Sigcomm: G: Behind the Curtain – The importance of replica selection in next generation networks

John P. Rula Fabián E. Bustamante
Northwestern University

1. PROBLEM & MOTIVATION

Smartdevices are becoming the primary or only Internet point of access for an ever larger fraction of users. Nearly a quarter of current web traffic is mobile, and recent industry studies have estimated a fourfold increase in global mobile data traffic by 2018, mainly driven by data demands and the growing number of smart phones and tablets [7].

Content delivery networks (CDNs) are responsible for delivering most of today's Internet data. CDNs replicate popular content on servers worldwide and redirect users to "nearby" replica servers on demand. The Domain Name System (DNS) is instrumental in this process since CDN redirection, and thus the performance of content delivery, is typically based on the location of users' DNS resolvers [21].

Considering the importance of content and the critical role of DNS for both name resolution and localization in today's networks, it is somewhat surprising how little is known about the infrastructure and configuration of cell network DNS and its impact on content distribution. The 2011 study of Xu et al. [21] is today's most comprehensive analysis of (the US) cellular network infrastructure, combining data from DNS logs, smartphone users and server logs. For the radio technologies in their study – 3G UTMS and EVDO – the authors point out the dominant role of radio latency and the limited number of network ingress points. They conclude that, in that setting, the best option for content providers is to locate servers close to these ingress points and that, given the restricted routing in these cellular networks, choosing content servers based on local DNS servers is sufficiently accurate.

Independent of replica placement, current client redirection policy (i.e., replica selection) in mobile content delivery networks relies on similar heuristics to those used in the wired Internet, mapping clients' IPs of their local-DNS servers (LDNS) to a network region [12, 13, 20], or through the use of anycast IP addresses.

1.1 Next-Generation Cellular Networks

The recent growth of 4G access technologies [7, 11], such as LTE, radically changes the scene. Around the world service providers are busy rolling out 4G networks to meet users' increasing demand for faster, higher bandwidth connections. The most recent CISCO VNI report estimates that by 2018, the majority of North America devices and connections will have 4G capability. While 4G will be 15% of world-wide connections then, these connections will be responsible for 51% of traffic. When compared with 3G networks, 4G LTE presents a significantly different network, with a radically larger number of ingress points, and offers much lower radio access latency and variance. *We show that these changes make accurate content replica selection critical to the performance of end users in cellular networks.*

Compared with 2G/3G networks, 4G LTE offers significantly higher performance, with order of magnitude improvements on radio access latency (5-10ms) [1], and infrastructure differences with a rapidly growing number of gateways, from the handful reported in Xu et al. [21] to several 10s as found by Zarifis [22]. These changes – shorter access times and significantly more options to route clients' requests to/from replicas – are rapidly turning replica selection in a key determinant of mobile content delivery performance [16, 22].

We observe a wide range of performance diversity of content replicas seen by each client. We measured the HTTP time-to-first-byte to replicas for a collection of websites. By aggregating all replica servers seen by clients in our experiments, and taking the average client latency to each replica, we found the lowest latency replica for each website. All other replicas are shown as a percentage difference between their latency and the "best" seen replica. The cumulative distribution of these ratios is shown in Fig. 1.

While the degree of replica differential performance varies per operator, we consistently found replica latency increases ranging from 50% to 100% in all networks. In an extreme case, we observed replicas with 400% increased latency over the closest observed replica server for clients in Sprint.

1.2 Cellular Network Opacity

Firewall and NAT policies of cellular operators prohibit external entities like CDNs from probing clients or infrastructure in their network. Table 1 lists the reachability of cellular network DNS servers seen externally throughout our experiments. Of the 4 major US providers, only T-Mobile and Verizon servers responded to pings, and no servers on any operator profiled responded to traceroute probes. This opaqueness of cellular networks limits the amount of client network localization information available to CDNs. *We show that this opacity leads to poor replica selection policies of CDNs since they are unable to adapt to the changing nature of cellular networks.*

Even if these probes were allowed they would be ineffectual. Balakrishnan et al. [3] showed the failure of IP-based identification and geolocation in cellular networks, due in part to the ephemeral

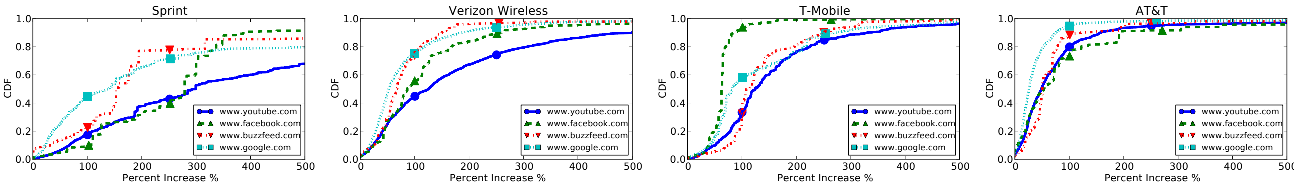

Figure 1: Performance comparison of HTTP time-to-first-byte for replica servers selected for mobile clients. We see cellular users directed toward a wide range of replica servers, with certain servers seeing 400% increased HTTP latency over the closest observed replica.

