A longitudinal study of IP Anycast

Danilo Cicalese
Telecom ParisTech
danilo.cicalese@enst.fr

Dario Rossi
Telecom ParisTech, Université Paris Saclay
dario.rossi@enst.fr

ABSTRACT

IP anycast is a commonly used technique to share the load of a variety of global services. Since more than one year, leveraging a lightweight technique for IP anycast detection, enumeration and geolocation, we perform regular IP monthly censuses. While this paper provides a brief longitudinal study of the anycast ecosystem, we make all our datasets (raw measurement from PlanetLab and RIPE Atlas), results (monthly geolocated anycast replicas for all IP/24) and code available to the community.

1 INTRODUCTION

IP anycast is an important building block of the current Internet, primarily used to share load [28] of a variety of global services – from DNS, to DDoS protection, to CDNs and content distribution, to even BitTorrent trackers and Internet radios [35].

Knowledge of IP anycast is instrumental not only for characterization, troubleshooting and infrastructure mapping [4] but also for security-related tasks such as censorship detection [55]. Yet, detailed knowledge and understanding of IP anycast in the scientific literature is generally limited to one or few deployments [13–16, 21, 24, 32, 44, 48, 49, 58, 62]. Fewer studies provide a broad spatial viewpoint [18, 50] and even fewer a temporal view [61]. The work focuses on *a broad and longitudinal view of anycast evolution*, that to the best of our knowledge has not appeared yet.

This paper is built on our own previous work [18, 20]. Shortly, [20] introduces a methodology that is able to (i) assert whether an IP is anycast, (ii) enumerate the replicas and (iii) geolocate them. It uses a set of latency measurement from a distributed set of vantage points with known location towards the same IP target. Our previous work [18] applies this methodology at scale, geolocating all the replicas for all IPv4 anycast at IP/24 level, through four censuses that refer to the same snapshot in time (March 2015). In this work, we extend findings in [18] along the temporal dimension, providing an analysis of monthly snapshots collected over more than a year-long period. Summarizing our main contributions:

- we conduct monthly anycast censuses at IP/24 level from distributed PlanetLab nodes, and conduct additional measurement from RIPE Atlas;
- we run our anycast geolocation algorithm [20] to build snapshots of anycast at IP/24, BGP announcement and AS levels, that we export as interactive tables and maps;
- based on these monthly censuses, we provide the first birdeye view of IPv4 anycast both from a spatial and a temporal viewpoint.

The rest of this paper puts this work in perspective with related effort (Sec.2), then describe our campaign (Sec.3) and comment our

main findings (Sec.4). To empower the community with the current state of anycast, as well as to enable further studies, we make all our raw dataset, results and code available at [5].

2 BACKGROUND

Anycast server enumeration and geolocalization is part of a broader effort from the research community to geographically map the Internet infrastructure and identify the various components of the physical Internet [25], possibly at scale, that we overview in the following .

Infrastructure mapping.. Techniques that are designed for application-level anycast are not applicable with IP-level anycast. There are only a handful techniques exist that allow to detect, enumerate or geolocate IP anycast replicas. Database-based techniques, that are unreliable with IP unicast [56], fail by definition with IP anycast, since they report a single geolocation per IP. Further, mapping techniques that exploit the EDNS-client-subnet (ECS) extension [17, 59] fail with anycast. Techniques relying on speed-of-light violation from ICMP measurements and BGP feeds [50] limitedly allow to detect anycast, but fail to provide replica geolocation. Techniques based on DNS queries of special class (CHAOS), type (TXT), and name (host-name.bind or id.server) provide reliable enumeration [31] but are DNS-specific and thus unsuitable to cover all services. While latency-based IP unicast geolocation [30, 37] is well understood, triangulation techniques do not apply in case of anycast, so that to the best of our knowledge, our previous technique [20] is the first and only to provide accurate geolocation of anycast replicas by only leveraging protocol-agnostic delay information.

While it is outside the scope of this work to recall the technique in details [20], to make this paper self-contained it is sufficient to state that the technique builds on inferring IP anycast by detecting speed-of-light violations via latency measurements: i.e., as packets travel slower than speed of light, an US and EU host probing the same target cannot *both* exhibit excessively low latency (e.g., few milliseconds) toward the *same* target, as this would violates physical laws. While this observation is not new [50], our iGreedy technique does [20] phrase the problem in terms of finding the maximum number of vantage points that are in such violation: by definition, these vantage points all contact a different anycast replica of the same IP target t. By extension, the location of a vantage point i that is found violating the speed-of-light constraint assists in geolocating the replica of t contacted by i: by definition, this replica is contained into a circle centered in the vantage point i and that stretches by at most the distance that the probe packet can have traveled during $RTT_{i,t}/2$. It turns out that, with the goal of attaining city-level precision, a very simple yet very good criterion is to choose the position of the most inhabited city as likely location of the t anycast replica. This follows from the fact that the decision to add an anycast replica, follows from the goal of ameliorating

performance for a large fraction of users, which live in large cities (interestingly, this was already used to bias geolocation of *unicast* addresses[29]). Overall, the technique has high recall (i.e., over the 75% of replicas are detected) and precise geolocation (i.e., over the 75% matches at city-level, and the average error in the remaining erroneous case is 384 Km).

