

A Case for Faster Mobile Web in Cellular IPv6 Networks

Utkarsh Goel
Montana State University
utkarsh.goel@montana.edu

Moritz Steiner
Akamai Technologies, Inc.
moritz@akamai.com

Mike P. Wittie
Montana State University
mwittie@cs.montana.edu

Martin Flack
Akamai Technologies, Inc.
mflack@akamai.com

Stephen Ludin
Akamai Technologies, Inc.
sludin@akamai.com

ABSTRACT

The transition to IPv6 cellular networks creates uncertainty for content providers (CPs) and content delivery networks (CDNs) of whether and how to follow suit. Do CPs that update their CDN contracts to allow IPv6 hosting achieve better, or worse performance in mobile networks? Should CDNs continue to host mobile content over IPv4 networks, or persuade to their CP customers the performance benefits of IPv6 content delivery?

In this paper we answer these questions through a comprehensive comparison of IPv4 and IPv6 mobile Web performance in cellular networks in the US from the point of view of Akamai's content delivery infrastructure. Our data show that IPv6 hosting outperforms legacy IPv4 paths in mobile Web. Our analysis leads to clear recommendations for CPs to transition to IPv6-hosted mobile Web. Finally, we propose new mechanisms, through which CDNs can safely transition mobile content to IPv6-enabled servers for improved content delivery.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Distributed Networks, Wireless Communication; C.2.5 [Local and Wide-Area Networks]: Internet; C.4 [Performance of Systems]: Performance Attributes, Measurement Techniques

Keywords

Mobile Web; performance; cellular; IPv6; IPv4; CDN; measurement; DNS

1. INTRODUCTION

Despite many years of research to improve Web performance in mobile and wireless networks, users remain dissatisfied with lengthy webpage load times [56]. As Internet Service Providers (ISPs) upgrade their network infrastructure from IPv4 to IPv6, understanding the performance of mobile content delivery in cellular IPv6 networks is crucial. In this study, we take a novel approach to characterize the dynamically changing IPv6 ecosystem from the point of view of Akamai's content delivery infrastructure

for cellular networks [51]. We argue that, unlike PlanetLab and Amazon EC2 datacenters [2, 15], Akamai's content delivery servers are so deeply deployed inside several cellular ISPs' networks that the end-to-end communication between mobile devices and Akamai's servers need not, strictly speaking, touch the wired public Internet outside the cellular network. As a result, Akamai's unique content delivery infrastructure enables us to view the end-to-end cellular ecosystem between mobile devices and cellular gateways and evaluate how content is delivered over cellular IPv6 networks from the perspective of content providers (CPs), ISPs, and other content delivery networks (CDNs) [6, 16].

CPs, such as Facebook and others, care about the experience of users with their respective applications. To deliver application content from datacenters to users in a timely manner, CPs make contractual agreements with CDNs to ensure content has high availability, is secure, and is delivered to users through low latency connections. Some CPs sign contracts with CDNs for content delivery only over a cellular ISP's IPv4 network, while other CPs sign contracts for content delivery over IPv6 networks. Although, CPs are aware that ISPs are deploying IPv6 in their networks and CDNs are offering IPv6 hosting of mobile content, CPs remain uncertain whether upgrading to IPv6 will improve or worsen the performance of mobile content delivery.

CDNs, such as Akamai and others, care about the performance of content delivery to their mobile users. Although CDNs strive to upgrade their infrastructure to overcome IPv4 address scarcity and to make content available over IPv6 networks [10, 27, 29, 50, 55], one of their goals is to ensure that the performance of content delivery over an ISP's IPv6 network is as good as over that ISP's IPv4 network. CDNs generally act as a surrogate infrastructure for its many CPs, and for stability, reliability, and contractual implications, the configurations for each CP are often only changed with respective permission granted from the CP. Thus, adoption of new content delivery techniques, such as IPv6, is often a multi-year process as its performance implications become better-understood in real-world conditions over time.

As cellular network operators adopt IPv6 addressing to resolve the challenges imposed by IPv4 address scarcity [29], the research community has shown interest in understanding different IPv6 deployment strategies within ISPs worldwide, and the adoption rate of IPv6 among mobile users and content providers [27, 50, 58]. Consequently, CPs and CDNs remain reluctant to embrace a wide-scale transition to IPv6, as they remain unaware of performance of IPv6 networks deployed by cellular carriers.

In this paper, we take a novel approach to investigate and expose performance of IPv6 and IPv4 ecosystems from a CDN's perspective – in between cellular ISPs and Content Providers. Our goal is to improve awareness within different networking communities

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MobiCom'16, October 03-07, 2016, New York City, NY, USA

© 2016 ACM. ISBN 978-1-4503-4226-1/16/10...\$15.00

DOI: <http://dx.doi.org/10.1145/2973750.2973771>

about performance benefits (if any) of serving mobile content over IPv6, as opposed to IPv4. Our work precisely describes the current role of a CDN in connecting mobile users to content servers in the evolving cellular ecosystem in the US. To the best of our knowledge, our work is the most detailed investigation to compare mobile Web performance over IPv4 and IPv6 networks. We classify the four major contributions of this work as follows:

Dataset Richness: We conducted a large scale, comprehensive study to measure IPv6 performance in four major cellular carriers in the US to compare its native and NAT64/DSLite deployments. Using Akamai CDN infrastructure, we collected a rich dataset consisting of millions of data points of measured IPv6 and IPv4 performance, during the months of January - August in 2015.

Measurement: Our study investigates IPv6 performance across multiple factors that influence Web performance on cellular networks.

- We compare IPv6 and IPv4 networks through 1) round trip time between clients and CDN servers; 2) time to resolve domain names from cellular DNS; and 3) webpage load time.
- We extend Akamai's Real User Monitoring System (RUM) [13] to accurately extract Web performance metrics for mobile content hosted on IPv6-enabled content servers in US cellular networks.

Inferences Drawn: Our experience with Akamai's content delivery infrastructure shows that IPv6 networks deployed by cellular ISPs outperform their IPv4 networks.

- Our analysis includes recommendations for CPs to host mobile content on IPv6 for improved user experience.
- We also recommend that CDNs deliver mobile content over IPv6 to avoid in-path middleboxes for IP address translation deployed by cellular carriers.
- And finally, we suggest cellular network operators upgrade their network infrastructure to support IPv6, instead of continuing to deploy legacy IPv4 technologies in their network.

Problems Discovered: During our study, we discovered the following three problems related to how IPv6 content is delivered in cellular networks. We also propose several solutions we adopted to address these problems.

- We discovered that the DNS lookup process takes longer on Android devices with IPv6 capability than Android devices with IPv4-only capability. The lookup time is high because IPv6-capable devices wait for both `Type A` and `Type AAAA` (pronounced 'quad A') DNS queries to finish before establishing a TCP connection. Because of higher DNS lookup time for IPv6-capable clients, we observe that IPv6 clients in the Sprint network often experience slower webpage loads when connecting to IPv4 servers, than IPv4-only clients connecting to same IPv4 servers.

Solution: To address this problem, we made four recommendations to the Google Android team to reduce long Round Trips Times (RTTs) to cellular resolvers. First, if the DNS lookup process on IPv6-capable Android devices could be modified to send `AAAA` and `A` DNS queries in parallel, lookup times could potentially become twice as fast. Second, we suggested that additional speedup in DNS lookup could be achieved by letting client browsers indicate to the mobile OS that they do not need to wait to get the `A` lookup back if they get an `AAAA` answer in DNS response. Third, in cases

where mobile clients are on an IPv6-only network, the mobile OS could be modified to only send DNS `AAAA` queries, instead of both `AAAA` and `A`. Finally, the existence of `A` or `AAAA` only answers could be cached per-name for some time on the device, not withstanding the caching of the actual answers, which would enable the device to make a smarter query the following time that an IP address is needed.

- In the case of T-Mobile (and applicable to other major IPv6-only networks worldwide), we discovered that when IPv6-capable clients resolve an IPv4-only domain name, the cellular DNS introduces an extra round trip to the DNS Authorities. We discuss details in Section 7.

Solution: We develop and prototype ONETRIP, a technique for DNS Authorities to eliminate the extra round trip in DNS lookups when resolving IPv4-only domains. Through in-lab simulations we show that DNS lookup times reduce significantly when DNS Authorities use ONETRIP. We also experimentally verify that in T-Mobile's production cellular network, ONETRIP maintains end-to-end connectivity.

- In a cellular network outside of the US, IPv6 packets were being routed via the US, which resulted in latency of end-to-end connections on IPv6 to be higher than IPv4 by 200 ms. While it is possible that IPv6 latency may be higher than IPv4 in some networks, we argue that it could be due to misconfigured routing policies.

