

Verfploeter: Broad and Load-Aware Anycast Mapping

ISI-TR-719

24 May 2017

Wouter B. de Vries¹

Ricardo de O. Schmidt^{1,2}

Wes Hardaker³

John Heidemann³

Pieter-Tjerk de Boer¹

Aiko Pras¹

1: University of Twente

2: SIDN Labs

3: USC/Information Sciences Institute

ABSTRACT

IP anycast provides DNS operators and CDNs with automatic fail-over and reduced latency by breaking the Internet into *catchments*, each served by a different anycast site. Unfortunately, *understanding* and *predicting* changes to catchments as sites are added or removed has been challenging. Current tools such as RIPE Atlas or commercial equivalents map from thousands of vantage points (VPs), but their coverage can be inconsistent around the globe. This paper proposes *Verfploeter*, a new method that maps anycast catchments using active probing. *Verfploeter* provides around 3.8M virtual VPs, 430× the 9k physical VPs in RIPE Atlas, providing coverage of the vast majority of networks around the globe. We then add load information from prior service logs to provide calibrated predictions of anycast changes. *Verfploeter* has been used to evaluate the new anycast for B-Root, and we also report its use of a nine-site anycast testbed. We show that the greater coverage made possible by *Verfploeter*'s active probing is necessary to see routing differences in regions that have sparse coverage from RIPE Atlas, like South America and China.

1. INTRODUCTION

IP anycast allows an Internet service operator to provide services such as DNS or HTTP content delivery from multiple sites that are, usually, physically distributed [2]. Anycast can reduce latency [26, 39] and blunt Distributed Denial-of-Service (DDoS) attacks by spreading traffic across different sites and providing greater aggregate capacity than any one site might [31]. IP anycast is used by (almost) all the Root DNS operators [2, 37], commercial DNS service providers and operators of top- and second-level domains. It is also used by multiple Content Distribution Networks (CDNs), including Microsoft/Bing [18], Verizon/Edgecast, and others.

Anycast operates by deploying servers at different *sites*¹, and having each site announce the same IPv4 or

¹ Anycast documents sometimes use the term *instance*, but that term can apply to both sites or individual servers. We avoid the term “instance” because it is ambiguous when anycast sites have multiple servers, as is often the case. Similarly, RFC 4786 [2] uses the term *node* for what we call a

v6 prefix using the Border Gateway Protocol (BGP—the standard protocol for inter-AS routing). All networks that receive these routes from BGP select their topologically closest site, as defined by the BGP metrics and policies. These networks define the *catchment* of that anycast site. Users in each catchment send queries to the site to which they are routed, and because routing is pre-computed, there is no query-time cost for site-selection.

Understanding anycast catchments is important for performance (both throughput and latency), defending against DDoS, and managing filtering. Anycast operators engineer their deployments to minimize latency to users, and to spread load over multiple sites. Some anycast systems employ tens or hundreds of anycast sites to minimize latency to users [10, 39], although evaluating how well that goal has been achieved requires ongoing observation [9, 39]. Operators also often balance load across multiple sites to manage capacity [11, 18], particularly for CDNs providing large volumes of content. Load balancing across anycast sites is particularly important to manage DDoS attacks [31], where matching attack traffic to capacity or isolating attack traffic to certain catchments are essential tools.

In addition to performance, anycast catchments can interact with country-specific policies for content filtering. Filtering of DNS to implement national-specific policies is not uncommon [3, 20, 22, 46], so it is important that policies and catchments align. Two examples of Root DNS service show cases where there is a mismatch. In 2010, the catchment I-Root DNS service's site in Beijing expanded outside China, imposing China's censorship policies outside its borders [28]. More recently, at the beginning of 2017, the catchment for a K-Root anycast site in Iran was seen outside that country, inconsistent with the policies of the K-Root operators and the hosts of that site [1, 27].

The challenge in managing anycast is that BGP routing is not always what one would expect. Absent other policies, BGP defines nearness in terms of AS-hops, but one AS hop across an organization with a global network; we avoid node because it often refers to specific servers.

work (such as AT&T, Tata, and NTT) can have very different latency than one AS hop across a small ISP. In addition, the trend of a flatter Internet [25] means that AS hops provide coarser control than it did in the past. More importantly, BGP policy controls allow ISPs to manipulate routing for business reasons; policy controls are widely used to do traffic management.

Current approaches to manage anycast catchments use one-off active measurements [11], platforms for active measurement such as RIPE Atlas [12, 39], commercial services (for example, [40]), and analysis of anycast service logs [21]. While these approaches have provided insight, and RIPE Atlas [36] and commercial services are in wide use, even the largest services have relatively small numbers of vantage points (from hundreds to 10,000 or so), and it is unclear how these measurement systems relate to actual operational traffic. Analysis of anycast service logs offer an accurate representation of actual load, but require the anycast service to be in operation and active use.

The overall contribution of this paper is to provide a new approach to mapping anycast catchments (§3) that has been validated through real world ground truth. This approach provides *broad coverage* and can be combined with traffic history to provide *estimated load*, providing *operational value* for one anycast service, and an approach that can be used by others. The insight in our new measurement approach is to use *active probing using the anycast service itself* and we can use *historical traffic to predict future load*. In §5 we show that active probing allows coverage of the ping-responsive Internet, currently about 4M /24 networks, providing 430× more information than current public measurement platforms. By contrast, coverage from existing platforms scale relative to the ability to deploy physical devices or virtual machines, both of which are limited.

The second contribution of this work is to use Verfploeter to examining the operational catchment for B-Root and to study anycast in Tangled, a nine-site anycast testbed (§6). B-Root deployed anycast only recently (May 2017), and *our approach actually contributed to the success of this planning and deployment*. Analyzing this active network deployment allows us to compare the predictive capability of our approach to prior approaches such as RIPE Atlas. Evaluation of our testbed lets us test a larger anycast deployment (nine sites compared to B-Root’s two sites). Our approach provides a new way to evaluate anycast stability with much broader coverage than recent studies [44].

We have made a complete version of Verfploeter available as open source at <https://github.com/woutifier> and <https://ant.isi.edu/software/lander/>. We have released all the data used in this paper (except *LN-4-12*, which is not ours); see citations in Table 1 and Table 2.

2. RELATED WORK

There have been several prior approaches to measure anycast catchment using a variety of techniques.

Use of Open Resolvers: Early work used Open DNS Resolvers in combination with PlanetLab and Ntalyzr to map catchments of anycast services [17]. While Open Resolvers provided a broad view at the time of their study (300k VPs), they are being steadily shut down out of concerns about their use in DNS amplification attacks [29]. While open resolvers offered a very large set of vantage points, it is still smaller than the method we propose that uses ping-responsive networks,

Measurement Platforms: The most common method of assessing anycast is to use public or private measurement platforms that offer physical or virtual VPs around the Internet. RIPE Atlas [36] and PlanetLab [32] are both openly available and widely distributed, and a number of commercial platforms are also available. Systems we are aware of range from hundreds to around 10k VPs.

Several studies, both by others and us, have used measurement platforms to study anycast [1, 9, 11, 12, 17, 28, 31, 39]. As pre-deployed measurement platforms these systems are available and can measure anycast services externally (without requiring support from the service operator). The main weaknesses of these systems are that they are slow and expensive to grow, and deployment is often skewed relative to the population of Internet users. This skew has been noted in many prior studies and was recently studied explicitly [8].

Traffic and Log Analysis: Anycast operators have always been able to assess current anycast performance by analyzing their own traffic and server logs. Recent work examined the Microsoft Bing CDN [11] and a variety of other CDNs [21]. As the service operator, log analysis requires no external measurements and can cover the entire service. While important, analysis of existing services can only study the *current* deployment—it requires active use by a large number of users and cannot directly support pre-deployment planning. Second, log files may be unavailable due to privacy concerns, cost of storage or retrieval, or concerns about performance impact on operational services.

