

Anycast on the Move: A Look at Mobile Anycast Performance

Sarah Wassermann
Inria Paris
sarah.wassermann@inria.fr

John P. Rula
Northwestern University
john.rula@eecs.northwestern.edu

Fabián E. Bustamante
Northwestern University
fabianb@eecs.northwestern.edu

Pedro Casas
AIT Vienna
pedro.casas@ait.ac.at

Abstract—The appeal and clear operational and economic benefits of anycast to service providers have motivated a number of recent experimental studies on its potential performance impact for end users. For CDNs on mobile networks, in particular, anycast provides a simpler alternative to existing routing systems challenged by a growing, complex, and commonly opaque cellular infrastructure. This paper presents the first analysis of anycast performance for mobile users. In particular, our evaluation focuses on two distinct anycast services, both providing part of the DNS Root zone and together covering all major geographical regions. Our results show that mobile clients tend to be routed to suboptimal replicas in terms of geographical distance, more frequently while on a cellular connection than on WiFi, with a significant impact on latency. We find that this is not simply an issue of lacking better alternatives, and that the problem is not specific to particular geographic areas or autonomous systems. We close with a first analysis of the root causes of this phenomenon and describe some of the major classes of anycast anomalies revealed during our study, additionally including a systematic approach to automatically detect such anomalies without any sort of training or annotated measurements. We release our datasets to the networking community.

I. INTRODUCTION

The impressive growth in mobile devices and use has led to an unprecedented amount of cellular traffic. The number of mobile subscriptions has grown rapidly in just a few years, surpassing 7.8 billion in the third quarter of 2017 [1]. Today, users spend most of their time browsing on their mobile phones, more than on any other device. In the United States, for instance, the average smartphone user spends 2x–3x more hours (87 hours per month) on her mobile device than she does on desktop machines (34 hours per month) [2].

Most of this content is delivered to end users by content delivery networks (CDNs). CDNs deploy servers around the world and redirect clients to nearby replicas to improve performance and reliability. The process of replica selection – the mapping of each client to a close replica – is key to CDN performance. Most commonly, it has relied on a DNS-based approach pioneered by Akamai.

Traditional replica selection systems, however, are being challenged by the rapid growth, increased complexity, and common opacity of cellular infrastructure [3]. Anycast offers mobile CDNs an alternative approach.

This work has been carried out while Sarah Wassermann was an MSc. student at the University of Liège.

With IP anycast, services advertise a single IP address from many physical locations (anycast sites) and clients’ requests are directed, based on BGP routing policies, to a “nearby” replica [4]. The approach is being used for redirecting clients in a range of applications, from naming (e.g., root servers and top-level domain resolvers) to content delivery. By letting BGP control request routes, anycast routing obviates the need for fine-grained infrastructure or client information. BGP routing also offers some degree of robustness, adapting to changes in service or network availability, and allows for some policy control. The benefits of anycast to service providers have motivated a number of recent experimental studies on its potential performance impact for end users. Prior analyses have shown that anycast routing can be suboptimal [5], unstable [6], [7], and seemingly chaotic [8], [9], as routing policies have not only technical motivations, but could be dictated by political or commercial reasons. Routing changes can silently shift traffic from one site to another with a consequent loss of shared state and potential performance impact [10]. Yet, despite the growing dominance of mobile Internet access, *no previous studies have evaluated the effectiveness of anycast for mobile users, including its routing behavior and performance.*

In this paper, we take a first look at anycast performance for mobile users. In particular, our evaluation focuses on two distinct anycast services, K- and F-Root DNS, each providing part the DNS Root zone and together covering all major geographical regions. Similarly to previous work [9], [11], [12], we use the geographic distance as our evaluation metric. Geographic distance is a useful metric to estimate RTT [13], and IP-to-location mapping techniques such as GeoPing [14] and Constraint-Based Geolocation (CBG) [15] rely on the correlation between network delay and geographic distance. In the specific case of wireless networks hosting users on the move, we consider geographic distance to be a more suitable proximity metric than delay. Indeed, delay is influenced by multiple factors which can vary greatly, even if the user does not move much. Note that, while we are interested in the impact and behavior of delay towards user-requested content, geographic distance provides a more stable ground truth to use as baseline for the comparison of measurements. *We show that mobile clients tend to be routed to suboptimal replicas in terms of geographical distance, more frequently while on a cellular connection than on WiFi, with a significant impact on perceived service performance.*

In the following section, we provide background on cellular networks and anycast routing and review prior work. We describe our methodology and datasets for measuring anycast routing for mobile end users in Section III, and present our findings in Section IV. We further investigate the causes of aberrant anycast behavior for mobile users in Section V and present a systematic approach to identify such anomalies in an unsupervised manner through clustering techniques in Section VI. Finally, we conclude in Section VII.

II. ANYCAST AND MOBILE USERS

Previous efforts have focused on techniques for characterizing anycast deployment [16], enumerating [17]–[19], and geolocating servers based on latency measurements [17], [18].

Cicalese et al. [16] presented the first Internet-wide anycast census and showed that major players in the Internet ecosystem are relying on anycast, even though only a small part of the IPv4 space has so far been anycasted.