Solution: Based on our findings, Akamai's network team is actively working with that cellular carrier to resolve misconfiguration in its IPv6 routing.

The rest of the paper is organized as follows. In the next section, we offer a discussion on how IPv6 is deployed by different cellular ISPs in the US. An overview of different IPv6 deployment strategies will support our measurement techniques and research findings. In Section 3, we describe our data collection methodology. In Section 4, 5, and 6, we investigate the component differences of Web performance in IPv6 and IPv4 networks through measuring the round trip latency between clients and CDN servers, the DNS lookup time, and the webpage load time. In Section 7, we introduce ONETRIP as a technique for DNS Authorities to eliminate unnecessary round trips from DNS lookups in IPv6-only networks. In Section 8, we discuss related work. Finally, we conclude in Section 9.

2. OVERVIEW OF IPv6 DEPLOYMENT IN CELLULAR NETWORKS

A cellular carrier that supports IPv6 addressing must provide a way for its IPv6 devices to connect with IPv4 and dual-stacked (both IPv4 and IPv6 addresses) Internet servers. Cellular network architectures are influenced by a variety of factors such as the capabilities of the existing infrastructure hardware to support IPv6 addressing, urgency to upgrade networks to IPv6, number of users with IPv6-capable devices, number of available IPv4 addresses, etc., which result in ISPs taking different approaches to IPv6 deployment [18]. In this section we provide an overview of how the different cellular carriers in the US have upgraded their IPv4 networks to provide IPv6 addressing to their users, based on the publicly available information from those carriers.

T-Mobile is an IPv6-only network for all phones with support for 4G4XLAT, which includes all phones with Android version 4.3 and above [24]. For older versions of Android, as well as iPhone, Blackberry, and Windows devices, T-Mobile is an IPv4-only




Figure 1: IPv6 infrastructure deployment in different cellular carriers in the US.

network. As a result, IPv6 devices in T-Mobile network always transmit IPv6 packets, whereas, IPv4 devices always transmit IPv4 packets. Thus, the choice of addressing (IPv4/IPv6) for devices is based on whether the device supports 464XLAT.

In Figure 1(a) we depict the high level infrastructure of the IPv6 network deployed by T-Mobile. We show that when an IPv6 device communicates with an IPv6 server the packets are routed directly to the server through the IPv6 network without any NAT-based stateful middleboxes (Steps 1 and 2). However, when an IPv6 device communicates with an IPv4 server, the IPv6 packets generated by the device are routed to a stateful NAT 64 middlebox (Steps 3 and 4). The NAT 64 middlebox translates IPv6 packets to IPv4 address family and forwards the translated packets to the IPv4 server (Step 5). The NAT 64 middlebox also converts reply IPv4 packets (from the IPv4 server) to IPv6 packets (forwarded to the mobile device), as shown in Steps 5, 4, and 3.

In summary, there are two ways in which T-Mobile routes packets from IPv6 devices: 1) via IPv6 network with no stateful middleboxes, and 2) via NAT 64 middlebox, where IPv6 network is used between clients and NAT 64 and IPv4 network is used between NAT 64 and IPv4 servers.

Verizon Wireless (Verizon) provides both IPv6 and IPv4 addressing to all of its devices connected to its LTE network [17]. For devices with no IPv6 support or not connected by LTE, Verizon provides only IPv4 addressing – resulting in the devices using Verizon’s IPv4 network. Thus, in the Verizon network, the choice of addressing on the device is based on whether the device is connected to the LTE network [54].

In Figure 1(b), we depict a high level infrastructure of Verizon’s IPv6 network. Similarly to T-Mobile, when an IPv6 device communicates with an IPv6 server, the packets are routed to the server through Verizon’s IPv6 network without any stateful NAT middleboxes (Steps 1 and 2). However, when an IPv6 device communicates with an IPv4 server, the device uses the Gateway-Initiated Dual-Stack Lite (DS Lite), a piece of software installed on-device, to encapsulate IPv4 packets inside IPv6 headers (Step 3) [22, 35]. The encapsulated packets are then forwarded to a middlebox that decapsulates the packet contents (Steps 4 and 5). Finally, the IPv4 packets are forwarded to the IPv4 server (Step 6). One of the benefits of using a gateway initiated DS Lite is that encapsulating IPv4 packets in an IPv6 header allows Verizon to use IPv6 addressing for routing packets within the cellular network.

In summary, there are two different ways in which Verizon routes packets from IPv6 devices: 1) via IPv6 network with no stateful middleboxes, and 2) via IPv4-in-IPv6 tunnels using DS Lite software on the phone.

AT&T Mobility (AT&T) and Sprint provide both IPv6 and IPv4 addressing to only some of their IPv6-capable devices [3, 19]. For other IPv6-capable and IPv4-only devices, both these networks provide only IPv4 addressing. Therefore, the choice of addressing

for devices in these networks is neither dependent on 464XLAT nor on the cellular technology and is rather likely to be decided by the carrier’s respective network configurations.

In Figure 1(c), we depict a high level infrastructure of the IPv6 network deployed by AT&T and Sprint. We show that when an IPv6 device communicates with an IPv6 server, the packets are routed directly to the server, through the IPv6 network without any stateful middleboxes (Steps 1 and 2) [3, 19]. However, when an IPv6 device communicates with an IPv4 server, unlike in T-Mobile and Verizon networks, the packets route through the IPv4 network, which consists of several stateful NAT middleboxes, such as NAT 44 and NAT 444, to convert private IPv4 addresses to public IPv4 addresses (Steps 3, 4, and 5).

In summary, there are two different ways in which both AT&T and Sprint route packets from IPv6 devices: 1) via IPv6 network with no stateful middleboxes, and 2) via IPv4 network with several stateful NAT middleboxes.

3. DATA COLLECTION METHODOLOGY

Recent studies on measuring CDN adoption rate among websites show that out of the most popular 1 K, 10 K, and 100 K websites listed on Alexa [14], 77%, 35%, and 19% are hosted on different CDN infrastructures, respectively [4]. Further, a study on understanding the CDN market share indicates that Akamai CDN infrastructure leads in delivering content for majority of the popular websites, including several e-commerce, media, government, news, and social media websites [6][16]. Specifically, out of the most popular 1,000, 10,000 websites listed on Alexa, as well as, top 500 websites listed on Fortune [5], Akamai delivers content for over 23%, 16.4%, and 32% websites, respectively [9]. Based on these results, we believe that our dataset on cellular network performance collected by globally distributed CDN servers of Akamai is representative of mobile Web performance in general.

We now shift our focus to organize our measurement data collection to accurately represent performance of native IPv6, legacy IPv4, NAT 64, and DS-Lite sessions. First, we provide an overview of how we collect performance data from client devices and CDN servers. Next, we discuss techniques used to filter data generated by a number of independent sources, such as client’s browser caching of content, Web proxies in the cellular network, and mobile device’s operating system. Finally, we describe how we sanitize our measurement data to only contain RTTs, DNS lookup times, and webpage load times for pages loaded over end-to-end (E2E) IPv6 and IPv4 sessions between clients and CDN servers, as well as pages loaded via NAT 64 middleboxes and DS-Lite.

Experimental Setup: To compare the Web performance perceived by end-users on IPv6 and IPv4 networks, we use Akamai’s RUM system [13], as depicted in Figure 2. Akamai’s Web servers inject JavaScript (RUM Javascript) into a small fraction of user requests for some of the customer-websites hosted on Akamai infrastruc-




Figure 2: Akamai’s RUM interactions with client’s browser.

ture (Steps 1 and 2). The injected JavaScript uses the browser-exposed Navigation Timing API to capture the time to resolve domain names, time to establish TCP connections, and webpage load time, among several other metrics [11]. The JavaScript then sends the collected timing data in the form of a `RUMBeacon` to a dual-stacked RUM server after the page load completes (Step 3). The RUM server then sends the data in the `RUMBeacon` to the RUM database (Step 4). Finally, in Step 5, the Web server augments the data into the database with the TCP latency estimated by the Web server that served the webpage, the publicly routable IP addresses of the CDN server and the client, indicator of whether the webpage was available via IPv4-only, or dual-stacked GET requests, an indicator of whether the webpage was requested over IPv4 or IPv6, an indicator of whether the `RUMBeacon` was submitted over IPv4 or IPv6, and the cellular ISP name to which the client’s IP address belongs as determined by Akamai’s *EdgeScope* [1].

