

Anycast Polarization in The Wild

ASM Rizvi^{1,2}, Tingshan Huang², Rasit Esrefoglu², and John Heidemann¹

¹ University of Southern California / Information Sciences Institute

² Akamai Technologies

Abstract. IP anycast is a commonly used method to associate users with services provided across multiple sites, and if properly used, it can provide efficient access with low latency. However, prior work has shown that *polarization* can occur in global anycast services, where some users of that service are routed to an anycast site on another continent, adding 100 ms or more latency compared to a nearby site. This paper describes the causes of polarization in real-world anycast and shows how to observe polarization in third-party anycast services. We use these methods to look for polarization and its causes in 7986 known anycast prefixes. We find that polarization occurs in more than a quarter of anycast prefixes, and identify incomplete connectivity to Tier-1 transit providers and route leakage by regional ISPs as common problems. Finally, working with a commercial CDN, we show how small routing changes can often address polarization, improving latency for 40% of clients, by up to 54%.

Keywords: Anycast · Polarization · BGP.

1 Introduction

Anycast is a routing approach where *anycast sites* in multiple locations announce the same service on the same IP prefix. First defined in 1993 [26], today anycast is used by many DNS services and Content-Delivery Networks (CDNs) to reduce latency and to increase capacity [34,11,12,8]. With sites at many Points-of-Presence (PoPs), clients can often find one physically nearby, providing low latency [30,17]. With a large capacity of servers that are distributed across many locations, anycast helps handle Distributed-Denial-of-Service (DDoS) attacks [23,28].

With an anycast service, each client is associated with one of the anycast sites. For IP anycast, neither clients nor the service explicitly choose this association. Instead, each anycast site announces the same IP prefix and BGP Internet routing selects which site each client reaches, defining that site’s *anycast catchment*. While BGP has flexible policies [3], these do not always optimize for latency or consider server performance [29,19].

Polarization is a pathology that can occur in anycast where all traffic from an ISP ignores nearby, low-latency anycast sites and prefers a distant, high-latency site [2,19,17,24]. Polarization can add 100 ms or more to the round-trip time if traffic goes to a different continent. Polarization may send traffic to a specific

anycast site which may result in an imbalanced load distribution and make that specific site more vulnerable to DDoS attack.

Polarization may happen for different reasons—incomplete transit connections, preference for a specific neighbor, and for unexpected route propagation. A prior study showed peering from the hypergiant picked one global exit from their corporate WAN, and it did not use the many local anycast sites for their .nl TLD [24]. This prior study showed the existence of polarization in two anycast services and explored some root causes, but it did not explore all known *root causes* of polarization, nor does it evaluate these causes over *many services* to understand how widespread this problem is. In this paper, we close this gap by providing a longitudinal analysis of polarization in the anycast services.

This paper makes three contributions. First, we describe two key reasons for polarization in known anycast prefixes and show how they can be observed remotely in real-world anycast prefixes (§3 and §4).

Second, we look for polarization in 7986 known anycast prefixes, finding it in at least 2273 (28.5%) showing that polarization is a common problem (§5.1). We show incomplete connectivity of Tier-1 providers (§5.4) and unwanted route leakage by regional ASes (§5.5) are the key reasons behind polarization. Our measurements show that polarization can have a large latency cost, often adding 100 ms latency for many users (§5). Our analysis only uses known IPv4 anycast prefixes. Our methods of classifying polarization problems apply to IPv6 as well, and IPv6 is an important area of future study because of its growth [16,9] and the potential different routing policies.

Finally, we show how a commercial CDN provider uses traffic engineering techniques to address polarization (§6) for their DNS service. We show small routing changes may significantly improve the overall anycast performance, and that community strings are an essential tool when other routing changes do not work. We demonstrate that simple routing changes can produce a 54% improvement in the mean latency, improving performance for 40% of all clients (§6.3).

Anonymization and Terminology: We anonymize all the names of the anycast service providers to emphasize technical issues rather than provider identity. In this paper, an anycast *service* is an application (such as DNS or web hosting) provided by an anycast *provider* using one or more anycast *prefixes*. A provider may operate multiple services (perhaps DNS for different companies, or optimized to different regions), and each service may be provided by one or more anycast prefixes. Each anycast prefix is typically announced by several sites.

2 Related Work

Extensive prior work has studied, with focuses on anycast topology, efficient routing, performance improvement, and DDoS mitigation.

Anycast topology: Different organizations design their anycast services in different ways. Anycast services have topological differences in number of anycast sites [33], number of providers [21], and regional or global deployment [35]. These topological differences affect the performance of anycast services. In this paper,

we show how topological differences with Tier-1 and regional ASes have impacts on anycast polarization. We show how the same topology with different routing configurations can mitigate polarization problems.

Anycast latency: Providing low latency is one of the goals of anycast services, multiple studies focus on anycast performance [19,33,6]. Prior studies described the importance of number of sites [30], upstream transit providers [21], selection of paths [19], stability in path selection [33], the impacts of polarization [24], and path inflation [17] over anycast latency. Polarization can increase latency for many clients [2,24], a problem sometimes described as path inflation [19]; the cost of poor site selection can be large [1,6,19]. Although a prior study shows the cost of polarization [24] in two services, to our knowledge we are the first to examine polarization across thousands of anycast prefixes.

Anycast performance improvement: Prior studies showed different possible ways to improve the performance of an anycast service. Li et al. proposed to use BGP hints to select the best possible routing path [19]. Removing peering relationship [24] and selective announcement [21] also helped others with performance improvement. In our study, we use multiple traffic engineering techniques to improve the performance of anycast services.

Traffic engineering is used in other previous studies for other purposes like load balancing [27,3,14], traffic shifting [7,32,5], DDoS mitigation [18,28], and for blackhole routing in IXPs and ISPs [10,15]. In our study, we show multiple traffic engineering techniques are required and BGP community strings are essential to improve performance in cases when other traffic engineering methods do not work.

3 Defining Anycast Polarization and its Root Causes

Recent work defined polarization [24], long a bane of anycast services (for example, [2]). We next give our definition of polarization, and add to prior work with a characterization of the two primary root causes of polarization. These steps pave the way for our measurement of polarization (§4) and evaluation of its consequences in the wild (§5).

3.1 Defining Polarization

Polarization is when an anycast client is served by a distant site for some specific prefix, even though there is a nearby site that could provide lower latency. We focus on polarization continent-by-continent. By a distant site, we mean an anycast site that is typically in a different continent, compared to a nearby site on the same continent.

We choose two thresholds to define distant and nearby sites. The thresholds are chosen to show the approximate delay between continents, with T_{lo} at 50 ms defining a nearby site, and T_{hi} of 100 ms to define a distant site.

Lower than T_{lo} latency guarantees that at least some VPs have access to a nearby anycast site, while the high thresholds, T_{hi} shows that other VPs miss this site and reach a distant site. This combination of VPs with lower than

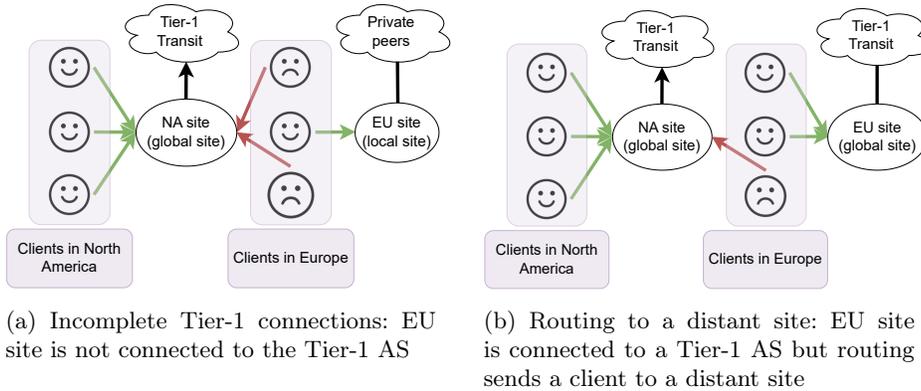


Fig. 1: Two scenarios of multi-pop backbone problems (each face represents a client in a different ASN)

T_{lo} and higher than T_{hi} latency indicates polarization. Latency that is between T_{lo} and T_{hi} may or may not represent a distant site on another continent, and thus we do not consider them for polarization. We also ignore high latency on continents when the anycast prefix has no sites there.

Many anycast sites are *global*, willing to serve any client. However, some larger anycast services have sites that are *local*, where routing is configured to provide service only for users in their host ISP (and sometimes its customers). We ignore classification of sites as local since prefixes that have both local and global sites typically are large enough to have multiple global sites on each continent. We focus on continent-level latency increases of T_{hi} or more, so this simplification has minimal change to our results. (Other work considers “optimal” latency, but we consider differences of a few milliseconds to be operationally unimportant.) In addition, it is often difficult to correctly identify local sites from only third-party observation. Next, we show two different types of polarization problems.

3.2 The Multi-PoP Backbone Problem

The *multi-PoP backbone problem* happens when a backbone network exists, and has points-of-presence (PoPs) in many places around the world, but that backbone forwards all traffic to a limited number of anycast sites that are not geographically distributed. We consider this case as polarization when some clients at least on some parts of the backbone could get lower latency while some other clients have to go to a distant site through the same backbone.