Performance Analysis of DNS Services: There have been a number of analysis of root DNS service, both pre-anycast [19] and with anycast for latency [9, 13, 17, 26, 39] and DDoS [31].

To the best of our knowledge, our paper is the first to present this ICMP-based anycast catchment determination approach. Further, we do not know of any larger scale catchment measurement with open datasets against a real-world anycast deployment.

3. Verfploeter: GLOBAL PROBING

Our approach has components to map anycast catch-

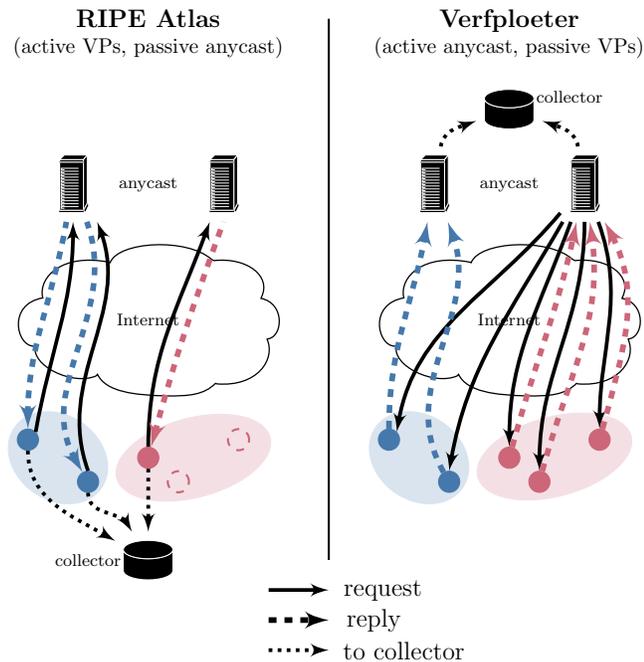


Figure 1: Traditional catchment mapping from active VPs using a testbed like RIPE Atlas (left); and using Verfloeter with queries originating in the anycast system (right).

ments for a large fraction of prefixes in the Internet, and to estimate load from each of these prefixes.

3.1 Mapping Anycast Catchments

Traditional approaches to measuring anycast catchments use many vantage points (VPs) around the Internet; each VP queries the anycast service to determine its catchment. In prior work for DNS, the VPs are typically RIPE Atlas probes [35, 39], and the queries use DNS TXT records, with the special CHAOS network type, and the name “hostname.bind” [45], or the newer NSID option in DNS [4]. One can augment these methods with traceroutes to detect possibly spoofed replies [17]. All of these approaches require deployment of active probes around the Internet. The largest studies we know of use between 9000 and 10000 VPs, all the active VPs in RIPE Atlas.

Our insight is that we do not need control over active vantage points if we can *solicit* messages from around Internet that will identify their catchment. Rather than handle both queries and responses from the VPs, we instead generate queries that cause the VPs to *respond and reply to the anycast system*; we define these as **passive VPs**. If we can capture traffic at all anycast catchments, we can determine the catchment of each VP that responds. In effect, we shift the active side that generates and receives queries from the VP to the anycast network itself.

Figure 1 compares these methods. On the left, traditional mapping has queries (black arrows) send from VPs into the anycast system. On the right, we send queries *from* the anycast network block (defined by the source address), *to passive VPs* in most /24 IPv4 networks. Their replies return to the site for their catchment, even if it is not the site that originated the query.

In Verfloeter, our *queries* are ICMP Echo requests (pings), sent using a custom program, soliciting ICMP Echo reply responses. Queries are sent from a designated measurement address that *must be in the anycast service IP prefix*. Unlike traditional catchment mapping, it is not the reply payload that indicates catchment, but instead *the catchment is identified by the anycast site that receives the reply*.

Our *passive VPs* are any computers in the Internet that reply to pings. We use a recent ISI IPv4 hitlist [16]. In principle, we could ping every IPv4 address to get complete coverage from all addresses that reply. We use hitlists instead because they provide representative addresses for each /24 block that are most likely to reply to pings, and with one address per /24 block, we can reduce measurement traffic to 0.4% of a complete IPv4 scan. (We select /24s as the smallest routable prefix in BGP today, since anycast depends on BGP.)

We must capture traffic for the measurement address with our response collection system. We can capture traffic at the routers (without having a computer at the address), or by running computers that capture traffic on the address itself. These captures must happen concurrently at *all anycast sites*. We have three different *response collection systems*: first is a custom program that does packet capture and forwards responses to a central site in near-real-time. For B-Root, we collect replies with LANDER [24], an existing packet capture system that collects data continuously. Third, we have also used tcpdump directly to capture traffic specifically for the measurement address.

Finally, we copy all responses to a central site for analysis. We do this manually, or with a custom program that forwards traffic after tagging it with its site.

Our approach to catchment mapping requires active participation at all anycast sites—it requires cooperation of the anycast operator, but it does not require additional Internet-wide infrastructure (such as distributed VPs). Fortunately, anycast operators are strongly motivated to understand their systems. These trade-offs are the opposite of traditional anycast mapping, which requires active VPs but not support of the target anycast system.

Note that we do not predict BGP routing and the resulting catchments, *we measure actual deployment*. To predict possible future catchments from different policies, one must deploy and announce a test prefix, and measure the resulting routes and catchments (assuming

the test prefix is subject to the same policies as the regular prefix in production). Fortunately anycast providers often announce both a /24 prefix with anycast and as the covering /23 prefix (this approach protects against corner cases with some routing policies [14]). The non-operational portion of the /23 could serve as the test prefix.

3.2 Load Estimation

Planning anycast deployment is more than just mapping catchments—different services can experience very different loads, depending on the distribution and usage patterns of its client base. We therefore build load estimates for each network block (/24 prefix) that accesses a service, so we can calibrate the loads that will be generated by a given catchment.

We assume operators collect query logs for their systems and using the recorded historical data we estimate future loads. (For example, all root operators collect this information as part of standard RSSAC-002 performance reporting [38].) We consider three types of load: queries, good replies, and all replies. Queries represent incoming load on the servers, while replies are the results. Query packet load counts may differ from replies if response rate limiting is used to blunt DNS amplification attacks [41]. We separate out good replies from all replies because of the large fraction of queries to non-present domains in root-server traffic (first observed in 1992 [15] and still true today); operators may wish to optimize for volume or for good replies.

In principle, we can estimate load over any time period. Practically, we compute it over one day, and look at overall traffic using hourly bins.

4. MEASUREMENT SETUP AND DATASET

Using the proposed ICMP-based method, Verfploeter, we measure the catchment of two anycast services, B-root and an anycast testbed (Tangled), from more than 6.4M VPs (IP addresses). We add geolocation information for these blocks using MaxMind [30]. Accuracy of this geolocation is considered reasonable at the country level [33]. We also use Route Views and RIPE RIS data to determine the AS number for each scanned IP address and the prefixes that are announced by each AS.

Data cleaning: We remove from our dataset the duplicate results, replies from IP-addresses that we did not send a request to, and late replies (15 minutes after the start of the measurement). Further investigation on the causes leading to the removed data are out of the scope in this paper.

