

The Cloud that Runs the Mobile Internet: A Measurement Study of Mobile Cloud Services

Foivos Michelinakis^{1,2,4}, Hossein Doroud², Abbas Razaghpanah³,
Andra Lutu⁴, Narseo Vallina-Rodriguez^{1,5}, Phillipa Gill⁶, Joerg Widmer¹
¹IMDEA Networks Institute, ²Universidad Carlos III de Madrid, ³Stony Brook University,
⁴Simula Research Laboratory, ⁵ICSI, ⁶University of Massachusetts

Abstract—Mobile applications outsource their cloud infrastructure deployment and content delivery to cloud computing services and content delivery networks. Studying how these services, which we collectively denote Cloud Service Providers (CSPs), perform over Mobile Network Operators (MNOs) is crucial to understanding some of the performance limitations of today’s mobile apps. To that end, we perform the first empirical study of the complex dynamics between applications, MNOs and CSPs. First, we use real mobile app traffic traces that we gathered through a global crowdsourcing campaign to identify the most prevalent CSPs supporting today’s mobile Internet. Then, we investigate how well these services interconnect with major European MNOs at a topological level, and measure their performance over European MNO networks through a month-long measurement campaign on the MONROE mobile broadband testbed. We discover that the top 6 most prevalent CSPs are used by 85% of apps, and observe significant differences in their performance across different MNOs due to the nature of their services, peering relationships with MNOs, and deployment strategies. We also find that CSP performance in MNOs is affected by inflated path length, roaming, and presence of middleboxes, but not influenced by the choice of DNS resolver.

I. INTRODUCTION

Mobile app developers have a wealth of tools and techniques at their disposal that help them decrease the amount of time and effort required to develop, deploy, and maintain their apps. One particularly powerful and very common technique is to use a variety of third-party online services such as on-demand cloud computing platforms (*e.g.*, Amazon Web Services) and content delivery networks (*e.g.*, Akamai) in their apps. This technique makes it easier and more efficient to deploy mobile apps at a global scale by shifting the burden of managing and maintaining server infrastructure from app developers to these *Cloud Service Providers* (CSPs). As a result, characterizing the performance of these services in the wild is critical in understanding, and ultimately reducing, the technological gap between the limited capabilities of the mobile Internet infrastructure and the performance requirements of current and future mobile apps [1].

While there have been quite a few studies on the implications of newer protocols such as QUIC on the performance of mobile apps [2], there is a notable lack of a systematic study on the performance of third-party cloud computing and content delivery services used by mobile apps. In fact, the relationships between CSPs, app developers, and Mobile Network Operators (MNOs) are tangled, and decisions made by each entity can

have a significant and far-reaching impact on the ecosystem as a whole. For example, MNOs may peer with popular CSPs to reduce their costs and improve performance for their users, thereby placing these CSPs and the apps that use them at an advantage over others.

In this paper, we empirically analyze the web of relationships between mobile apps, CSPs, and MNOs. In particular, we aim to answer the following questions:

- Which are the most dominant CSPs enabling the mobile Internet?
- How well are these CSPs interconnected with MNOs at a topological level?
- What is the performance of these services (*i.e.*, as perceived by end-users) when accessed from commercial MNOs?

We conduct the first comprehensive study of its kind, combining different measurement techniques and vantage points to fully capture the synergies between the entities forming this complex ecosystem. As a starting point, we use traffic logs that we collected with Lumen Privacy Monitor [3], a mobile privacy and transparency tool. Lumen’s rich traffic logs allow us to accurately identify the most prevalent CSPs providing on-line infrastructure to 8,281 mobile apps in the wild. Then, we run a purpose-specific month-long measurement campaign using the MONROE platform for mobile broadband measurements [4] to capture the interactions between ten commercial MNOs from four European countries and the most popular CSPs, as well as their transport- and application-layer performance. Specifically, we focus on analyzing the effect of replica selection, the role of the DNS subsystem, and the impact of in-path TCP splitting proxies, as well as routing- and peering-level effects on transport-layer performance. Our study also includes mobile subscriptions roaming internationally.

Our analysis reveals that six CSPs— Amazon Web Services (AWS), Google, Facebook, Akamai, Amazon CloudFront, and Highwinds — provide infrastructure and online support to 85% of the apps that we measure with Lumen. We capture interesting multi-CSP strategies that 687 second level domains (15% of domains) use to increase their geographical coverage and reliability (Section IV). We also track the integration and collaboration strategies between the top CSPs identified through Lumen and the MNOs available in the MONROE platform (Section V). In particular, Akamai’s strategic al-

liances with multiple MNOs stand out. The varying degrees of collaboration between MNOs and CSPs translates into notable performance differences, which we actively measure and analyze in the same section. Namely, the tight integration between MNOs and CSPs results in lower latency and connection times: Google’s relationship with various MNOs has resulted in 15% lower connection times on average compared to other similarly performing CSPs. We detect various levels of EDNS adoption among the studied operators, which, however, does not seem to have any significant impact on performance. Finally, international roaming may add significant delays, especially in the case of well provisioned CSPs, defeating their attempts to put content close to the user.

II. BACKGROUND AND RELATED WORK

The line distinguishing a CDN from a cloud computing provider can be blurred at times: third-party service providers may simultaneously offer cloud computing and CDN services using the same domain names and IP blocks. Due to this classification challenge, in this study we analyze them together using the term *Cloud Service Providers (CSP)*.

CSP deployment strategies: CDNs and cloud services may follow different strategies to deploy their servers at a global scale. Cloud services like AWS leverage a reduced number of datacenters located in strategic locations. Instead, most CDNs deploy thousands of caches and proxies as close to the end-user as possible to minimize the end-to-end latency. For example, Akamai’s infrastructure controls more than 233,000 servers in over 130 countries and 1,600 networks [5], while AWS operates just 44 large-scale datacenters located in 16 geographic regions [6]. Large CDNs may also host their services in IP blocks owned by MNOs and fixed-line ISPs who may also commercialize their own CDNs and cloud solutions (e.g., Level 3 and TeliaSonera). This state of affairs makes it difficult to attribute a given domain name or IP address to a particular CSP, as we will discuss in Section V.

