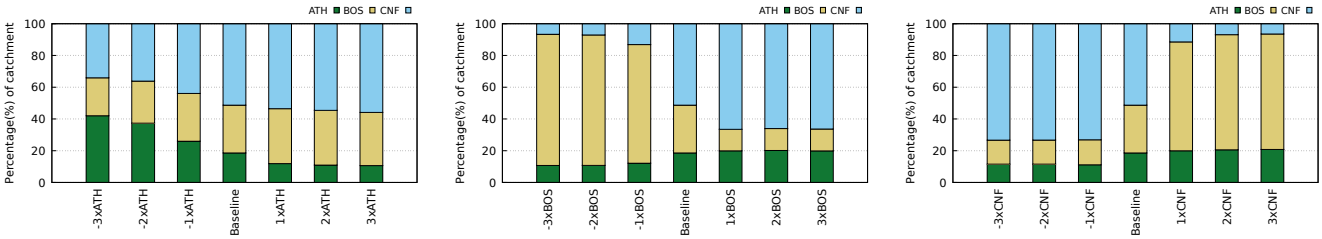

(a) ATH site. **(b) BOS site.** **(c) CNF site.**

Figure 11: `Peering`: Impact of path prepending in catchment distribution with ATH, BOS and CNF sites on 2020-05-30.

Site	Transits	IXP peers	Options
MIA	1	0	1
LHR	1	0	1
SYD	1	0	1
CDG	1	0	1
LAX	1	0	1
IAD	1	0	1
POA	2	218	218

Table 6: Selected sites from `Tangled` along with the number of transits, IXP peers, and propagating ASes considering IPv4 BGP sessions on 2020-02-27.

or BOS, most traffic goes away from these sites. Prepending one time from ATL leaves no traffic in ATL, and prepending one time from BOS cuts traffic from 42% to 14%. Even with negative prepending, BOS can get over 90% traffic, and ATL can get nearly 90% traffic. So, BOS and ATL can cut or gain almost all the traffic with positive and negative prepending.

MSN receives a small fraction of traffic in the baseline, and some blocks are always "stuck" at MSN. With two negative prepends, MSN receives only 27% catchment, however, with the third negative prepend, MSN receives almost 80% catchment. This slow and sudden increase in the catchment suggests us why we need a BGP "playbook" for an anycast setup.

A.3 Communities coverage on `Tangled`

`Tangled` supports selective announcement to the transit providers and IXP peers that a site connects to. Depending on the number of connected transits and IXP peers, we have different number of community options at each site.

In Table 6 we show all the sites of `Tangled` that we use. Most of them just have one transit provider and no direct connection to IXP, where POA is an exception. On this site we have two distinct transit providers and a direct connection to local IXP. While IXP says to provide selective communities, we did not explore this in detail.

A.4 More Sites in `Tangled`

We next want to confirm that increasing the number of sites in `Tangled` shows the similar results that we get in §7.2.

Like `Peering`, we take 3, 5 and 7 sites in `Tangled` testbed (Figure 13). As the overall capacity increases with more number of sites, in `Tangled` also we can see the baseline traffic is reduced in each site as the traffic spreads out. LHR gets 55, 31 and 25% traffic when there are 3, 5 and 7 sites.

With more sites, as there are more capacity in other sites, one site can cut almost all of its traffic. For example, with 3 prepends, LHR can cut all of its traffic when there 7 sites which is not possible when there are 3 or 5 sites.

We can see a "shadowing" instance when we have a 7-site testbed. After adding ENS and POA, we can see that all traffic from LAX site disappears (Figure 13). We believe Cogent traffic now shifts from LAX to POA, and academic traffic shifts from LAX to ENS. We explain this issue in §7.2 when IAD shadows LAX.

Like `Peering`, adding more sites can create new options where the shifted traffic goes. For example, with 3 sites, LHR traffic goes to MIA when we make prepending. But when we have 7 sites, a significant amount of LHR traffic goes to POA—POA traffic increases from 26% to 40% when we prepend from LHR. Hence, it is necessary to keep a "Playbook" to see the traffic distribution after a BGP change.

APPENDIX B SIDE-EFFECTS OF CATCHMENT MANIPULATION

As BGP manipulation steers traffic among anycast sites, it might redirect clients to a distant site causing an increased path length and Round Trip Time (RTT). Next, we want to measure these side-effects in our playbook to understand how they change client latency. For short DDoS-events, latency may be secondary, but some DDoS events last for days and so latency may be relevant.