By backbone, we mean both the Tier-1 providers (for example, as reported in customer-cone analysis [4]), and hypergiants (for example, Microsoft or Google), since both operate global backbone networks and have many PoPs. Organizations with large backbone networks often peer in multiple Points of Presence (PoPs) around the globe. They often have many customers, either by providing transit service, or operating large cloud data centers, or both. Connectivity of an

anycast network with these backbones is important since these backbones carry a significant amount of user traffic.

Polarization with multi-PoP backbones can occur for two reasons. First, although both the anycast prefix and backbone have many PoPs, if they share or peer in only a few physical locations, traffic may be forced to travel long distances, creating high latency.

Second, even when backbones and the anycast prefix peer widely, the backbone may choose to route traffic to a single site. This scenario was described when both the Google and Microsoft backbones connected to `.nl` DNS prefixes, with global Google traffic going to Amsterdam and ignoring anycast sites in North America [24].

Figure 1 shows two examples of multi-PoP backbone problems. Often *local anycast sites* are deployed in certain ASes for the benefit of that ISP’s customers. Such sites may serve a few customers of the host AS, or a few of its customers, but if they are not connected to an IXP or a transit provider, they have limited scope. Given no-valley routing policies [13], these sites are not widely visible. These sites will not be the preferred sites, even by the clients on the same continent. Many anycast providers design these sites to serve local clients. Figure 1a shows a multi-PoP backbone problem, where the site in Europe is a local anycast site, meaning it has only a few private peers, is not connected to a popular IXP, and does not have any transit connections. On the other hand, the site in North America is a global site connected to a Tier-1 provider. Since the Tier-1 AS has a missing connectivity in Europe, we call this problem as *incomplete Tier-1 connection*. Due to incomplete Tier-1 connection, some clients from Europe will go to North America (marked by two sad faces), where the Tier-1 AS is connected, resulting in two different latency levels inside Europe.

Multi-PoP backbone polarization due to backbone routing choices is shown in Figure 1b. In this scenario, a Tier-1 provider or hypergiant connects multiple sites in Europe and North America. However, due to routing preference, we can see a client from Europe is going to North American site (marked by a sad face).

3.3 The Leaking Regional Routes Problem

Our second class of polarization problems is *leaking regional routes*. In this scenario, the anycast operator peers with a regional AS at some location, but that peering attracts global traffic and so incurs unnecessarily high latency. By regional ASes, we mean non-Tier-1 ASes that purchase transit from another provider.

This scenario causes polarization because of the *prefer-customer* routing policy common in many ASes. Because the regional AS purchases transit, presumably from a Tier-1 AS, that transit provider will prefer to route to the regional network and its customer, *over* any other anycast sites. Often this preference will influence all customers of the Tier-1, and most Tier-1 ASes have many customers. In addition, by definition, a regional AS has a limited geographic footprint, so its connectivity does not offset the shift of the Tier-1’s customer cone. Thus when anycast prefixes peer with a regional network can have global effects.

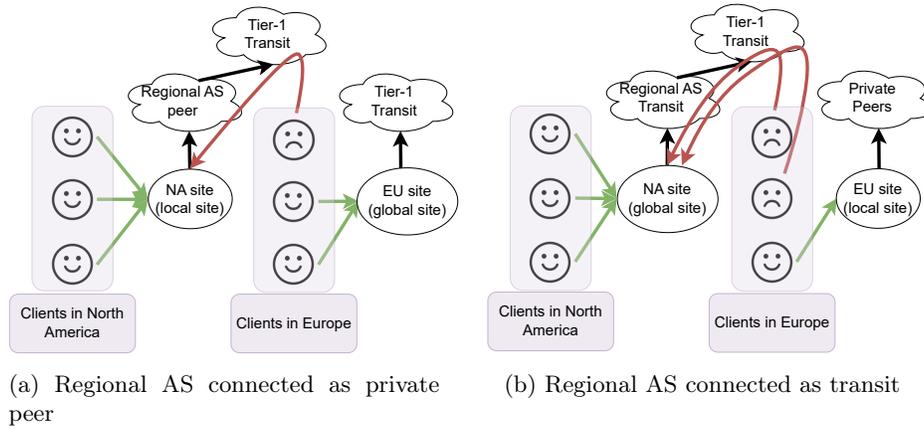


Fig. 2: Regional leakage problem

Regional ASes may be connected to the anycast prefix as a private peer or as a transit provider. As a private peer, the regional ASes are not expected to propagate the anycast prefix to their Tier-1 upstream. However, sometimes a regional AS with private peering may violate this assumption and propagate the anycast prefixes to its Tier-1 transit provider. We can see a route leakage event in [Figure 2a](#) where a regional AS peer propagates routes to its Tier-1 transit and brings traffic from Europe (as shown by the sad client in Europe).

The regional ASes may also serve as the anycast service’s transit provider at this site. As a transit provider, these regional ASes rely on their upstream Tier-1 ASes to propagate their customer prefixes. We can see such polarization in [Figure 2b](#), where a regional AS is connected as a transit in North America and propagates anycast prefix to its upstream Tier-1 provider. The upstream Tier-1 transit is globally well-connected, and attracts traffic from Europe (illustrated by two sad faces in Europe).

4 Detecting and Classing Polarization in the Wild

To meet our goal to study polarization in the wild, we must take third-party observations that can detect polarization and its root causes.

Our measurement approach has three steps: First, we use prior work that identified anycast prefixes [31] to get a list of /24 anycast prefixes to study. Second, we test each /24 anycast prefix for polarization by measuring latency from many locations with RIPE Atlas. Finally, for prefixes that demonstrate polarization, we take traceroutes and use what they find to identify root causes for polarization. [Figure 3](#) shows the steps to find polarization problems and their causes; we confirm specific cases when we can reach operators in [§5](#), and examine how one operator can improve cases in [§6](#).

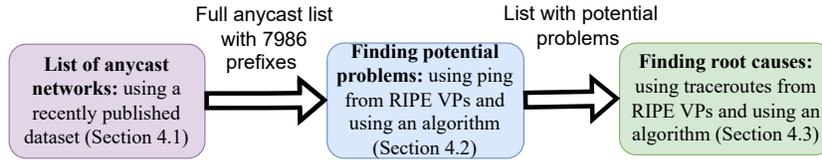


Fig. 3: Steps to find polarization problems and root causes

4.1 Discovering Anycast Prefixes

We first need to find anycast prefixes to search. Fortunately, prior work has developed effective methods to discover anycast prefixes [31]. We directly use the results of their work and begin with their list of 7986 anycast prefixes with 3 to 60 anycast sites distributed around the world.

We evaluate each of these known anycast prefixes in our study. However, we expect that the chance of polarization is low for prefixes with many anycast sites (more than 25); with many sites, there is often one nearby. For prefixes with few sites (less than 5), the sites are sparsely distributed and many clients will observe high latency, even without polarization (§5.3).

4.2 Finding Potential Polarization

To find anycast polarization, we ping from the RIPE Atlas Vantage Points (VPs) to the known anycast prefixes in September 2023 and look for latency variability within the VPs of a continent. We use 100 RIPE Atlas VPs—72 worldwide VPs, and to ensure global coverage, we pick 7 VPs from each Asia, South America, Africa, and Oceania continents to analyze the latency from the RIPE Atlas VPs to the anycast destinations. We selected the country from each continent, and the VPs were assigned randomly. Our selection results in a diverse list of VPs from approximately 50 countries and 85 different ASNs. We create a separate measurement for each anycast prefix with worldwide VPs, which results in 7,986 different PING measurements with different numbers of source countries, and AS numbers.

Using 100 RIPE Atlas VPs, we list potential polarization problems. By using a bigger list of RIPE Atlas VPs, one may observe even more polarization problems (we may miss identifying local sites with a small number of probes, and thus a potential polarization). However, these 100 RIPE Atlas VPs covering all the continents are sufficient for our analysis to determine the causes of polarization.

Based on the ping latency measurements, as we mentioned in §3.1, we detect polarization as when some VPs get good latency ($T_{lo} < 50 ms$) and other nearby VPs see bad latency ($T_{hi} > 100 ms$).

Filtering to reduce false positives: Next, we show how we reduce the number of false positives for more accurate results. We may identify false polarizations because of the regional anycast prefixes—covering only one or two continents like Edgio [35]. Also, the initial anycast prefix list may not be 100% accurate because

of the reassignment of addresses or falsely identified anycast prefixes. To ensure global coverage, we pick VPs from all the continents to verify the prefix is a meaningful anycast prefix to find polarization. We evaluate the anycast prefixes where some VPs from at least three continents get a low latency ($< T_{lo}$). Low latency from at least three continents ensures global coverage of the anycast sites.

We acknowledge that some prefixes may occasionally experience large latency due to either polarization or other network situations. However, prior studies showed anycast catchment normally remains stable [28]. Also, we took ping and traceroutes measurements on different days.

4.3 Finding Root Causes

Next, we describe the measurement methods that indicate each of the root causes for polarization that we identified in §3. We take traceroutes to each anycast prefixes with potential polarization and examine the penultimate AS hop as seen from different VPs.