4.1 B-Root

We validate the proposed methodology by providing a detailed view of the catchment of one of the DNS root-

Id	Service	Method	Start	Dur.
<i>SBA-4-20</i>	B-Root	Atlas	2017-04-20	8 m
<i>SBA-4-21</i>	[34]		2017-04-21	8 m
<i>SBA-5-15</i>			2017-05-15	10 m
<i>SBV-4-21</i>	B-Root	Verf-	2017-04-21	20 m
<i>SBV-5-15</i>	[23]	ploeter	2017-05-15	20 m
<i>STA-2-01</i>	Tangled [42]	Atlas	2017-02-01	10 m
<i>STV-2-01</i>	Tangled	Verf-	2017-02-01	10 m
<i>STV-3-23</i>	[43]	ploeter	2017-03-23	24 h

Table 1: Scans of anycast catchments for B-Root and our testbed (Tangled).

Id	Service	Date	Site	Queries	
				q/day	q/s
<i>LB-4-12</i>	B-Root [6]	2017-04-12	LAX	2.34G	27.1k
<i>LB-5-15</i>	B-Root [7]	2017-05-15	both	2.20G	25.4k
			LAX	1.78G	20.6k
			MIA	0.407G	4.71k
<i>LN-4-12</i>	NL ccTLD	2017-04-12		redacted	

Table 2: Datasets used to study load (IPv4 UDP queries only).

servers. B-root is the most recent root letter to make the change from unicast to anycast. B-Root deployed anycast at the beginning of May, 2017 [5], adding a site in Miami to its original site in Los Angeles (Table 3).

B’s new deployment of anycast makes it an interesting analysis target. Unlike the other DNS Roots, B does not have a history of anycast deployment to guide its choices (although of course it draws on experience of other anycast deployments).

Dataset: We study B-Root catchments using several scans using both RIPE Atlas and Verfploeter, as shown in Table 1. We estimate B-Root load using two day-long datasets listed in Table 2. As a baseline we use data from DITL 2017, taken Wednesday, 2017-04-12 (UTC), before B-Root was using anycast. We then test against Thursday, 2017-05-15 (UTC), after B-Root anycast was well established.

4.2 Anycast Testbed

We augment our measurements of B-Root with mea-

Service	Location	Host	Upstream
B-Root	US, Los Angeles	USC/ISI	AS226
	US, Miami	FIU/AMPATH	AS20080
Tangled	AU, Sydney	Vultr	AS20473
	FR, Paris	Vultr	AS20473
	JP, Tokyo	WIDE	AS2500
	NL, Enschede	Univ. of Twente	AS1103
	UK, London	Vultr	AS20473
	US, Miami	Florida Int. Univ.	AS20080
	US, Washington	USC/ISI	AS1972
	BR, Sao Paulo	Florida Int. Univ.	AS1251
	DK, Copenhagen	DK-Hostmaster	AS39389

Table 3: List of anycast sites used in our measurements.

surements of our anycast testbed, *Tangled*. This testbed has 9 sites around the world: 5 sites in Europe, 2 in the USA, and 3 other sites spread across Asia, Oceania and South America (Table 3). *Tangled* allows us to study how a larger amount of sites interact, and to perform experiments which we cannot do in an operational anycast service. We use it to understand anycast instability and ASes that appear in multiple catchments (§6).

Limitations Three of the testbed sites share a common ISP, which might impact the overall catchment. The anycast site in São Paulo has all its traffic routed via the same link as the site in Miami, which might cause announcements from São Paulo to be hidden. Finally, the connectivity at the site in Japan is such that it does not attract much traffic since announcements from other sites are almost always preferred over it. Prior to the measurement, the connectivity of each site was validated individually by announcing our prefix from that location only. Such limitations are not particular to our testbed as similar features can also be observed in public anycast services [39].

Dataset As shown in Table 1 we measured the catchment using both Verfploeter and Atlas on Wednesday, 2017-02-01 (UTC). We also determined the catchment of *Tangled*, using only Verfploeter, every 15 minutes during a 24 hour period starting 2017-03-23 10:57 UTC, for a total of 96 measurements. In total we collected 342,604,759 ICMP replies, of which 324,675,876 remained after cleaning.

For each measurement we transmitted one ICMP packet to each of the 6.4M IPs from the hitlist, at a rate of 10k/second to prevent overloading networks or network equipment. Each measurement round took approximately 10.5 minutes to complete. A unique identifier in the ICMP header was used in every measurement round to ensure data set separation.

5. ANALYSIS OF THE Verfploeter MECHANISM

In this section we examine the Verfploeter measurement method. We show the broader coverage of Verfploeter compared to RIPE Atlas, and how catchment mapping from Verfploeter can be combined to historic traffic load to accurately predict load at individual anycast sites.

5.1 Utility: Operational Evaluation of Anycast Catchments

A long-standing goal of anycast mapping is to assess load balancing and routing problems [9, 39]. We next look at B-Root’s anycast distribution. Deployed recently in May 2017, it has only two sites, but we are able to deploy Verfploeter on it.

We have measured the geographic footprint of B-Root with RIPE Atlas (Figure 2a) and Verfploeter (Fig-

ure 2b). These maps highlight a couple of important differences between these measurement methods.

First, Verfploeter *has much broader coverage*: Atlas coverage is good in Europe and reasonable in North America, but sparse elsewhere and almost absent in China. Verfploeter provides good coverage for most of the populated globe. Second, even where coverage is good, Verfploeter *provides far more numerous observations*—the scale of Figure 2b is 1000× greater than Figure 2a.

These differences are particularly important for examination of B-Root catchments in South America and China. The broader coverage is important to understand, for example, a hypothetical B-Root site selection in China: Atlas cannot comment, but Verfploeter shows most of China selects the MIA site.

The denser coverage in South America also helps highlighting the impact of B-Root’s hosting ISPs. B-Root’s ISP in MIA (AMPATH) is very well connected in Brazil and Argentina, but does not have direct ties to the west coast of South America. This difference shows in the wider use of the MIA site in Brazil, and less use of it in Peru and Chile.

B-Root’s goal in measuring anycast is to understand routing choices; we return to this question in §6.1.

5.2 Utility in Mapping Multi-Site Anycast

B-Root shows the benefits of increased number of VPs with Verfploeter, but we would like to understand how the different approaches work on anycast deployments with more sites. We therefore turn to *Tangled*: an anycast testbed designed and deployed by us (§4.2).

Figure 3 maps the catchments of *Tangled* with Atlas and Verfploeter. Again, outside of Europe, the greater density of coverage of Verfploeter provides clear qualitative differences between the two maps. For example, the IAD site (dark yellow) shows up prominently across North America with Verfploeter, but with Atlas, CDG and ENS seem to serve that region. We also see very different mixes of sites in Australia. And only Verfploeter provides coverage of China.

The key result from these graphs is that *Verfploeter coverage tracks the Internet as a whole*, not just where physical VPs can be placed. We quantify this difference in the next section.

5.3 Greater Coverage in Verfploeter

In §5.1 and §5.2 we showed how the greater coverage in Verfploeter reveals aspects of B-Root and our testbed *Tangled* that would otherwise be missed. This coverage is possible because Verfploeter’s virtual VPs only require a computer that responds to ICMP, instead of physically deployed devices (Figure 1); this way we can cover millions of /24s.