Anycast characterisation. Research on anycast has so far prevalently focused either on architectural modifications [13, 32, 33, 45] or on the characterization of existing anycast deployments. Overall, a large fraction of these studies quantify the performance of anycast in current IP anycast deployments in terms of metrics such as proximity [13, 14, 21, 49, 58], affinity [13–16, 48, 49, 58], availability [14, 44, 48, 58], and load-balancing [14]. Interestingly, while the body of historical work targets DNS, more recent work [32, 48] has tackled investigation of anycast CDN performance (e.g., client-server affinity and anycast prefix availability for the CacheFly CDN). More recently, [24] investigate DNS root servers, outlining a rule of thumb to determine the right number of anycast replicas, whereas [61] investigates affinity of DNS root servers over a period of two weeks in two different years. We are not aware of any other studies presenting a more systematic temporal analysis than [61], and clearly none targeting a larger spatial set than DNS root servers.

In [18] we leverage measurement infrastructures, namely Planet-Lab and RIPE Atlas, to perform Internet-scale census of IP anycast, by actively probing all /24 subnets and geolocating anycast replicas, finding that only a tiny fraction (0.03%) of IP/24 are anycast – i.e., it appears that finding anycast deployments is like a finding a needle in the IP haystack. At the same time, by *actively* probing these anycast targets, we also unveil that all major Internet players do use anycast and that a wide variety of services are used. We instead use a complementary approach in [35], where we *passively* inspect the anycast traffic at one specific DSLAM in EU, to assess their actual usage in real networks (users have a 50% chance to encounter one anycast instance in their daily activities, that even include radio streaming session last for hours, as well as BitTorrent anycast trackers).

However, while [18] and [35] present a very complete and detailed view of the *spatial* characteristics of anycast deployment, e.g. their geographical distribution, the services offered over anycast (active inspection) and their usage (passive monitoring), these studies represent a snapshot at a fixed point in time. This is orthogonal with respect to the focus of this work, that instead presents a longitudinal study of anycast, based on monthly censuses we run since December 2015.

Internet Censuses. In the past, several studies focus their attention on scaling active scanning techniques to provide broad spatial surveys of the Internet infrastructure. Given the lack of high-rate scanning tools (such as [27, 47, 51]) at that time, researchers have studied sample of the Internet-space [42] or have splitted the IPv4 space over multiple vantage points [1, 39, 40] or completed the scans in an extended period of time. Since 2006, authors [39] measures periodically the population of visible Internet edge hosts (at IP/32 level) from eight different vantage points in two different locations, providing an IPv4 hitlist (one likely alive IP/32 target per

IP/24). In 2008, authors in [22] scanned the Internet to find DNS servers that provide incorrect resolutions. In 2010, the IRLscanner tool allowed to scan the IP/32 Internet in about 24 hours, and results from 21 Internet-wide scans using 6 different protocols have then been presented in [47]. In 2012, the (highly discussed[46]) Carna Botnet [1] has used 420k insecure embedded devices to build a distributed port scanner to scan all IPv4 addresses using nmap [53].

In the recent years, the situation has drastically changed with the advent of new network scanner tools as ZMap [10] and Masscan [51], able to achieve scan rates in excess of 10 Mpps, which let a IP/32 scan complete in less than five minutes. This has led to a huge increase of sporadic and regular scans, including the malicious ones: as documented in [26], using a network telescope, authors detected over 10 million scans from about 1.5 million hosts during January 2014. These are mainly regular scans, with daily [3] or lower frequency [52, 57]. Despite only a tiny fraction of these scans target more than 1% of the monitored IPv4 address space, they generate the majority of the unsolicited traffic hitting the darknet. Anycast censuses, such as those we performed in [18], raise the spatial requirement to another level, since the same target needs to be actively probes from multiple vantage points: experimentally, several hundreds vantage points have been shown to provide a good geographical coverage [20]. As such, to the best of our knowledge we are not aware of any IPv4 anycast survey with the exception of [18].

3 MEASUREMENT CAMPAIGN

Platforms. Several measurement platforms[6, 9, 11, 36] exist that have different characteristics in terms of vantage points cardinality, AS diversity, geographic coverage and limits, in terms of probing traffic or rate[12]. In this paper we make use of two platforms, namely PlanetLab[6]and RIPE Atlas[9], that we select due to their complementarity. Specifically, PlanetLab does not enforces a specific probing limit, nor implement a credit system: we use it to perform exhaustive censuses to a large set of targets, that we expect to be mostly unicast. Conversely, RIPE Atlas has better coverage: we use it to refine the information concerning a specific subset of targets, i.e., those that were found to be anycast with PlanetLab.

Targets. For the selection of the targets, we rely on the USC/LANDER hitlist [41], providing a list of (likely alive) target IP/32 host per each /24 prefix. Every two months the list is updated, and so our target selection. We only consider hitlist IPs that have been successfully contacted (i.e., denoted by a positive score [41]), which leaves us about 6.3 millions potential targets (out of 14.7 millions). We argue that /24 is a reasonable granularity since (in principle) IP-anycast is subject to BGP prefix announcement that should not (but seldom are) more specific than /24. Additionally, while in principle we could use one IP/32 per each announced BGP prefix, [2] observes that prefixes longer than /24 have low visibility: as such, we limit the granularity to IP/24 level (i.e., one IP/32 per /24) targeting less than *0.4%*of the whole address space.

Vantage points. For the selection of the vantage points (VP), we proceed following the guidelines in [19]: in PlanetLab, where the total number of VPs is small, we simply select all the available; in
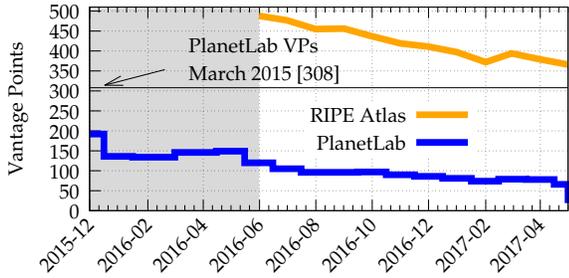
Figure 1: Measurement campaign: evolution of number of PlanetLab and RIPE Atlas VPs

Figure 2: Measurement campaign: boxplot of the number of responsive unique IPs/24 across all PlanetLab VPs.