CSP-MNO integration: The quality of experience (QoE) the end user perceives when connecting to CSPs may be determined by the underlying connectivity agreements between MNOs and CSPs [7, 8]. A recent crowd-sourcing measurement campaign [9] suggests that certain mobile domains perform poorly on many MNOs. App developers, MNOs and CSPs are increasingly engaging in new peering agreements [7, 10, 11] and initiatives to avoid such inefficiencies. Two example are Netflix’ Open Connect Initiative [12] and Akamai’s Accelerated Network Partner (AANP) program [13].

Multi-CDN strategies: CSP usage by mobile apps can be complex at times. Some apps combine several cloud services to perform specific operations – e.g., Netflix uses AWS for encoding their videos while using multiple CDN providers to enhance their resilience, coverage, and efficiency. This strategy is known as *Multi-CDN* [14]. Adhikari *et al.* [14] studied Netflix and Hulu’s multi-CDN strategies and their CDN selection algorithms. Their work concludes that considering network conditions in the CDN selection algorithm

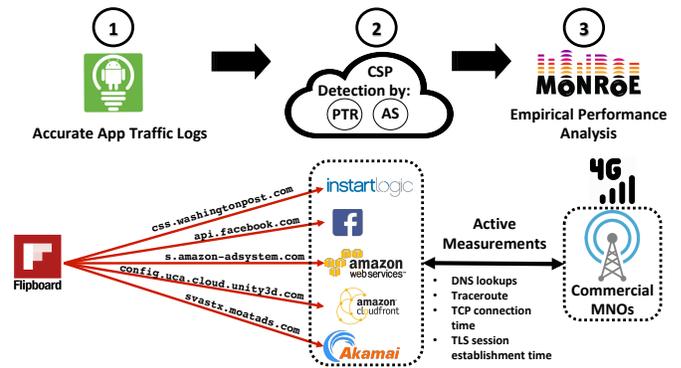


Fig. 1: Schema of our study methodology using a simplified case of the Flipboard app as a toy example. We followed three complementary steps in our study: 1) we analyze app traffic logs to identify the network domains reached by thousands of mobile apps (each red arrow represents a traffic flow to a domain); 2) we detect those domains hosted in CSPs; and 3) we actively measure the performance of CSP-hosted domains on the MONROE measurements platform.

or utilizing multiple CDNs simultaneously can improve the average available bandwidth by 12% and 50%, respectively.

Measuring CSP performance: CSPs’ performance in mobile networks may be affected by multiple factors such as radio link variability, the presence of in-path middleboxes [15], traffic shaping policies [16, 17], the behavior of the DNS resolver [18, 19], the peering relationships between cloud providers and MNOs [7, 8], and inflated network paths [20]. While previous research studies assumed that users are always paired with geographically close content replicas thanks to DNS-based geolocation techniques [21] and IP anycast, [22] showed that this assumption might not always be true due to inaccurate geolocation of mobile users resulting in sub-optimal server assignment. Furthermore, Rula *et al.* conducted a crowd-sourced measurement campaign to study DNS behavior in mobile networks and revealed that client-to-resolver inconsistencies make DNS-based solutions unsuitable for determining the location of clients in MNOs [18]. Recently, the Internet has witnessed the growth of new anycast-enabled CDNs (e.g., CloudFlare) to overcome these limitations. However, this strategy depends on the stability of the paths toward the nearest server.

III. METHODOLOGY AND DATASETS

We follow a multi-step research method to study mobile CSPs as we depict in Figure 1. In summary: *i)* we obtain accurate mobile traffic traces provided by thousands of users of the Lumen app to identify the set of network domains reached by thousands of mobile apps; *ii)* we rank each domain by their popularity, and identify those hosted on CSPs using a purpose-built CSP classifier; *iii)* we measure the performance of CSP-hosted domains on the MONROE platform and infer peering relationships between CSPs and MNOs. As our study is built upon real-world mobile traffic, we can comprehensively

study the most prevalent CSPs by running realistic active measurements on a set of representative CSP-hosted domains.

Toy example: the Flipboard app. Figure 1 provides a high-level description of our method using the Flipboard app as a toy example, depicting how our method helps us understand the relationships between mobile apps, domains (specifically, Fully Qualified Domain Names (FQDNs)), and CSPs hosting these domains. First, we use Lumen Privacy Monitor (Lumen) [3] to capture the different domains Flipboard connects to during normal operations (**Step 1** in Figure 1). In **Step 2**, we combine a number of techniques — including reverse DNS lookups, domain classification, and IP block analysis, among others — to identify which domains rely on cloud providers, and to determine the actual CSP providing support. This step allows us to know that the Flipboard app communicates with 5 different FQDNs, and that each contacted domain is hosted in a different CSP.

Mobile apps, including Flipboard, typically connect to third-party services for purposes of advertising and tracking [23], or to embed other services like online payment and weather reports [24]. These third-party services may also rely on CSPs for outsourcing their cloud infrastructure. For example, the Flipboard app leverages Facebook’s Graph API, which is hosted in Facebook’s own cloud infrastructure, for user login and possibly for advertising purposes. Armed with this FQDN-CSP mapping, we select a number of representative domains to perform active performance measurements on (*e.g.*, TCP connection time) using the MONROE platform (**Step 3**). We further describe each step and their relevant datasets in the following subsections.

A. Step 1. Collecting Accurate Traffic Logs

Lumen is a free Android tool for transparency and user control that captures, reassembles, and analyzes mobile app’s traffic flows on the device itself. Lumen operates as a middleware between apps and the network interface, and intercepts all network traffic locally and in user space using the Android VPN API. This allows Lumen to correlate traffic flows with disparate and rich contextual information available on the device. For example, Lumen matches DNS queries to outgoing flows and the app process owning the socket in order to obtain an accurate profile of a given app’s traffic patterns.

Lumen is publicly available to download from the Google Play Store [25], allowing us to crowd-source mobile traffic measurements at scale from all over the world. This feature makes Lumen a unique mobile vantage point to understand how mobile apps communicate with online services using real user input and network-stimuli and, therefore, the real interactions between mobile apps and the CSPs supporting them. Lumen’s global user base allowed the collection of a representative dataset accounting for over 5M anonymous network flows corresponding to over 8,000 different mobile apps reaching more than 18,000 FQDNs. In order to preserve user privacy, Lumen performs its flow processing and analysis

on the device, only sending anonymized data — no payload or user identifier is collected — to our servers¹.