Provider	Total	Ping	Traceroute
Sprint	20	0	0
Verizon	34	32	0
AT&T	47	3	0
T-Mobile	40	40	0

Table 1: Number of external DNS resolvers able to be reached externally by either ping or traceroute probes.

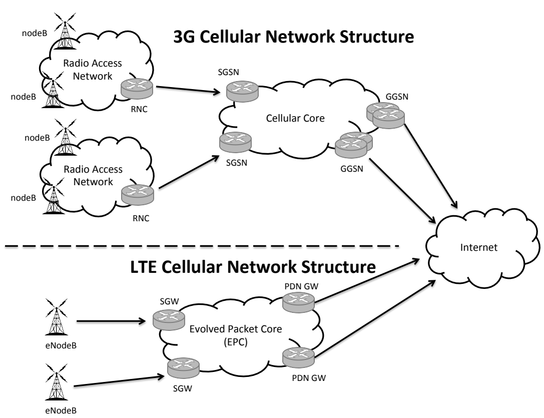

Figure 2: Network architecture changes cellular networks between 2/3G and LTE networks. LTE introduces a simpler, flatter network structure and an all-IP network.

and itinerant nature of mobile client's IPs – IPs assignment change rapidly and similar IPs are assigned to geographically distant devices. This dynamic assignment of client IP addresses is also shared by other features of cellular infrastructure including Client-LDNS mappings.

2. BACKGROUND & RELATED WORK

2.1 Background

In this section, we give an overview of current cellular infrastructure, the changes ongoing across cellular networks as they transition toward Long-Term Evolution (LTE) networks, and how these changes point toward the need for more intelligent replica selection for cellular devices.

2.1.1 Cellular Network Architecture

LTE has been growing rapidly since its entering the market in 2009. Service providers are busy rolling out 4G LTE networks to meet users' increasing demand for faster, higher bandwidth connections. LTE promises speeds up 150/75 Mbps of downstream/upstream throughput, significantly faster than what is possible in existing 3G networks. The 2014 CISCO VIN report estimates that, by 2018, the majority of North America devices and connections will have 4G capability. Transitioning to LTE technologies requires cellular operators to make substantial changes to their core networks, flattening their architectures and moving to an all-IP network. For example, LTE introduces an enhanced radio access component, the eNodeB, which removes the need for previous hierarchical structures such as the Radio Network Controller (RNC) by combining its functionality into a single node. These changes are illustrated in Figure 2. LTE also requires operators to switch over to the Evolved Packet Core (EPC), which requires an all-IP network [6], reducing the need for legacy, circuit based technologies.

2.1.2 Mobile Content Delivery

CDNs host and deliver the large majority of the mobile web content and, as in the wired Internet, most CDNs use the local DNS resolver (LDNS) of clients to locate them and find nearby replica servers for content delivery.

When a client requests an object hosted by a CDN, the client's local DNS resolver contacts the authoritative DNS (ADNS) of the domain name run by the CDN. The CDN uses the location of the client's DNS resolver as an approximate location for the client, and redirects the client to content servers nearby. In wired networks, this approach has been shown to be sufficiently accurate except when paired with certain ISP configurations or the use of public DNS services [15].

In their thorough discussion of the 2G/3G cellular network infrastructure, Xu et al. [21] suggest the placement of CDN replicas nearby cellular gateways given the constrained routing of traffic to and from mobile clients. Such placement of replica servers is necessary, if not sufficient for better end-to-end performance. For instance, if the mapping of clients to gateways is not consistent or related to the client's network location (potentially routing clients requests/response through the distant gateway), the end-to-end performance of clients would remain poor or inconsistent.

Current client redirection policy (i.e., replica selection) in mobile content delivery networks relies on similar heuristics to those used in the wired Internet, mapping clients' IPs of their local-DNS servers (LDNS) to a network region [12, 13, 20], or through the use of anycast IP addresses..

2.2 Related Work

Our work builds on the many previous efforts to explore mobile device performance within cellular networks, characterize client DNS infrastructure, and analyze the interplay between DNS and CDNs in replica selection. It represents the first analysis of cellular DNS infrastructure and its impact on content replica selection in 4G networks.