Note: Earlier CDN deployments only served some of the static webpage content. However, as the need for responsive Web performance increased, CPs moved their base pages and other page resources onto replica servers as well. DNS CNAMEing allows base page domain to resolve to a domain page hosted by a CDN [47].

Using Akamai’s RUM, in Figure 3 we show the round trip TCP latency estimated by Akamai servers when connecting to TCP terminating proxies deployed by cellular ISPs using measurement techniques developed by Goel *et al.* [37]. The height of each bar graph represents the median TCP latency, and the extreme ends of error bars represent the 25th and 75th percentile of the latency respectively. We show that the median latency between Akamai CDN servers and cellular gateways of T-Mobile and AT&T is only 2 ms and 4 ms, respectively. For Verizon and Sprint, the latency is less than 10ms. We argue that such a low latency is possible only when Akamai’s CDN servers are in extreme proximity with the cellular gateways, as opposed to significantly higher latency estimated by Amazon EC2 and PlanetLab datacenters in the US [28]. Therefore, our view of the IPv6 ecosystem using Akamai’s infrastructure allows us to isolate the performance differences if IPv4 and IPv6 within the cellular network (without introducing the confounding factor of the public Internet).

Data Sanitization: We filter our dataset to include performance numbers that pertain only to webpages loaded on Google Chrome browser on Android devices. We choose to consider measurement data for Android devices only as the IPv6 behavior in iOS devices used Happy Eyeballs to select between IPv4 and IPv6 addresses in ways that biased the measurements [45]. This behavior was subsequently improved in late 2015 with the release of iOS 9.

Recent study on detecting Performance Enhancing Proxies (PEPs) in cellular networks has revealed that cellular networks in the US do not use PEPs for HTTPS traffic [37]. Therefore, to




Figure 3: RTT from CDNs to cellular TCP Split proxies.

remove any influence of PEPs (in terms of Web content caching and TCP split connections) in our dataset, we consider latency for only HTTPS sessions. Latency for HTTPS sessions allows us to accurately estimate the latency between CDN servers and client devices and ensure that the estimated latency is **not** between servers and PEPs in cellular networks. In this work we focus on Web performance over native IPv6 and legacy IPv4 networks, and so eliminate factors, such as presence of PEPs that may confound our measurement data. Analysis of how PEPs in cellular networks impact Web performance is potential future work of this study.

Finally, to remove any influence of content caching in Web browsers on the measurement data, we consider data for only newly created TCP connections. During our data collection we identified that several page loads experienced TCP connection time of less than 20ms. Such a low TCP connection time is unjustified in a cellular network [12]. Instead, these short connection times are more likely due to the times it takes a browser to fetch a cached TCP connection. Thus, we extract data only for connections whose setup time is more than 20ms. We employ a similar technique to extract DNS resolution times that were resolved at the time of webpage load, thus eliminating the influence of any cached DNS entries in the browser.

Our sanitized dataset consists of measurement data for RTT and DNS lookup time for several million sessions between clients and Akamai CDN servers, and webpage load time from several hundred page loads.

Data Analysis: To record the latency over an IPv6 connection we use RTT to clients estimated by CDN servers for webpages requested over IPv6 network. To get the latency over connections via NAT 64 (in T-Mobile) or DS-Lite (in Verizon), or by IPv6 clients using IPv4 network (in AT&T and Sprint), we use latency estimated by CDN servers for connections, where webpages were requested over IPv4 network and the `RUMBeacon` was submitted over IPv6 – indicating that the webpage loaded on an IPv6 client. To record the latency over IPv4 connections, we use latency estimated by CDN servers for webpages requested over IPv4 and where `RUM beacon` was also submitted over IPv4 – indicating that the webpage is loaded on an IPv4 client. We apply similar techniques to extract webpage load time.

To record DNS lookup times that pertain to domain names resolved by IPv6 clients, we extract data points for clients on which either websites were loaded over IPv6, NAT 64, or DS-Lite connectivity, or the RUM beacon was submitted to the RUM server over an IPv6. Similarly, to get DNS lookup times that pertain to domain names resolved by IPv4 clients, we extract data points for clients on which either websites were loaded over IPv4 or the `RUMBeacon` was submitted to the RUM server over IPv4 network.




Figure 4: RTT distribution for different cellular carriers in the US.




Figure 5: 24-hour RTT distribution for different cellular carriers in the US.

4. ROUND TRIP LATENCY OVER IPv6 AND IPv4 CELLULAR NETWORKS

The Round trip time (RTT) between clients and servers plays an important role in influencing Web performance [38]. In this section, we investigate whether serving mobile content over IPv6 results in lower latency between mobile clients and content servers. We also investigate whether the performance of IPv6 network as well as the performance gap between IPv6 and IPv4 remains same at peak and non-peak traffic hours of a day.

In Figure 4, we show the overall distribution of RTT between clients and CDN servers over IPv6 and IPv4 networks of different cellular carriers, collected in five months of 2015. The solid CDF lines in these graphs show the RTT when IPv6 clients connect to IPv6 servers, over the IPv6 network. The dashed CDF lines show the RTT when IPv6 clients connect to IPv4 servers via NAT 64 middleboxes (in T-Mobile), via IPv4-in-IPv6 tunnel (in Verizon), or via the IPv4 network (in AT&T and Sprint). The dotted CDF lines show the RTT when IPv4 clients connect to IPv4 servers, over the IPv4 network. Additionally, in Figure 5, we show the RTT distribution for 24-hour period, averaged over two months (June and July in 2015).

In the case of T-Mobile in Figure 4(a), we observe that the RTT for sessions over IPv6 network is lower than the RTT over the IPv4 network. For example, for median and 80% of sessions, the RTT over IPv6 network is about 49% and 64% faster than RTT over IPv4 network, respectively. Even for sessions via NAT 64 middlebox, the IPv6 RTT is lower than RTT over IPv4 network. For example, for connections that go through NAT 64 middleboxes, the latencies for median and 80% of sessions are about 18% and 27% faster than RTT over the IPv4 network, respectively.

Further, to eliminate any effects of provisioning differences between IPv4 and IPv6 networks, we compare performance over IPv6 and IPv4 networks at peak and non-peak traffic hours in Figure 5(a). We observe that the RTT over T-Mobile’s IPv6 network outperforms latency over its IPv4 network at all times. In fact, NAT 64 sessions also experience lower latency than IPv4 sessions

at all times of the day. Although we see that the performance gaps between all three distributions is not the same at all times, native IPv6 connectivity always provides least possible latency among others. Based on our observations, we speculate that such differences in round trip latency in the T-Mobile network arise from the elimination of overhead of NAT middleboxes deployed in the IPv4 network to perform address translation of client sessions. With no middleboxes in the case of end-to-end IPv6 connectivity or one middlebox in the case of NAT 64 sessions, users experience lower latency.

In the case of Verizon in Figure 4(b), we observe that RTT for IPv6 sessions is similar to RTT for DS Lite sessions. We expect the two RTTs to be similar since in case of DS Lite (IPv4-in-IPv6 tunneling) all packets are sent from the device over the IPv6 network, resulting in similar RTT as end-to-end IPv6 sessions. The RTT over IPv4 network however is influenced by Carrier Grade NATs and Large Scale NATs and thus experience significantly higher latency than end-to-end IPv6 or IPv4-in-IPv6 tunneled sessions. For example, for median and 80% of sessions on Verizon, the RTT over IPv6 network is about 29% and 44% faster than RTT over its IPv4 network, respectively. Additionally, when comparing performance over 24-hour periods in Figure 5(b), we observe that IPv6 latency inside Verizon’s network outperforms latency over its IPv4 network in both peak and non-peak traffic hours. Based on our observations, we speculate that the reduced RTT over Verizon’s IPv6 network is due to two major factors: 1) no stateful middleboxes in its IPv6 network, and 2) the use of IPv6 connectivity only over LTE network, as opposed to use of IPv4 connectivity over 3G network.¹

Finally, in the case of AT&T and Sprint in Figures 4(c) and 4(d) respectively, we observe that RTT over IPv6 network is lower than RTT over their respective IPv4 networks, especially in the long tail. For example, for median and 80% of sessions on AT&T network

¹We could not disambiguate our measurement data specific to sessions over LTE and 3G networks as Akamai’s RUM uses JavaScript to collect client-side performance and at the time of our measurement (Jan - Aug 2015), Chrome browser did not capture the cellular technology over which the webpage is loaded [30].