Calder et al. [20] analyzed the performance of anycast in the Bing CDN. The authors showed that, although anycast performed well in general, about 20% of the clients were redirected to a suboptimal replica, with a negative impact on performance. Kuipers et al. [9] reported a related study on the performance of the K-Root DNS service using the RIPE Atlas framework [21]. Like Calder et al., they observed that anycast was suboptimal for multiple probes with, for instance, 46% of them being redirected to a suboptimal DNS server, both in terms of latency and geographical distance. Li and Spring [22] claimed that the main causes for this behavior are Tier-1 ASes, which forward almost all requests systematically to the same replica, irrespectively of the user’s location.

Addressing the flaws of anycast remains a major challenge. Schmidt et al. [5] analyzed the impact of the number of anycast sites on the performance of this paradigm in the DNS infrastructure. The authors concluded that increasing the number of sites does not solve the suboptimal mapping observed nor, counterintuitively, reduces clients’ latencies. In [8], Bellis et al. described their use of RIPE Atlas to detect and address issues related to F-Root, such as the redirection of requests to replicas in different continents (e.g., from a probe in Europe to a server in Atlanta).

Zarifis et al. [23] studied the impact of the Internet topology and routing on the performance perceived by mobile users, and revealed that a significant fraction of Internet paths are inflated, with a non-negligible performance penalty. The authors investigated the root causes for route inflation and found the lack of carrier ingress points to be one of the main reasons.

Our work builds on approaches and insights from many of these past efforts. *Yet, to the best of our knowledge, this is the first work focused on the performance of anycast in mobile networks.*

III. METHODOLOGY AND DATASETS

This section describes our datasets and measurement methodology for analyzing anycast from mobile devices. We focus our analysis on two major root DNS services: F-Root

and K-Root. We selected these services since they are both widely replicated, covering together all geographic areas (with approximately 60 servers for K-Root and nearly 140 for F-Root) and both with publicly available site locations and unicast IP addresses. The latter point is key to let us evaluate the performance of anycast routing relative to its “optimal” (in terms of unicast) site location.

We collected active measurements from geographically distributed clients on both cellular and WiFi networks from September 2016 to April 2017, using the ALICE engine¹. ALICE conducts mobile network measurements by executing a small, self-contained experiment script, run approximately every hour. In each experiment, for each of the two root DNS services, clients launched ping and traceroute measurements towards the root server’s anycast address, as well as towards five unicast addresses of the analyzed DNS service determined to be the geographically closest to the client. Target unicast addresses were selected based on the origin country of each client’s IP address, as reported by `whois`.

Note that, besides our previous justification for choosing geographical distance over latency as selection criteria, we consider the closest unicast replicas in terms of geographical distance and not latency also due to practical concerns. Indeed, we cannot know the closest unicast replicas in terms of latency without flooding the user’s connection with latency measurements towards all available unicast replicas before starting every new experiment, which would be impractical and highly invasive. Moreover, a previous study [9] shows that geographically distant replicas generally result in poorer latency, further supporting our methodology.

We collected data from mobile devices while clients were connected to either WiFi or cellular networks. We thus divide our data into a set containing experiments launched from cellular networks (CELL) and another one including the experiments issued from WiFi (WIFI). As measurements were collected opportunistically, based on connection availability and resource usage, the number of experiments in each of these sets is not the same.

CELL dataset. CELL includes more than 20,000 experiments, issued from 151 different clients. Our cellular users were scattered across nearly 40 different countries, with 70% of them located in the United States, Greece, Brazil, and France. Furthermore, the analyzed clients were hosted in a range of major ASes, including AS 29247 (GR-Cosmote, Greece), AS 26599 (Vivo, Brazil), and AS 22394 (Verizon Wireless, US).

WIFI dataset. WIFI encompasses three times as many experiments as CELL, issued from 251 clients. Clients were located in nearly 50 different countries around the world with 70% of them in Greece and the United States. Some of the most active clients (i.e., those that have launched the most experiments) were hosted in AS 6799 (OTENET-GR, Greece), AS 9121 (TTNet, Turkey), and AS 7922 (COMCAST, US).

¹<http://aqualab.cs.northwestern.edu/projects/261-alice>




Fig. 1: Travel distance from clients to anycast servers.

A total of 125 clients launched experiments from both cellular and WiFi networks.

We are making both the CELL and WIFI datasets publicly available on GitHub².

IV. FINDINGS AND OBSERVATIONS

In this section, we present key observations from our analysis of mobile anycast performance. In particular, our results suggest that mobile clients are most of the time routed towards suboptimal replicas in terms of geographical distance (Section IV-A), more frequently while on a cellular connection than on WiFi, and that the additional distance traveled by requests has a significant impact on performance (Section IV-B). We analyze the scope of this phenomenon (Section IV-C) and the existence of better alternatives (Section IV-D).

A. The Travel Distance Problem

Our study is partially motivated by anecdotal evidence of long geographic distances between mobile clients and their assigned anycast DNS servers. To characterize these distances, given that anycast IP addresses cannot be geolocated [24], we geolocate the penultimate hop on the path from a client to the anycast server and estimate its location using the Akamai EdgeScape service³. We use this and the client location to compute the geographic distance between a client and her assigned anycast server. For each experiment, the client recorded her anonymized geographic location to a 10 km² area. Figure 1 shows the distribution of these distances, in kilometers, for both WiFi and cellular users.