RIPE Atlas, where the number of VPs is large and due due to credits limit, we carefully select 500 VPs, making sure that each VP is far from the others by at least 200 km (roughly 2ms).

Fig. 1 shows the evolution of the number of available vantage points over time. In the PlanetLab case, the number of nodes available drastically decrease from 300 in March 2015 [18] to roughly 50 in May 2017. In the case of RIPE Atlas, the decrease is due to the fact that we launched a long-standing periodic measurement in June 2016 with an initial set of VPs, some of which later become unavailable. Interestingly, we will see that anycast results appears to be consistent despite this decrease.

Despite a handful of carefully selected vantage points [34] allow to *correctly detect* anycast deployments, it is clear that the shrinking size of the available PlanetLab VP it is not adequate to *thoroughly enumerate* all the locations of an anycast deployment. – for which RIPE Atlas measurements become necessary as we shall see.

Censuses. Anycast censuses require the same target to be probed from multiple vantage points: to limit the intrusiveness of our scans, and since we expect that changes in the anycast deployments happen at a low pace, we decide to run scans at a *monthly* frequency. Our first anycast censuses date back to March 2015 [18]. We then re-engineered our system and started to run monthly censuses from PlanetLab in December 2015. We kept tuning and improving system performance and reliability until June 2016, date at which we additionally started the measurements from RIPE Atlas. We opted to strip down as much as possible the information collected per VP, narrowing down to about 30MB per VP on average, so that the (compressed) raw PlanetLab measurement data hosted at [5] for all censuses amount to about 60GB. The RIPE Atlas measurement are publicly accessible via RIPE Atlas (measurement identifier are at [5]).

Fig. 2 shows the number of responsive unique IPs discovered collectively by all PlanetLab VP in each census, and the right y-axis report this number as the fraction of replies from the contacted targets. The shadowed part indicates the months where we were still updating the system. Notably, we slowed down the probe rate per VP to about 1,000 targets per second to comply with recommendations in [38]. As a beneficial side effect, the packet loss rate also decreased. We can see that the total number of unique IPs is always greater than the number observed by a single VP, and that
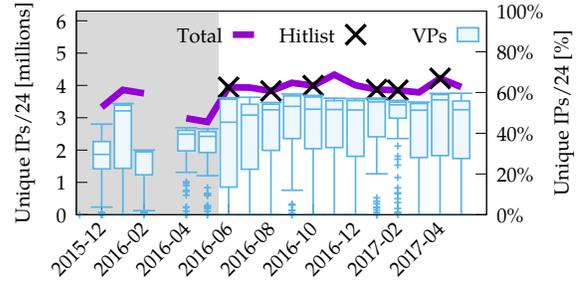
it fluctuates between the 2, 9 millions of May 2015 and 4, 3 millions of November 2016 (almost in pair with the expectations of other work [23, 63]). This number has increased since June 2016 when we started to regularly update the hitlist [41] (denote with crosses in the figure). The figure also reports the distribution of responsive IPs per vantage points (boxplots): recall varier widely per census, per VP, and over time, with some VPs able to collect only few hundred ICMP replies. Luckily, albeit the number of PlanetLab VP decreases, the median number of contacted target exceeds 3 millions.

4 RESULTS

This section provides a longitudinal view of anycast evolution. We report both a broad picture including all deployments (Sec.4.1), as well as a more detailed view by cherry-picking some representative ones (Sec.4.2). Without loss of generality, we refer to the last year worth of censuses collected between May 2016 and May 2017.

4.1 Broad view

Longitudinal view. First, we assess the extent of variability of anycast deployments. We start by considering an IP/24 granularity, and depict in Fig. 3 the evolution of the number of IP/24 anycast deployments, i.e., the number of deployments that have been found to be anycast by running iGreedy[20] over PlanetLab measurements. We recall that iGreedy requires to solve a Maximum Independent Set (MIS) optimization problem for each of the over 4 million responsive targets every month: the code available on GitHub [7] is able to complete the analysis of a census in few hours, which returns the set of geolocated replicas $\mathcal{G}_t$ for each responsive IP/24 target t. While full details of the geolocation for each target and over all months are available online as a Google-map interface[5], in this paper we limitedly consider the footprint $G_t = |\mathcal{G}_t|$ of the deployment, i.e., the number of distinct instances irrespectively of their location.

The figure shows that in our censuses, the number of anycast deployments has slightly increased in the last year, peaking in April 2017 at 4729 IP/24 belonging to 1591 routed BGP prefixes and 413 ASes. In the last six months, the number of anycast deployments has never dropped below 4500 while in June 2016, when we started the censuses regularly, we found only 4297 IP/24 (1507 routed BGP prefixes and 379 ASes). Compared to our previous results of March 2015 [18], this represents a 2,5-fold increase over a period of 2 years. This may be due to several reasons: part of the increase is rooted
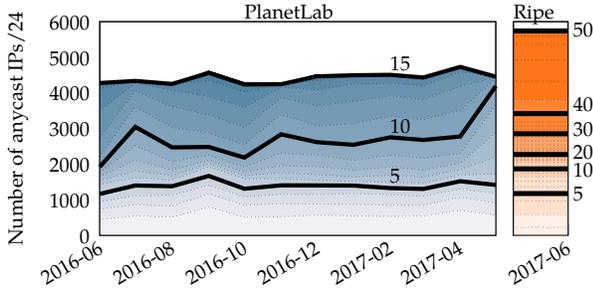
Figure 3: Broad longitudinal view of anycast evolution: Number of IP/24 anycast deployments (y-axis) and breakdown of their geographical footprint (heatmap and contour lines) in PlanetLab (left, over the last year) vs RIPE Atlas (right, last month)

.

in anycast adoption over time first, and another part is rooted in system improvements to reduce packet losses at PlanetLab monitors, which increases the recall. This also means that results in [18] were fairly conservatively assessing the extent of the anycast Internet.