Figure 6: DNS Lookup Time distribution for clients in different cellular carriers in the US.

in Figure 4(c), the RTT for IPv6 network is 17% and 24% faster than RTT over its IPv4 network respectively. Further, the RTT for sessions established by IPv6 clients with IPv4 servers is similar to RTT for sessions established by IPv4 clients with IPv4 servers. We expect the two RTTs to be similar because both AT&T and Sprint are dual-stacked networks and therefore both IPv6 and IPv4 clients must connect to IPv4 servers over their respective legacy IPv4 networks with NAT 44 and NAT 444 middleboxes – resulting in similar latency. Additionally, when comparing AT&T and Sprint’s IPv6 network performance with their respective IPv4 networks over 24-hour periods in Figures 5(c) and 5(d), we observe that IPv6 sessions on both AT&T and Sprint experience lower RTT than latency experienced by IPv4 sessions. Therefore, based on our observations we speculate that similarly to T-Mobile, IPv6 networks of both AT&T and Sprint outperform their respective IPv4 networks because there are no stateful middleboxes in their IPv6 networks.

Discussion: Although, we observe that providing mobile content over IPv6 offers reduced latency for end-users, we identified a cellular network outside the US where, during our study, RTT over IPv6 network was higher than RTT over its IPv4 network by almost 200ms. To investigate, we ran traceroutes from CDN servers to several IPv6 client IP addresses in that network and identified that IPv6 packets were being routed through another country, resulting in higher RTT over its IPv6 network. To investigate whether similar routing was applicable to IPv4 packets in that network at the time of our measurement, we ran traceroutes from the same CDN servers to IPv4 client IP addresses in that network and found that packets were **not** being routed via another country. While it is possible that IPv6 could be slower than IPv4 in some networks, it could be related to how IPv6 packets are forwarded on the Internet. Therefore, proximity of Akamai servers to cellular gateways eliminates the effects of misconfigured routing in the public Internet.

5. DNS LOOKUP TIME FOR IPv6 AND IPv4 CLIENTS

In addition to RTT, DNS lookup time is another important factor which influences the Web performance in cellular networks [64]. In this section, we measure the DNS lookup time for both IPv6 and IPv4 clients resolving dual-stacked domain names (domains which can be resolved to both IPv4 and IPv6 addresses). In Figure 6, we show the distribution of DNS lookup times when IPv6 and IPv4 clients resolve dual-stacked domain names. The dotted CDF lines represent the lookup time when domains are resolved for IPv6 clients, whereas, the solid CDF lines represent lookup time when domains are resolved for IPv4 clients.

In general, and contrary to the trends in the previous section, we see that the DNS lookup takes longer for IPv6 clients than IPv4 clients in T-Mobile, AT&T, and Sprint networks. However, for Verizon’s IPv6 and IPv4 clients, the DNS lookup times are

similar. DNS lookup times for IPv6 clients are influenced by their technique of DNS resolution. Client devices are unaware of whether content from a domain is available over IPv4 network or over IPv6 network. Therefore, clients must send both AAAA (IPv6) and A (IPv4) DNS queries to their local resolvers to resolve domain names. If an IPv6 address for the requested domain is available, Android clients prefer to connect with the IPv6 address, instead of the IPv4 address of the server [67]. Although the two DNS requests could be sent in parallel to reduce the time to perform the DNS lookups, we observe that regardless of the type of domain (IPv4-only or dual-stack), IPv6 Android clients always issue both AAAA and A DNS queries serially. Further, before returning the DNS response to the application the IPv6 clients wait until responses for both queries arrive, or the resolutions times out. Therefore, the DNS resolution on IPv6 Android clients require two round trips between clients and DNS server, whereas IPv4 Android clients wait for only one round trip for resolving the domain via type A query.

In the case of T-Mobile in Figure 6(a), we observe that the median DNS lookup time for IPv6 clients is 25.7% slower than IPv4 clients, because IPv6 clients wait for both type AAAA and A queries to finish, whereas IPv4 clients wait for only type A queries. However, for about 20% of DNS requests, IPv6 clients experience faster resolution time than IPv4 clients. Since T-Mobile’s IPv6 clients can only transmit IPv6 packets into its network, DNS lookups for IPv6 clients take place over T-Mobile’s IPv6 network. Our earlier observation from Figure 4(a) shows that RTT over T-Mobile’s IPv6 network is lower than IPv4 network, which likely helps about 20% of IPv6 DNS lookups to complete faster in spite of the additional RTT. Therefore, DNS lookups for some IPv6 clients outperform the lookup time for IPv4 clients, even though DNS lookup process for IPv6 clients waits for an additional DNS query to finish.

In the case of Verizon in Figure 6(b), we observe that about 60% of the DNS queries by IPv6 clients take same time as queries by IPv4 clients. Similarly to T-Mobile, IPv6 clients in Verizon network transmit IPv6 packets into the network, which results in DNS lookups over the IPv6 network. From Figure 4(b), we know that RTT over Verizon’s IPv6 network is significantly lower than its IPv4 network, therefore the DNS lookup time for IPv6 clients is similar to lookup time for IPv4 clients.

In the case of AT&T and Sprint in Figures 6(c) and 6(d) respectively, we observe that the median DNS lookup times for IPv6 clients are about 38% slower than lookup times for IPv4 client. Since both IPv4 and IPv6 clients in these networks use their respective IPv4 networks to send DNS queries to local resolvers and that RTTs for IPv4 packets sent by IPv6 and IPv4 clients are similar (from Figures 4(c) and 4(d)), IPv6 clients wait for responses for two DNS queries in serial, as opposed to IPv4 clients that wait for only one DNS lookup.

Discussion: We observe that the DNS lookup process takes longer for devices with IPv6 capabilities than devices with IPv4-only




Figure 7: Dual-Stack webpage PLT distribution in different cellular carriers in the US.

capabilities in AT&T, Sprint, and partly in T-Mobile networks. Following our observation, we argue that if the DNS lookup process on IPv6 capable Android devices could be modified to send AAAA and A DNS queries in parallel, lookup time could be significantly reduced. Additionally, if the client browser could indicate to the mobile OS that they do not need to wait to get the A lookup back if they get an AAAA answer in DNS response, the DNS lookup time can be further reduced. Further, for devices in the T-Mobile network (and other IPv6-only networks such as Orange Poland and SK Telecom, Telenor [50]), if the OS installed on the device could identify whether the connected network is IPv6-only, the DNS process could simply resolve a AAAA DNS query alone, instead of the current implementation where both type A and AAAA DNS requests are resolved. And finally, the existence of A or AAAA only answers could be cached per-name for some time, not withstanding the caching of the actual answers, which would enable the device to make a smarter query the following time that an IP address is needed. Such a network-specific DNS resolution process will allow clients connected to IPv6-only networks to speed up their DNS lookups. We made these recommendations to Google’s Android team to improve DNS resolution process.

6. PAGE LOAD TIME OVER IPv6 AND IPv4 NETWORKS

Interactive webpages often require multiple DNS lookups and many round trips between clients and servers to download Web objects onto the client’s browser. To investigate the overall impact on Web performance of IPv6’s low RTTs and high DNS lookup times, we compare the webpage load time (PLT) over IPv6 and IPv4 networks, using the browser’s Navigation Timing API. For this part of the study, we analyze measurement data for two types of webpages: 1) those available over both IPv6 and IPv4 networks (dual-stacked), and 2) those available over IPv4 network only. Our comparison of PLTs only includes page load requests, for which DNS lookups were performed by the client browser and the pages were loaded over newly established TCP connections. Our immediate goal is to eliminate any PLT data that includes DNS lookup from Web browser’s cache and reuse of an existing TCP connection.

To mitigate the influence of mobile hardware on PLT [61], we consider PLTs from only one specific device model for each cellular network. For each carrier we selected the device model for which Akamai’s RUM had the highest number of requests for downloading webpage over both IPv4 and IPv6 networks of the corresponding carrier. We select different devices in the investigated networks as device popularity depends on how an ISP promotes its services to their customers. In Table 1, we list the device models selected for each network from our dataset to allow potential future work to match the model, or consider our results in light of any significant hardware issues with a given device.

Table 1: Selected mobile device models with highest number of webpage load requests in different cellular networks.

Network	Device Model Name	Model ID
T-Mobile	Samsung Galaxy S5	SM-G900
Verizon	Samsung Galaxy S5	SM-G900V
AT&T	Samsung Galaxy S6 Edge	SM-G925
Sprint	Samsung Galaxy S6 Edge	SM-G925

Table 2: Details of IPv4-only webpages loaded over IPv4 networks of different cellular carriers.