To quantify the difference in coverage that is visible

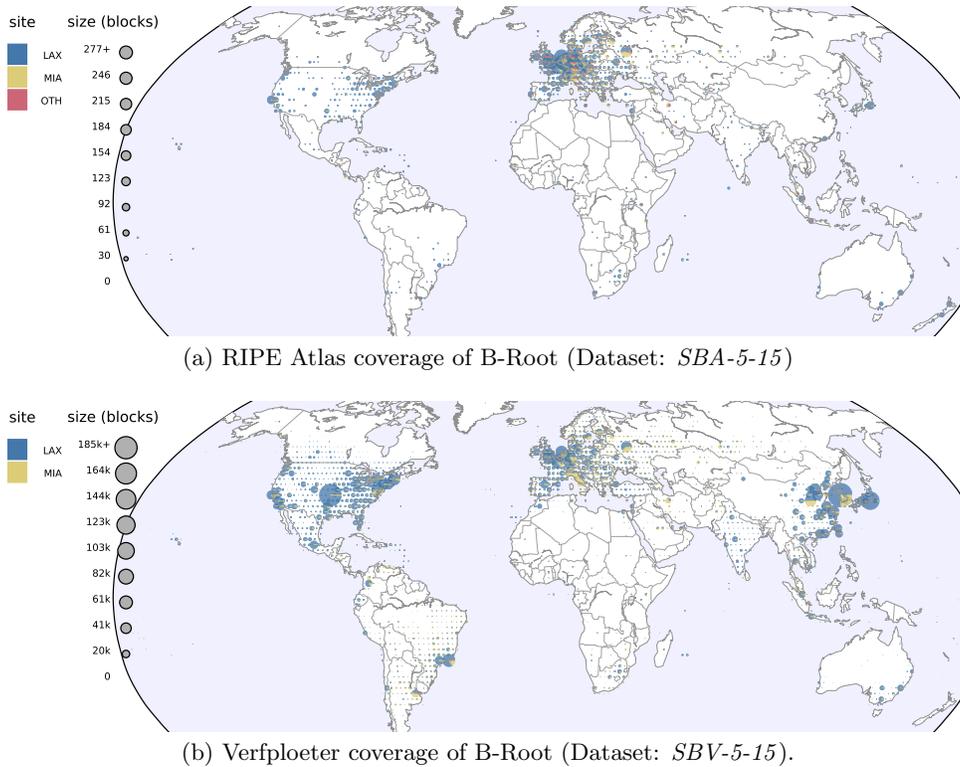


Figure 2: Geographic coverage of vantage points for RIPE Atlas and Verfloeter for B-Root, in two-degree geographic bins. The pie in each bin is colored by site (blue: LAX; yellow: MIA; red: other). Circle areas show number of address blocks (Verfloeter) or VPs (Atlas) at different scales.

	RIPE Atlas (VPs)	Atlas (/24s)	Verfloeter (/24s)
considered	9807	9083	6,877,175
non-responding	455	406	3,090,268
responding	9352	8677	3,786,907
no location	0	0	678
geolocatable	9352	8677	3,786,229
unique		2079	3606300

Table 4: Coverage of B-Root from the perspective of the RIPE Atlas and Verfloeter measurement systems, measured in VPs (Atlas) or /24 blocks (both). (Datasets: *SBA-5-15*, *SBV-5-15*)

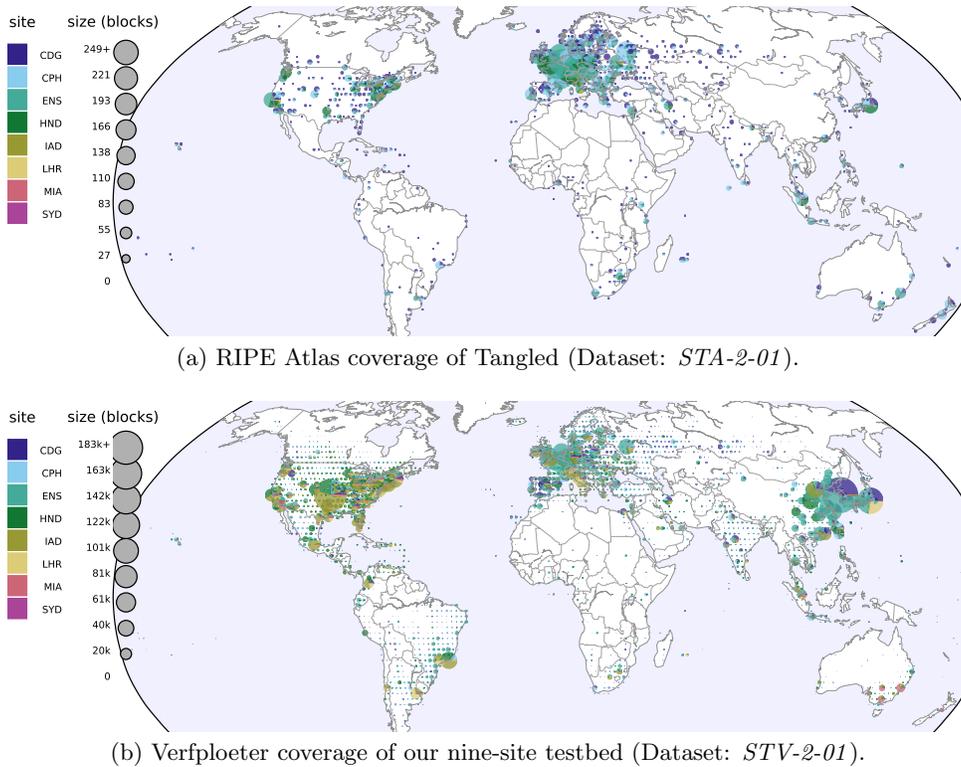
in Figure 2, Table 4 compares how many blocks the two measurement approaches see. For both systems we try to use all available VPs, but some VPs are unavailable: for Atlas, 455 VPs do not respond (within 406 blocks), presumably because they are temporarily down. For Verfloeter, about 3M ping targets do not reply, presumably because the target was temporarily down, or it was in a block of dynamic addresses and temporarily unused. If desired, both of these non-response rates could be reduced by retrying later, or with additional addresses for Verfloeter. All Atlas VPs have geolocation (set when the VP is registered), but we discard a few Verfloeter blocks (678) that we cannot geolocate.

The key result about coverage is that Verfloeter sees around $430\times$ more blocks than Atlas. Although Atlas finds a few unique blocks (presumably blocks that discard all pings), about 77% of Atlas blocks are also seen by Verfloeter, and Verfloeter sees around 3.61M additional blocks.

5.4 From Observations to Load

We next look at how well different measurement systems relate to actual load on an anycast service. It is well known that the distribution of RIPE Atlas reflects more about who RIPE interacts with than global Internet traffic—as an European project, and Europe being the main region of RIPE NCC operation, Atlas’ deployment is by far heavier in Europe than in other parts of the globe (and this is a well known shortcoming [8]). Our goal here is to *calibrate* different measurement systems to best match actual traffic.

Estimating Load: To estimate load on B-Root, we begin with our prediction about anycast catchments from Verfloeter, then we weight each /24 block by its known traffic load (§3.2). There are blocks for which we do not have anycast mapping, either because they do not reply to our probes, or because the specific address we chose to contact did not reply; these blocks



(a) RIPE Atlas coverage of Tangled (Dataset: *STA-2-01*).

(b) Verfloeter coverage of our nine-site testbed (Dataset: *STV-2-01*).

Figure 3: Catchments for Tangled from RIPE Atlas and Verfloeter. Circle areas show number of blocks (Verfloeter) or VPs (Atlas) at different scales; each is a pie chart with colors showing each site.

	Blocks		Queries	
	/24s	%	q/day	%
seen at B-Root	1,388,338	100%	2.19G	100%
mapped by Verfloeter	986,605	87.1%	1.80G	82.4%
not mappable	401,733	12.9%	384M	17.6%

Table 5: Coverage of Verfloeter from B-Root. (Dataset: *SBV-5-15*, *LB-5-15*.)

are mapped to the “other” anycast site. (Although we assume their traffic will go to our sites in similar proportion to blocks in known catchments.)

Figure 4a shows the result of this load prediction. It is useful to compare this estimate to Figure 2b, which counts /24 blocks that source traffic, and Figure 2a, which counts Atlas VPs.