Figure 1(a) presents the travel distance between clients and their K-Root anycast servers ($D_{client \rightarrow anycast}$). We see a significant difference between cell and WiFi: $D_{client \rightarrow anycast}$ is smaller than 4,000 kilometers for approximately 75% of




Fig. 2: Latency from clients to anycast servers.

the experiments carried out on a WiFi network, whereas this holds for merely 50% of the experiments issued from cellular clients. While we can observe the same phenomenon for F-Root on Figure 1(b), the differences are not as striking as for the K-Root service. The difference in number and geographic distribution between F- and K-Root services explains the measured differences in the distances traveled by a client request on a cellular or WiFi connection. Still, clients on cellular networks travel farther to their assigned replicas (with potential latency implications).

B. Impact on Latency

Does the additional distance traveled by users' requests result in worse user performance? We evaluate the impact of the distance between clients and their replicas in terms of latency. Figure 2 shows the distribution of their latencies. Indeed, the figure reveals that latencies are higher for the measurements carried out on cellular networks than for the ones performed on WiFi. Again, the difference is more pronounced for K-Root (Figure 2(a)) than for F-root (Figure 2(b)). A comparison of both figures shows that, while the distribution of latencies is similar for K-Root and F-Root on WiFi, we see significantly worse performance for K-Root compared to F-Root on cellular networks: 90% of the experiments output a RTT lower than 200 milliseconds for F-Root, but only 70% for K-Root DNS.

Other factors beyond distance traveled, such as signal strength and radio access latency, could result in performance degradations for clients on cellular networks. To understand the importance of geographic distance on observed latency, we look at the correlation between distance and minimum RTT for cellular measurements. Figure 3, a scatter plot of this correlation, clearly shows that, while not the only factor, the distance between cellular clients and their assigned anycast servers is indeed an important aspect to take into account when it comes to performance optimization. We observed the same

²<https://github.com/SAWassermann/mobile-anycast>

³<https://www.akamai.com/us/en/products/web-performance>

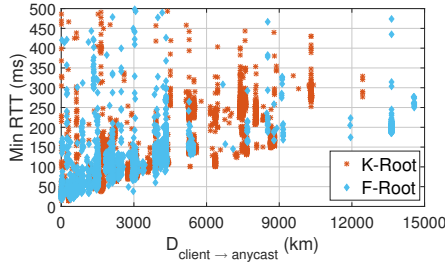


Fig. 3: Correlation between geographic distance and minimum RTT.

behavior for WiFi clients; the graph is not included due to space constraints.

C. Where Does the Travel Distance Problem Occur?

Is it possible that the observed issues are limited to some specific geographical regions or a few particularly mismanaged ASes? To explore this, we now focus our analysis on five ASes that provide the most measurements for clients in cellular networks. This leaves us with 2,310 measurements issued from 16 users for K-Root and 1,860 measurements retrieved from 15 mobile clients for F-Root DNS.

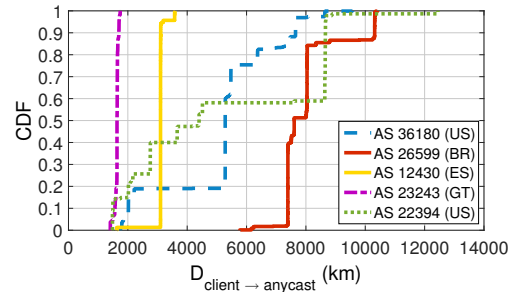
Figure 4 presents the distributions of the distances traveled to anycast servers for clients in these ASes. While we see large variability within and across ASes, particularly in US-based networks potentially covering larger geographic areas, there are no clear patterns to argue for the identified problem to be a region- or AS-specific issue.

As noted before, F-Root DNS replicas seem to be closer to their clients, regardless of the network. Indeed, for 80% of the launched measurements towards F-Root anycast servers, $D_{client \rightarrow anycast}$ is less than approximately 4,300 kilometers, while this distance is larger than 8,500 kilometers for 20% of the measurements issued towards K-Root servers.


D. Do Cellular Clients Have a Closer Option?

Another possible explanation for the long distances would be that these cellular clients are simply located far away from all the available replicas. Further analysis shows that this is not necessarily the case. We compute the distance between cellular clients and both their assigned anycast server ($D_{client \rightarrow anycast}$) and their geographically closest unicast replica ($D_{client \rightarrow unicast}$). We compute the *additional distance traveled by anycast* as the difference between these two distances (i.e., $D_{client \rightarrow anycast} - D_{client \rightarrow unicast}$), which we denote by δ_{client} .

Figure 5 presents the additional distance traveled by anycast requests, δ_{client} , for the measurements launched from the five ASes providing the most measurements for K- and F-Root DNS. We can easily infer from this graph that our clients are most of the time routed to a suboptimal replica and that this issue does not seem to be specific to one AS or region. Nevertheless, comparing Figures 5(a) and 5(b) leads us to conclude that anycast routing is significantly worse for K-Root than for F-Root, even though the mappings for F-Root are far from optimal.




(a) $D_{client \rightarrow anycast}$ observed for K-Root.