Network	Webpage	#Object	Size
T-Mobile	Clothing	444	0.70 MB
Verizon	Internet Retailer	165	1.30 MB
AT&T	Home Improvement	63	0.67 MB
Sprint	Internet Retailer	165	1.30 MB

Further, to accurately characterize the performance of webpage loads over IPv6 and IPv4 networks, we consider PLTs for only one Web URL loaded on one specific device model for each network. Similarly to our choice of device model, we selected the URL for which Akamai’s RUM received the highest number of requests from the selected device model. The selected page has 94 embedded Web objects in total and 0.83 MB in size. To prevent the influence of any in-network HTTP caches on PLTs, the page was loaded over Secure HTTP (HTTPS). Further, we identified different IPv4-only webpages for different carriers that had the most number of page load requests from the selected device. We list the details of IPv4-only URLs selected in each network for performance comparison in Table 2 respectively.

IPv6 webpage load time: In Figure 7, we show the distribution of page load time of one dual-stacked webpage loaded by IPv6 clients over IPv6 networks and loaded by IPv4 clients over IPv4 network. The solid CDF lines show PLTs when the page was loaded over IPv6 network. The dotted CDF lines show PLTs when the page was loaded over IPv4 network. In general, we observe that for all four US carriers, the PLTs of pages loaded by IPv6 clients over IPv6 networks are lower than PLTs of the same pages loaded by IPv4 clients over the respective carrier’s IPv4 networks. Further, despite DNS lookup times being higher for IPv6 clients, we observe that PLTs are lower for IPv6 clients loading pages over IPv6 network. We argue that DNS lookup times for IPv6 clients influence the overall page load time by just one extra round trip and that the actual benefits of faster IPv6 network are observed when multiple Web objects are loaded over the IPv6 network in several round trips.

In case of T-Mobile in Figure 7(a), we observe that for median and 80% of page loads by IPv6 clients, the PLTs over IPv6 network are 9% and 14% faster than PLTs over T-Mobile’s IPv4 network. In Figure 7(b), we observe similar reductions in PLTs for pages




Figure 8: IPv4-only webpage PLT distribution in different cellular carriers in the US.

loaded by Verizon’s IPv6 clients over Verizon’s IPv6 network. Specifically, we show that the median and 80% of the PLTs by IPv6 clients over Verizon’s IPv6 network are 48% and 64% faster than PLTs over its IPv4 network, because of the significant differences in RTTs between Verizon’s IPv6 and IPv4 networks as shown in Figure 4(b). For AT&T and Sprint as well, we observe that PLTs are lower over the IPv6 network.

IPv4 webpage load time: In Figure 8, we show distribution of PLTs of an IPv4 webpage, loaded by IPv6 and IPv4 clients. The dashed CDF lines show PLTs when IPv6 clients load an IPv4 webpage via NAT64 (in T-Mobile), via IPv4-in-IPv6 tunnel (in Verizon), or via IPv4 network (in AT&T and Sprint). The dotted CDF lines show PLTs when IPv4 clients load an IPv4 webpage over the IPv4 network. In general, we observe that T-Mobile and Verizon IPv6 clients experience reduced PLTs for IPv4 webpages, due to reduced RTT when using NAT 64 in T-Mobile (as shown in Figure 4(a)) and the use of LTE and IPv6 network in Verizon (as shown in Figure 4(b)). For example, in Figure 8(b), we show that the median and 80% of the IPv4 page loads in Verizon, are about 49% and 67% faster than page loads over IPv4 network. In the case of IPv6 clients in AT&T and Sprint network in Figures 8(c) and 8(d) respectively, we observe that IPv4 webpage PLTs are similar (in AT&T), or occasionally slower (in Sprint) than PLTs experienced by IPv4 clients. We expect the two PLTs to be similar since both IPv6 and IPv4 clients use the IPv4 network to load IPv4 websites. However, in some cases we expect the PLTs by IPv6 clients to be slower than PLTs by IPv4 clients, since such clients wait for an additional DNS query for each domain in the webpage.

Although we only show performance gains with IPv6 for only one URL loaded on one type of device model, we also looked at other URLs loaded from other device models as well. We identified that IPv6 connectivity also improves the page load time of other URLs, though due to space constraints we do not show PLT distributions for other URLs and devices.

Discussion: Based on our findings on webpage load times, we show that for IPv6 clients in T-Mobile and Verizon networks, both IPv4-only and dual-stacked websites will load faster than IPv4 clients loading the same websites over the carrier’s respective IPv4 network. Although we observe that DNS lookups are slower for IPv6 clients in T-Mobile, AT&T, and Sprint networks, we observe their impact on the overall PLT over these networks to be minimal, since the RTT over IPv6 network is lower than RTT over IPv4 networks, which accounts for most of the round trips when loading a webpage [38].

For dual-stacked networks, such as AT&T and Sprint, IPv6 clients in such networks experience faster page loads for only dual-stacked websites. Further, websites available over IPv4 only may occasionally experience poor performance on IPv6 clients,




Figure 9: Sequence of how IPv4-only domains are resolved for IPv6-clients in IPv6-only networks.

likely due to slower DNS lookups that lower IPv6 latency cannot offset. Finally, we argue that since users are likely to visit websites from different networks, providing mobile content over IPv6 will not hurt the Web performance and in fact in some cases the Web performance will improve if pages are served over IPv6.

7. DNS LOOKUPS IN T-MOBILE’S IPv6-ONLY NETWORK

As cellular carriers upgrade their network infrastructure to IPv6 to mitigate IPv4 address scarcity, we find that the DNS protocol starts to introduce avoidable performance overhead. We observe an extra round trip in cellular ISPs, such as T-Mobile, Orange Poland, SK Telecom, and others [50], which use IPv6-only addressing in their network [24]. In effect, ISPs that fully embrace IPv6 technology are penalized with slower domain name resolution.

We illustrate the problem scenario in Figure 9, which shows a sequence of DNS messages exchanged between a user device, T-Mobile’s DNS 64 server, and a DNS Authority. An IPv6-only environment requires a mobile client to send a AAAA DNS query to the cellular DNS server to resolve an IPv4-only domain (Step 1). The DNS request then travels to the DNS Authority (Step 2). The DNS Authority replies with NOANSWER flag in the DNS response (Step 3), because an IPv6 address is not available for the IPv4 domain in question. Instead of returning the DNS response with NOANSWER back to the client, T-Mobile’s DNS server sends a subsequent A DNS query for the same domain (Step 4), to which the DNS Authority replies with an IPv4 address (Step 5). After receiving the IPv4 address, T-Mobile’s DNS server synthesizes an IPv6 address corresponding to the IPv4 address and sends the synthesized address to the client in response to the client’s DNS request (Step 6). At this point, the mobile client is unaware whether the returned IPv6 address is a synthesized address, or a real address returned by the DNS Authority. Since T-Mobile employs NAT 64 middleboxes to translate synthesized IPv6 addresses back to real IPv4 addresses (Figure 1(a)), the synthesized address that the clients receive does not alter their end-to-end connectivity [20, 24].

Thus, if IPv6 addresses are available for a domain, the DNS

lookup will finish in Step 3. However, when a domain has only IPv4 addresses available, the DNS lookup requires an extra round trip between the cellular DNS and the DNS Authority (Steps 4 and 5) [24]. The latency of this extra round trip could significantly influence the PLT when cellular DNS servers are not in proximity to DNS Authorities and the webpage requires multiple DNS lookups.

To address this overhead in DNS lookup process, we design ONETRIP – a technique for DNS Authorities to eliminate the extra round trip from DNS lookup process in IPv6-only mobile networks. We show that for mobile clients in IPv6-only networks, a DNS Authority can proactively synthesize IPv6 addresses from IPv4 addresses, for all domains it holds mappings between domains and IPv4 addresses. The DNS Authorities could then reply with a synthesized IPv6 address, instead of a NOANSWER, to any IPv4-only domain name lookup from IPv6-only networks.

7.1 ONETRIP’s Approach

To synthesize an IPv6 address from an IPv4 address, similarly to how T-Mobile’s DNS 64 servers synthesize [24], ONETRIP requires two types of datasets that identify the /64 prefix used to synthesize an IPv6 address by cellular DNS servers. The first are mappings between client and cellular DNS server IP addresses. The second are mappings between cellular DNS server IP addresses and the /64 prefix used by the NAT 64 middlebox to which their clients connect. The two datasets collectively allow ONETRIP to map a client IP address to /64 prefix, used by the NAT 64 middlebox associated with each DNS (DNS 64) server. ONETRIP makes these mappings available to DNS Authorities so that when a cellular DNS server sends a request to the Authority, the Authority already knows (based on cellular DNS IP in the DNS query) the /64 prefix of the NAT 64 middlebox to which the client is connected. The Authority then uses the corresponding /64 prefix to synthesize the IPv6 address from the IPv4 address associated with the IPv4-only domain name in question.