Several recent efforts have looked at the performance characteristics of mobile clients within cellular networks. Sommers et al. [19] compared the performance of mobile devices when they were in cellular networks or connected to 802.11 WiFi. Nikravesh et al. [14] looked at the longitudinal overall performance of mobile network performance from end devices, while Zarifis et al. [22] analyzed the increased latency experienced by mobile clients. With the exception of Zarifis et al., no paper considered the location of content replicas in their performance analysis.

Other work has investigated the impact of edge network effects such as client inaccessibility and IP address shuffling, and their impact on web services. Balakrishnan et al. [3] looks at the consistency of IP addresses within cellular networks as well as their geo-location accuracy. Casado et al. [4] studied how the opacity of edge networks and their stationarity of IP addresses can adversely affect network services. We continue this work with our observations of client LDNS resolver inconsistency, and its ill suited position as a client locator for CDNs.

Recently, Schomp et al. [17] and Alzoubi et al. [2] have performed large scale studies characterizing the structure and performance of client side LDNS resolvers. Our work is a natural complement as their investigation was limited to wired and broadband networks, because cellular network infrastructure can only be discovered from the end devices within the particular operator network given common firewall policies.

Krishnamurthy et al. [12] performed one of the earliest studies of content delivery networks and evaluated the effectiveness of replica selection for client performance. Our work naturally extends this, updating analysis with 4G cellular networks.

The effectiveness of DNS servers for content replica selection has been extensively explored before (e.g., [13, 18]). However most previous studies have been done under the assumption that client to resolver mappings remained constant. Huang et al. further investigated the replica selection mechanisms of a major CDN [10]. More recently Otto et al. [15] looked at the impact of using remote public DNS on replica selection for broadband connections. Here the authors show that using a client's ISP DNS servers yield the best performance a majority of the time, with public DNS services incurring significant performance of selected replicas. Our work shows that many times, the performance of replicas chosen by cellular operator DNS are in many cases no better than the publicly chosen ones, and offer worse performance in almost 25% of cases. The behavior of cellular DNS creates an entirely different environment for content replica selection, requiring its own set of localization techniques aside from client resolver.

Our work is closest to Xu et al. [21], which looked at the role cellular network structure plays in content placement. While the authors conducted initial investigations on DNS and content placement in cellular networks, our study includes longitudinal data from clients which allows us to monitor changes in DNS configuration from mobile end hosts. In addition, their work was conducted prior to major deployments of 4G technologies such as LTE which significantly decrease radio access latency and variance, as well as prior to the expansion of cellular network ingress points. Therefore, replica selection in modern networks represents a more important piece of mobile user performance than previously cited.

3. UNIQUENESS OF APPROACH

This work approaches the issue of client performance in cellular networks from looking at the differences in the structure and operation of cellular networks and how this affects the operation and performance of network services such as CDNs. Our investigation into DNS behavior allowed us to overturn common assumptions about client localization heuristics.

This work was the first to discover and analyze how the impact of dynamic infrastructure services such as DNS affect the performance of content delivery networks for cellular clients. Our approach is also unique in its analysis of longitudinal data from cellular clients, which allowed us to observe the spatial and temporal changes in behavior of cellular networks. In addition to cellular DNS, we are currently exploring the dynamic assignment of clients to packet gateways in cellular networks, and its implications for content delivery.

4. RESULTS & CONCLUSIONS

Many of the largest CDNs rely on DNS for client location/replica selection in cellular networks. Our results show that cellular DNS makes a poor indicator for client location, due to the opaqueness of cellular resolvers to external services, and the inconsistency between mobile clients and their visible LDNS resolvers. CDNs typically aggregate client resolvers behind traceroute divergence points and map clients based on measurements to these points. Unlike the majority of DNS resolvers, cellular DNS resolvers are unable to be probed and measured by CDNs. The inability to traceroute the cellular DNS resolvers invalidates this approach.

4.1 Client-LDNS Inconsistency

Our analysis revealed inconsistency between client and resolver mappings in all cellular providers investigated. These inconsistencies are not limited to the external resolver IP address, but include the /24 prefix of the external address. Figure 3 shows the LDNS pairs observed by each device over time, enumerated based on the order of appearance in our measurements.

The temporal stability of mapping between clients and external resolvers varies across carriers. Sprint and Verizon clients, for instance, show relatively stable mappings while the mappings for the remaining carriers appeared to be very unstable. Unstable mappings are not all the same, however, as can be seen when contrasting T-Mobile and AT&T. In the first two, changes in resolver IP addresses are typically accompanied by changes in the resolvers /24 prefix.