(b) $D_{client \rightarrow anycast}$ observed for F-Root.

Fig. 4: $D_{client \rightarrow anycast}$ in top 5 ASes (w.r.t. number of measurements).



(a) δ_{client} observed for K-Root.



(b) δ_{client} observed for F-Root.

Fig. 5: δ_{client} in top 5 ASes (w.r.t. number of measurements).

Particularly clear examples for the K-Root problem are clients residing in the United States, but being routed towards anycast servers located in London and Tehran. This is even more surprising as there are at least seven K-Root servers distributed across the United States, with four hosted on the East Coast and two on the West Coast. Finally, we observe that, in each of the examined ASes, some mobile clients are redirected to the nearest anycast F-Root server, while the ideal case never occurs for the K-Root service.

V. WHY TRAVELING SO FAR? A FIRST LOOK INTO CELLULAR ANOMALIES

In the following paragraphs, we explore some of the causes of anycast routing anomalies for cellular clients. In particular, we focus on the measurements with a value of δ_{client} higher than 1,000 kilometers. This corresponds to more than 80% of the experiments for K-Root and to 70% of the experiments for F-Root. We refer to these measurements as *suboptimal anycast measurements*.

We first compare the lengths of the AS-paths leading from the cellular clients, on the one hand, towards the anycast-assigned replicas and, on the other hand, towards the nearest unicast servers. We then explore whether the announcement policies of different root nodes – whether they are locally or globally announced – affects anycast quality. Finally, we identify three classes of anomalies including: (i) distant client packet gateways, (ii) poor anycast routing within Tier-1 networks, and (iii) improper routing within cellular networks.

AS-path comparison. We analyze the lengths of the AS paths between cellular clients and their assigned anycast server, and between those clients and the geographically closest unicast replica. We find that, when anycast servers are geographically farther away than unicast servers, they still tend to be closer in terms of number of traversed ASes as one would expect with anycast routing. The paths client \rightarrow anycast server are shorter than the paths client \rightarrow unicast replica in nearly 40% of the measurements towards K-Root and in more than 75% towards F-Root. While more than 80% of the paths to F-Root anycast servers have a length of at most six AS hops (more than 60%, considering K-Root), approximately 55% of unicast paths for both DNS services are longer than six hops.

Root announcement policies. We also analyze the node types of both DNS services. While all but one K-Root servers are global, there are only four global servers for F-Root (one in the Netherlands and three in the United States). We find that clients from Singapore, Australia, the United States, and the Dominican Republic are routed towards global nodes in the United States, with a significant latency penalty for most of them, despite the presence of local nodes in their geographical surroundings. Note that, in cases where both a local and global node reside in the same city, we are not able to distinguish between them from traceroutes alone.

Further investigation into anomalous anycast routing for our cellular clients reveals three classes of anycast performance problems. While by no means exhaustive of all problems encountered by cellular clients, these common problems appear to be either unique, or significantly more common, to cellular with respect to fixed-line clients.

A. Distant Packet Gateway

As previously reported in [25], a cellular client’s packet gateway (PGW) largely determines the client’s network position and locality, since all client traffic is routed through that PGW. This implies that client packets first encounter Internet routing once they are beyond the PGW and thus, depending on the relative distance between a client’s PGW and the assigned

Client Location: Boston, MA				PGW Location: Los Angeles, CA			
Anycast Destination				Nearest Unicast Destination			
Hop	IP	ASN		Hop	IP	ASN	
1-5	Private addresses	-	-	1-5	Private addresses	-	-
6	208.185.160.182	6461	Los Angeles, CA	6	208.185.160.182	6461	Los Angeles, CA
7	208.185.160.181	6461	Los Angeles, CA	7	208.185.160.181	6461	Los Angeles, CA
8	64.125.28.41	6461	San Jose, CA	8	64.125.28.45	6461	San Jose, CA
9	64.125.30.230	6461	San Jose, CA	9	64.125.31.15	6461	San Jose, CA
10	64.125.31.218	6461	Denver, CO	10	63.146.26.253	209	-
11	64.125.29.18	6461	Chicago, IL	12	63.234.254.142	209	Salt Lake City, UT
12	64.125.29.208	6461	New York, NY	15	209.33.214.250	11071	St. George, UT
13	64.125.29.127	6461	London, UK	Round-Trip Time: 160 ms			
14	64.125.30.237	6461	London, UK	<div style="border: 1px solid black; padding: 5px;"> <ul style="list-style-type: none"> - Distant client PGW - Tier-1 routing issues </div>			
15	64.125.31.193	6461	London, UK				
16	64.125.27.50	6461	London, UK				
17	213.161.79.50	6461	London, UK				
18	193.0.14.129	25152					

Round-Trip Time: 258 ms

Fig. 6: Distant client PGW location & Tier-1 routing issues.




Fig. 7: Distance client \rightarrow assigned PGW in top 5 ASes (w.r.t. number of different clients).

server, the *wrong* assignment of a client to a distant PGW can be the determinant factor of anycast routing performance. Knowing the location of a client’s PGW can therefore greatly aid in diagnosing anomalous anycast routing.