Fig. 3 additionally encodes, as a heatmap, the estimated geographical footprint G_t, where deployments are ranked from bottom to top in ascending size (equivalently, darker colors). A few contour lines indicate the number of cumulative deployments having no more than 5, 10 or 15 replicas. Interestingly, Fig. 3 shows that, despite a shrinking number of PlanetLab VPs, the number of anycast IP/24 remains steady over time. Particularly, the number of deployments having few replicas (e.g., 5 or less) remains flat over time, hinting to the fact that the geographical coverage of PlanetLab is still enough to correctly detect most

Yet, as previously observed, the shrinking number of PlanetLab VPs surely affects the completeness of the replica enumeration. We thus complement PlanetLab censuses with a refinement campaign from RIPE Atlas, which is also reported in Fig. 3: during June 2017, we target all IP/32 that have been found to be anycast in PlanetLab during the previous year. Out of the overall 5841 IP/24s, approximately 300 were not reachable in June 2017 and 5105 IP/24s are confirmed to be still anycast. Particularly, we used 500 RIPE Atlas VPs, i.e., about one order of magnitude more than PlanetLab, which ensures a good geographic coverage. Thus, while PlanetLab may provide a rather conservative lower bound of the actual footprint for a target t, we expect $G_t^{RIPE} > G_t^{PL}$. Fig. 3 confirms these expectations: in several cases, the number of anycast instances discovered in RIPE Atlas doubles with respect to PlanetLab, and the maximum number exceeds 49 replicas (18 in PlanetLab). Overall, according to RIPE Atlas, half of the deployments are in more than 5 different locations, but only few of them have more than 35 locations (including DNS root servers, Verisign, Microsoft, WoodyNet and Cloudflare). Consider also that as a consequence of the drop in the number of PlanetLab VPs in the the last months, the largest footprint measurable from PlanetLab drops as well (notice the sharp increase for deployments with at most 10 replicas). This confirms that PlanetLab remains useful for anycast detection, but also that RIPE Atlas becomes necessary for enumeration and geolocation,
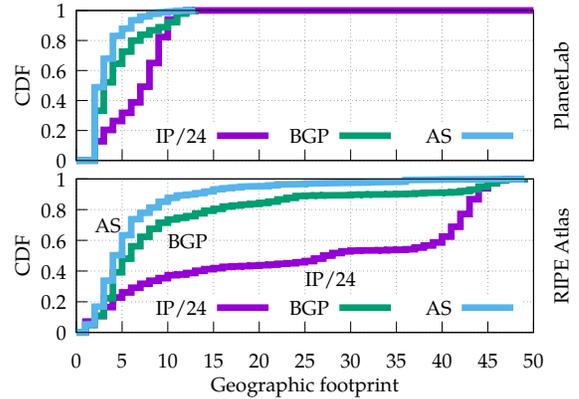

Figure 4: Distribution of the geographic footprint of anycast deployment at IP/24 (G_t), BGP-announced prefix (G_B) and AS level (G_A). Results from PlanetLab (top, all months) vs RIPE Atlas (bottom, last month).

reinforcing the need for a more systematic coupling of PlanetLab and RIPE Atlas measurement.

Aggregation level. Clearly, while we operate censuses at IP/24 level, it is then possible to aggregate the information at BGP or AS level. Denoting with $\mathcal{S}_x$ the set of IP/24 included in a BGP-announced prefix (or an AS) x, we can define the *spatial IP footprint* as $S_x = |\mathcal{S}_x|$. By extension, we can define the BGP-level (AS-level) geographic footprint G_B by considering only the largest IP/24 in the prefix $G_B = \max_{t \in \mathcal{S}_B} G_t$ (or AS $G_A = \max_{t \in \mathcal{S}_A} G_t$). To perform this aggregation step, for each month in the census, we retrieve the AS and prefix information using all the RIPE-RIS and RouteViews collectors with BGPStream [54], and cross-validate the information using the TeamCymru IP to ASn service [60].

The different viewpoints are illustrated in Fig.4 that reports for PlanetLab (top, all months) vs RIPE Atlas (bottom, last month) the cumulative distribution function of the geographic footprint at IP/24, BGP-announced prefix and AS levels. The geographic footprint per-IP/24 vs per-BGP/AS varies widely, which is due to the fact that the spatial distribution is highly skewed, so that ASes making use of a large number of IP/24 to be over-represented. Particularly, while more than 50% AS (75% of BGP announced prefixes) make use of a single anycast IP/24, about 10% ASes (BGP prefixes) hosts more than 10 anycast IP/24, topping to 384 (for 104.16.0.0/12) and 3016 (for AS13335). Since all three level of aggregation have relevance to give an unbiased picture of Internet anycast, we make available monthly snapshots with IP/24, BGP and AS aggregations as tabular data [5], which is also browsable online with a Slickgrid interface.