The most striking operational difference between measurements of blocks and actual load estimates is that load seems to *concentrate* traffic in fewer hotspots. This outcome should not be surprising: DNS is a common service operated by most ISPs with a local recursive resolver. Thus an ISP with users spread over a large region may still send all NS traffic through recursive resolvers at a few data centers.

Second, Verfloeter can only map blocks that respond to our probes. Table 5 shows coverage as seen from B-Root’s traffic logs, showing that there are a large num-

Date	Method	Measurement	% LAX
2017-04-21	Atlas	967 VPs	68.8%
2017-05-15		9,682	82.4%
2017-04-21	Verf-	4.069M /24s	82.4%
2017-05-15	ploeter	3.923M	87.8%
2017-05-15	+ load	n/a q/day	81.6%
2017-05-15	Act. Load	2.188G q/day	81.4%

Table 6: Quantifying differences B-Root anycast with different measurement methods and times.

ber of blocks (about 12.9%) that are not mapped. Figure 4a plots the load from these blocks in red, showing that most are in Korea, with some in Japan and central and southeast Asia. In §5.5 we show that these missing blocks do not alter our predictions.

Finally, we see that load is higher in some regions than the number of blocks would suggest, particularly in India. This difference may be explained by many users using relatively few IP blocks in these areas (that is, a great deal of network address translation).

Quantifying Differences from VPs to Blocks to Load: While Figure 2 and Figure 2b show visual differences, we turn to Table 6 to quantify those differences and their impact on assessment of catchment sizes in B-Root. When we compare Atlas, Verfloeter, and Verfloeter with load, we see very different mea-

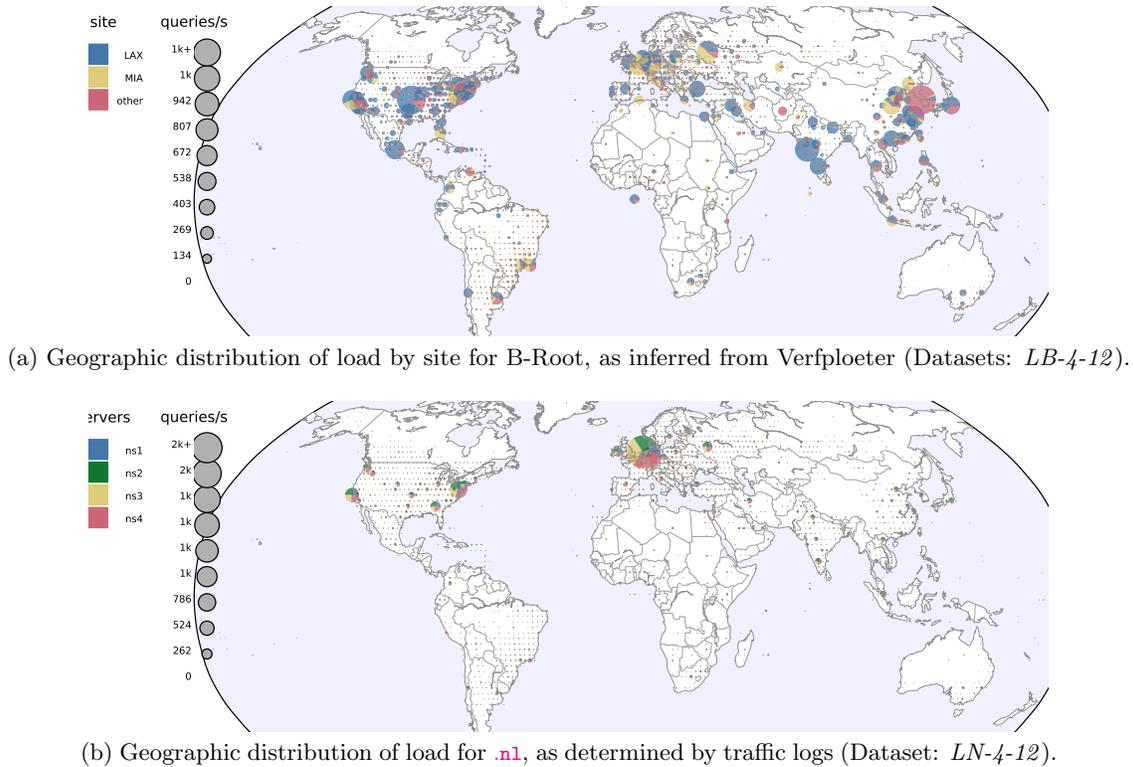


Figure 4: Measured DNS traffic over geography for B-Root and *.nl*.

surements (thousands of VPs, millions of blocks, or billions of queries per day). The different weighting factors result in different fractions of traffic between the LAX and MIA sites, as shown in the “% LAX” column. In §5.5 we will compare these values to measured load to see which is most accurate, but next we see how these changes will be even larger for DNS services with less even global load.

Uneven Load: Load for B-Root is global and largely follows the distribution of Internet users, so Figure 4a has only moderate differences from Figure 2b.

Other DNS systems are more regional. Figure 4b shows load for four of the *.nl* nameservers, the country domain for the Netherlands. This plot does not include data from their two nameservers that use anycast, so it may under-represent global traffic, but we know it captures at least half of all global traffic to this domain.

Unlike B-Root, we see that the majority of traffic to *.nl* is from Europe and the Netherlands. There is also significant traffic from the U.S. and some global traffic. With this type of client distribution, calibrating the measured catchment using load information is critical.

5.5 Using Verfploeter to Predict Load

We next examine how accurately Verfploeter’s load modeling can *predict future load*. Our goal is to determine how much unmappable blocks (§5.4) affect ac-

curacy, and how much does routing and load shift over time. A study of long-term predictions will require more experience with Verfploeter, but we address the basic accuracy question here.

We study the accuracy of load predictions with Verfploeter by analyzing what network blocks B-Root sees traffic from that Verfploeter has found to be unmappable by examining the DNS network load at B-Root on 2017-05-15 (Dataset: *LB-5-15*) and the Verfploeter analysis performed on the same day (Dataset: *SBV-5-15*). Recall from Table 6 that although Verfploeter finds 87.8% of network blocks reach the LAX, the load prediction is that 81.6% of traffic should go to LAX, and that that prediction does not consider blocks that send traffic to B-Root but do not respond to Verfploeter (12.9% from Table 5).

Predicted vs. Measured Load: The last line of Table 6 shows the *actual* load, as measured at all B-Root sites on 2017-05-15. At 81.4%, we see our prediction using same-day Verfploeter and load is only quite close (81.6%). We conclude that the blocks that are unobservable by Verfploeter do *not* have significant effects on our overall load estimate. Although they account for 17.6% of queries (Table 5, and the red slices in Figure 4a), the fraction of traffic that goes to each B-Root site appears to follow the ratio seen in measured blocks.

Our second observation is that *our load-weighted pre-*

distributions are very close to observed load. Verfloeter without load adjustment is further off, with 87.8% of blocks going to LAX. We conclude that weighting by load is important. Surprisingly, Atlas estimates, at 82.4%, are actually closer than Verfloeter if Verfloeter is not load-weighted.

The key take-away of this result is that with load-weighted Verfloeter *preliminary results suggest it is possible to make reasonably predictions about future anycast deployments* by measuring the deployment on a test network and predicting traffic mapping with recent load data. We hope to expand these results beyond B-Root as ongoing work.

Long-duration predictions: Finally, we can also look at long-duration prediction. We performed a similar prediction analysis in advance of the B-Root deployment using the Verfloeter data gathered on 2017-04-12 and network traffic from 2017-04-12. We see a fairly large shift in blocks between these dates, with Verfloeter shifting from 82.4% to LAX in April to 87.8% in May. By weighting the *SBV-4-21* Verfloeter dataset from the B-Root test prefix with the *LB-4-12* measured load, we find that the predicted DNS request load arriving at LAX is 76.2%. This is significantly less than the 81.6% measured load in *LB-5-15*, which highlights the discrepancy between shifts in routing over one month between the *SBV-4-21* and *SBV-5-15* dataset collection periods.