We investigate the consequences of this phenomenon in terms of latency. We find that for about 35% of the considered anycast measurements, the clients’ request spends more than 50% of its time on the way towards the PGW. In particular, for 45% of the traceroutes, the clients’ probe needs more than 70 milliseconds to reach the assigned packet gateway. These latencies are far from being negligible and suggest that the routing within the users’ cellular network could be significantly enhanced.

A striking example is depicted in Figure 6. These paths correspond to a client located in Boston (MA), but whose PGW is situated in Los Angeles (CA). While the closest unicast replica is hosted in St. George (UT), the client’s anycast request is routed to London, with a large latency penalty of almost 100 milliseconds. In this scenario, the journey of the packet is not only costly in terms of kilometers and latency due to the suboptimal geolocation mapping between the client and an anycast server, but also because the packet needs to travel more than 4,000 kilometers to reach the PGW.

A more detailed analysis of our suboptimal anycast measurements shows that, in 40% of the cases, the concerned clients are located more than 1,000 kilometers away from their assigned PGW, and, for approximately 10%, this distance is larger than 3,000 kilometers. Unfortunately, even when the

PGW is far away from a cellular client, it does not mean that the assigned anycast replica is near the packet gateway. Indeed, for the users presenting a distance to their PGW larger than 1,000 kilometers, their packet still has to travel more than 2,000 kilometers to reach its destination in more than 65% of these measurements.

We take a closer look at the top five ASes in terms of number of different clients suffering from suboptimal anycast measurements. Figure 7 depicts the distribution of the distances between the clients and their PGW for each investigated anycast case. Our results show that distant PGWs do not seem to be a systematic issue: while this distance varies widely in ASes 21928 (T-Mobile, US), 22394 (Verizon Wireless, US), and 26599 (Vivo, Brazil), it is relatively stable and small for ASes 12361 (Vodafone-Panafon, Greece) and 23243 (Comcel, Guatemala). Specifically, for all the considered measurements launched in ASes 12361 and 23243, the cellular clients present a distance to their PGW smaller than 700 kilometers, which is the case for less than 1% of the measurements seen in AS 26599. For AS 21928, we note that 60% of the measurements have been carried out while clients were very close to their assigned PGW (between 30 and 60 kilometers), while the remaining 40% have been launched when they were farther away from it (up to 3,000 kilometers).

B. Tier-1 Routing

We observe that many of the instances of poor anycast performance occur when anycast paths traverse Tier-1 transit networks. While Tier-1 routing problems are not solely specific to cellular networks [22], we found them to be much more prevalent on suboptimal cellular paths than on WiFi ones. We found that 64% of the paths leading towards the assigned anycast replica traverse at least one Tier-1 AS, while this is the case for 73% of the paths towards the geographically closest unicast server.

With large Tier-1 networks, we see clients often routed to the same anycast replica, regardless of where clients enter these networks. Many problematic cases we investigated appear to be caused by packets remaining in the transit network until routed to a distant destination. We find this behavior in Tier-1 networks with varying levels of consistency. For example, clients entering AT&T (AS 7018) split anycast destinations between sites in Reno (NV) and London. We also see clients routed through AboveNet (AS 6461) being consistently routed to K-Root sites in London. An observed example corresponds to a cellular client residing in Los Angeles and whose PGW is in the same city. Even though the geographically closest unicast server is in Reno, her packet is routed to a server situated in London. Analyzing the corresponding traceroutes reveals that more than 50% of the hops in the path leading to the anycast server lie in AboveNet, a well-known Tier-1 AS. As opposed to the anycast path, the path connecting the cellular user to her nearest unicast server exits the AS 6461 fairly quickly: less than 25% of the traceroute hops are in this Tier-1 AS. As a consequence, the latency towards the anycast server is much higher than the one

Client Location: Sao Paulo, BR				PGW Location: Sao Paulo, BR			
Anycast Destination				Nearest Unicast Destination			
Hop	IP	ASN	Location	Hop	IP	ASN	Location
1-2	Private addresses	-	-	1-2	Private addresses	-	-
3	177.79.213.38	26599	Sao Paulo, BR	3	177.79.213.38	26599	Sao Paulo, BR
4	187.100.51.69	27699	Sao Paulo, BR	7	187.100.86.164	27699	Sao Paulo, BR
5	187.100.193.254	27699	Sao Paulo, BR	8	187.100.49.25	27699	Sao Paulo, BR
6	187.100.192.246	27699	Sao Paulo, BR	9	187.100.49.17	27699	Sao Paulo, BR
7	213.140.39.70	12956	-	10	5.53.0.189	12956	Natal, BR
8	176.52.250.246	12956	-	12	94.142.97.9	12956	Salvador, BR
9	195.22.199.89	6762	Miami, FL	15	5.53.3.43	12956	-
10	195.22.209.107	6762	London, UK	10	94.142.119.168	12956	Santos, BR
11	195.66.224.183	8881	London, UK	11	94.142.127.69	12956	Buenos Aires, AR
12	193.0.14.129	25152	London, UK	12	213.140.35.83	12956	Buenos Aires, AR
Round-Trip Time: 271 ms				13	176.52.248.199	12956	Buenos Aires, AR
- Improper routing within operator's network.				14	5.53.1.149	12956	Buenos Aires, AR
				19	200.40.98.29	6057	Montevideo, UY
				20	200.40.98.27	6057	Montevideo, UY
				21	200.7.84.36	28000	Montevideo, UY
				22	179.0.156.11	28000	Montevideo, UY
				Round-Trip Time: 119 ms			

Fig. 8: Improper cellular network routing.

towards the closest unicast replica. In this scenario, the seen traceroutes show a similar behavior as the ones in Figure 6.