Next, we describe ONETRIP’s approach to collect the above mappings using Akamai’s global infrastructure of content delivery.

Collecting DNS IP – NAT 64 Prefix mappings: Using the crowd-sourced dataset collected by Netalyzr in over 11 months in 2013-2014 [44], we extracted mappings between cellular DNS server IP addresses and the /64 prefixes used by NAT 64 middleboxes deployed in T-Mobile’s network [62]. The Netalyzr mobile application sends a AAAA DNS request to a cellular DNS server for resolving an IPv4-only domain name whose DNS Authority replies with the IP address of the DNS server that made the request. From the DNS response received by the client, we extract the IP address of the cellular DNS server by converting the last (least significant) 32 bits of the synthesized IPv6 address format from HEX to DEC. We also extract the /64 prefix from the first (most significant) 64 bits of the synthesized IPv6 address.

Collecting Client IP – DNS IP mappings: For the purpose of this work, we modified our measurement system to make clients resolve a unique hostname that allows us to map the client IP addresses to the cellular DNS IP addresses in a manner similar to the above.

Finally, from the above two mappings we generate the associations between client IP addresses and /64 prefixes that enable DNS Authorities to identify the NAT 64 middlebox, to which the clients connect. These mappings are also useful when DNS requests use the EDNS0 extension (containing the IP address of the mobile client) for CDNs to perform server selection based on the client IP instead of the DNS server IP [25]. Specifically, when ISP resolvers send IPv6 client addresses in the DNS request, DNS Authorities search for the NAT 64 address, to which the

client is currently connected, and then perform mapping based on the location of NAT64 address in the network. DNS Authorities may also combine NAT64 location information with the (limited) information about the location of client’s IPv6 address.

Additionally, the mappings between client IP and NAT64 prefixes enable Web/proxy servers to reduce latency for HTTP transactions, as opposed to using synthetic IPv6 addresses only for the purposes of DNS lookups. Specifically, servers that currently embed static IPv4 addresses in HTTP headers or payload [24], can now embed synthesized IPv6 addresses to sidestep 464XLAT software on resource limited mobile devices. This procedure reduces latency perceived by end users because the 464XLAT software on the client first converts the IPv4 address (returned in HTTP response) into an IPv6 address. Next, IPv6 packets are forwarded to the NAT64 middlebox that converts them back to IPv4 packets for the IPv4 host. By allowing servers to embed synthetic IPv6 addresses in HTTP headers, or payloads, latency introduced by 464XLAT software on the client is eliminated.

Discussion: Based on the mappings we collect between clients and DNS server IP addresses, we find that each DNS IP is used by thousands of mobile clients to resolve domain names. Therefore, ONETRIP’s implementation on DNS Authorities can improve DNS lookup time for thousands of cellular clients for every IPv4-only domain name in resolution. We also observe from the data collected by Netalyzr that the DNS server address and /64 prefix mappings were stable and that there were no changes over a period of 11 months in 2013-2014. We argue that such demonstrated historical mapping stability supports ONETRIP’s approach to reliably synthesize IPv6 addresses on the DNS Authorities, however there is a serious risk that doing this mapping outside of operator control could impair their ability to operate their network. The ONETRIP mechanism is however a good way to identify the performance improvements available by improving the DNS protocol.

7.2 Speeding DNS Lookups with ONETRIP

We discussed ONETRIP’s approach for DNS Authorities to eliminate the extra round trip during DNS resolutions. However, it still remains unclear whether T-Mobile’s cellular DNS would honor, or reject, the IPv6 addresses synthesized by an Authority outside of T-Mobile’s network. To answer this question, we configured IPv6 clients in T-Mobile’s network to send AAAA DNS requests to their respective cellular DNS servers. We also setup a DNS Authority, outside of T-Mobile’s network, maintaining DNS records for a domain name with only IPv4 address. Using the two mappings collected by ONETRIP, we configured the DNS Authority to reply with a synthesized IPv6 address to the DNS request for the IPv4-only domain name, instead of a NOANSWER.

Our experiments show that T-Mobile’s DNS servers do not discard any DNS replies with IPv6 addresses synthesized by a DNS Authority outside of the T-Mobile’s network. Further, the mobile clients successfully connect to the IPv4 servers, using the synthesized IPv6 address.

Finally, we perform several experiments to understand the performance improvements that ONETRIP brings to DNS content delivery when DNS Authorities are hosted at different proximity from cellular DNS servers. Specifically, our evaluation of ONETRIP is based on several possible latency values that exist between cellular DNS servers and DNS Authorities. We expect geographically-distributed DNS Authorities to have lower latency to cellular DNS servers than a centralized DNS Authority. For example, a previous study has shown that Internet backbone round-trip time is around 82 ms between the East and West coast in the US [60]. The same study also provides latencies between




Figure 10: Reduction in DNS Lookup time when using ONETRIP on DNS Authority.

different regions in the US and the EU. We use these latency values as representative of latency between cellular DNS servers and DNS Authorities hosted at these different locations.

A recent study by Varvello *et al.* shows that the 25th, 50th, and 75th percentile of websites (out of the top 9000 Alexa Websites that support HTTP/2) consists of over 12, 18, and 30 unique domain names, respectively [63]. We therefore argue that depending on the number of unique domain names in a given webpage and the latency between cellular DNS servers and DNS Authorities, the webpage load time could be significantly improved through ONETRIP’s removal of one round trip per resolution.

In Figure 10, using simulations we show the absolute reduction in DNS lookup time when using ONETRIP on DNS Authorities. To measure the effectiveness of ONETRIP, we perform in-lab simulations to emulate the behavior of T-Mobile’s DNS 64 servers, where we perform DNS lookups from an emulated cellular DNS issuing name resolutions to a DNS Authority. On the x-axis, we represent five different latency values between cellular DNS and DNS Authority. On the y-axis, we show the latency reduction in the DNS lookup time achieved when ONETRIP is used on the DNS Authority.

In general, we see that as the network latency between cellular DNS and Authority increases, the gains increase as well because ONETRIP shaves off the extra round in the DNS lookup process. The overall gain further increases when the number of unique domain names associated with different Web objects embedded in a webpage increases, because ONETRIP reduces the latency for resolving each of these unique domain names. For example, when the latency between cellular DNS server and DNS Authority is about 82 ms and the number of unique domain names (that needs to be resolved from that Authority) on a webpage is 18, ONETRIP reduces the overall DNS lookup time by about 1.4 seconds ($18 * 82$ ms). The latency reduction shown in Figure 10 is the maximum saving ONETRIP offers when all domain name lookups are critical for webpage rendering; the minimum benefit is just one RTT saved on the lookup of the base page. In general, as ONETRIP reduces the number of round trips between the cellular DNS and DNS Authorities from two to one, ONETRIP offers a reduction of about 50% in the DNS lookup time between cellular DNS servers and DNS Authorities in IPv6-only network environments for IPv4-only content.

7.3 Discussion on ONETRIP’s Approach

ONETRIP relies on identifying the /64 prefix associated with the NAT 64 middlebox, to which a client connects. Similarly, previous studies have developed several techniques to detect the presence of NAT 64 middleboxes in the network [32], including techniques to resolve an IPv4-only domain name with a AAAA DNS request and checking if an IPv4 answer is available [59]. Other studies have developed techniques to identify IPv4 addresses from syn-

thesized IPv6 packets [21]. Some other techniques showcase new extensions to DNS protocol and a new Resource Record that DNS servers can adopt to let clients know what the original IPv4 address is for the domain name in question [23, 43, 66]. Previous studies also investigated application layer protocols such as STUN to detect the NAT 64 prefix [57]. A study by Ding *et al.* offers a detailed comparison of different techniques used to identifying NAT 64 prefix in IPv6 networks [33]. The same study also showcase how the EDNS0 extension could be used by DNS resolvers to send NAT 64 prefixes in the DNS request, together with the the technique to calculate the prefix by resolving an IPv4-only domain name.

ONETRIP, in contrast to the previous techniques, is unique in that it enables DNS Authorities to detect and effectively use NAT 64 prefixes to reduce DNS lookup time, without the need of any support from the cellular ISPs, allowing measurement of the performance difference. Our motivation for ONETRIP is to eliminate the extra round trip present in DNS lookups in IPv6-only networks, with the goal of faster mobile Web for the end-users. Similar round trip elimination has also been proposed in QUIC [40] and TCP Fast Open [53]. Although for eliminating the extra round trip, we also recommend the use of the EDNS0 option in the DNS query for the cellular DNS servers to pass along the NAT 64 prefix that they are using, such that in the absence of an AAAA record DNS Authorities could synthesize one from the A record using the contents of the EDNS0 option. However, we argue that support for such an EDNS0 option may not be appealing for some cellular ISPs as it introduces additional operational overheads in ISPs’ functionality. Therefore, we designed ONETRIP for DNS Authorities maintained by Content Providers, Content Delivery Networks, and any independent server operator, which they can implement without any support from cellular ISPs. A strictly better option for Content Providers is to make their content available over IPv6 as this reduces the round-trip by making the AAAA record available. Finally, although ONETRIP does not address performance overhead introduced by Android DNS lookup process when IPv6 clients wait for both AAAA and A replies (as discussed in Section 5), our recommendations to Google’s Android team would address the issue of sequential lookups.