This shift suggests that the accuracy of load estimates depends on how old the data is.

We know that routing changes in the Internet over time [9]; this early result suggests some care must be taken with long-duration predictions. We are currently studying them more systematically.

6. RESULTS: UNDERSTANDING ANYCAST WITH Verfloeter

We next use Verfloeter to explore three questions about anycast. These questions have each been raised in prior work; here we use Verfloeter to revisit them, both to show its utility and to refine these prior results.

6.1 Use of AS-Prepending in B-Root

An important operational question for B-Root is understanding how to balance load between sites. Although both sites are able to handle normal traffic, DNS operators need to shift load during emergencies like DDoS attacks [31]. They may also want to control load during regular operation, perhaps because different sites have cost structures that are traffic-sensitive.

We used RIPE Atlas and Verfloeter to investigate the use of AS Prepending to adjust the catchment of a test prefix on B’s sites. Figure 5 shows how the distribution changes as AS prepending is applied between the two sites, as measured with both methods. (Note

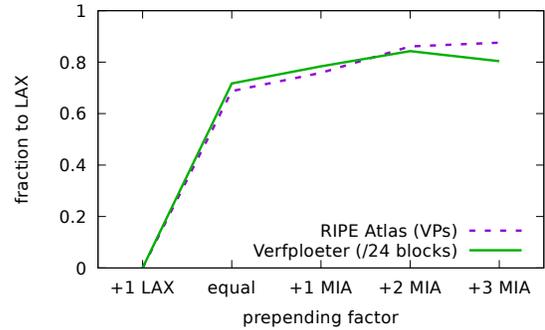


Figure 5: Split between MIA and LAX in VPs for Atlas and /24s for Verfloeter. (Dataset: *SBA-4-20*, *SBA-4-21*, *SBV-4-21*.)

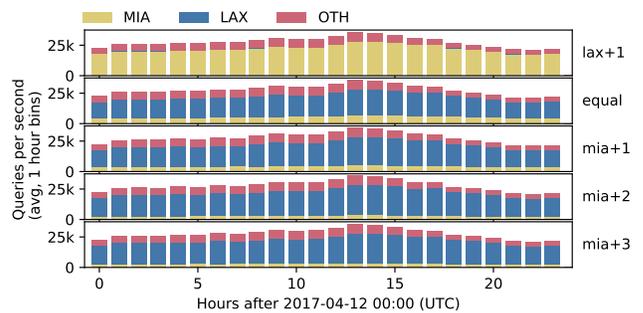


Figure 6: Predicted load for B-Root with multiple AS prepending combinations; catchment data from Verfloeter with load (Datasets: *SBV-4-21*, *LB-4-12*).

that the units for each measurement is different: RIPE Atlas is measured in VPs, and Verfloeter is measured in /24 blocks.) By default, with no prepending, 74% of Atlas VPs arrive at LAX, while Verfloeter shows that 78% of responsive /24 prefixes will arrive at LAX.

These results show that both measurement systems are useful to evaluate routing options. With only two sites, either measurement method seems sufficient for rough analysis. We expect the greater precision of Verfloeter will be important with more sites, or with more subtle methods of route control (for example, use of BGP communities to re-balance some traffic).

We next use our load predictions to see how load shifts at different prepending values over the course of a day. Figure 6 shows load predictions using catchment data from Verfloeter combined with DITL data of B-Root (2017-04-12). The top graph shows nearly all traffic going to the MIA site, since LAX’s BGP announcement includes an “AS prepending” of one (and the small share of load, OTH, that is not mappable by Verfloeter). When LAX and MIA announce routes without prepending, most of the traffic load shifts to LAX (second graph from top-down). The last three

graphs show the results of prepending MIA’s BGP announcement by up to 3 times, resulting in an increasing traffic share shifting to LAX. However, even by announcing our prefix with 3 times our AS at MIA (MIA+3), we still see a small fraction of traffic being mapped to MIA.

6.2 Discovering Divisions Within ASes

Prior work has often assumed that anycast catchments align with ASes, thus one VP can represent where the load of the entire AS goes. While generally true for smaller ASes, this assumption is less likely to hold for large, multi-national ASes where different parts of the AS may be served by different anycast sites. Such large ASes are likely to have geographically distributed peering locations and so may prefer to direct some of their users to different anycast sites to reduce service latency.

This high density of VPs in Verfploeter allows us to test this assumption by looking for differences in anycast catchments that occur *within* individual ASes. We first remove those VPs from the dataset that show instability (see §6.3), to prevent unstable routing from being classified as a division within the AS. Without removing these VPs we observe approximately 2% more divisions (e.g., ASes which are served by more than one site). We count the number of sites that are seen (from different VPs) within a single AS, in a single measurement round.

In total, we see multiple sites from 7,188 ASes, or approximately 12.7% of all ASes that were announcing at least a single prefix at the time of the measurement. Note that this is a lower-bound, using a larger and/or more diverse anycast service we might be able to determine a higher, and more accurate, percentage of ASes that are split into multiple individually routed parts.

Routing policies are a likely cause for these divisions. And, as routing on the Internet is largely determined by BGP, we show the number of prefixes that are announced via BGP by an AS versus the number of sites that it sees in Figure 7. Indeed, those ASes that announce more prefixes tend to see a higher amount of sites from their network.

In Figure 8 we show the number of sites that are seen from announced prefixes, grouped by prefix length. VPs in prefixes longer than a /15 are mapped to a single site in most cases, and up to 10% of the prefixes have VPs mapped to two different sites. Larger prefixes are often divided further.

These results show that, in order to get a complete view of the catchment, in many cases you need *more* than a single VP per AS. The AS-coverage metric that is often used for measurement platforms underestimates the amount of information that is missed.

6.3 Stability of Anycast for Clients

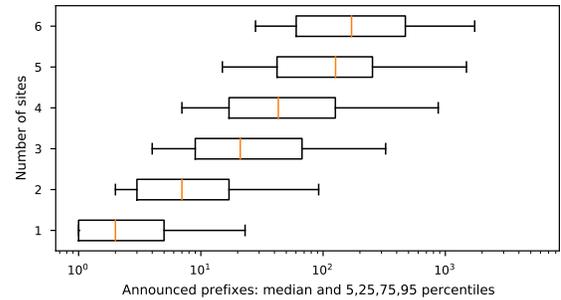


Figure 7: The number of sites that are seen from an AS versus the median amount of prefixes that are announced by those ASes. (Dataset: *STV-3-23*.)

A long-term concern with anycast is how *stable* the mapping from a client to a site is [44]. Since TCP connections require shared state at both ends, if users switch anycast sites within the lifetime of a TCP connection, that connection will break and need to be restarted. The existence of multiple successful CDNs that use IP anycast (including Bing, Edgecast, and Cloudflare) suggest that anycast is almost always stable, but recent work has suggested that anycast may be persistently unstable for a tiny fraction of (user, service) combinations (less than 1%) [44]. From a service operators point of view, it is interesting to know if a single measurement can be representative for a longer time, or if the catchment is continuously in flux.

Verfploeter allows us to revisit this question from Tangled to many VPs. We measured the global catchment of our testbed every 15 minutes for a day (96 observations). We categorize the responses (or non-responses) into 4 groups: *stable*, VPs that maintain the same catchment across measurements; *flipped*, VPs that change catchment, with responses sent to a different anycast site than the prior measurement; *to-NR*, VPs that switched to “not responding” in the current measurement; and *from-NR*, VPs that started responding in the current measurement. We do not count VPs that remain non-responsive after being counted as *to-NR*.