B. Step 2. Mapping FQDNs to CSPs

Identifying the synergies between mobile apps, FQDNs, and CSPs is a challenging problem. To tackle this problem, our approach focuses on the 18,000 FQDNs available in the Lumen dataset. We retrieve and analyze the PTR records (if available) associated with each IPv4 and IPv6 address by running reverse DNS lookups to identify whether a FQDN is hosted on a given CSP. This allows us to map FQDNs to CSPs using CDNFinder’s PTR to CSP mapping [26]. However, CDNFinder’s mapping does not include marginal CDNs like CDNetworks as well as pure cloud service providers like AWS or Claranet. Moreover, we could only retrieve PTR records for 62% of the total IP addresses present in the Lumen dataset. In order to overcome these limitations and increase CDNFinder’s coverage, we take the following steps:

- 1) We run a semi-supervised PTR classification by searching for strings that may suggest CSP-related operations like `cdn` and `host` on the PTR records.
- 2) To identify CSP-related PTR records that are absent in CDNFinder’s mapping, we implement a semi-supervised PTR classifier that leverages public domain classifiers, specifically McAfee’s [27] and OpenDNS’ [28] domain classifiers. To that end, we first extract the most common categories assigned by the aforementioned domain classifier services to well-known CSP-related PTR records (namely “Internet Services” and “Content Server”). Then, we check if any of the PTR records that we obtain through our reverse DNS lookups fall in any of these categories. Unfortunately, this approach introduces false positives as third-party in-app services like ad networks may be classified as “Internet Services” too. The sheer size of the PTR records impedes our ability to sanitize all of them manually so we limit our manual inspection to PTR records associated with 248 popular FQDNs.
- 3) For each IP address associated with Lumen’s FQDN entries, we run WHOIS queries and retrieve the information on registrant organization and listed email addresses. We browse the website of the email address domain to check if the organization offers any CSP-related products.
- 4) To identify CSPs in IP addresses that do not have PTR records associated with them, we leverage the organization name string as present on AS-level records. This analysis allows us to infer the presence of CSPs for 37% of FQDNs without PTR records and to increase the identification coverage of FQDNs associated with CSPs by 14%.

Combining these four techniques allows us to create a mapping of 194 second-level PTR records associated with 125 third-party CSPs, of which only 43 were initially present in CDNFinder. We made our PTR-CSP mapping open to the public [29].

¹Our institutional IRB classifies this project as “non-human research subject” as we analyze the behavior of software, and not people.

TABLE I: List of MNOs per country. MNOs listed in bold are roaming internationally (home country code in brackets).

Country	MNOs
Norway	Telenor, Telia, Telia (SE)
Italy	Vodafone, Wind, TIM
Spain	Yoigo, Orange, Vodafone (IT)
Sweden	Telia, Telenor, 3

Summary: Our extended FQDN-CSP mapping — both for IP addresses as well as PTR-records — and their associated AS numbers allow us to measure the prevalence of each CSP in the mobile ecosystem, reveal instances of multi-CDN strategies, analyze CSP peering relationships with MNOs, and compile a set of representative domains to empirically measure on the MONROE platform (**step 3**). Due to time and technical restraints, we limit our active measurements campaign to a subset of 1,334 FQDNs hosted by the 6 most prevalent CSPs across apps: Amazon Web Services (AWS), Google, Facebook, Akamai, Amazon CloudFront, and Highwinds.

C. Step 3. Empirical Performance Analysis

In order to assess the performance of CSP-hosted domains, we run a dedicated measurement campaign on the MONROE platform [4], the first open access measurement platform for independent and large-scale experimentation on commercial MNOs. MONROE consists of programmable nodes spread across several European countries, each one multi-homed to three MNOs. For this study, we use nodes in 4 countries — including SIMs performing international roaming — as listed in Table I. We benefit from MONROE’s openness and capabilities to run long-lived active measurements in realistic but controlled scenarios in commercial MNOs. Namely:

- **DNS test:** We run DNS lookups for each target mobile domain. We compare the response provided by the default MNO DNS resolver with Google’s and OpenDNS public resolvers. This allows us to compare the quality of the responses and identify possible DNS-level inefficiencies during the replica selection. All of our DNS queries include the EDNS flag [21] as it may be used by CSPs like EdgeCast and Amazon services to locate the end user and improve the quality of replica selection [30]. If we do not receive a reply, as it occurs for Vodafone (IT) possibly due to the presence of an in-path DNS proxy filtering those requests, we repeat the request without the flag. We detect support for this functionality based on whether the DNS response includes an ECS option with the client’s subnet [21].
- **Traffic performance test:** We measure the TCP connection and TLS session establishment time, if applicable, towards the resolved IP addresses. We open both TCP and TLS connections over TCP ports 80 and 443 by making an HTTP GET request for the `favicon.ico` object. The presence of the object in the server is not relevant as the handshakes are triggered regardless of its existence.

TABLE II: Top 5 FQDN by app penetration.

App (%)	SLD	IP(%)	CSP(s)
29	googlesyndication.com	0.8	Google
28	doubleclick.net	1.4	Google
27	facebook.com	0.6	Facebook, Akamai
26	crashlytics.com	1.4	AWS
25	googleadservices.com	0.5	Google

- **Network topology test:** For each resolved IP address, we run UDP traceroutes to study CSP-MNO peering and topological relationships.

We run the aforementioned experiments in isolation from other experiments continuously over 4 weeks, from April 5, 2017 until May 6, 2017. The combined results of the three tests for a given FQDN produce a “sample”. The measurements are run continuously, and the time period between our samples varies between 4 and 24 hours. Thus, our experiments may cover several instances across time of day and various radio conditions. Nevertheless, we run the experiments sequentially to guarantee similar network conditions across FQDNs. We do not measure metrics such as Time to First Byte [31] and download speed as they are more likely to be affected by server-side artifacts, which is beyond the scope of this study. The measurement code and dataset are publicly available [29] to satisfy the reproducibility principle.