Further analysis highlights that AboveNet is by far the most often encountered Tier-1 AS. Indeed, this Tier-1 is seen in 55% of the anycast traceroute paths traversing at least one of these ASes, while the second-most frequently observed one (AS 12956, Telxius) in only 24% of these traceroutes. When looking at the cases for which we have a δ_{client} between 5,000 and 10,000 kilometers, AboveNet is still the most popular Tier-1 (appearing in 40% of the traceroutes with Tier-1 hops), but ASes such as 7018 (AT&T) and 2914 (NTT Communications) are traversed more frequently with respect to the cases with smaller values of δ_{client} . In addition to that, we were curious about the impact of Tier-1s on latency. Our results reveal that anycast is suboptimal in terms of latency with respect to the closest unicast replica in more than 70% of the cases when the packet traverses at least one Tier-1 AS on its way to the assigned anycast server.

C. Improper Cellular Network Routing

We found cases where paths leaving towards the anycast and nearest unicast servers diverge immediately at the IP level after exiting a client's PGW, even though they both end up in the same AS for the next IP hop. We discovered that 17% of our analyzed cases suffer from improper routing within the cellular network. However, this issue does not seem to be tied to specific ISPs or regions, as we detected it in multiple ASes located across Europe and the United States.

Figure 8 illustrates this phenomenon. In the figure, the IP paths diverge right after they exit the client's PGW. These diverging paths remain in the operator's same AS (AS 27699), and even have the same next hop AS. In this scenario, the client's packet reaches the unicast replica significantly faster than the assigned anycast server; we observe an impressive latency difference of 151 milliseconds. While we cannot know what caused this exact instance, we saw for several operators multiple transit providers connected to, or very near, cellular network PGWs. As we noted in the previous class, the choice

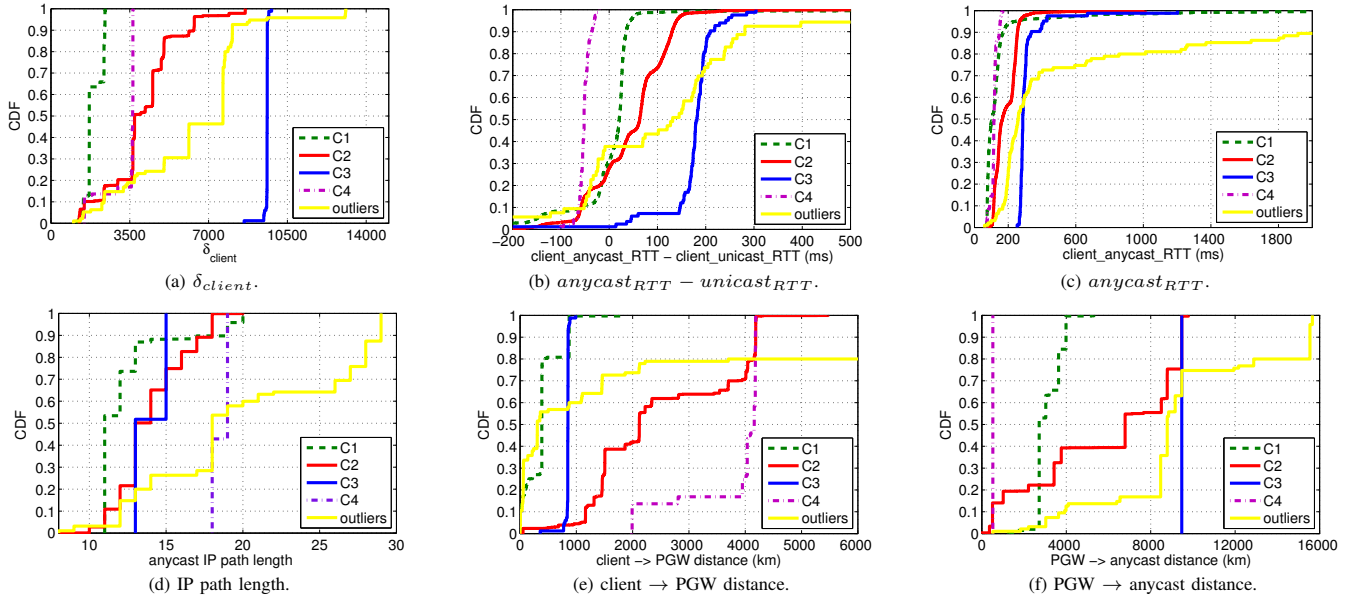


Fig. 9: Characterization of the suboptimal anycast measurements using DBSCAN clustering. Highly suboptimal anycast measurements (as compared to unicast) are located in clusters C_2 , C_3 and within the outliers.

of transit can play a large role in the behavior of anycast routes. This is especially true if the provider is a Tier-1 network, as many commonly are for large cellular networks.