Figure 9 shows the results of one day of these measurement. Because the fractions of stable and flipping are so different, we break the graph into three sections. We see that the *catchment is very stable* across the measurement rounds, with a median of 3.54M (about 95%) VPs always replying and maintaining their prior catchment. The fraction of VPs that fluctuate between responsive and non-responsive states is small across all 96 measurements. A median of 89k (about 2.4%) VPs changed from responsive to non-responsive between measurements; and a median of 89k (about 2.4%) VPs changed from non-responsive to responsive. Note that fluctuating and flipping VPs are not necessarily always the same ones.

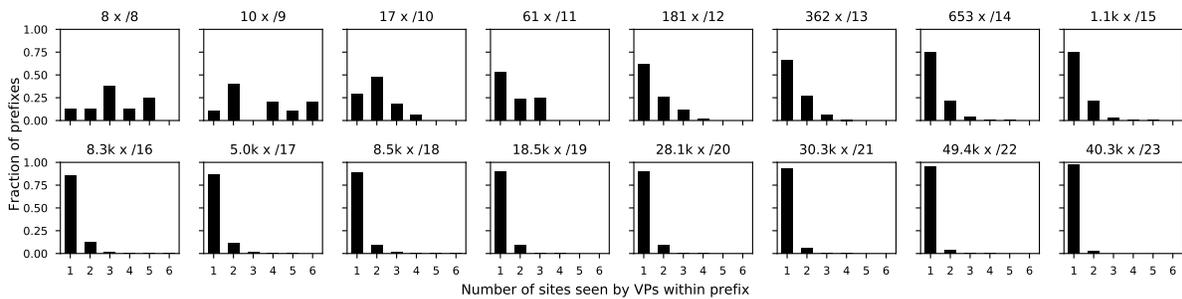


Figure 8: The number of sites that are seen for each prefix as announced in BGP. (Dataset: *STV-3-23*.)

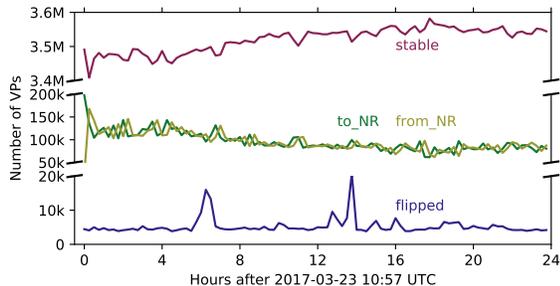


Figure 9: Stability over 24 hours. One datapoint per 15 minutes. If a VP stopped responding then it counts as to_NR, if it started responding it counts as from_NR. (Dataset: *STV-3-23*.)

	AS	IPs (/24s)	Flips	Frac.
1	4134 CHINANET	47,963	257,915	0.51
2	7922 COMCAST	3,933	19,133	0.04
3	6983 ITCDELTA	1,372	15,403	0.03
4	6739 ONO-AS	849	13,347	0.03
5	37963 ALIBABA	2,493	10,988	0.02
	<i>Other</i>	43,388	188,630	0.37
	Total	108,493	505,416	1.00

Table 7: Top ASes involved in site flips. (Dataset: *STV-3-23*.)

Across the measurement period, we also see a median of 4.6k (about 0.1%) VPs change catchment (the blue line in Figure 9). All these VPs are located within 2809 ASes. Table 7 shows that 63% of the flipping VPs are part of only 5 ASes; and 51% are within AS 4134 (Chinanet) alone. These shares suggest that potentially instabilities in the catchment are reserved to a small number of ASes; and that the large number of flips within AS 4134 seems to be a result of routing conditions within the AS itself. Catchment flips can be caused by changes in routing policies or link state, and frequent flipping can be caused by load balanced links. These results confirm prior results taken with RIPE Atlas [44], but with a much larger set of vantage points.

7. CONCLUSIONS

The key result of this paper is to show that Verfloeter allows measurements of anycast catchments across millions of networks in the Internet. Verfloeter allows us to see 430× more network blocks than RIPE Atlas, a widely used, large-scale platform for active measurements.

Such measurements are important for operating anycast services (§5.1), and more important as anycast services grow in number of sites (§5.2). With large DNS and CDN anycast networks using hundreds or thousands of sites, catchment mapping with broad coverage (§5.3) is increasingly important, particularly since regular catchment evaluation is necessary to avoid performance errors [9, 39].

Furthermore, the combination of historic traffic load and catchment mapping (§5.4) can provide a predictive tool for anycast operation (§5.5). The broad coverage of Verfloeter allows us to identify individual networks that are very likely to be the source of larger amounts of traffic.

We have used Verfloeter to understand the new B-Root anycast system (§6.1), evaluate split catchments in large ASes (§6.2), and confirm prior results in anycast stability with a larger dataset (§6.3).

Acknowledgments

We thank Moritz Müller (SIDN), Giovane Moura (SIDN) and Anna Sperotto (U. Twente) for their valuable input and feedback.

Work at U. Twente was partially supported by SURFnet Research on Networks, SAND (<http://www.sand-project.nl>) and NWO DAS (<http://www.das-project.nl>).

Wes Hardaker’s work in this paper is partially supported by USC as part of B-Root research activity.

John Heidemann’s work in this paper is partially sponsored by the Department of Homeland Security (DHS) Science and Technology Directorate, HSARPA, Cyber Security Division, via BAA 11-01-RIKA and Air Force Research Laboratory, Information Directorate (agreement FA8750-12-2-0344) and via contract number HHSP233201600010C. The U.S. Government is authorized to make reprints for governmental purposes notwithstanding any copyright. The views contained herein are those of the authors and do not necessarily represent those of NSF, DHS or the U.S. Government.