Data sanitation: We leverage metadata provided by the MONROE nodes to avoid bias introduced by uncontrolled changes in the wireless technology coverage (see Section VI for further details) while we use packet captures to ensure that the measurements run by higher layer tools are accurate. We also remove measurements that may be affected by MNOs enforcing volume caps which may inflate latency. After sanitizing our dataset, we obtain a set of 173,679 valid samples.

IV. CSP PREVALENCE ON MOBILE APPS AND SERVICES

In this section we analyze the prevalence of CSPs among mobile apps in order to identify the main players supporting the mobile Internet and multi-CDN strategies. 55 of the CSPs that we identified in the previous section are present in the Lumen dataset and 85.2% of the apps connect to at least one of them. However, such a high prevalence is not necessarily a consequence of app-developer decisions. As we can see in Table II, advertising-related FQDNs have the highest app penetration [24, 32]. Just the domain `googlesyndication.com`, hosted by Google on its own datacenters, is present in over 29% of the apps in our dataset.

Figure 2 ranks CSPs by the number of mobile apps connecting to them, also showing the percentage of FQDNs and second-level domains that they support. We report CSP prevalence by apps, by FQDNs and second-level domains (SLDs in short) to give a sense of both app and domain-level usage. The figure reveals a clear power-law distribution. While 49 CSPs (*e.g.*, Purepeak, not shown in the figure) have a marginal presence as they receive connections from less than 1% of the apps, six CSPs play a central role on this market, being associated with 85.05% of the apps. AWS

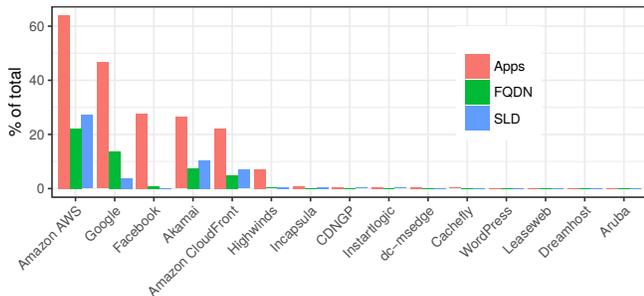


Fig. 2: Top-15 CSPs prevalence by app, FQDN and SLD.

TABLE III: Multi-CSP strategies by FQDN and SLD.

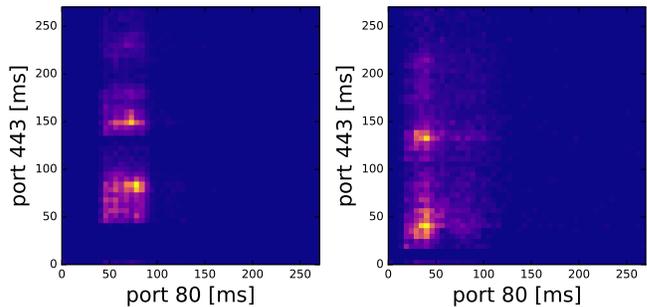
# of CSP	1	2	3	≥ 4
FQDN(%)	97.4	2.5	0.1	0.0
SLD(%)	84.9	13.6	1.3	0.2

and CloudFront, both owned by Amazon, are the most used CSPs by mobile on-line services, supporting 27% and 6% of SLDs, respectively. Other CSP services like Facebook are easily found across mobile apps as many of them integrate Facebook services, including advertising, through the Facebook Graph API. However, Facebook also leverages Akamai’s infrastructure as well as its own, which is only open to affiliate companies like Instagram. Google, instead, has opened their infrastructure to third parties with the Google App Engine service.

Finally, we leverage our PTR-CSP mapping to identify instances of multi-CSP strategies on a per-FQDN and per-SLD basis. According to our results (Table III), only 3% of the analyzed FQDNs and 15% of SLDs use at least 2 CSPs. After carefully inspecting such FQDNs, we can conclude that they are associated with large companies such as Samsung, Adobe, Chartboost, Unity or Facebook among others. This observation suggests that multi-CSP strategies are specific to large companies — probably because of cost-related reasons — despite their performance and reliability benefits.

V. CSP PERFORMANCE AND INTEGRATION WITH MNOS

This section aims to analyze the actual performance of CSPs in the MONROE platform. To that end, we study first the presence of in-path middleboxes on MONROE’s MNOS given that their presence can bias CSP performance measurements (Section V-A). In Section V-B, we characterize the DNS infrastructure deployed by our tested MNOS and their support for EDNS, a DNS extension used by many CSPs to correctly locate the end-user. Third, we study the CSPs’ performance in the MONROE platform (Section V-C) following the methodology described in Section III-C, further analyzing in detail the effect of topological and peering relationships between MONROE MNOS and our six target CSPs (Section V-D) on TCP and TLS connection establishment. We conclude with an analysis of the impact of international roaming on TCP and TLS performance (Section V-E).



(a) Location: Spain (roaming)

(b) Location: Italy

Fig. 3: Heatmap of TCP connection times over ports 80 and 443 for a Vodafone Italy SIM when roaming (left) and when connecting from the home network (right). Lighter colors indicate more repetitions.

A. In-path Middleboxes

MNOs may deploy Performance Enhancing Proxies (PEP) to optimize mobile traffic performance [15] using techniques like TCP splitting. However, TCP-splitting proxies can introduce bias in our measurements, as the TCP connection time obtained is to the proxy rather than to the final endpoint (*i.e.*, the CSP). In order to identify such scenarios, we run the Netalyzr network troubleshooting tool[33] directly on the MONROE nodes [15, 34]. Netalyzr only revealed a TCP-splitting proxy for Vodafone Italy in TCP port 80. Its presence is also visible in Figure 3, where Vodafone Italy’s TCP connection time over port 80 is much lower than on port 443 for most of our measurements. MNOs do not deploy TCP-splitting proxies on port 443 as their presence can interfere with mobile apps’ securing TLS flows [35]. Due to the above, we focus our TCP performance analysis only on TCP port 443, a middlebox-free path for all our MNOs.