Finally, as a rough measure of persistence of individual anycast deployments, Fig.5 depicts a breakdown of the number of months that they are present in our censuses at IP/24, BGP or AS levels. Notice that over 45% of anycast ASes (60% of anycast IP/24) consistently appear in our measurements for the whole year and 70% AS (78% IP/24) appear at least 6 months. Only less than 10% deployments are seen only once.
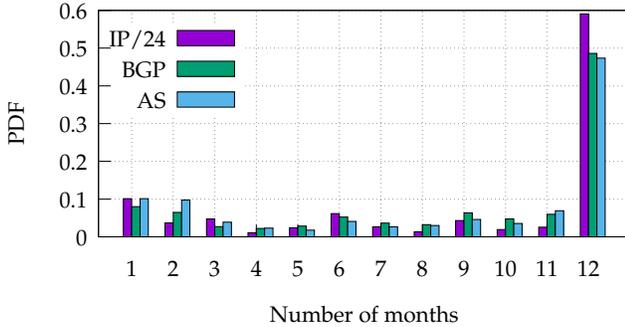
Figure 5: Anycast deployments stability over time: number of censuses that the IP/24, BGP prefix or AS is present over the one year observation period.

4.2 Focused view

Top-10 deployments. We now provide a more detailed view of a few selected ASes out of the 566 in our censuses. Particularly, Tab. 1 reports detail concerning the top-10 deployments (company name and type, AS number and the number of BGP prefixes announced by that AS), the spatial footprint (i.e, the number S_A of IP/24 per AS and its temporal variability) and the geographical footprint (i.e, the number G_A of distinct replicas and its temporal variability). To compactly represent the size of a deployment, we report the maximum number S_A^+ of observed anycast IP/24 over the last year for that AS, as well as the $G_A^+ = \max G_A^{RIPE}(t)$ maximum number of locations observed from RIPE Atlas. Selection in Tab. 1 reports the top-5 in terms of S_A^+ spatial footprint (top) and the top-5 for G_A^+ geographic footprint (bottom).

Considering the spatial footprint IP/24, Cloudflare (AS13335) has a leading role: it is present in all the censuses with over 3 thousands IP/24 belonging to about 200 announced prefixes (mainly /20 but also less specific prefixes, as a /12 or a /17), and we did not observe significant variation over time. Furthermore, as confirmed from RIPE Atlas, the deployment has an heterogeneous geographical footprint, with some /24 having only 10–15 instances, while in the majority of the cases the /24 appear at over 40 distinct locations. Notice that, this would had been unnoticed if we had performed censuses at a different granularity (i.e., one IP/32 per BGP prefix as opposite as to one IP/32 per IP/24 in that prefix). Few other companies have over 100 anycast IP/24 prefixes in our censuses. For instance, Google (AS15169) had a 3-fold increase in the number of IPs/24 in the last year, from 130 IPs/24 announced mainly by /16 in June 2016, to 330 IP/24 announced also by 190 new /13 prefixes in March 2017. The majority of Google /24 have instances at more than 30 locations. Opposite behavior are also possible: for instance, Fastly (AS54113), a Content Delivery Network shrunk its spatial footprint, the majority of which belong to an IP/16 and regularly appear in all our censuses. Interestingly, as early depicted in Fig. 3, the overall aggregate of all anycast deployments (i.e., the number of anycast /24 in the Internet and their geographical breakdown) is stable despite the variability of the individual deployments.

The common way to deploy anycast is to announce an IP prefix from multiple points using the same AS [43], that we refer to Single

Table 1: Focused view on footprint variability of top-5 spatial (top) and top-5 geographical (bottom) deployments

Deployment footprint:				Spatial		Geographical	
Company	AS	Type	BGP	S_A^+	CV_S	G_A^+	CV_G
Cloudflare	13335	CDN	206	3016	0.04	49	0.07
Google	15169	SP	16	524	0.38	30	0.08
Afilias	12041	TLD	218	218	0.15	6	0.10
Fastly	54113	CDN	34	175	0.09	20	0.07
Incapsula	19551	DDoS	146	146	0.23	15	0.17
Cloudflare	13335	CDN	206	3016	0.04	49	0.07
L root	20144	DNS	1	1	0	47	0.13
F root	3557	DNS	2	2	0	40	0.19
Woodynet	42	TLD	132	133	0.02	39	0.12
Verisign	26415	Reg.	2	2	0	36	0.20

Origin AS (SOAS). Another way is to announce the IP prefix using multiple ASes, usually referred as Multiple Origin ASes (MOAS) prefixes. In our dataset, we identified hundred IPs/24 as MOASes, that are commonly announced by few siblings, i.e., different ASes belonging to the same organization that announce the same prefix. However we spot cases where the number of ASes is greater than 10: for instance, we find that Verisign announces MOASes with 17 different ASes in the range of AS36617-AS36632; similarly, the Registry and DNS company AusRegistry, announces MOASes with 13 different ASes.

To compactly represent the temporal variability of spatial footprint, we use the coefficient of variation, computed as the ratio of the standard deviation over the average number of anycast IP/24 per month $CV_S = \text{std}(S_A(t))/\mathbb{E}[S_A(t)]$, From Tab. 1, we can see that for deployments that have large spatial footprint (top), the variability CV_S can be important (e.g., Google or Incapsula), hinting to deployments that have grown (or shrunk) significantly. Conversely, among deployments with large geographical footprint, several have a very small spatial footprint ($S_A^+ \geq 2$) and exhibit no variation $CV_S = 0$.

Finally, as simple indicator of geographical footprint variability we compute $CV_G = \text{std}(G_A^{PL}(t))/\mathbb{E}[G_A^{PL}(t)]$ from PlanetLab measurements. Notably, we expect part of the variability to be due to measurement imprecision: e.g., shrinking number of VPs, packet losses and increased delay, can lead to underestimate the number of distinct locations. Yet, as it can be seen from Tab. 1, we find that the geographical variability is lower than the spatial one: this is reasonable since, while spatial variability hints to configuration changes in software, the geographical one hints to physical deployments of new hardware.