8. RELATED WORK

IPv6 adoption and deployment challenges: Previous studies have measured the adoption rate of IPv6 across different ISPs worldwide and indicate a significant growth in IPv6 traffic over the recent years [7, 27, 41, 50, 58]. While understanding the adoption of IPv6 is important, several communities (including network operators, CDNs, and content providers) have shown interest in discussing challenges faced by mobile ISPs and content providers to adopt IPv6 in their infrastructures in a panel at the @Scale

conference held in 2015 [10]. Several case studies conducted at Akamai and Fortinet provide experiences with IPv6 deployment on a global scale [49] and the current state of the IPv6 in terms of information security [8, 36].

IPv6 performance measurement: A study by Dhamdhare *et al.* investigates the impact of BGP route changes on the performance of IPv6 networks [31]. Plonka *et al.* investigate the flow bit rates of IPv4 and IPv6 traffic at different times of the day [52]. A study by Donley *et al.* investigates the impact of several NAT middleboxes on IPv4 latency and offers a performance comparison between IPv4 and native IPv6 connectivity in wired networks [34]. Other studies investigate throughput, RTT, packet loss, and hop counts provided by IPv6 networks [46, 48, 65, 68, 69].

Addressing IPv6 connectivity issues: At the 2015 @Scale conference, a representative from Verizon Wireless suggested that browsers should not fallback on IPv4 when IPv6 connectivity is slow, because a fallback to IPv4 can mask critical performance issues related to IPv6 connectivity [10]. Collectively, several mobile operators suggested that application developers should support IPv6 connectivity to help kickstart a transition from IPv4 to IPv6. This suggestion complements Apple's recent announcement of IPv6 support in all iOS9 applications [45].

Redirecting clients from broken IPv6 links: Several studies suggest methods to improve Web performance by selectively handing out answers to AAAA queries based on the performance of the current IPv6 connectivity is behind the client's resolver network [26, 39, 42, 49].

Our work, in contrast to these studies, focuses on understanding how different IPv6 deployment strategies in cellular networks influence the mobile Web performance, from the perspective of CDNs deployed in the middle of cellular networks and content providers. Based on our study, we argue that hosting mobile content on IPv6-enabled networks and content servers is another direction that CDNs and content providers could adopt to improve the end-user experience.

9. CONCLUSIONS

Content Providers (CPs) and Content Delivery Networks (CDNs) are not fully aware of the differences in mobile Web performance in cellular IPv6 and IPv4 networks. In this paper, we provide our experience with the changing IPv6 ecosystem in major cellular networks from the perspective Akamai's global infrastructure for content delivery. We perform extensive measurement to understand how different IPv6 technologies deployed by cellular carriers impact mobile Web performance. Our results indicate that cellular IPv6 networks outperform their legacy IPv4 networks. We argue to CPs and CDNs that the transition from IPv4 to IPv6 is another milestone for improving mobile Web performance.

ACKNOWLEDGMENTS

We thank Ruomei Gao, Ajay Kumar Miyapuram, Erik Nygren, Kanika Shah, our shepherd, Matt Welsh, and the anonymous reviewers for providing us constructive feedback. We also thank National Science Foundation (NSF) for supporting this work through grants NSF CNS-1555591 and NSF CNS-1527097.

10. REFERENCES

- [1] Akamai EdgeScape. http://uk.akamai.com/dl/brochures/edgescape_service_description.pdf, Mar. 2002.
- [2] PlanetLab. <https://www.planet-lab.org/>, Jun. 2007.
- [3] Sprint to Participate in World IPv6 Day. <http://newsroom.sprint.com/news-releases/sprint-to-participate-in-world-ipv6-day.htm>, Apr. 2011.
- [4] Dyn Research: CDN Adoption By The Numbers. <http://dyn.com/blog/dyn-research-cdn-adoption-by-the-numbers/>, Jun. 2014.
- [5] List of Fortune 500 companies and their websites. <http://www.zyxware.com/articles/4344/list-of-fortune-500-companies-and-their-websites>, Jul. 2014.
- [6] Akamai Facts and Figures. <https://www.akamai.com/us/en/about/facts-figures.jsp>, Aug. 2015.
- [7] Akamai's State of the Internet. <http://www.wsta.org/wp-content/uploads/2013/11/Q1-2015-SOTI-Security-Report-Low-Res.pdf>, 2015.
- [8] Akamai's State of the Internet/Security. <https://www.stateoftheinternet.com/downloads/pdfs/2015-q1-state-of-the-internet-report.pdf>, 2015.
- [9] Alexa and Fortune 500 CDN Marketshare. <http://blog.cloudharmony.com/2015/03/cdn-marketshare-alexa-fortune-500.html>, Mar. 2015.
- [10] Facebook: IPv6 is here and You're Hurting your Users. https://www.youtube.com/watch?v=_7rcA1bvzVY, Sept. 2015.
- [11] Navigation Timing. <http://w3c.github.io/navigation-timing/>, Aug. 2015.
- [12] NSF Workshop on Achieving Ultra-Low Latencies in Wireless Networks. <http://inlab.lab.asu.edu/nsf/files/WorkshopReport.pdf>, Mar. 2015.
- [13] Real User Monitoring. <https://www.akamai.com/us/en/resources/real-user-monitoring.jsp>, Aug. 2015.
- [14] Alexa: Your Complete Web Analytics Toolkit. <http://www.alexa.com/>, Jan. 2016.
- [15] Amazon EC2 - Virtual Server Hosting. <https://aws.amazon.com/ec2/>, Feb. 2016.
- [16] Akamai. Our Customers. <https://www.akamai.com/us/en/our-customers.jsp>, Jun. 2015.
- [17] APNIC 34. IPv6 at Verizon Wireless. https://www.hpc.mil/images/hpcdocs/ipv6/vzw_apnic_13462152832-2.pdf, Aug. 2012.
- [18] J. Arkko and F. Baker. Guidelines for Using IPv6 Transition Mechanisms during IPv6 Deployment. <https://tools.ietf.org/html/rfc6180>, May 2011.
- [19] AT&T Public Policy Blog. It's World IPv6 Day. <http://www.attpublicpolicy.com/administration/it's-world-ipv6-day/>, Jun. 2011.
- [20] M. Bagnulo, P. Matthews, and I. van Beijnum. Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers. <https://tools.ietf.org/html/rfc6146>, Apr. 2011.
- [21] C. Bao, C. Huitema, M. Bagnulo, M. Boucadair, and X. Li. IPv6 Addressing of IPv4/IPv6 Translators. <https://tools.ietf.org/html/rfc6052>, Oct. 2010.
- [22] Ben Parker. IPv6 Transition for VZW. https://sites.google.com/site/ipv6implementors/2010/agenda/14_Parker_VerizonWireless.pdf, Jun. 2010.
- [23] M. Boucadair and E. Burgey. A64: DNS Resource Record for IPv4-mapped IPv6 Address. <https://tools.ietf.org/html/draft-boucadair-behave-dns-a64-01>,