VI. TOWARDS AUTOMATIC CHARACTERIZATION OF ANYCAST CELLULAR ANOMALIES

We devote the last section to further investigate the aforementioned suboptimal anycast measurements (i.e., with δ_{client} higher than 1,000 kilometers) and corresponding anomalies. We take an exploratory approach, and use machine-learning techniques to provide first steps in the automatic identification and characterization of anycast anomalies in cellular networks. Given the general lack of ground truth regarding the nature and root causes of the anomalies, we perform an unsupervised analysis, relying on clustering techniques.

We take two different approaches: firstly, we use the well-known DBSCAN clustering algorithm [26] to identify homogeneous groupings and outliers within the suboptimal measurements. Then, we apply a more advanced, DBSCAN-based algorithm we have previously conceived in [27] to automatically spot out the most relevant anomalies within the measurements, without requiring any sort of training or labeled data.

A. DBSCAN-based Analysis

By applying DBSCAN to the suboptimal anycast measurements, we are able to identify four consistent clusters $C_{i=1..4}$ (silhouette scores above 0.7) and a group of *outliers*, i.e., measurements not belonging to any of the four clusters. Following the results obtained in [27], we use auto-calibration for DBSCAN parameters ρ (distance search coefficient) and λ (minimum connected region size) [26], taking $\rho = \alpha \times n$ ($\alpha = 0.01$) and $\lambda = \beta \times \bar{d}$ ($\beta = 0.2$), where n is the

number of measurements and \bar{d} the average distance between all the different pairs of measurements. Each measurement is described by a vector of 12 features: δ_{client} , client to anycast RTT (ARTT) and distance (CAD), client to closest unicast RTT (URTT), client to PGW RTT (CPRTT) and distance (CPD), PGW to anycast RTT and distance, anycast IP- and AS-path length (IPPL and ASPL), number and fraction of Tier-1 hops in the anycast path.

Figure 9 presents an overall characterization of the results. According to Figures 9(b) and 9(c), anomalous or highly suboptimal anycast measurements as compared to unicast are located in clusters C_2 , C_3 and within the outliers, representing altogether more than 80% of all the suboptimal measurements. Indeed, despite the higher geographical distance of the anycast replicas, C_1 and C_4 measurements tend to have a similar performance between anycast and unicast, and both clusters consist of measurements with comparable latencies, trading closer PGWs with farther anycast replicas between clusters.

Previous anomaly *types* (e.g., distant PGWs and routing/path-inflation issues) are clearly observed in C_2 , C_3 , and the outliers. Outliers systematically correspond to poor anycast performance as compared to unicast, with large latencies, far away replicas with long paths (potentially linked to routing and path inflation issues), PGW selection at distant locations, and very high differences in terms of geographical distance to selected anycast versus unicast replicas. C_3 measurements are characterized by large δ_{client} values and much higher latencies as compared to unicast replicas. Finally, C_2 measurements are more spread and correspond to mostly worse-than-unicast performance scenarios, also showing far located PGWs and anycast replicas, but resulting in RTTs below 250 milliseconds for more than 85% of the measurements.

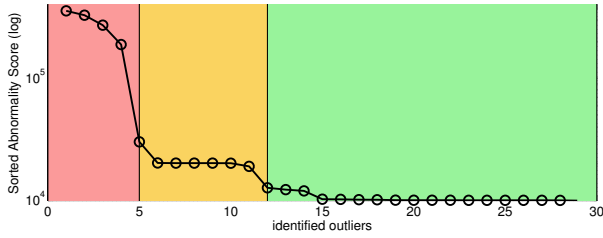


Fig. 10: Anomalies detected by sub-space clustering.

B. Anomaly Detection with Sub-Space Clustering

To conclude, we explore the possibility of automatically spotting out the aforementioned anomalies to better support the root cause analysis process. In particular, we apply a fully unsupervised approach for anomaly detection (UAD) previously conceived in [27] to the set of suboptimal anycast measurements. In a nutshell, the UAD algorithm uses DBSCAN to cluster the measurements by relying on different sub-spaces of the complete feature space, and accumulates a weighted distance between outliers and clusters on these sub-spaces to compute an abnormality score; the higher this score, the more different (i.e., anomalous) is the corresponding measurement from the rest. DBSCAN parameters are set through the same auto-calibration approach described before. We refer the reader to [27] for further details on the algorithm.

Figure 10 depicts the resulting abnormality scores obtained by UAD, sorted in descending order. The number of identified outliers corresponds to approximately 1% of the suboptimal anycast measurements. There are different regions clearly visible in the ranking, with knees and breaks showing different levels of abnormality. To shed some light on the detected anomalies, Table I reports the corresponding feature values for the top-10 ranked anomalies.

The first two anomalies are characterized by very large RTTs to both anycast and unicast replicas; anomalies with IDs 4, 6, 7, 8 correspond to similar measurements, characterized by a very high δ_{client} . Interestingly, anomalies 5 and 9 correspond to some of the anomalies manually studied in Section V (cf. Figures 6 and 8). Remaining anomalies 3 and 10 correspond to far located anycast replicas and reasonably close PGWs, suggesting anomalies linked to routing. In all cases, the performance of anycast in terms of latency is far worse than that of unicast, with a δ_{client} higher than 2,000 kilometers.