Temporal variability. We now inspect temporal variability at a finer grain. We start by depicting in Fig.6 the temporal evolution of the spatial footprint, normalized over the maximum observed for that deployment, i.e., $S_A(t)/max_t S_A^+(t)$. The picture includes deployments in the anycast top-5 (Afilias, Google) as well as anycast deployments of other key Internet players (Microsoft, Akamai, Netflix, Windstream). Evolutions represent a sample of what can be found in our censuses: for instances the two ASes reported in the
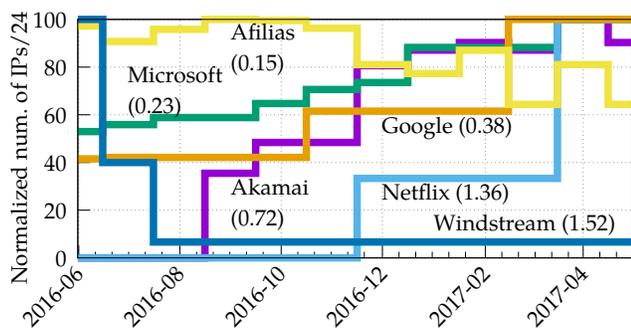
Figure 6: Spatial footprint evolution: Number of IPs/24 for selected anycast AS deployments (PlanetLab).

Figure 7: Geographical evolution: Number of locations for selected deployment (RIPE Atlas)

picture owned by Akamai (AS21342) and Neflix (AS40027) start being announced as anycast during our observation period and either systematically (Akamai) or abruptly (Neflix) increase the amount of responsive anycast IP/24 over time. Google (AS15169) and Microsoft (AS8068) both have a sizeable presence at the beginning of the observation period, with roughly 50% of the IP/24 already in use, and roughly double the amount of IP/24 used at the end of the period in a smooth (Microsoft) or abrupt (Google) fashion. Finally, close to the beginning of our observation period, Windstream drastically reduces its anycast spatial footprint, keeping just a single anycasted IP/24. While these observations have anecdotal value, and cannot explain the reason behind changes in the deployment, they however confirm that anycast deployment have a rather lively temporal evolution, the extent of which is captured by the coefficient of variation.

It is intuitive that the number (and location) of vantage points upper-bounds (and constrains) the number of anycast instances that can be found. Given the slow but steady decrease of the PlanetLab VPs, we unfortunately do not deem PlanetLab measurement reliable in assessing, at a fine grain, the geographic growth (which can be underestimated) or shrink (which can be due to VP decrease) of anycast deployments. We thus decided to regularly monitor anycast prefixes using 500 Ripe Atlas VPs. We picked targets from two key CDN players, namely Cloudflare (8 different IP/20) and Fastly (5 different IP/24). As per Tab. 1, Cloudflare is the top-1 player over all anycast, and given its sheer footprint, we do not expect it to further grow: especially, the number of locations it appears already significantly exceeds (by more than a factor of 4×) the number[1] suggested in [24]. As such, we use Cloudflare as a litmus paper for our measurement. Our selection of Fastly is then motivated by the fact that despite it appears in the top-5 and so we expect it to steadily appear in our measurement, it has 1/10 the spatial footprint and 1/2 the geographic footprint of Cloudflare: so it not only has room to grow, but also possibly has the money necessary for the investment.

In the case of Cloudflare, Fig. 7 shows that, as expected the number of instances is stable , so that fluctuations are only measurement
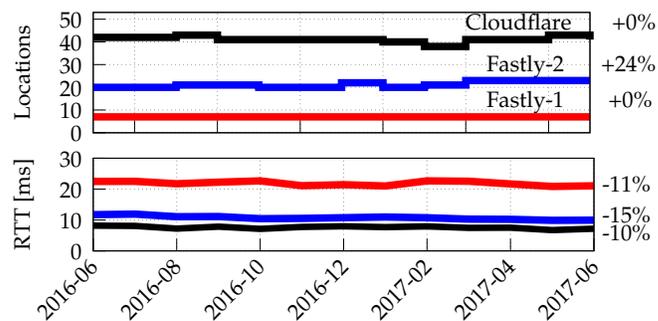
artifacts. In the case of Fastly, we observed a general growth from 19 locations in June 2016 to 24 in June 2017 (except for a stable IP/24 with only 7 different locations). For references purposes, 33 locations are mentioned in [8], which corresponds to a 73% recall, in line with expectations for the iGreedy methodology [19]). Latency measurements are shown in the bottom part of the figure. We can observe a 10% of latency reduction also for stable deployments (Fastly-1 and Cloudflare) and it thus to be imputed to other causes (e.g., increased peering connectivity). At the same time, at least for Fastly, it appears that increasing the number of instances reduces the average latency toward our RIPE Atlas probes by an additional 5%, and *halves* the 95th percentile (from 137ms in June 2016 to 68ms in June 2017), unlike expectations [24].

5 CONCLUSIONS

Internet anycast is important building block of the current Internet: the study of deployments and their evolution is useful to enrich our understanding of Internet operations. In this longitudinal study, we learn that anycast detection (important for censorship studies[55]) is reliable in spite of varying (and especially diminishing) vantage points. We additionally see that anycast *spatial footprint* evolves significantly for individual deployments, though it remains steady in the aggregate. PlanetLab censuses can reliably measure this variability. However, while we gather that anycast *geographical footprint* evolves, a more tight coupling with RIPE Atlas would be needed, due to decreasing PlanetLab VPs, to accurately track the state of anycast Internet at replica level (e.g., monthly detection from PlanetLab, followed from a refinement of geolocation for detected anycast deployments).