Finding penultimate AS hop and its type: At first, we find out the penultimate AS hop in the AS path to the destination. We observe the whole AS path, and pick the AS that is present just before the anycast prefix’s AS. Multiple routers before the final destination may represent the same penultimate AS hop.

After getting the penultimate AS hop, we determine the type of the penultimate AS hop. We consider CAIDA top 10 ASes as the Tier-1 ASes (AS3356: Lumen/Level-3, AS1299: Arelion/Telia, AS174: Cogent, AS6762: Telecom Italia, AS2914: NTT America, AS6939: Hurricane Electric, AS6461: Zayo Bandwidth, AS6453: TATA Communications, AS3257: GTT Communications, and AS3491: PCCW Global) [4], and the others are regional ASes. Hypergiants have heavy outbound traffic. We consider AS15169: Google, AS8075: Microsoft, AS40027: Netflix, AS16509: Amazon or AS32934: Facebook [25] as hypergiants. We classify them by their AS numbers.

Finding multi-pop backbone problems: Multi-pop backbone problems happen when Tier-1 providers or hypergiants are either partially connected or route traffic to a distant site. We identify the partial connectivity when we find VPs with Tier-1 AS or hypergiant in the penultimate AS hop in their paths to a distant site, and when no other VPs from the same continent go to a nearby site using the same penultimate AS. To find routing problems to a distant site, we check for VPs from the same continent that get both good and poor latency with the same Tier-1 AS or hypergiants in the AS path. We use our latency thresholds (§4.2) to understand whether VPs are going to a nearby site or a distant cross-continent site. This method cannot find poor routing instances when all the VPs are going to a distant site even though they have nearby sites. These nearby sites cannot be identified by the traceroutes since their routing sends all the VPs to the distant site. We must talk to the operators to learn such poor routing cases.

Finding leaking regional problem: We identify a leaking regional problem when a regional AS with a smaller customer cone attracts a lot of traffic to a specific anycast location.

The regional AS can be connected as a transit or as a private peer that is leaking routes to its upstreams. To identify private peer leaking routes, we search for other Tier-1 ASes that are present in the penultimate hops and connected in multiple locations. The existence of other penultimate Tier-1 ASes indicates the possible real transit providers. When there is another Tier-1 transit, the regional AS should emphasize on the regional propagation of the announcements. To find regional ASes that are connected as transits, we check for other Tier-1 ASes in multiple locations. If we find other Tier-1 ASes, then the regional AS is possibly connected as a private peer but leaking routes to its upstream.

This approach can predict *possible* regional route leaking, but it may also happen that these regional ASes are connected as transit providers. Knowing the exact peering relationship is only possible when we can talk to the operators as the inferred peering relationships from the public databases are not always accurate, and corner cases like route leakage may infer wrong peering relationships [20]. Hence, to validate actual route leakage we must talk to the operators.

4.4 Finding Impacts

After getting the polarization problems and their root causes, we want to see the impacts of polarization. We find out the penultimate AS hop that is common in the paths to the distant site. We measure the median and 95th percentile latency from a continent. Polarization results in inter-continental traffic, and its impact is expected to show in high 95th percentile latency. Using the difference between the median and 95th percentile latency, we show the impact of polarization.

Next, we show that polarization problems are common and they have a significant impact over anycast performance.

5 Measurement Results and Impacts of Polarization

We next show how common polarization is in known anycast prefixes, and how polarization affects service performance.

We validate results for all prefixes with operators we could reach.

5.1 Detecting Polarization in Anycast Prefixes

Our first goal is to understand how common polarization is. Following our methodology in §4, we first examine all known anycast prefixes for polarization. We see that of the 7986 examined prefixes, about 28% show potential polarization (Table 1).

In this paper, we focus on the 626 anycast prefixes that show polarization problems for clients in Europe and North America. We focus on these anycast prefixes to detect root causes because these continents have mature Internet

service and multiple IXPs, so they are locations where polarization can occur and we believe we have cases where actions can be taken to resolve the issue. We leave the study of polarization on other continents as future work, since polarization that occurs due to incomplete in-country peering will be addressed as the domestic Internet market matures and interconnects. (An anycast prefix that peers with one AS cannot avoid polarization with other ISPs in the same country if there is incomplete domestic peering.)

Since we only focus on the EU and NA continents to find the root causes behind polarization using traceroutes (§4.3), we select 46 European VPs and 20 North American VPs. We choose 2-3 VPs each from 20 different European countries to make 46 European VPs, and we select 15 VPs from the USA and 5 VPs from Canada to make 20 North American VPs. These probes are chosen randomly from a specific country (we select the country and probes are chosen randomly), and they cover approximately 55 different ASes from these countries. We take traceroute measurements from these VPs to each of 626 anycast prefixes. To find the catchment, we use the penultimate routing hop’s geo-location reported by RIPE or a CDN’s internal tool for IP to geo translation. This location may not be the true catchment, but this location is sufficient to find out the cross-continent traffic to the destination. We confirm the catchments of 18 anycast prefixes where we had the chance to contact the operators.

5.2 Detecting Root Causes

Given potential polarization in known anycast prefixes, we then apply root cause detection (§3) to these prefixes.

We find multi-pop backbone problems in 376 anycast /24 prefixes out of 626 prefixes (Table 1). We observe a Tier-1 AS in the penultimate AS hop for these 376 anycast prefixes. Among these multi-pop backbone problems, our methodology finds 218 instances where a Tier-1 provider is incompletely connected. We suspect this behavior when nearly all VPs from a continent have catchment in a distant site through the same Tier-1 AS. We also find 158 cases when VPs route to a distant site. We suspect this event when some VPs from a continent experience good latency (<50 ms) while others observe poor latency (>100 ms), and when these VPs utilize the same Tier-1 AS in the penultimate AS hop.

Our methodology shows 233 cases where a regional AS leaks routes. Among these, in 177 cases we find other Tier-1 ASes in the penultimate AS hop. As there are other Tier-1 ASes in the path, we suspect these Tier-1 ASes are the real transits, and the regional ASes are possibly leaking routes. In 56 other instances, we find no other Tier-1 ASes in the path. We suspect a regional AS is connected as a transit in these 56 cases.

We contacted the operators of 18 of these 626 cases. The operators confirmed all these polarization events.

An anycast service provider may have multiple /24 prefixes, and because of their topological similarity we find polarization problems in many of their /24 prefixes. Table 2 shows top providers who have polarization problems in many of their /24 anycast prefixes. We can see some providers have polarization in almost all of their anycast /24 prefixes.

Category	Count	%	Confirmed
Known anycast prefixes	7986	100	
No observed potential polarization	5713	72	
Potential polarization	2273	28	
In continents outside EU and NA	1647	20	
In EU and NA	626	8	18
No class found (a)	161	2	
Only multi-pop backbone problem (b)	232	3	
Only regional leakage (c)	89	1	
Both classes (d)	144	2	
Multi-pop backbone problem (b+d)	376	5	9
Incomplete Tier-1 connections	218	3	9
Routing to a distant site	158	2	0
Leaking regional problem (c+d)	233	3	9
Leakage by regional	177	2	6
Leakage by regional transits	56	1	3

Table 1: Detected polarization and inferred root causes.

Provider	Potential problems	Total anycast
Anon-DNS-2	214	216
Anon-DNS-3	94	159
Anon-CDN-1	20	47
Anon-DNS-4	12	75
Anon-DNS-5	9	9

Table 2: Top anycast prefixes with potential polarization problems

5.3 Impacts of the Number of Sites on Polarization

Does polarization correlate with the total number of anycast sites? Prior work [30] suggested that 12 sites can provide good geographic latency, but those results assume good in-continent routing. Does that assumption hold in practice?

To answer this question, we explore the relationship between the number of anycast sites and the degree of polarization. To get the number of anycast sites, we utilize the count reported by the recent study [31] that we used to get the anycast prefixes.

Figure 4 shows how much polarization occurs relative to the number of anycast sites. We group anycast prefixes by number of sites into bins of 5 (so the first bin is 5 sites or less, the next is 6 to 10, etc.). The number of sites in each part of the graph varies and is shown on the top of each bin, but is always at least 52 prefixes, and often hundreds. For each bin, we show the percentage of prefixes that appear polarized.

We see some polarization in prefixes of many different number of sites. But for Europe and North America, polarization is most common in prefixes with 30 or fewer sites (the left part of Figure 4a). For other continents, some polarization

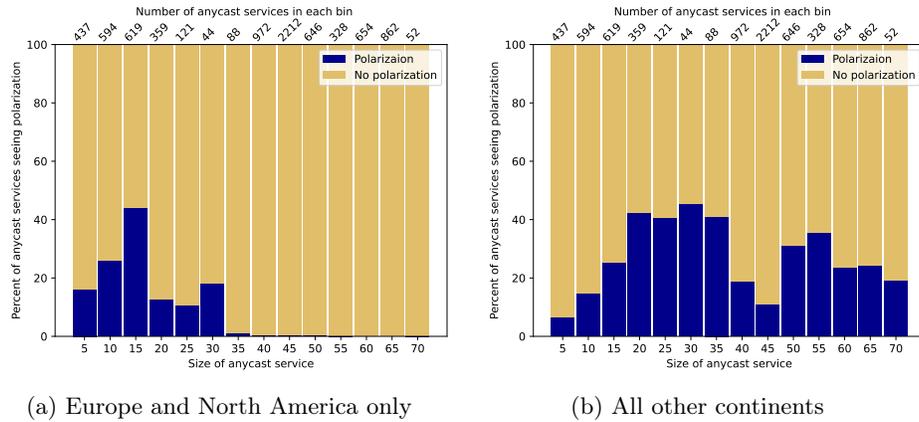


Fig. 4: Percent of anycast prefixes that see polarization, grouped in bins by number of anycast sites per anycast prefix.

occurs regardless of how many sites the prefix has. In [Figure 4b](#), we can see polarization even for prefixes with 30 to 70 sites.