- Oct. 2010.
- [24] C. Byrne. 464XLAT: Breaking Free of IPv4. https://conference.apnic.net/data/37/464xlat-apricot-2014_1393236641.pdf, Feb. 2014.
- [25] F. Chen, R. K. Sitaraman, and M. Torres. End-User Mapping: Next Generation Request Routing for Content Delivery. In *ACM Sigcomm CCR*, Apr. 2015.
- [26] L. Colitti. Broken IPv6 Clients. <https://sites.google.com/site/ipv6implementors/2010/agenda/BrokenIPv6clients.pdf>, Jun. 2010.
- [27] L. Colitti, S. H. Gunderson, E. Kline, and T. Refice. Evaluating IPv6 Adoption in the Internet. In *Passive and Active Measurement*, Apr. 2010.
- [28] N. Craig. A simple and rough comparison of Akamai and Cloudfront CDN's. <https://thedotproduct.org/a-simple-and-rough-comparison-of-akamai-and-cloudfront-cdns/>, Mar. 2014.
- [29] J. Curran. ARIN IPv4 Free Pool Reaches Zero. <https://www.arin.net/announcements/2015/20150924.html>, Sept. 2015.
- [30] M. CÃaçeres, F. J. Moreno, and I. Grigorik. Network Information API. <https://w3c.github.io/netinfo/>, Dec. 2015.
- [31] A. Dhamdhare, M. Luckie, B. Huffaker, k. claffy, A. Elmokashfi, and E. Aben. Measuring the Deployment of IPv6: Topology, Routing and Performance. In *ACM IMC*, Nov 2012.
- [32] A. Y. Ding, J. Korhonen, T. Savolainen, Y. Liu, M. Kojo, S. Tarkoma, and H. Schulzrinne. Reflections on Middlebox Detection Mechanisms in IPv6 Transition. In *IAB Workshop on Stack Evolution in a Middlebox Internet (SEMI)*, Jan. 2015.
- [33] Y. Ding, T. Savolainen, J. Korhonen, and M. Kojo. Speeding up IPv6 transition: Discovering NAT64 and Learning Prefix for IPv6 Address Synthesis. In *IEEE International Conference on Communications (ICC)*, June 2012.
- [34] C. Donley. IPv6-IPv4 Performance Comparison: The Effect of NAT. https://www.nanog.org/sites/default/files/tuesday_general_donley_ipv6_performance.17.pdf, 2013.
- [35] F. Brockners, S. Gundavelli, S. Speicher and D. Ward. Gateway-Initiated Dual-Stack Lite Deployment. <https://tools.ietf.org/html/rfc6674>, Jul. 2012.
- [36] Fortinet. IPv6: Network Security and the Next Generation of IP Communication. <http://www.fortinet.com/sites/default/files/solutionbrief/SG-IPV6.pdf>, 2015.
- [37] U. Goel, M. Steiner, M. P. Wittie, M. Flack, and S. Ludin. Detecting Cellular Middle-boxes using Passive Measurement Techniques. In *Passive and Active Measurements Conference (PAM)*, 2016.
- [38] U. Goel, M. P. Wittie, and M. Steiner. Faster Web through Client-assisted CDN Server Selection. In *IEEE International Conference on Computer Communications and Networks*, Aug. 2015.
- [39] B. I. T. A. Group. IPv6 AAAA DNS Whitelisting. http://www.bitag.org/documents/BITAG_TWG_Report-DNS_Whitelisting.pdf, Sept. 2011.
- [40] J. Iyengar and I. Swett. QUIC: A UDP-Based Secure and Reliable Transport for HTTP/2. <https://tools.ietf.org/html/draft-tsvwg-quic-protocol-00>, Jun. 2015.
- [41] R. Kisteleki. Measuring IPv6 usage at web clients and DNS resolvers. https://sites.google.com/site/ipv6implementors/2010/agenda/07_robert-kisteleki-measure-v6.pdf, Jun. 2010.
- [42] E. Kline. IPv6 Whitelist Operations. https://sites.google.com/site/ipv6implementors/2010/agenda/IPv6_Whitelist_Operations.pdf, Jun. 2010.
- [43] J. Korhonen and E. T. Savolainen. EDNS0 Option for Indicating AAAA Record Synthesis and Format. <https://tools.ietf.org/html/draft-korhonen-edns0-synthesis-flag-02>, Feb. 2011.
- [44] C. Kreibich, N. Weaver, B. Nechaev, and V. Paxson. Netalyzr: Illuminating the Edge Network. In *ACM SIGCOMM Conference on Internet Measurement*, Nov. 2010.
- [45] P. Lakhera. Your App and Next Generation Networks. <https://developer.apple.com/videos/wwdc/2015/?id=719>, June. 2015.
- [46] Y.-N. Law, M.-C. Lai, W. L. Tan, and W. C. Lau. Empirical Performance of IPv6 vs. IPv4 under a Dual-Stack Environment. In *IEEE International Conference on Communications (ICC)*, May 2008.
- [47] P. Mockapetris. Domain Names - Implementation and Specification. Nov. 1987.
- [48] A. M. Moreiras. Performance comparison between IPv4 and IPv6 on the Internet. <http://www.cepro.br/pub/CEPTRO/ArquivoNoticiaIPv6Performance/ipv6xipv4performance-napla.pdf>, May 2011.
- [49] E. Nygren. Implementing IPv6 on a Global Scale: Experiences at Akamai, Dec. 2011.
- [50] E. Nygren. Four years since World IPv6 Launch: entering the mainstream. <https://blogs.akamai.com/2016/06/four-years-since-world-ipv6-launch-entering-the-mainstream.html>, Jun. 2016.
- [51] E. Nygren, R. K. Sitaraman, and J. Sun. The Akamai Network: A Platform for High-Performance Internet Applications. In *ACM SIGOPS Operating Systems Review*, Vol. 44, No.3, July 2010.
- [52] D. Plonka and P. Barford. Assessing performance of Internet services on IPv6. In *IEEE Symposium on Computers and Communications (ISCC)*, July 2013.
- [53] S. Radhakrishnan, Y. Cheng, J. Chu, A. Jain, and B. Raghavan. TCP Fast Open. In *ACM CoNEXT*, Dec. 2011.
- [54] B. Reed. LTE devices must support IPv6, says Verizon. <http://www.networkworld.com/article/2257267/lan-wan/lte-devices-must-support-ipv6--says-verizon.html>, Jun. 2009.
- [55] P. Richter, M. Allman, R. Bush, and V. Paxson. A Primer on IPv4 Scarcity. In *Sigcomm CCR*, Apr. 2015.
- [56] V. Rideout and V. S. Katz. Opportunity for all? Technology and learning in lower-income families. http://www.joanganzcooneycenter.org/wp-content/uploads/2016/01/jgcc_opportunityforall.pdf, Jan. 2016.
- [57] J. Rosenberg, R. Mahy, P. Matthews, and D. Wing. Session Traversal Utilities for NAT (STUN). <https://tools.ietf.org/html/rfc5389>, Oct. 2008.
- [58] P. Saab. Facebook V6 World Congress 2015. <https://www.youtube.com/watch?v=An7s25FSK0U>, Mar. 2015.
- [59] T. Savolainen and J. Korhonen. Discovery of a Network-Specific NAT64 Prefix using a Well-Known Name. <https://tools.ietf.org/html/draft-savolainen-heuristic-nat64-discovery-01>, Feb. 2011.

- [60] N. Smith. Latency between Amazon Web Services' Regions. <http://www.nsmith.net/articles/2011-08/latency-between-amazon-web-services-regions/>, Aug. 2011.
- [61] M. Steiner and R. Gao. What slows you down? Your network or your device? *arXiv:1603.02293*, Mar. 2016.
- [62] N. Vallina-Rodriguez, S. Sundaresan, C. Kreibich, N. Weaver, and V. Paxson. CRAWDAD dataset icsi/netalyzer-android (v. 2015-03-24). Downloaded from <http://crawdad.org/icsi/netalyzer-android/20150324/middleboxes>, Mar. 2015. traceset: middleboxes.
- [63] M. Varvello, K. Schomp, D. Naylor, J. Blackburn, Alessandro Finamore, and K. Papagiannaki. Is The Web HTTP/2 Yet? In *Passive and Active Measurements Conference (PAM)*, Mar. 2016.
- [64] A. Vulimiri, P. B. Godfrey, R. Mittal, J. Sherry, S. Ratnasamy, and S. Shenker. Low Latency via Redundancy. In *ACM CoNEXT*, Dec 2013.
- [65] Y. Wang¹, S. Ye, and X. Li. Understanding Current IPv6 Performance: A Case Study from CERNET. In *Symposium on Computers and Communications (ISCC)*, 2005.
- [66] D. Wing. Learning the IPv6 Prefix of a Network's IPv6/IPv4 Translator. <https://tools.ietf.org/html/draft-wing-behave-learn-prefix-04>, Oct. 2009.
- [67] D. Wing and A. Yourtchenko. Happy Eyeballs: Success with Dual-Stack Hosts. <https://tools.ietf.org/html/rfc6555>, Apr. 2012.
- [68] R. Yasinovskyy, A. Wijesinha, R. Karne, and G. Khaksari. A comparison of VoIP performance on IPv6 and IPv4 networks. In *IEEE/ACS International Conference on Computer Systems and Applications (AICCSA)*, May 2009.
- [69] X. Zhou, M. Jacobsson, H. Uijterwaal, and P. Van Mieghem. IPv6 Delay and Loss Performance Evolution. *Int. J. Commun. Syst.*, 21(6), June 2008.