Finally, by closely monitoring a few deployments with RIPE Atlas, we gather that anycast deployment that already have a large geographical footprint, apparently benefit (in that their service latency decreases) from further growing the deployment beyond sound rules of thumb [24], which requires more systematic investigation. We believe that making this knowledge, and especially datasets and tools, available to the scientific community contributes to enrich the Internet map [4] along another dimension.

---

[1]"After carefully studying four very different anycast deployments, we claim that 12 anycast sites would be enough for good overall latency. Two sites per continent, in well chosen and well connected locations, can provide good latency to most users." [24]

REFERENCES

[1] http://internetcensus2012.github.io/InternetCensus2012/paper.html. (2012).
[2] https://labs.ripe.net/Members/emileaben/has-the-routability-of-longer-than-24-prefixes-changed. (2013).
[3] Shadowserver open resolver scanning project. dnsscan.shadowserver.org. (2013).
[4] labs.ripe.net/Members/emileaben/measuring-ixps-with-ripe-atlas. (2015).
[5] http://www.enst.fr/~drossi/anycast. (2017).
[6] https://www.planet-lab.org. (2017).
[7] https://github.com/TeamRossi/anycast-census. (2017).
[8] Fastly location. https://www.fastly.com/network-map/. (2017).
[9] Ripe Atlas. https://atlas.ripe.net. (2017).
[10] David Adrian, Zakir Durumeric, Gulshan Singh, and J Alex Halderman. 2014. Zippier ZMap: Internet-Wide Scanning at 10 Gbps.. In WOOT.
[11] Ark. http://www.caida.org/projects/ark/. (2007).
[12] V. Bajpai and J. Schonwalder. 2015. A Survey on Internet Performance Measurement Platforms and Related Standardization Efforts. IEEE Communications Surveys & Tutorials (2015).
[13] Hitesh Ballani and Paul Francis. 2005. Towards a Global IP Anycast Service. In ACM SIGCOMM.
[14] Hitesh Ballani, Paul Francis, and Sylvia Ratnasamy. 2006. A measurement-based deployment proposal for IP anycast.. In ACM IMC.
[15] Biet Barber, Matt Larson, and Mark Kosters. Traffic Source Analysis of the J Root Anycast instances. 39th Nanog. (2006).
[16] Peter Boothe and Randy Bush. DNS Anycast Stability: Some Early Results. 19th APNIC. (2005).
[17] Matt Calder, Xun Fan, Zi Hu, Ethan Katz-Bassett, John Heidemann, and Ramesh Govindan. 2013. Mapping the expansion of Google's serving infrastructure. In ACM IMC.
[18] Danilo Cicalese, Jordan Auge, Diana Joumblatt, Timur Friedman, and Dario Rossi. 2015. Characterizing IPv4 Anycast Adoption and Deployment. In ACM CoNEXT.
[19] Danilo Cicalese, Diana Joumblatt, Dario Rossi, Marc-Olivier Buob, Jordan Augé, and Timur Friedman. 2015. Latency-Based Anycast Geolocalization: Algorithms, Software and Datasets. http://www.enst.fr/ drossi/dataset/anycast/anycast-techrep.pdf. In Tech. Rep.
[20] Danilo Cicalese, Diana Joumblatt, Dario Rossi, Marc-Olivier Buob, Jordan Auge, and Timur Friedman. 2016. Latency-Based Anycast Geolocalization: Algorithms, Software and Datasets. IEEE JSAC (2016).
[21] Lorenzo Colitti. Measuring anycast server performance: The case of K-root. Nanog. (2006).
[22] David Dagon, Niels Provos, Christopher P Lee, and Wenke Lee. 2008. Corrupted DNS Resolution Paths: The Rise of a Malicious Resolution Authority.. In NDSS.
[23] Alberto Dainotti, Karyn Benson, Alistair King, Bradley Huffaker, Eduard Glatz, Xenofontas Dimitropoulos, Philipp Richter, Alessandro Finamore, and Alex C Snoeren. 2016. Lost in Space: Improving Inference of IPv4 Address Space Utilization. IEEE J-SAC (2016).
[24] Ricardo de Oliveira Schmidt, John Heidemann, and Jan Harm Kuipers. 2017. Anycast Latency: How Many Sites Are Enough?. In PAM.
[25] R. Durairajan, S. Ghosh, X. Tang, P. Barford, and B. Eriksson. 2013. Internet Atlas: A Geographic Database of the Internet. In HotPlanet.
[26] Zakir Durumeric, Michael Bailey, and J Alex Halderman. 2014. An Internet-Wide View of Internet-Wide Scanning.. In USENIX Security Symposium. 65–78.
[27] Z. Durumeric, E. Wustrow, and J. Alex Halderman. 2013. ZMap: Fast Internet-wide Scanning and Its Security Applications. In USENIX Security Symposium.
[28] Robert Engel, Vinod Peris, Debanjan Saha, Erol Basturk, and Robert Haas. 1998. Using IP Anycast For Load Distribution And Server Location. In Proc. IEEE GI.
[29] B. Eriksson, P. Barford, J. Sommers, and R. Nowak. 2010. A Learning-based Approach for IP Geolocation. In PAM.
[30] Brian Eriksson and Mark Crovella. 2013. Understanding geolocation accuracy using network geometry. In IEEE INFOCOM.