B. DNS infrastructure

DNS Proxies: As for stateful TCP traffic, MNOs may also deploy DNS proxies to gain full control over user’s traffic [15]. Their presence can interfere with CSP’s performance by altering both DNS queries and responses. DNS resolvers deployed by all Swedish carriers perform cache delegation of DNS records, a common practice across MNOs [36]. Furthermore, TIM Italy proxies DNS traffic and performs DNS wildcarding (*i.e.*, they resolve non-existing names). Finally, Vodafone (IT) seems to actively block all DNS requests towards any DNS resolvers containing the EDNS flag. This behavior could be caused either by a DNS proxy or a DNS-aware firewall. Consequently, we cannot study the impact of the EDNS extension on Vodafone (IT).

EDNS support: Both Google’s Public DNS and OpenDNS publicly claim to support the EDNS `edns-client-subnet` flag. We study whether DNS resolvers — including the default DNS resolver provided by the MNO as well as Google’s and OpenDNS resolvers — support EDNS by checking whether the ECS option is included in the response, as specified by the RFC 7871 [21].

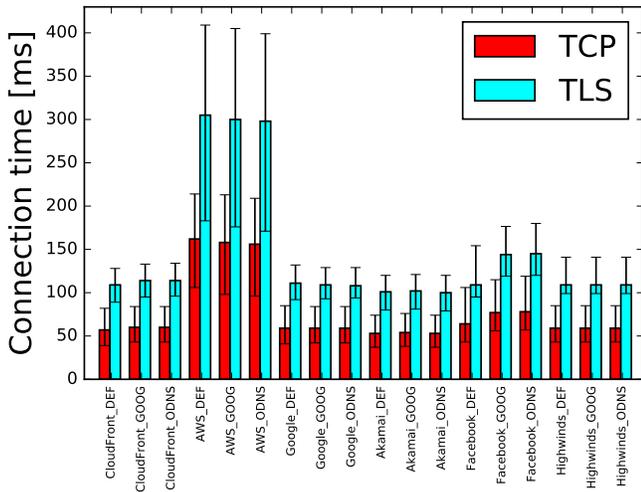


Fig. 4: Median values of TCP connection time and TLS handshake duration for $\langle \text{CSP} \rangle \langle \text{DNS resolver} \rangle$ combinations. Error bars represent the 25th and 75th percentile.

According to our results, only Google’s public DNS seems to completely follow the RFC 7871 recommendations: the ECS option is absent from OpenDNS responses despite supporting EDNS [30]. This observation leads us to believe that OpenDNS may have their own interpretation of the standards. The only instance where the flag is absent from the response, when contacting Google is for TIM (IT), probably due to an in-path DNS proxy. Additionally, we check both the EDNS buffer size and reply size. For all the Norwegian operators as well as Telia SE, DNS replies are limited to 512 bytes, which is an indication that EDNS is not supported.

As opposed to public DNS resolvers, MNO’s recursive DNS resolvers may not need to provide EDNS support to help the authoritative name server locate the end-user. MNOs may assign the same public IP address space both to end-users and recursive DNS resolvers. To study the feasibility of this, at least for the MNOs under consideration, we record the public IP of the user by sending traffic to a machine we own. Then, we get the public IP of the default DNS resolver of the operator, by resolving the URL whoami.akamai.net, which returns the IP from which Akamai servers receive the DNS request. Our experiment shows that all our studied MNOs use different address spaces for their DNS infrastructure and their subscribers, hence making the localization of the user based on the public IP address of the resolver impossible.

C. CSP Performance

Once we understand the DNS infrastructure of each MONROE MNO, we study the “quality” of the DNS responses provided by each DNS service and its impact on TCP and TLS connection time. In only 45% of the tests, the three resolvers return IPs that belong to the same set of $\setminus 24$ subnets. However, only 2% of the responses provided by the third-party DNS providers point to machines hosted in a different country than the one in the responses of the default MNO resolver. In that

case, the increase of the median handshake duration is at most 2ms and 9ms for TCP and TLS, respectively.

Figure 4 shows the TCP and TLS connection time across all our MNOs, grouped by the DNS resolver being used. As we can see, the three resolvers perform similarly at the transport level regardless of EDNS support. MONROE nodes connecting to IPs provided by the default DNS resolver have marginally better TCP connection times than those connecting to servers proposed by OpenDNS and Google’s public DNS (around 3ms). Facebook is the only CSP exhibiting significantly better performance when using the MNO’s default resolver in all MNOs but in the case of TIM (IT) where all the DNS responses perform similarly because of its in-path DNS proxy. For the other MNOs, the responses provided by the default DNS resolver are 10ms faster at the TCP level compared to Google’s and OpenDNS’ responses.

Figure 4 also reveals performance differences across CSP services, namely due to the scale and coverage of their infrastructure. As we can see, Amazon CloudFront (a CDN provider with a vast infrastructure) and Amazon AWS performance differences are remarkable. A possible explanation for that is that CloudFront is a dedicated CDN provider service, with replicas in multiple locations, whereas AWS is a cloud computing platform with just a few data centers in Europe. In fact, among the many FQDNs studied, we can identify that many Amazon AWS customers do not leverage AWS’s global infrastructure and decide to host their services entirely in US data centers, hence further inflating the delays.

D. CSP-MNO Integration

In this section, we study the topological relationships (*e.g.*, peering) and the geographical distribution of CSPs’ data centers and MNOs’ Points of Presence (PoP) in order to identify their effect on TCP and TLS connection time. We use MONROE’s traceroute measurements to characterize the interconnection between MNOs and CSPs. We retrieve the Autonomous Systems (ASes) that advertise in BGP the most specific network prefix covering the IP of each hop in the traces we collect.

Given the difficulty to accurately measure the geographic distance between an MNO and a CSP without insider knowledge, we define the following distance metrics:

- **Country distance:** This metric counts the number of unique countries traversed by a traceroute probe from the MONROE vantage point to the target CSP. We retrieve country-level information by mapping each hop’s IP along the data path to a country code using MaxMind’s free GeoIP service. Consequently, MaxMind’s accuracy [37] constrains our analysis accuracy.
- **Organization distance:** This metric reflects the number of unique organizations traversed by a traceroute probe from the MONROE vantage point to the target domain. For each hop’s IP address along the traceroute data path, we retrieve the AS using the most specific prefix advertised in BGP. Then, we use CAIDA’s AS-to-Organization mapping

TABLE IV: Median and standard deviation values of the organization and country distance per CSP when aggregating all the MNOs that we measure in the MONROE platform. We target six main CSPs we previously identified in the analysis of the Lumen dataset.