We conclude that prefixes with many sites generally get good routing in continents with mature Internet markets. However, routing is more challenging, as shown by greater polarization, for prefixes with only a few sites, and when operating globally. In mature Internet markets, high rates of AS interconnectivity decrease risks of polarization. However, outside EU and NA, anycast is more difficult to deploy well, likely because of poor local inter-AS connectivity means some local customers cannot use a local anycast site.

Next, we show examples of such connectivity issues that cause polarization and their impacts on latency.

5.4 Impacts of multi-pop backbone problems

We next look at the problem of multi-PoP backbones ([§3.2](#)). We examine several examples of this in real-world anycast prefixes in [Table 3](#) and show how it can result in extra latency of 100 ms or more due to inter-continental traffic (shown in [Figure 5](#)). We show the number of /24 prefixes for each of the anycast services, however, the problems that we show in [Table 3](#) demonstrate a problem for a specific /24 prefix of that anycast service.

5.4.1 An incomplete Tier-1 connection in Anon-CDN-2 We find an anycast prefix of Anon-CDN-2 where a fraction of traffic from Europe was going to the US. We believe this event is an example of polarization for incomplete Tier-1 connections. We find 40 out of 46 VPs (87%) of the European VPs remain within the continent, and 6 out of 46 VPs (13%) of the other VPs end up in San Jose, USA using a fixed Tier-1 AS.

We find European sites are globally well connected. We observe a Tier-1 AS (AS3356 - Level-3) connected to a European site. We believe the European site is

Provider	No. of /24 prefixes	Reason of the problem	Common AS to the distant site	Source cont.	Med.	95th	Example
Anon-CDN-2	14	Incomplete Tier-1	AS1299	EU NA	12 25	173 88	Poland to USA
Anon-Cloud-1	1	Incomplete Tier-1	AS3356, AS6939	EU NA	51 55	158 162	Canada to Germany
Anon-CDN-6	67	Incomplete Tier-1 (peers as transits)	AS6762	EU NA	16 39	122 141	Greece to USA
Anon-CDN-3	9	Exceptional incomplete Tier-1	AS6453	EU NA	32 25	113 39	Germany to USA
Anon-CDN-1	67	Leaking regional transit	AS1273	EU NA	21 88	53 150	USA to UK
Anon-DNS-1	18	Leaking regional peer	AS4826	EU NA	174 37	314 214	USA to Australia
Anon-CDN-5	67	Leaking regional peer	AS7473	EU NA	39 24	248 252	USA to Singapore
Anon-CDN-7	67	Leaking regional peer (merging org)	AS209	EU NA	29 26	84 65	Finland to USA
Anon-DNS-2	216	Leaking regional Incomplete Tier-1	AS4637 AS1299	EU NA	165 81	301 254	Canada to Tokyo

Table 3: Polarization in real-world anycast prefixes

a global site attracting VPs from many locations. As proof, we find that African VPs have catchments in Europe. European sites also increase their connectivity by having many private peers connecting most small ASes within the continent. Even with this well-connectivity, 13% of the European VPs have catchment in the US. As a result, we believe this is an example of polarization.

The VPs that are going to North America have AS1299 (Arelion Sweden AB) in common within their paths. We believe AS1299 is working as a transit although we do not know the contract type between Anon-CDN-2 and AS1299 for that anycast site. Based on the other traceroutes, we did not find any path that remains within Europe through AS1299. That means AS1299 has a missing connection in Europe. Figure 5a shows NA site is incompletely connected to AS1299 and causes polarization for European clients.

Impacts: The impact of this cross-continent traffic has a significant impact over latency. As an example, we find a VP from Poland goes to San Jose, USA. The cross-continent traffic results in a high 95th percentile latency. While the median is only 12 ms, we get the 95th percentile latency is 173 ms (Table 3).

5.4.2 Multiple incomplete Tier-1 in Anon-Cloud-1 Next, we show an example from Anon-Cloud-1 where an anycast prefix uses two different Tier-1 ASes in two different continents.

We find two global sites in Europe and North America connected through two Tier-1 ASes as shown in Figure 5b. We observe traffic going to Dallas, USA using AS3356 (Level-3), and to Frankfurt, Germany through AS6939 (Hurricane Electric). Both these ASes are Tier-1 ASes and have a wide range of connectivity. We are unsure about their contract type with Anon-Cloud-1. They can be

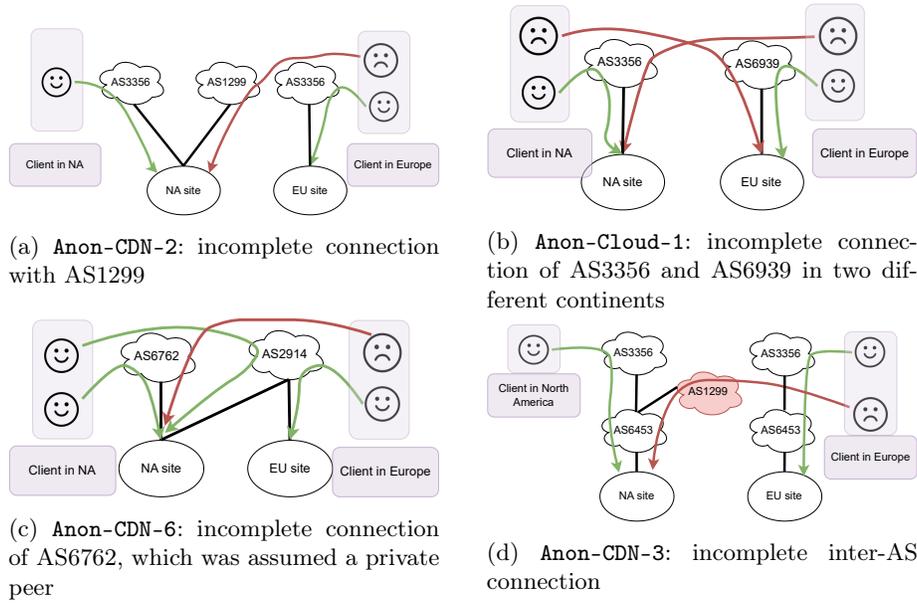


Fig. 5: Real-world multi-pop backbone problems

incomplete transit connections, or one of these Tier-1 ASes has a private peering relationship with **Anon-Cloud-1**, but works like a transit.

The traceroutes exhibit the preferred route to reach an anycast site. Of course, a client may have multiple paths to reach European and North American sites via two Tier-1 ASes. However, going to a distant site through a specific Tier-1 AS means that the client has a preference for that particular AS, and that AS may not have any connectivity within the client’s continent.

Impacts: When we have two big Tier-1 ASes have incomplete transit connections in two continents, we observe a significant fraction of cross-continent traffic. We observe 40% VPs from Europe and North America end up in a different continent. The cross-continent traffic using AS3356, and AS6939 add over 100 ms of latency. The cross-continent traffic makes 3×95^{th} percentile latency compared to the median latency (Table 3). This example shows how bad the impact can be when we have two Tier-1 providers connected incompletely in two continents. As an example of cross-continent traffic, we observe traffic from Denmark goes to Dallas, USA, and traffic from Canada goes to Germany.

5.4.3 Incomplete Tier-1: peers working as transits in Anon-CDN-6 We find another case where a Tier-1 AS has incomplete connections with **Anon-CDN-6**. We contacted the operators, and they confirmed their peering relationship with this Tier-1 AS. But that Tier-1 AS has a huge customer cone, and their policy is to propagate the routes to their global customer cone.

In this event, `Anon-CDN-6` has a site in Miami, USA where the Miami location is connected to AS6762 (Telecom Italia) as a private peer. The operators confirmed that AS2914 (NTT America) is the real transit provider since it is connected to three continents—Asia, Europe, and North America. AS6762 is only connected in one location. The anycast operators did not expect the prefix to be propagated out of the continent by AS6762 since it is connected as a private peer. But in reality, it was propagated to other ASes connected in other continents. From [Figure 5c](#), we can see that a client from Europe has catchment in North America via AS6762.

Impacts: AS6762 propagates the anycast prefix out of the continent. As a result, we found 5 out of 46 VPs (10%) from Europe have catchment in Miami, USA, and experienced over 100 ms of extra latency.