[31] Xun Fan, John S. Heidemann, and Ramesh Govindan. 2013. Evaluating anycast in the domain name system.. In IEEE INFOCOM.
[32] Ashley Flavel, Pradeepkumar Mani, David A Maltz, Nick Holt, Jie Liu, Yingying Chen, and Oleg Surmachev. 2015. FastRoute: A Scalable Load-Aware Anycast Routing Architecture for Modern CDNs. In USENIX NSDI.
[33] Michael J. Freedman, Karthik Lakshminarayanan, and David Mazières. 2006. OASIS: Anycast for Any Service. In USENIX NSDI.
[34] Gregoire Gallois-Montbrun and Victor Nguyen. Defining a minimum vantage point selection to detect BGP Hijack with iGreedy. http://perso.telecom-paristech.fr/ drossi/teaching/INF570/projects/2015-paper-3.pdf. (2015).
[35] Danilo Giordano, Danilo Cicalese, Alessandro Finamore, Marco Mellia, Maurizio Munafo, Diana Joumblatt, and Dario Rossi. 2016. A First Characterization of Anycast Traffic from Passive Traces. In TMA.
[36] Vasileios Giotsas, Amogh Dhamdhere, and Kimberly C Claffy. 2016. Periscope: Unifying looking glass querying. In PAM. Springer, Cham, 177–189.
[37] B. Gueye, A. Ziviani, M. Crovella, and S. Fdida. 2004. Constraint-based geolocation of internet hosts.. In ACM IMC.
[38] Hang Guo and John Heidemann. 2017. Detecting ICMP Rate Limiting in the Internet. Technical Report.
[39] J. Heidemann, Y. Pradkin, R. Govindan, C. Papadopoulos, G. Bartlett, and J. Bannister. 2008. Census and Survey of the Visible Internet. In Proc. ACM IMC.
[40] Nadia Heninger, Zakir Durumeric, Eric Wustrow, and J Alex Halderman. 2012. Mining Your Ps and Qs: Detection of Widespread Weak Keys in Network Devices.. In USENIX Security Symposium.
[41] hitlist. Internet Addresses Hitlist Dataset (20140829/rev4338). Provided by the USC/LANDER project (http://www.isi.edu/ant/lander). (2006).
[42] Ralph Holz, Lothar Braun, Nils Kammenhuber, and Georg Carle. 2011. The SSL landscape: a thorough analysis of the x. 509 PKI using active and passive measurements. In Proc. ACM IMC.
[43] Xin Hu and Z Morley Mao. 2007. Accurate real-time identification of IP prefix hijacking. In Security and Privacy, 2007. SP'07. IEEE Symposium on.
[44] Daniel Karrenberg. Anycast and BGP Stability: A Closer Look at DNSMON Data. 34th Nanog. (2005).
[45] Dina Katabi and John Wroclawski. 2000. A Framework for Scalable Global IP-anycast (GIA). In ACM SIGCOMM.
[46] Thomas Krenc, Oliver Hohlfeld, and Anja Feldmann. 2014. An internet census taken by an illegal botnet: a qualitative assessment of published measurements. ACM CCR (2014).
[47] Derek Leonard and Dmitri Loguinov. 2013. Demystifying internet-wide service discovery. IEEE/ACM Transactions on Networking (2013).
[48] Matt Levine, Barrett Lyon, and Todd Underwood. Operational experience with TCP and Anycast. 37th Nanog. (2006).
[49] Ziqian Liu, Bradley Huffaker, Marina Fomenkov, Nevil Brownlee, and Kimberly C. Claffy. 2007. Two Days in the Life of the DNS Anycast Root Servers.. In PAM.
[50] Doug Madory, Chris Cook, and Kevin Miao. Who are the anycasters. Nanog. (2013).
[51] masscan. https://github.com/robertdavidgraham/masscan. (2013).
[52] Ariana Mirian, Zane Ma, David Adrian, Matthew Tischer, Thasphon Chuenchujit, Tim Yardley, Robin Berthier, Joshua Mason, Zakir Durumeric, J Alex Halderman, and others. 2016. An Internet-Wide View of ICS Devices. In IEEE PST.
[53] Nmap. https://nmap.org. (1997).
[54] Chiara Orsini, Alistair King, Danilo Giordano, Vasileios Giotsas, and Alberto Dainotti. 2016. BGPStream: a software framework for live and historical BGP data analysis. In Proc. ACM IMC. ACM.
[55] Paul Pearce, Roya Ensafi, Frank Li, Nick Feamster, and Vern Paxson. 2017. Augur: Internet-Wide Detection of Connectivity Disruptions. In IEEE SP.
[56] Ingmar Poese, Steve Uhlig, Mohamed Ali Kaafar, Benoit Donnet, and Bamba Gueye. 2011. IP Geolocation Databases: Unreliable? ACM SIGCOMM CCR (2011).
[57] Christian Rossow. 2014. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In NDSS.
[58] S. Sarat, V. Pappas, and A. Terzis. 2006. On the Use of Anycast in DNS. In ICCCN.
[59] Florian Streibelt, Jan Böttger, Nikolaos Chatzis, Georgios Smaragdakis, and Anja Feldmann. 2013. Exploring EDNS-client-subnet Adopters in Your Free Time. In ACM IMC.
[60] TeamCymru. http://www.team-cymru.org/Services/ip-to-asn.html. (2007).
[61] Lan Wei and John Heidemann. 2017. Does Anycast Hang up on You?. In TMA.
[62] Jing'an Xue, David Choffnes, and Jilong Wang. 2017. CDNs Meet CN. In IEEE Access.
[63] Sebastian Zander, Lachlan LH Andrew, and Grenville Armitage. 2014. Capturing ghosts: Predicting the used IPv4 space by inferring unobserved addresses. In Proc. ACM IMC.