CSP	Performance Tests #	Country dist. median (std)	Org. dist. median (std)
CloudFront	408,537	3 (1.6)	2 (1.0)
AWS	400,240	2 (1.3)	2 (1.0)
Google	69,839	2 (0.8)	2 (0.7)
Akamai	67,375	2 (0.7)	2 (0.7)
Facebook	9,731	3 (1.1)	3 (0.9)
Highwinds	8,984	3 (0.9)	3 (0.7)

TABLE V: The effect of organization and country distance on TCP connection time [median (std)].

CSP	MNO	Organization dist.	Country dist.	TCP conn. time [ms]
Akamai	Telia NO	2 (0.3)	2 (0.4)	42.0
	Telenor NO	3 (0.9)	2 (1.0)	65.0
	Telia SE	3 (0.6)	2 (0.4)	52.0
	TIM IT	2 (0.3)	2 (0.5)	30.0
	Orange ES	2 (0.5)	2 (0.7)	39.0
	Yoigo ES	2 (0.7)	2 (0.5)	38.0
Google	Telia NO	3 (0.0)	2 (0.0)	43.0
	Telenor NO	2 (0.4)	2 (0.4)	75.0
	Telia SE	3 (0.7)	2 (0.3)	57.0
	TIM IT	2 (0.5)	2 (0.5)	30.0
	Orange ES	2 (0.4)	3 (1.2)	56.0
	Yoigo ES	2 (0.4)	3 (0.8)	41.0

dataset [38] to identify the parent organization for each IP address².

Large Internet entities such as Tier-1 ISPs and CSPs may use multiple AS numbers and yet advertise various IP blocks with the same origin AS [39]. This makes identifying the actual entity behind a given IP block and its usage harder. For example, Telia Sonera – an ISP offering both fixed and mobile connectivity that owns multiple AS numbers – advertises its reachability information in BGP, including its mobile subscribers, using the AS number associated to its Tier-1 ISP.

Results: Table IV presents the median country and organization distance for each CSP across all our MONROE nodes. We identify organizations and countries along the path based on the IPs we see in the traceroute data. In general terms, Akamai and Google have the smallest distance metrics, presumably because of the extensive use of peering and the presence of caches within MNO networks. Table V shows the impact of distance values on the TCP connection time for a selection of MNOs and CSPs. When aggregating all results across all MNOs, we can see that most flows cross 2 or 3 organizations on average at most. The following paragraphs discuss specific MNO cases.

Telia (SE, NO): Both Telia and Telenor operate in Sweden and Norway. The median value organization distance between Telia and the majority of our target CSPs is 3, regardless of

the operating country. The only exception is Akamai for Telia NO, where the median value goes down to 2 organizations. From our MONROE measurements, we find that Telia (SE) (AS3301) reaches over 85% of the Akamai services directly through its Swedish parent organization, Telia Company (AS1299), which also acts as its Internet transit provider [40], hence possibly inflating the network path. Similarly, we find that Telia (NO) (AS12929) always routes its traffic through Telia Company (AS1299), which has over 1,500 customer networks and 50 peers. The ongoing strategic alliance between Akamai and Telia [41] allows the CDN-MNO collaboration to improve the end-user experience in a cost-effective manner. We measure a median TCP connection time of 52ms and of 42ms for Telia SE and Telia NO to Akamai, respectively. According to [40], Telia Company registers as a peer for Google Inc. (AS15169), explaining the organization distance of 3 to Google from Telia NO and Telia SE. The TCP connection time we measure towards Google servers is similar to Akamai’s prior values: 43ms from Telia NO and 57ms from Telia SE, respectively.

Telenor (SE, NO): Telenor Sweden’s median distance towards each one of the six CSPs is 3, while for Telenor NO the distance varies in median value with Google having the smallest value (2 organizations) and Akamai – a median value of 3 organizations. Telenor NO (AS8786) belongs to the Telenor group and is registered for mobile operations in Norway. AS8786 always routes its traffic using the parent company AS2119. Similarly, Telenor SE also depends on the same AS2119 to reach targets. Our traceroute measurements reveal a high diversity of AS-level paths when reaching Akamai target servers with a median organization distance of 3. In 20% of these paths, we observe IP address blocks belonging to the Amsterdam Internet Exchange (AMS-IX). This suggests that Telenor (NO) leverages its peering connections in the Netherlands at AMS-IX to reach content hosted by Akamai. This result does not suggest an ongoing alliance between Telenor and Akamai. This observation translates in performance degradation: 65ms in median value from Telenor NO to Akamai (Table V) In Telenor SE we also identify that 27% of the requests have a country distance of 4 or higher, while for the rest of the MNOs this is less than 2%.

TIM (IT): Telecom Italia (AS6762) (TIM) is one of the largest operators in the world, with peering connections to all the other Tier-1 ISPs [40]. Thanks to its dense global interconnection and large user-base, CSPs such as Akamai entered into partnerships agreements with TIM to optimize content delivery and increase the quality of experience of content consumers [42]. We note that when breaking down the organization distance (Table IV) for TIM on the six different CSPs, only Facebook and Highwinds have a median distance value of 3 organizations, while for the rest we find a median distance value of 2 organization. The tight integration of TIM with Akamai and Google translates into low TCP connection times: 30 ms both for Akamai and Google.

Yoigo (ES) and Orange (ES): The two MNOs that we

²The same organization may own several AS numbers, thus we refrain from using the AS path length as a distance metric.