5.4.4 Exceptional incomplete Tier-1: incomplete inter-AS connections We find a case when we observe a polarization incident even with complete Tier-1 connections. We find this case using manual observation of the traceroutes when we could not find a proper classification of the polarization problem.

Traceroutes to `Anon-CDN-3` show that AS6453 (Tata Communications) as a penultimate AS hop in many different paths to Europe and North America. This behavior proves AS6453 is working as a transit connected in both Europe and North America with `Anon-CDN-3` ([Figure 5d](#)). So, we define this connectivity as a complete transit connection. Even with this complete connection, we observe some European VPs have catchments in North America, which results in polarization and high latency.

Many traceroutes show Tier-1 ASes like AS3356 and AS1299 just before the penultimate hop AS6453, as we can see from [Figure 5d](#). The presence of these Tier-1s two hops back from the anycast site suggests that this anycast prefix mostly relies on the Tier-1 ASes to get into AS6453 and then to `Anon-CDN-3`. We find the VPs from Europe with AS1299 in the paths end up in North America through AS6453. We suspect AS1299 is connected to AS6453 only in North America, or for `Anon-CDN-3`, AS1299 has a preference for the North American connection. On the contrary, we observe the opposite for AS3356. AS3356 has connectivity with AS6453 in both continents and as a result, we are not observing any performance issues.

Impacts: We find 8 out of 46 VPs (18%) of the European VPs choose a path through AS1299 to go to North America to connect to AS6453 and `Anon-CDN-3`. While the median and 95th percentile latency have a small difference for the North American VPs; because of this polarization, European VPs observe over $3\times$ 95th percentile latency ([Table 3](#)).

5.5 Impacts of Leaking Regional Problems

We next turn to leaking regional ([§3.3](#)), our second class of polarization problems. We find several cases where regional ASes (non-Tier-1 ASes) send a great deal of traffic to a distant site. These regional ASes purchase transit from a Tier-1, and so polarization results when their transit-providing Tier-1 adopts a prefer-client routing policy. Alternatively, these regional ASes are private peers but

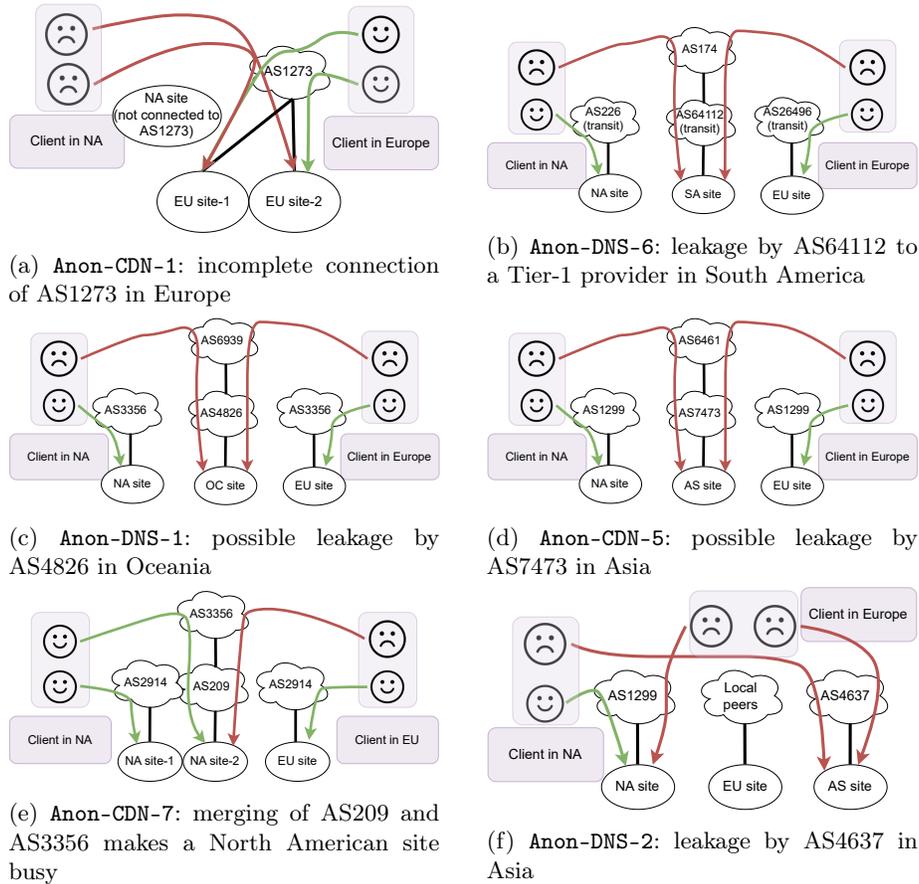


Fig. 6: Real-world leaking regional problems

make unwanted route propagation to their upstream. **Figure 6** shows the leaking regional problems.

5.5.1 Anon-CDN-1: Leaking by regional transits with multiple connectivity We find a polarization instance because of the leaking of an anycast prefix by a regional AS in one of **Anon-CDN-1** anycast prefixes. We confirmed this event with the anycast operators. The operators informed us that the regional AS is connected as a transit provider.

In this polarization problem, we find a significant portion of traffic from North America ends up in Europe using AS1273 (Vodafone). Contacting **Anon-CDN-1** operators about this issue, we confirmed that AS1273 is indeed connected as a transit for their anycast prefix. We did not find any VPs having AS1273 in the path that stays within North America (**Figure 6a**). We believe AS1273 is only connected in Europe, not in North America, resulting in incomplete re-

gional transit connections and polarization. We also confirmed this finding with `Anon-CDN-1`, and they informed us that AS1273 is connected to multiple countries in Europe.

Impacts: Since AS1273 is only connected in Europe as a transit, we find around 13 out of 20 VPs (64%) of the North American VPs end up in Europe. Only 36% VPs stay within North America. As an example of cross-continent traffic, we find a VP from the USA that goes to London, UK. As a result, we find both median and 95th percentile latency are high for this polarization ([Table 3](#)).

5.5.2 Anon-DNS-6: Leaking by a regional transit in South America A global DNS provider (`Anon-DNS-6`) peers in several locations. A South American site peers with a regional network (AS64112, PIT-Chile) who purchases transit from a Tier-1 provider (AS174, Cogent). Because of AS174’s prefer-customer routing policy, peering with the regional network causes all customers of the Tier-1 provider in North America and Europe to go to South America ([Figure 6b](#)).

Impact: Cross-continent routing adds 100 ms or more latency, so `Anon-DNS-6` instead prevents route announcements to this Tier-1 from this site. With limited routing, this anycast site is unavailable to other regional networks in South America that can be reached via transit.

Better solutions to this problem is for `Anon-DNS-6` to peer with all regional networks, or that they influence routing inside the Tier-1 AS, neither of which is easy.

5.5.3 Anon-DNS-1: Possible route leakage by regional AS in Oceania We identify polarization in `Anon-DNS-1` anycast prefix due to route leakage by a regional peer.

We suspect this anycast prefix has route leakage because many VPs have a non-Tier-1 AS (AS4826: Vocus Connect, Australia) as the penultimate AS hop in all the paths to the distant site, and some other VPs have a Tier-1 AS (AS3356) in multiple AS paths ([Figure 6c](#)). Our assumption is that AS3356 is the actual transit, and AS4826 is connected as a peer but it does not work as a peer, rather it leaks `Anon-DNS-1` routes to its peers and transits. We find AS4826 propagates its routes to many different peers including a Tier-1 AS (AS6939). Since AS6939 is well-connected to the rest of the world, it brings a large fraction of VPs to AS4826.

We must talk to the operators to know whether AS4826 is really a private peer or a transit provider. However, having a transit connection with a smaller AS is highly unlikely since this anycast network (/24 prefix) has connectivity with another Tier-1 AS with a large geographic presence.

Impacts: We observe severe polarization due to this route leakage by AS4826, where AS4826 is only connected in Australia (as shown in [Figure 6c](#)). Since AS4826 is propagating its routes to AS6939, and since AS6939 is a well-connected Tier-1 AS, traffic from all over the world ends up in Australia. We find a case where traffic from the Ashburn, USA connects to AS6939 in Los Angeles, USA, then travels to AS4826 in San Jose, USA, and then ends up in Sydney, Australia,

resulting in over 200 ms of latency. From Europe, we find a case when a VP from Switzerland travels to Ashburn, USA, Los Angeles, USA, San Jose, USA, and then Sydney, Australia, which takes over 300 ms latency.

5.5.4 Anon-CDN-5: Route leakage by a regional AS in Asia We have already shown how a possible route leakage brings traffic to South America and Oceania. Next, we show a similar case in Asia with a different autonomous system.

We find another polarization event due to route leakage by AS7473 (Singapore Telecom). We confirm this event with Anon-CDN-5 that AS7473 is connected with Anon-CDN-5 as a private peer but it propagates Anon-CDN-5 prefix to a big Tier-1 AS (AS6461 - Zayo Bandwidth) as shown in Figure 6d. Since AS6461 is connected heavily with the rest of the world, we find many VPs from Europe and North America end up in Singapore. From the traceroutes, we find a Tier-1 AS (AS1299) is the real transit connected to different locations worldwide.