8. REFERENCES

- [1] ABEN, E. DNS Root Server Transparency: K-Root, Anycast and More. Web, Jan. 2017. RIPE NCC.
- [2] ABLEY, J., AND LINDQVIST, K. Operation of Anycast Services. RFC 4786, December 2006.
- [3] ANONYMOUS. The collateral damage of Internet censorship by DNS injection. *ACM Computer Communication Review* 42, 3 (July 2012), 21–27.
- [4] AUSTEIN, R. DNS Name Server Identifier (NSID) Option. RFC 5001, August 2007.
- [5] B-ROOT. B-Root Begins Anycast in May. <http://root-servers.org/news/b-root-begins-anycast-in-may.txt>, Apr. 2017.
- [6] B-ROOT. DITL 2017 for B-Root. dataset DITL_B_Root-20170407 at <https://ant.isi.edu/traces/>, 2017. also available through <https://www.dns-oarc.org/>.
- [7] B-ROOT. Load for 2017-05-15 for B-Root. dataset Load_B_Root-20170515 at <https://ant.isi.edu/traces/>, 2017.
- [8] BAJPAI, V., ERAVUCHIRA, S. J., AND SCHÖNWÄLDER, J. Lessons Learned from using the RIPE Atlas Platform for Measurement Research. *ACM Computer Communication Review (CCR)* 45, 3 (July 2015), 35–42.
- [9] BELLIS, R. Researching F-root anycast placement using RIPE Atlas. RIPE Labs: https://labs.ripe.net/Members/ray_bellis/researching-f-root-anycast-placement-using-ripe-atlas, Oct. 2015.
- [10] CALDER, M., FAN, X., HU, Z., KATZ-BASSETT, E., HEIDEMANN, J., AND GOVINDAN, R. Mapping the expansion of Google’s serving infrastructure. In *Proceedings of the ACM Internet Measurement Conference* (Barcelona, Spain, Oct. 2013), ACM, pp. 313–326.
- [11] CALDER, M., FLAVEL, A., KATZ-BASSETT, E., MAHAJAN, R., AND PADHYE, J. Analyzing the Performance of an Anycast CDN. In *Proceedings of the ACM Internet Measurement Conference* (Tokyo, Japan, Oct. 2015), ACM, pp. 531–537.
- [12] CICALESE, D., AUGÉ, J., JOUMLATT, D., FRIEDMAN, T., AND ROSSI, D. Characterizing IPv4 Anycast Adoption and Deployment. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies (CoNEXT)* (2015), pp. 16:1–16:13.
- [13] CICALESE, D., JOUMLATT, D., ROSSI, D., BUOB, M.-O., AUGE, J., AND FRIEDMAN, T. A Fistful of Pings: Accurate and Lightweight Anycast Enumeration and Geolocation. In *IEEE Conference on Computer Communications (INFOCOM)* (April 2015), pp. 2776–2784.
- [14] COLITTI, L. K-root anycast deployment. Presentaiton at RIPE-52, April 2006.
- [15] DANZIG, P. B., OBRACZKA, K., AND KUMAR, A. An analysis of wide-area name server traffic: A study of the Domain Name System. In *Proceedings of the ACM SIGCOMM Conference* (Jan. 1992), pp. 281–292.
- [16] FAN, X., AND HEIDEMANN, J. Selecting representative IP addresses for Internet topology studies. In *Proceedings of the ACM Internet Measurement Conference* (Melbourne, Australia, Nov. 2010), ACM, pp. 411–423.
- [17] FAN, X., HEIDEMANN, J., AND GOVINDAN, R. Evaluating anycast in the Domain Name System. In *Proceedings of the IEEE Infocom* (Turin, Italy, Apr. 2013), IEEE, pp. 1681–1689.
- [18] FLAVEL, A., MANI, P., MALTZ, D. A., HOLT, N., LIU, J., CHEN, Y., AND SURMACHEV, O. FastRoute: A Scalable Load-Aware Anycast Routing Architecture for Modern CDNs. In *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI)* (2015), pp. 381–394.
- [19] FOMENKOV, M., K. C. CLAFFY, HUFFAKER, B., AND MOORE, D. Macroscopic internet topology and performance measurements from the DNS root name servers. In *Proceedings of the USENIX Large Installation Systems Administration Conference* (San Diego, CA, USA, Dec. 2001), USENIX, pp. 221–230.
- [20] GILL, P., CRETE-NISHIHATA, M., DALEK, J., GOLDBERG, S., SENFT, A., AND WISEMAN, G. Characterizing web censorship worldwide: Another look at the opennet initiative data. *ACM Transactions on the Web* 9, 1 (Jan. 2015).
- [21] GIORDANO, D., CICALESE, D., FINAMORE, A., MELLIA, M., MUNAFÀ, M., JOUMLATT, D. Z., AND ROSSI, D. A first characterization of anycast traffic from passive traces. In *IFIP workshop on Traffic Monitoring and Analysis (TMA)* (April 2016), pp. 30–38.
- [22] GRUBB, B. 8.8.8.8: The four digits that could thwart Australia’s anti-piracy, website-blocking regime. *Sydney Morning Herald* (June 24 2015).
- [23] HARDAKER, W. Verfploeter measurements of B-Root. USC-LANDER dataset B-Root_Verfploeter-20170421 and B-Root_Verfploeter-20170515, <https://ant.isi.edu/datasets/>, 2017.
- [24] HUSSAIN, A., BARTLETT, G., PRYADKIN, Y., HEIDEMANN, J., PAPADOPOULOS, C., AND BANNISTER, J. Experiences with a continuous network tracing infrastructure. Tech. Rep. ISI-TR-2005-601, USC/Information Sciences Institute, Apr. 2005.
- [25] LABOVITZ, C., IEKEL-JOHNSON, S., MCPHERSON, D., OBERHEIDE, J., AND

- JAHANIAN, F. Internet inter-domain traffic. In *Proceedings of the ACM SIGCOMM Conference* (New Delhi, India, Aug. 2010), ACM, pp. 75–86.
- [26] LIANG, J., LI, J. J. H. D. K., AND WU, J. Measuring query latency of top level DNS servers. In *Proceedings of the Passive and Active Measurement Workshop* (Hong Kong, China, Mar. 2013), Springer, pp. 145–154.
- [27] MADORY, D. Iran Leaks Censorship via BGP Hijacks. Web, Jan. 2017. Dyn.
- [28] MADORY, D., POPESCU, A., AND ZMIJEWSKI, E. Accidentally Importing Censorship - The I-root instance in China. Presentation, June 2010. Renesys Corporation.
- [29] MAUCH, J. Open resolver project. Talk at DNS-OARC, May 2013.
- [30] MAXMIND. MaxMind GeoLite2. <http://dev.maxmind.com/geoip/geoip2/geolite2/>.
- [31] MOURA, G. C. M., DE O. SCHMIDT, R., HEIDEMANN, J., DE VRIES, W. B., MÜLLER, M., WEI, L., AND HESSELMAN, C. Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event. In *Proceedings of the ACM Internet Measurement Conference* (Nov. 2016), pp. 255–270.
- [32] PLANETLAB. <https://www.planet-lab.org/>.
- [33] POESE, I., UHLIG, S., KAAFAR, M. A., DONNET, B., AND GUEYE, B. IP geolocation databases: Unreliable? *SIGCOMM Comput. Commun. Rev.* 41, 2 (Apr. 2011), 53–56.
- [34] RIPE. RIPE Atlas measurements of B-Root. <https://atlas.ripe.net/measurements/>, measurement IDs 10310, 8310594, 8312460, 8312974, 8313009, 8313262, 2017.
- [35] RIPE NCC. DNSMON. <https://atlas.ripe.net/dnsmon/>.
- [36] RIPE NCC STAFF. RIPE Atlas: A Global Internet Measurement Network. *The Internet Protocol Journal* 18, 3 (Sept. 2015), 2–26.
- [37] ROOT DNS. <http://www.root-servers.org/>.
- [38] RSSAC. RSSAC advisory on measurements of the root server system. Tech. Rep. RSSAC002v3, ICANN, June 2016.
- [39] SCHMIDT, R. D. O., HEIDEMANN, J., AND KUIPERS, J. H. Anycast latency: How many sites are enough? In *Proceedings of the Passive and Active Measurement Workshop* (Sydney, Australia, Mar. 2017), Springer, pp. 188–200.
- [40] THOUSAND EYES. DNS monitoring service. Service description at <https://www.thousandeyes.com/solutions/dns>, 2017.
- [41] VIXIE, P., AND KATO, A. DNS referral response size issues. Work in progress (Internet draft draft-ietf-dnsop-respsize-14, May 2012).
- [42] DE VRIES, W. RIPE Atlas measurements of Tangled. Tangled_Atlas-20170201 <https://traces.simpleweb.org/> and RIPE Atlas measurement IDs 7794356-7794364, 2017.
- [43] DE VRIES, W. Verfloeter measurements of Tangled. dataset Tangled_Verfloeter-20170201 and Tangled_Verfloeter-20170323, <https://traces.simpleweb.org/>, 2017.
- [44] WEI, L., AND HEIDEMANN, J. Does anycast hang up on you? In *IEEE International Workshop on Traffic Monitoring and Analysis* (Dublin, Ireland, July 2017), IEEE, p. to appear.
- [45] WOOLF, S., AND CONRAD, D. Requirements for a Mechanism Identifying a Name Server Instance. RFC 4892, June 2007.
- [46] YURDAM, S. Using graffiti, Turks share tips for getting around Twitter ban. *The Observers* (21 March 2014).