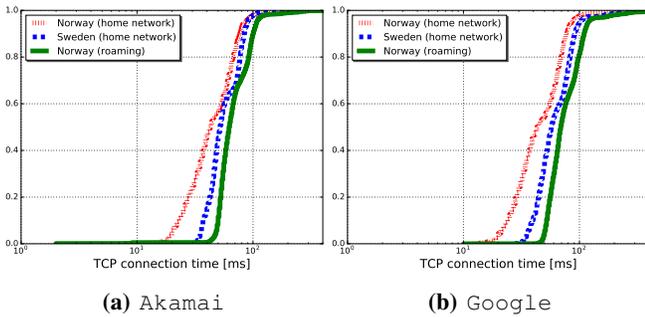


Fig. 5: Effect of roaming on TCP handshake over Telia.

measure in Spain, Yoigo and Orange, present similar distances to the two CSPs listed in Table V. Our traceroute experiments show that Yoigo (AS16299) relies exclusively on its two transit providers to reach popular content, namely Telia Company (AS1299) and Orange Spain (AS12715). Similarly, Orange ES relies on its two providers, Orange S.A. (AS5511) and Level 3 (AS3356) to reach both Akamai and Google services. Both MNOs have Tier-1 ISPs as providers, and leverage the latter’s dense interconnection with Google and Akamai to ensure good performance for their customers. Our measurements report a median TCP connection time of 39ms from Orange ES to Akamai and of 38ms from Yoigo ES to Akamai. The TCP connection time from both MNOs is slightly higher towards Google (Table V), which may be caused by a higher median country distance.

E. International Roaming

When a SIM card is in international roaming state, MNOs can either forward their traffic to the home network before reaching the Internet (*i.e.*, home routing) or use the host MNO infrastructure (*i.e.*, local breakout) [15]. Most MNOs decide to implement home routing so that they can keep control over their subscribers’ traffic at the expense of inflated path length.

Our traceroute analysis reveals that our MNOs implement home routing roaming as the number of hops to a given target remains the same – due to the presence of a transparent tunnel – but path latency increases. Figure 5 presents the CDF of TCP connection time of all the successful connections over ports 80 and 443 in logarithmic scale for Telia SIM cards. We compare the performance for the three Telia SIM cards in our MONROE nodes: two locally connected in Norway and Sweden, and a third Telia Sweden SIM card roaming on a Norwegian MONROE node. According to our results, TCP connection time takes on average 20ms more on roaming devices than on those connecting directly to the local network (Figure 3) because of the country distance inflation. The home routing roaming approach, by its nature, defeats the purpose of CDNs placing content close to the user. Roaming users do not benefit from existing peering agreements between the host network and large CSPs. For this reason, when the target server belongs to a well-provisioned CSP— possibly better peered — this impact is greater, compared to CSPs present in a few locations. For example, CloudFront and Google

services over the network of Telia have at least 20% delay inflation. In contrast, the delay inflation for AWS under the same conditions is usually below 15%. We also identify a clear performance degradation in the Vodafone (IT) SIM card roaming in the Spanish node. It is visible in Figure 3b, where the point cloud is located further away from the axis (higher delay) compared to the local connection.

VI. STUDY LIMITATIONS

App representativity: Our mobile app sample is limited to the traffic logs obtained from our Lumen users. Nevertheless, as discussed in our previous work [24], we consider the apps in our dataset to be representative of those used by average mobile users from all over the world: 48% of the apps in our records have more than 1M installs while 71% of the apps listed on the Google Play Top-50 charts for USA, Spain, Germany, India and UK are also present in our dataset.

CSP representativity: Because of limited testing cycles on the MONROE platform, our study focuses on the six most representative CSPs across our sample of mobile apps. However, these CSPs are likely the better peered ones with large MNOs due to their popularity. Finally, we intentionally execute measurements towards FQDNs rather than towards specific CSPs to analyze realistic domains and characterize DNS-level artifacts. Unfortunately, this is skewing our number of samples towards the most popular CSPs.

Cellular technology: We only consider measurements run only over LTE due to its rapid adoption rate and its low latency radio link. Including 3G and 2G cellular technologies in our studies could bias our empirical results due to the significant differences in the radio access link.

Active measurements: We execute our active measurements in real networks against real systems. As a result, the scale and accuracy of our results may be limited by a number of factors beyond the scope of this study. We cannot fully control aspects such as CSP load balancing mechanisms and server load, cellular network behavior, network load and congestion, and radio link stability which may influence and introduce bias in our results.

Given the aforementioned limitations, our goal is not drawing conclusive causal relationships between MNOs and CSPs, but providing a first study of this complex ecosystem to motivate further research. To that end, we made public our data and measurements scripts so that other researchers can continue, extend and improve our work.

VII. CONCLUSIONS

In this paper, we performed the first holistic analysis of the complex ecosystem formed by mobile applications, cloud service providers (CSPs), and Mobile Network Operators (MNOs). We aimed to comprehensively characterize their relationships and dynamics and measured their performance with dedicated active measurements. We leveraged accurate traffic fingerprints from thousands of mobile apps that we collected through crowd-sourcing with Lumen [3]. This data allowed us

to *i*) identify the most relevant CDNs and cloud providers for mobile traffic; *ii*) map their connectivity with relevant European MNOs; and *iii*) measure their performance using the MONROE platform [4]. Our results show a significant reliance of apps on mobile CSPs with the major CSPs being used by 85% of the apps. We reported path inflation (*e.g.*, due to poor peering relationships and roaming) and presence of middle-boxes (*e.g.*, in-path DNS proxies) which can significantly impede CSP performance, but we saw no noticeable difference in performance metrics when using different DNS resolvers or enabling the EDNS parameters. Our active measurement dataset, the code for the measurement experiments and the CSP mapping tool are publicly available [29].

ACKNOWLEDGMENTS

This work has been supported by the European Union H2020-ICT grants 644399 (MONROE) and 688421 (MAMI), by the Madrid Regional Government through the TIGRE5-CM program (S2013/ICE-2919), the Ramon y Cajal grant from the Spanish Ministry of Economy and Competitiveness RYC-2012-10788 and the NSF Awards: CNS-1740895, CNS-1350720, CNS-1719386 and CNS-1564329. The opinions expressed and arguments employed reflect only the authors' view. The European Commission is not responsible for any use that may be made of this research paper. Part of this work was carried out while Foivos Michelinakis was visiting KAIST.