Impacts: We find 7% VPs from Europe and 6% VPs from North America end up in Singapore, even if this anycast network has multiple presence in those continents. While VPs that stay within the continent observe low latency, due to this long path to Singapore, we observe over 200 ms difference between the median and 95th percentile latency (Table 3).

5.5.5 Regional route leakage: a special case when organizations merge

We find a polarization problem in one of CDN prefixes (Anon-CDN-7) due to the merging of two organizations.

We find this case by looking over the traceroutes from the VPs to a distant site for a specific prefix of Anon-CDN-7. We find “AS3356 AS209” in many of the paths that show bad latency to a distant site (Figure 6e). In that particular anycast prefix, AS209 (Century Link) is connected only in one location. Since AS209 propagates its routes to AS3356 (Level-3), and since AS3356 is a big Tier-1 AS with global connectivity, we observe a significant fraction of cross-continent traffic to the distant site where AS209 is connected. Contacting the CDN, we learn that AS209 is connected as a private peer, and so it should not propagate the CDN prefix to a big Tier-1 provider. We *suspect* that this issue occurs because of the merging of AS209 and AS3356 [22]. We confirmed this incident with the operators of Anon-CDN-7.

Impacts: Anon-CDN-7 network (announced by a /24 prefix) peers with AS209 in Sterling, USA location. Since AS209 propagates its route to AS3356, we find VPs from Europe and even from Asia have catchment in the USA (Figure 6e), resulting in over 200 ms of latency.

5.6 Combination of Problems

We find several Anon-DNS-2 prefixes where we believe multiple connectivity problems exist that cause severe polarization.

Anon-DNS-2 has multiple anycast prefixes that are announced from different locations connected to different transits and peers. We are certain that Anon-DNS-2 has anycast sites at least in North America, Europe, and Asia since

some of the VPs from these continents experienced good latency. However, the big Tier-1 ASes like AS1299 or AS174 is only connected in North America. Based on our traceroutes, we also find European and Asian sites have peers that are not big Tier-1 ASes. With a connectivity like this, we expect most North American VPs will stay within North America because of the Tier-1 AS, and some European and Asian VPs will go to North America because of the incomplete transit connectivity.

In reality, we observe cross-continent traffic in different directions. We find the peers (AS4637 Telstra Global) in Asia leaks routes (or working as a transit) to its Tier-1 providers. Since that Tier-1 provider is well-connected, we find North American traffic goes to Asia (left side unhappy user goes to Asia in [Figure 6f](#)). We also find Asian VPs going to North America since North American sites are connected to Tier-1 ASes. We find only a few European VPs stay within the continent because of their local peers, but a significant portion moves to North America and Asia (shown in [Figure 6f](#) by two unhappy European users). This cross-continent traffic results in increased latency, in some cases they add more than 200 ms of latency.

6 Improvement by Routing Nudges

We already show polarization exists in many different anycast prefixes and that it can have a significant impact on performance. However, often small changes in routing can address polarization, even without adding new peering relationships with other ASes. We find many potential polarization problems as mentioned in [Table 1](#). In this section, we show two examples of how `Anon-CDN-4` improves anycast performance in two anycast systems. While we would love to have improvements for more anycast systems, we could only obtain before-and-after results for changes by the operators of this CDN.

6.1 Anycast Configuration

`Anon-CDN-4` uses anycast for their DNS services. To ensure reliability, `Anon-CDN-4` uses multiple anycast /24 prefixes for DNS. These anycast prefixes are announced from around 268 sites located in 93 cities distributed around the world. Each site has multiple machines to serve the client load. Each site has different upstream connectivity, with different numbers of peers, network access points, and transit providers. `Anon-CDN-4` uses multiple Tier-1 ASes as the transits for their anycast prefixes (/24s). These transits differ in each anycast prefix.

We show improvements in two anycast prefixes after making routing changes by `Anon-CDN-4`. The first anycast prefix has a presence in 15 cities covering 9 countries. AS2914 (NTT America) provides transit connectivity and is connected in multiple geographic locations covering Asia, Europe, and North America. There are other peers and network access points connected the anycast sites. The second anycast prefix covers 17 cities in 11 countries. AS1299 (Arelion Sweden) provides transit for this anycast prefix covering Europe and North America continents.

6.2 Routing problems

The two anycast prefixes mentioned in §6.1 have different routing problems. In the first anycast prefix, `Anon-CDN-4` encounters two different problems: multi-pop backbone problems (§3.2), and leaking regional problems (§3.3). The second anycast prefix only has a leaking regional problem (§3.3).

First anycast prefix has both multi-pop backbone and leaking regional problems: The first anycast prefix has multi-pop backbone problems with multiple Tier-1 ASes. `Anon-CDN-4` connects to AS1299 (Arelion Sweden) and AS3356 (Level-3) as private peers in Dallas, USA. We also find AS6762 (Telecom Italia) as a private peer connected in Virginia, USA, and Milan, Italy. `Anon-CDN-4` operators expected these peers to confine their announcements within a smaller customer cone. But we find that these peers propagate `Anon-CDN-4` routes to the rest of the world. We suspect these peers treat `Anon-CDN-4` as their customers since they are also connected as a transit for other anycast prefixes of the same anycast service.

In the second problem of the first anycast prefix, `Anon-CDN-4` connects to AS209 (Lumen) as a private peer in Virginia, USA location, and propagates routes to AS3356. Since AS3356 is well-connected to the rest of the world, `Anon-CDN-4` observes cross-continent traffic to Virginia, USA.

Second anycast prefix has a leaking regional problem: In the second problem, `Anon-CDN-4` has a regional private peer (AS7473) connected in Singapore, leaks their routes to other upstream Tier-1 ASes. As a result, `Anon-CDN-4` observes cross-continent traffic from other continents to Singapore.

6.3 Solving Problems

`Anon-CDN-4` solves these performance issues by changing their routing configuration.

Solving two problems in the first anycast prefix: To solve the two problems in the first anycast prefix, `Anon-CDN-4` stops announcing to the peers that were causing the polarization problem. `Anon-CDN-4` blocks announcements to each of the Tier-1 private peers (AS1299, AS3356, and AS6762), and to AS209 to prevent the propagation of routes to AS3356.

Solving leaking regional problem in the second anycast prefix: For the second anycast prefix, `Anon-CDN-4` tries two things. Since many local VPs were taking advantage by using AS7473, `Anon-CDN-4` realized that blocking AS7473 may result in even worse performance overall. That is why, instead of blocking the announcement to AS7473, `Anon-CDN-4` takes a more cautious action. In one change, they prepend twice from Singapore location so that fewer VPs end up in Singapore. In another change, they use community strings to tell AS7473 to keep the `Anon-CDN-4` prefix within Asia and Oceania regions. Only the second change results in better performance overall which we will describe next. This example shows the importance of having multiple routing configurations for traffic engineering, and using the one that results in best performance.