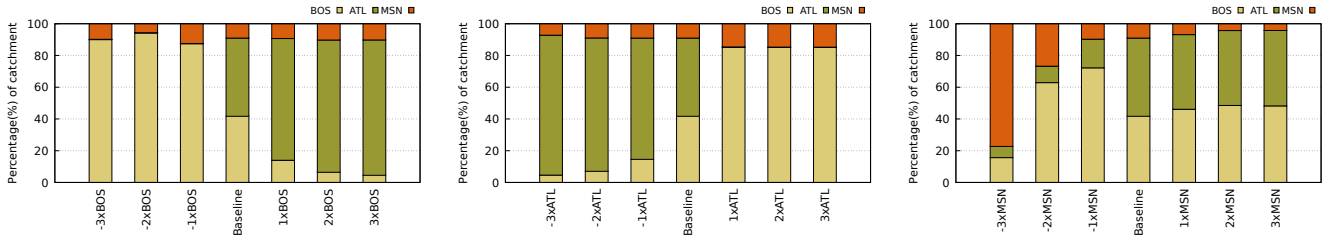

(a) BOS site. **(b) ATL site.** **(c) MSN site.**

Figure 12: `Peering`: Impact of path prepending in catchment distribution with BOS, ATL and MSN sites on 2020-05-29.

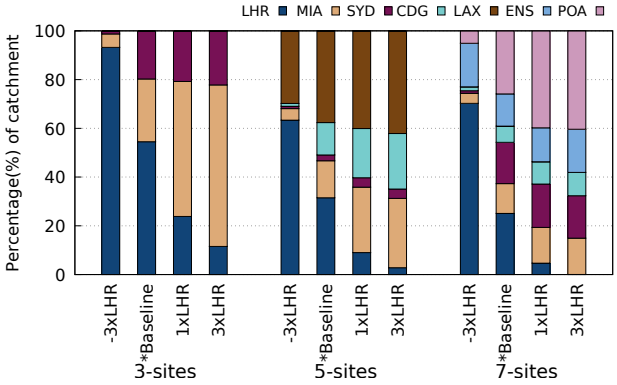

Figure 13: `Tangled`: Impacts of changing the number of anycast sites.

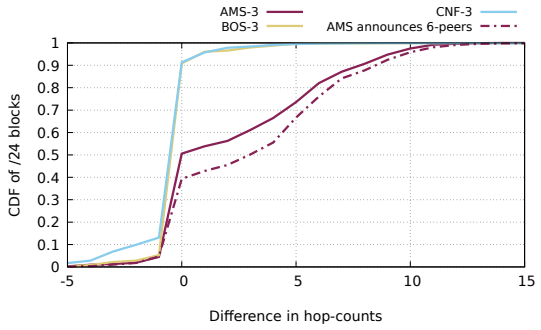

Figure 14: Changes in number of router hops after BGP manipulation

B.1 Changes in Number of Router Hops

BGP manipulation changes the catchment of the clients. Clients might need to go to a distant site. We investigate how the router distance changes when we manipulate BGP.

To know the changes in the hop count, we measure the Time To Live (TTL) value for each of the incoming packets at each site. Different targets set different initial TTL values, and each router decreases the TTL value by one when it forwards a packet. From the difference between the initial and received TTL values at an anycast site, we can determine the number of routing hops from the targets to the anycast sites. As different targets set different initial TTL values, we do not guess the initial values. We use the difference between the received TTL values before and after a BGP manipulation to get the change in the hop count. We admit that a particular target sometimes sets different initial TTL values at different site due to the possible presence of a NAT, however, this does not affect our result as they are very small in number.

Next, we observe how the hop count changes when we manipulate BGP from different sites. We consider the maximum three prepended paths from AMS, BOS and CNF sites to observe the maximum possible distorted scenario with path prepending. We also consider an extreme condition with community strings where we announce only to 6 peers of AMS site. From Figure 14, we can see that over 40% prefixes travel the similar or even less hop distance (when the difference in hop-count is 0 or negative). It is also apparent that prepending from a well-connected site like AMS can cause around 50% prefixes to choose a longer path, and 25% prefixes travel more than 5 router hops. When we prepend from a relatively less visible site like BOS or CNF, around 10% prefixes find a better route in terms of hop-count, and around 80% prefixes travel the same hop distance. In summary, a significant amount of prefixes see the same hop distance even after three prepends or announce only to 6 peers, and prepending from a well-connected site has the worst impact over the router hop distance of the clients.

B.2 Changes in RTT

Next, we want to see how the RTT changes when we manipulate BGP from different sites. We use Verfploeter [18], but we extended it to measure the RTT.

Variation in RTT follows the similar pattern like the variation in hop count. When we prepend from a well-connected site like AMS, we see an increased RTT—over 60% prefixes

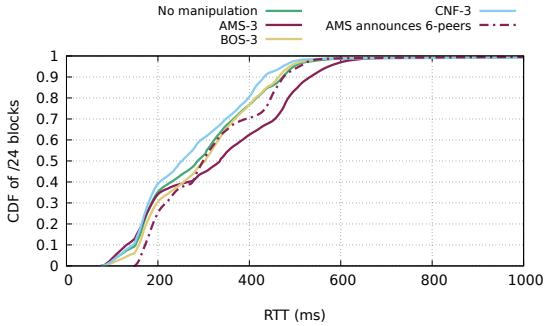

Figure 15: RTT variation after BGP manipulation

get an increased RTT (difference between the no manipulation and AMS-3 lines in Figure 15). When we announce only to 6 peers from AMS site, over 50% prefixes see an increased RTT. We find the opposite with a relatively not well-connected site—improved or mostly similar RTT when we prepend from BOS or CNF.