REFERENCES

- [1] E. Nygren Et al., "The Akamai network: a platform for high-performance internet applications," *ACM SIGOPS*, 2010.
- [2] U. Goel Et al., "HTTP/2 Performance in Cellular Networks," in *ACM MobiCom*, 2016.
- [3] A. Razaghpanah Et al., "Haystack: A multi-purpose mobile vantage point in user space," *arXiv preprint arXiv:1510.01419v3*, 2016.
- [4] Ö. Alay Et al., "Experience: An Open Platform for Experimentation with Commercial Mobile Broadband Networks," in *Proc. ACM Mobicom*, 2017.
- [5] "Akamai Facts & figures," 2017, <https://www.akamai.com/uk/en/about/facts-figures.jsp>.
- [6] "AWS Global Infrastructure," <https://aws.amazon.com/about-aws/global-infrastructure/>.
- [7] B. Frank Et al., "Pushing CDN-ISP Collaboration to the Limit," *ACM SIGCOMM CCR*, 2013.
- [8] I. N. Bozkurt Et al., "Why is the internet so slow?!" in *Proc. PAM*. Springer, 2017.
- [9] D. Wu Et al., "Mopeye: Opportunistic monitoring of per-app mobile network performance," in *Proc. USENIX ATC*, 2017.
- [10] I. Castro Et al., "Remote peering: More peering without internet flattening," in *Proc. ACM CoNEXT*, 2014.
- [11] A. Dhamdhere and C. Dovrolis, "The Internet is flat: modeling the transition from a transit hierarchy to a peering mesh," in *Proc. ACM CoNEXT*, 2010.
- [12] "Netflix Open Connect." [Online]. Available: <https://openconnect.netflix.com>
- [13] "Akamai Network Partnerships." [Online]. Available: <https://www.akamai.com/uk/en/products/network-operator/akamai-network-partnerships.jsp>
- [14] V. K. Adhikari Et al., "Measurement study of Netflix, Hulu, and a tale of three CDNs," *IEEE/ACM ToN*, 2015.
- [15] N. Vallina-Rodriguez Et al., "Beyond the radio: Illuminating the higher layers of mobile networks," in *ACM MobiSys*, 2015.
- [16] "During Netflix money fight, Cogent's other big customers suffered too." [Online]. Available: <https://arstechnica.com/information-technology/2014/11/during-netflix-money-fight-cogents-other-big-customers-suffered-too/>
- [17] A. Molavi Kakhki Et al., "Identifying traffic differentiation in mobile networks," in *Proceedings of ACM IMC*, 2015.
- [18] J. P. Rula and F. E. Bustamante, "Behind the curtain: Cellular dns and content replica selection," in *Proc. ACM IMC*, 2014.
- [19] J. S. Otto Et al., "Content delivery and the natural evolution of DNS: remote dns trends, performance issues and alternative solutions," in *Proc. ACM IMC*, 2012.
- [20] F. Zarinni Et al., "A first look at performance in mobile virtual network operators," in *Proc. IMC*. ACM, 2014.
- [21] C. Contavalli Et al., "Client Subnet in DNS Queries," RFC 7871 (Informational), Internet Engineering Task Force, May 2016. [Online]. Available: <http://www.ietf.org/rfc/rfc7871.txt>
- [22] S. Narayana Et al., "Distributed wide-area traffic management for cloud services," in *ACM SIGMETRICS*, 2012.
- [23] N. Vallina-Rodriguez Et al., "Tracking the Trackers: Towards Understanding the Mobile Advertising and Tracking Ecosystem," in *Proc. DAT Workshop*, 2016.
- [24] A. Razaghpanah Et al., "Apps, Trackers, Privacy, and Regulators: A Global Study of the Mobile Tracking Ecosystem," in *Proc. NDSS*, 2018.
- [25] ICSI, "Lumen Privacy Monitor," 2016, <https://play.google.com/store/apps/details?id=edu.berkeley.icsi.haystack>.
- [26] "CDNfinder by CDNPlanet," <https://www.cdnplanet.com/tools/cdnfinder/>.
- [27] McAfee, "Trusted Source," <http://www.trustedsource.org/>.
- [28] OpenDNS, "Domain Tagging," <https://domain.opendns.com>.
- [29] "Cloudmap Project," 2018, <http://wireless.networks.imdea.org/cloudmap-project>.
- [30] "edns-client-subnet participants," <http://www.afasterinternet.com/participants.htm>.
- [31] E. Halepovic et. al, "Can You GET Me Now?: Estimating the Time-to-first-byte of HTTP Transactions with Passive Measurements," in *Proc. ACM IMC*, 2012.
- [32] N. Vallina-Rodriguez Et al., "Breaking for Commercials: Characterizing Mobile Advertising," in *Proc. ACM IMC*, 2012.
- [33] ICSI, "Netalyzer," <http://netalyzer.icsi.berkeley.edu/>.
- [34] C. Kreibich Et al., "Netalyzer: illuminating the edge network," in *Proc. ACM IMC*, 2010.
- [35] A. Razaghpanah Et al., "Studying TLS Usage in Android Apps," in *Proc. ACM CoNEXT*, 2017.
- [36] M. Almeida Et al., "Dissecting DNS Stakeholders in Mobile Networks," in *Proc. ACM CoNEXT*, 2017.
- [37] Y. Shavitt and N. Zilberman, "A geolocation databases study," *IEEE J-SAC*, vol. 29, no. 10, 2011.
- [38] CAIDA/UCSD, "Mapping Autonomous Systems to Organizations: CAIDA's Inference Methodology." [Online]. Available: <https://www.caida.org/research/topology/as2org/>
- [39] J. P. Rula Et al., "Cell spotting: studying the role of cellular networks in the internet," in *Proc. ACM IMC*, 2017.
- [40] CAIDA/UCSD, "AS Relationships Dataset." [Online]. Available: <http://as-rank.caida.org/>
- [41] T. Sonera, "Telia and Akamai announce strategic relationship to deliver enhanced web services throughout Europe." [Online]. Available: <https://www.teliacompany.com/en/news/press-releases/2000/1/telia-and-akamai-announce-strategic-relationship-to-deliver-enhanced-web-services-throughout-europe/>
- [42] Akamai Press Release, "Akamai And Telecom Italia Enter Into Partnership To Offer Content Delivery And Web Optimization Solutions." [Online]. Available: <https://www.akamai.com/us/en/about/news/press/2015-press/akamai-and-telecom-italia-enter-into-partnership-to-offer-content-delivery-and-web-optimization-solutions.jsp>