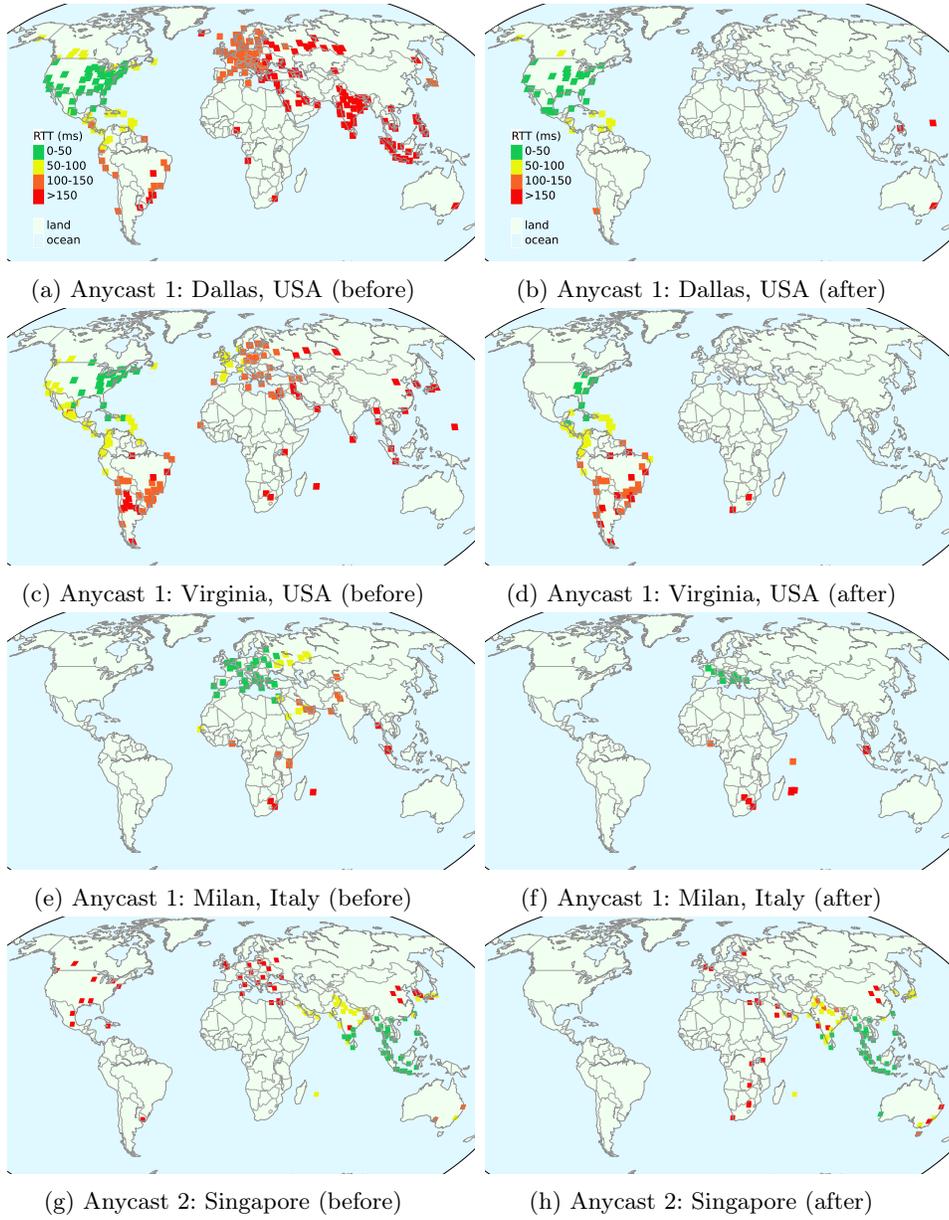


Fig. 7: Changes in Anycast catchment for a anycast site due to a routing change

Measuring performance, before and after: We use Anon-CDN-4’s internal measurement system, observing latency from about 2300 global VPs. These vantage points have a global coverage with 74 African, 839 Asian, 772 European, 315 North American, 90 Oceanian, and 171 South American VPs. We do not use RIPE Atlas VPs from this point onwards.

6.3.1 Changes in the catchments After making the routing changes, Anon-CDN-4 observes significant changes in the catchment distribution. We show the catchment distribution based on the Anon-CDN-4’s internal measurement of the catchments. We examine the VPs going to an anycast site before and after the routing change. In [Figure 7](#), we visualize the geolocations of the VPs to an anycast site, along with the measured latency. Each point on the map represents a VP and the colors represent the latency level from that VP to the anycast site.

Catchment changes in the first anycast prefix: For the first one, Anon-CDN-4 blocks their private peering with AS1299 and AS3356 from Dallas.

The topmost left graph ([Figure 7a](#)) shows the catchment before making the change. As we can see many VPs from Europe and Asia (shown by red dots) end up in Dallas, USA. As a result, these VPs experience bad latency (over 100 ms). Traceroutes confirm cross-continent VPs using AS1299 and AS3356 to reach Dallas, USA. After blocking the announcement to AS1299 and AS3356, we observe no cross-continent VPs from Europe and Asia ([Figure 7b](#)). We can only see VPs from the US have catchment in Dallas, USA, and experience better latency (less than 50 ms).

Anon-CDN-4 also blocks private peers AS209 and AS6762 from Virginia, USA. With these two ASes, Virginia, USA was receiving traffic from different continents ([Figure 7c](#)). After blocking the announcement, Virginia, USA site receives traffic mostly from North and South America ([Figure 7d](#)). We also observe a less number of cross-continent VPs when we block AS6762 from Milan, Italy ([Figure 7e](#) and [Figure 7f](#)).

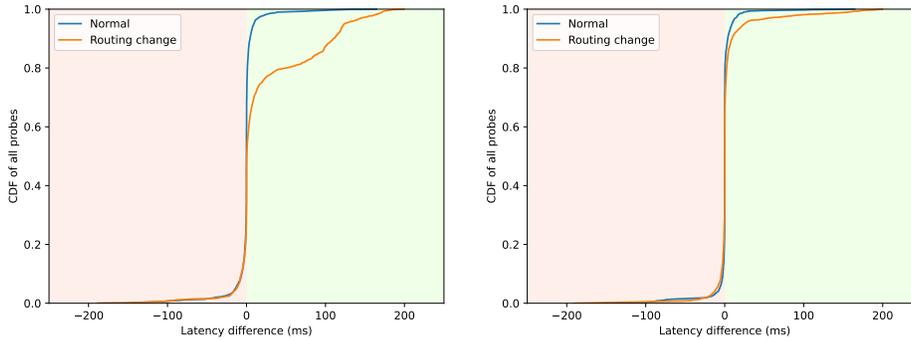
Catchment changes in the second anycast prefix: In the second anycast prefix, the Singapore site was receiving traffic from other continents ([Figure 7g](#)). Anon-CDN-4 uses community strings to keep the announcement propagation within the Asia and Oceania continents. As a result, we can see less number of VPs to Singapore from other continents ([Figure 7h](#)).

6.3.2 Impacts over Performance We have shown above the catchment changes after deploying the new routing configurations. Next, we show how the performance changes.

Improvement in the first anycast prefix: After the changes in the first anycast prefix, we find most improvement in the Europe, Asia, and North America continents ([Table 4](#)). We find 54.6%, 23.4%, and 23.0% improvement in mean latency in these continents, respectively. Many European and Asian VPs were going to Dallas, USA, and Virginia, USA. After the new announcement, we do not see this cross-continent traffic. Even if we block announcement from the

Anycast prefix	Routing changes	Improvement(%)					
		Africa	Asia	Europe	North America	Oceania	South America
Prefix-1	Blocking AS1299 and AS3356 from Dallas, USA, AS209 and AS6762 from Sterling, USA, and AS6762 from Milan, Italy	5.4	23.4	54.6	23.0	2.2	-0.15
Prefix-2	Announcement to AS7473 only within Asia Pacific and block others (using community strings)	10.5	12.3	34.8	19.7	2.3	1.9
Prefix-2	Announcement to AS7473 with two prepending	3.8	-9.6	6.5	6.4	6.3	3.8

Table 4: Continent-wise improvement in latency by routing changes



(a) Prefix-1: blocking in Dallas, USA, Virginia, USA, and Milan, Italy

(b) Prefix-2: changes in Singapore - announcing only within Asia Pacific

Fig. 8: CDF of all the VPs wrt latency difference (ms)

North American sites, we observe 23.0% improvement among the North American VPs. This is because many North American VPs from the West Coast had catchments to the Dallas and Virginia sites. After the routing change, the traffic starts going to their nearby locations. The performance in the other continents remains mostly stable.

Improvement in the second anycast prefix: After the change in the second anycast prefix using community strings, we find the most improvement in the European and North American VPs. From Table 4, we can see that European VPs observe 34.8% improvement, and North American VPs observe 19.7% improvement. This is because several European and North American VPs had catchments in Singapore. After the new announcement, we do not observe this

type of cross-continent traffic. We observe improvement in other continents as well.

New latency distribution in both anycast prefixes: Since Anon-CDN-4 blocks routing announcements to different peers in different geo-locations, some VPs may observe worse performance who are dependent on the peers that Anon-CDN-4 blocks. However, if there are nearby sites, then the VPs may be redirected to the nearby sites after the routing changes, and still observe good latency. To know how many VPs are getting worse latency after the routing change, we show a CDF graph with respect to latency difference in Figure 8. We measure the latency decrease for each VP after the routing change. A positive difference indicates an improved performance (light green region in Figure 8), and a negative difference indicates a degraded performance (light red region in Figure 8). Since latency may vary slightly between two measurements, the graphs show a normal difference line (blue lines) to show the regular latency variance without any routing changes. The orange lines show the latency decrease after the routing changes.

In the first anycast prefix, we can see 40% of the VPs get lower latency after the routing change (Figure 8a). On the other side, we can see the blue and orange lines overlap, which indicates no significant number of VPs gets worse latency. For the second anycast prefix, we can see around 15% of the VPs observe lower latency (Figure 8b), when most other VPs observe regular differences (blue and orange lines overlap).

6.3.3 Community strings are important Anon-CDN-4 also attempts to use path prepending at their Singapore location to stop getting cross-continent traffic to Singapore. (Since many local VPs are dependent on AS7473 to reach Singapore site, Anon-CDN-4 did not want to fully block the announcement through AS7473 fully.) Table 4 shows the outcome after they use path prepending for two times. Even after twice prepending from Singapore, they could only reduce 6.5% mean latency in Europe, and 6.4% mean latency in North America. At the same time, the mean latency becomes 9.6% worse in Asia. Path prepending is an available tool for traffic engineering—anycast operators can make path prepending without requiring support from their upstream providers. However, as we can see from this result that path prepending may not always be useful.

On the other hand, when we restrict the announcements only within Asia and Oceania continents using community string, we observe significant performance improvement for all the continents (Table 4). We recommend the anycast operators to have transits with providers that support community strings.

Anon-DNS-6 has also explored the use of community strings to adjust their routing. However, while community-strings are supported at most commercial sites, support at non-commercial (research and education) sites is less uniform. Without pre-deployed community strings, routing changes typically require custom tickets. Standardization of community strings across all sites would avoid this cost.