VII. CONCLUSIONS

In this paper, we presented findings from the first look at anycast performance for mobile users. Using data collected from a crowdsourced platform, we showed that mobile clients are very frequently mapped to a geographically suboptimal anycast replica, with a significant impact on performance in terms of latency. We found that the long distances between clients and their assigned anycast server are not due to a lack of better, closer unicast replicas, and that this phenomenon is not bound to specific regions or particular ASes. Exploring root causes highlighted three classes of anycast anomalies, namely distant client packet gateways, poor anycast routing

TABLE I: Top 10 detected anomalies.

ID	δ_{client} (km)	ARTT (ms)	URTT (ms)	IPPL	CAD (km)	CPD (km)
1	2,387	6,828	2,431	12	3,871	138
2	5,022	1,638	1,428	12	8,040	1,457
3	5,027	229	97	12	8,042	1,454
4	13,097	273	95	14	14,582	1,084
5	3,714	271	119	12	7,392	2,118
6	13,073	250	77	14	14,567	1,099
7	13,084	256	76	14	14,573	1,093
8	13,063	261	90	14	14,559	1,108
9	4,519	258	160	18	5,281	4,185
10	9,597	322	138	13	10,316	848

within Tier-1 networks, and improper routing out of cellular networks. We finally presented a clustering-based analysis of the suboptimal measurements, including a fully unsupervised, automatic approach to identify the most critical cases. In addition to that, we release our datasets to the networking community. In ongoing work, we are exploring some of these issues, including patterns of cellular specific anycast problems, their role on CDN request routing, and the impact of these pitfalls on clients' quality of experience.

REFERENCES

- [1] Ericsson, "Ericsson Mobility Report – On the Pulse of the Networked Society," Tech. Rep., November 2017.
- [2] Ofcom, "The Communications Market Report," Tech. Rep., December 2016.
- [3] J. P. Rula and F. E. Bustamante, "Behind the Curtain: Cellular DNS and Content Replica Selection," in *IMC*, 2014.
- [4] J. Abley and K. Lindqvist, "Operation of Anycast Services," Internet Requests for Comments, RFC Editor, RFC 4786, December 2006. [Online].
- [5] R. Schmidt et al., "Anycast Latency: How Many Sites Are Enough?" in *PAM*, 2017.
- [6] P. Bret et al., "TCP over IP Anycast – Pipe Dream or Reality?" 2015. [Online].
- [7] J. Hiebert et al., "Determining the Cause and Frequency of Routing Instability with Anycast," in *AMS*, 2006.
- [8] R. Bellis, "Researching F-root Anycast Placement Using RIPE Atlas," 2015. [Online].
- [9] J. Kuipers, "Analysing the K-root Anycast Infrastructure," in *25th Twente Student Conference on IT*, 2016.
- [10] L. Wei and J. Heidemann, "Does Anycast Hang up on You?" in *TMA*, 2017.
- [11] Z. Li et al., "Longitudinal Analysis of Root Server Anycast Inefficiencies," 2017.
- [12] Z. Liu et al., "Two Days in the Life of the DNS Anycast Root Servers," in *PAM*, 2007.
- [13] B. Huffaker et al., "Distance Metrics in the Internet," in *IEEE ITS*, 2002.
- [14] V. N. Padmanabhan and L. Subramanian, "An Investigation of Geographic Mapping Techniques for Internet Hosts," in *SIGCOMM*, 2001.
- [15] B. Gueye et al., "Constraint-Based Geolocation of Internet Hosts," in *IMC*, 2004.
- [16] D. Cicalese et al., "Characterizing IPv4 Anycast Adoption and Deployment," in *CoNEXT*, 2015.
- [17] D. Cicalese et al., "A Fistful of Pings: Accurate and Lightweight Anycast Enumeration and Geolocation," in *INFOCOM*, 2015.
- [18] —, "Latency-based Anycast Geolocalization: Algorithms, Software and Datasets," Tech. Rep., 2015.
- [19] D. Madory et al., "Who are the Anycasters?" in *NANOG 59*, 2013.
- [20] M. Calder et al., "Analyzing the Performance of an Anycast CDN," in *IMC*, 2015.
- [21] RIPE NCC Staff, "RIPE Atlas: a Global Internet Measurement Network," *Internet Protocol Journal*, September 2015.
- [22] Z. Li and N. Spring, "Tier-1's break Anycast DNS," in *AIMS*, 2017.
- [23] K. Zarifis et al., "Diagnosing Path Inflation of Mobile Client Traffic," in *PAM*, 2014.
- [24] D. McPherson et al., "Architectural Considerations of IP Anycast," Internet Requests for Comments, RFC Editor, RFC 7094, January 2014. [Online].
- [25] Q. Xu et al., "Cellular Data Network Infrastructure Characterization and Implication on Mobile Content Placement," in *SIGMETRICS*, 2011.
- [26] M. Ester et al., "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise," in *KDD*, 1996.
- [27] P. Casas et al., "Unsupervised Network Intrusion Detection Systems: Detecting the Unknown without Knowledge," *Comput. Commun.*, vol. 35, no. 7, pp. 772–783, April 2012.