7 Conclusion

This paper proposes a way to discover and resolve the polarization problems in anycast services. We evaluate nearly 7,986 anycast prefixes, and show that the polarization problem is common in the wild Internet. We present our method to classify the polarization problems. We demonstrate two different classes of polarization problems. Our evaluation shows that the causes for both are common in known anycast prefixes. Polarization can take the clients to a distant site, often in a different continent, when the clients have a nearby anycast site. Because of the cross-continent routes, clients can observe over 100 ms of extra latency. We show some polarization problems can be solved using traffic engineering techniques. Small changes in the routing policy can improve the latency for many VPs. We also show network operators should have multiple traffic engineering techniques, including BGP community strings, to improve the performance of their anycast service.

Acknowledgments: ASM Rizvi and John Heidemann’s work was partially supported by DARPA under Contract No. HR001120C0157. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DARPA. John Heidemann’s work was also partially supported by the NFS projects CNS-2319409, CRI-8115780, and CNS-1925737. ASM Rizvi’s work was begun while on an internship at Akamai. We thank the anonymous shepherd and reviewers for their input.

References

1. Ballani, H., Francis, P., Ratnasamy, S.: A measurement-based deployment proposal for IP anycast. In: Proceedings of the 6th ACM SIGCOMM conference on Internet measurement. pp. 231–244 (2006)
2. Bellis, R.: Researching F-root anycast placement using RIPE Atlas. ripe blog https://labs.ripe.net/Members/ray_bellis/researching-f-root-anycast-placement-using-ripe-atlas (Oct 2015), https://labs.ripe.net/Members/ray_bellis/researching-f-root-anycast-placement-using-ripe-atlas
3. Caesar, M., Rexford, J.: BGP routing policies in ISP networks. IEEE Network Magazine **19**(6), 5–11 (Nov 2005). <https://doi.org/http://dx.doi.org/10.1109/MNET.2005.1541715>
4. CAIDA: AS rank. <https://asrank.caida.org/> (2020), [Online; accessed 12-Oct-2021]
5. CAIDA: CAIDA UCSD BGP community dictionary. <https://www.caida.org/data/bgp-communities/> (2020), [Online; accessed 12-Oct-2021]
6. Calder, M., Flavel, A., Katz-Bassett, E., Mahajan, R., Padhye, J.: Analyzing the performance of an anycast CDN. In: Proceedings of the 2015 Internet Measurement Conference. pp. 531–537 (2015)
7. Chandra, R., Traina, P., Li, T.: BGP communities attribute. Tech. Rep. 1997, RFC Editor (1996), <https://www.rfc-editor.org/rfc/rfc1997.txt>

8. Cicalese, D., Augé, J., Joumblatt, D., Friedman, T., Rossi, D.: Characterizing ipv4 anycast adoption and deployment. In: Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies. pp. 1–13 (2015)
9. Cloudflare: Ipv6 adoption. <https://radar.cloudflare.com/reports/ipv6/> (01 2024)
10. Dietzel, C., Feldmann, A., King, T.: Blackholing at IXPs: On the effectiveness of DDoS mitigation in the wild. In: International Conference on Passive and Active Network Measurement. pp. 319–332. Springer (2016)
11. Fan, X., Heidemann, J., Govindan, R.: Evaluating anycast in the domain name system. In: 2013 Proceedings IEEE INFOCOM. pp. 1681–1689. IEEE (2013)
12. Flavel, A., Mani, P., Maltz, D., Holt, N., Liu, J., Chen, Y., Surmachev, O.: Fastroute: A scalable load-aware anycast routing architecture for modern CDNs. In: 12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15). pp. 381–394 (2015)
13. Gao, L.: On inferring autonomous system relationships in the Internet. ACM/IEEE Transactions on Networking **9**(6), 733–745 (Dec 2001). <https://doi.org/http://dx.doi.org/10.1109/90.974527>, <http://www-unix.ecs.umass.edu/~lgao/ton.ps>
14. Gao, R., Dovrolis, C., Zegura, E.W.: Interdomain ingress traffic engineering through optimized AS-path prepending. In: International Conference on Research in Networking. pp. 647–658. Springer (2005)
15. Giotsas, V., Smaragdakis, G., Dietzel, C., Richter, P., Feldmann, A., Berger, A.: Inferring BGP blackholing activity in the internet. In: Proceedings of the Internet Measurement Conference. pp. 1–14. ACM (2017)
16. Google: Google ipv6 statistics. <https://www.google.com/intl/en/ipv6/statistics.html/> (01 2024)
17. Koch, T., Li, K., Ardi, C., Katz-Bassett, E., Calder, M., Heidemann, J.: Anycast in context: A tale of two systems. In: Proceedings of the ACM SIGCOMM Conference. ACM, Virtual (Aug 2021). <https://doi.org/https://doi.org/10.1145/3452296.3472891>
18. Kuipers, J.H.: Anycast for DDoS. https://essay.utwente.nl/73795/1/Kuipers_MA_EWI.pdf (2017), [Online; accessed 12-Oct-2021]
19. Li, Z., Levin, D., Spring, N., Bhattacharjee, B.: Internet anycast: Performance, problems, and potential. pp. 59–73. Budapest, Hungary (Aug 2018). <https://doi.org/https://doi.org/10.1145/3230543.3230547>, http://www.cs.umd.edu/projects/droot/anycast_sigcomm18.pdf
20. Luckie, M., Huffaker, B., Dhamdhere, A., Giotsas, V., Claffy, K.: As relationships, customer cones, and validation. In: Proceedings of the 2013 conference on Internet measurement conference. pp. 243–256 (2013)
21. McQuistin, S., Uppu, S.P., Flores, M.: Taming anycast in the wild Internet. In: Proceedings of the Internet Measurement Conference. pp. 165–178 (2019)
22. Monroe, L.: Centurylink completes acquisition of level 3. <https://news.lumen.com/2017-11-01-CenturyLink-completes-acquisition-of-Level-3> (2017)
23. Moura, G.C.M., de O. Schmidt, R., Heidemann, J., de Vries, W.B., Müller, M., Wei, L., Hesselman, C.: Anycast vs DDoS: Evaluating the November 2015 root DNS event. In: Proceedings of the ACM Internet Measurement Conference (Nov 2016). <https://doi.org/http://dx.doi.org/10.1145/2987443.2987446>
24. Moura, G.C., Heidemann, J., Hardaker, W., Charnsethikul, P., Bulten, J., Ceron, J.M., Hesselman, C.: Old but gold: prospecting tcp to engineer and live monitor dns anycast. In: International Conference on Passive and Active Network Measurement. pp. 264–292. Springer (2022)

25. Munteanu, C., Gasser, O., Poese, I., Smaragdakis, G., Feldmann, A.: Enabling multi-hop isp-hypergiant collaboration. In: Proceedings of the Applied Networking Research Workshop. pp. 54–59 (2023)
26. Partridge, C., Mendez, T., Milliken, W.: Host anycasting service. Tech. Rep. 1546, RFC Editor (1993), <https://www.rfc-editor.org/rfc/rfc1546.txt>
27. Quoitin, B., Pelsser, C., Swinnen, L., Bonaventure, O., Uhlig, S.: Interdomain traffic engineering with BGP. *IEEE Communications magazine* **41**(5), 122–128 (2003)
28. Rizvi, A., Bertholdo, L., Ceron, J., Heidemann, J.: Anycast agility: Network playbooks to fight {DDoS}. In: 31st USENIX Security Symposium (USENIX Security 22). pp. 4201–4218 (2022)
29. Schlinker, B., Kim, H., Cui, T., Katz-Bassett, E., Madhyastha, H.V., Cunha, I., Quinn, J., Hasan, S., Lapukhov, P., Zeng, H.: Engineering egress with edge fabric: Steering oceans of content to the world. In: Proceedings of the Conference of the ACM Special Interest Group on Data Communication. pp. 418–431 (2017)
30. Schmidt, R.d.O., Heidemann, J., Kuipers, J.H.: Anycast latency: How many sites are enough? In: International Conference on Passive and Active Network Measurement. pp. 188–200. Sydney, Australia (Mar 2017), <https://www.isi.edu/%7ejohnh/PAPERS/Schmidt17a.html>
31. Sommese, R., Bertholdo, L., Akiwate, G., Jonker, M., van Rijswijk-Deij, R., Dainotti, A., Claffy, K., Sperotto, A.: Manycast2: Using anycast to measure anycast. In: Proceedings of the ACM Internet Measurement Conference. p. 456–463. IMC '20, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3419394.3423646>, <https://doi.org/10.1145/3419394.3423646>
32. Step, O.: BGP community guides. <https://onestep.net/communities/>, [Online; accessed 12-Oct-2021]
33. Wei, L., Heidemann, J.: Does anycast hang up on you? In: 2017 Network Traffic Measurement and Analysis Conference (TMA). pp. 1–9. IEEE, Dublin, Ireland (Jul 2017). <https://doi.org/https://doi.org/10.23919/TMA.2017.8002905>
34. Weiden, F., Frost, P.: Anycast as a load balancing feature. In: Proceedings of the 24th international conference on Large installation system administration. pp. 1–6. USENIX Association (2010)
35. Zhou, M., Zhang, X., Hao, S., Yang, X., Zheng, J., Chen, G., Dou, W.: Regional ip anycast: Deployments, performance, and potentials. In: Proceedings of the ACM SIGCOMM 2023 Conference. pp. 917–931 (2023)