To understand the potential impact of these inconsistencies consider that, as shown by Xu et al. [21], DNS resolvers tend to be clustered at egress points. A change of resolver can result in the association of a mobile client with a completely different (and distant!) egress point.

4.2 Cellular Replica Selection

We investigate the impact of the inconsistency of client LDNS resolvers by comparing the mapping of replica servers to LDNS resolvers. For the selected DNS resolver, we construct a map of $< replicaIP, ratio >$ pairs capturing, for each replica server, the server IP address and the fraction of time that replica was used:

$$replica\_map = < (ip_1, \frac{ip_1 seen}{total\_seen}), ..., (ip_n, \frac{ip_n seen}{total\_seen}) >$$

We then us *cosine similarity* [5, 20] to quantify the similarity of replica servers mapped to each DNS resolver. The cosine similarity between two vectors A and B quantifies the degree of overlap between two vectors by computing the dot product of the vectors and dividing by the product of their lengths:

$$cos\_sim = \frac{A \cdot B}{\|A\| \|B\|}$$

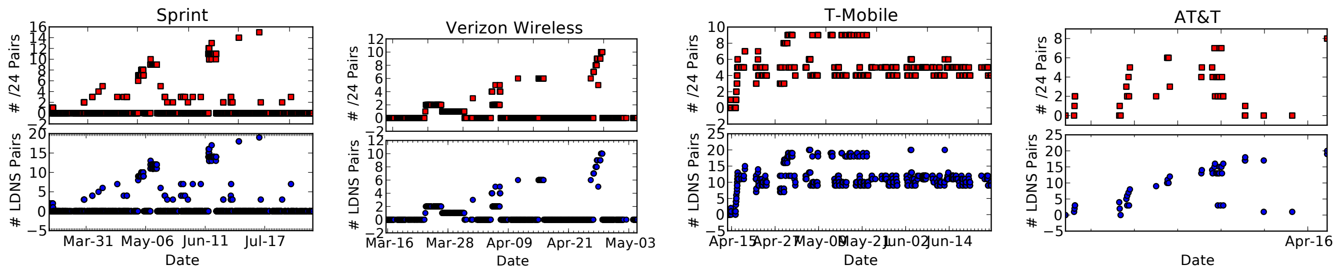

(a) Sprint (b) Verizon (c) T-Mobile (d) AT&T

Figure 3: Number of external resolvers observed by a client in each of the networks we looked at. Bottom: number of external resolver IP addresses. Top: number of unique /24 prefixes observed by resolvers. Client DNS resolvers change not just within localized clusters, but span multiple /24 prefixes over time.

The fact that CDNs aThis may explain the relatively small numbers of replica servers mapped to each cellular DNS resolver, particularly when compared to typical CDN-resolver mappings in wired networks.

Looking at the replica maps for each cellular operator and comparing cosine similarities, it appears that CDNs are grouping replica mappings by resolver /24 prefix. Figure 4 shows the cosine similarity (overlap) between replica sets for DNS resolvers in the same /24 prefix, and those in different prefixes. We see large degrees of independence between sets from differing /24 prefixes, with over 60% of sets having a cosine similarity of 0, meaning there is no overlap at all between replica vectors.

This high degree of replica set independence becomes a significant issue since, as we showed previously, cellular clients change LDNS resolvers frequently and across /24 prefixes potentially leading to large performance variability.

4.3 Other Heuristics for Replica Selection

In addition to client LDNS, other heuristics are possible for CDN replica selection including client IP address as well as the use of anycast addresses for replica servers. We evaluate the effectiveness of these alternative methods in this section.

Client IP. Sidestepping client-LDNS mapping, clients' IP could potentially be used for replica selection. The IETF proposed EDNS(0) Client Subnet (ECS) extension could be used to pass on a prefix of the client IP that a CDN could use for replica selection [8].

While the approach has been shown to be effective in wired networks [15], mobile client IPs have been previously shown to be quite ephemeral [3]. Using our dataset, Figure 5 illustrates the range of client IPs seen a by a single client in the T-Mobile network

As the figure shows, our example client is assigned over 60 unique IP addresses over a 3 month period. This pattern is not unique to T-Mobile but can be found in every one of the providers we investigated.

Anycast Addresses. While following a different model, for completeness we also evaluated the consistency of any cast address for mobile client as a few CDNs (e.g., EdgeCast) rely on replica selection.

Figure 6 illustrates the range of unicast IP addresses GoogleDNS requests were routed to over time, spanning multiple /24 prefixes. While seeing multiple IP addresses for GoogleDNS is not unusual when one considers that requests are load balanced across a number of machines, the presence of multiple /24 prefixes indicates requests are continually being routed to a geographically distinct data center according to their documentation [9].

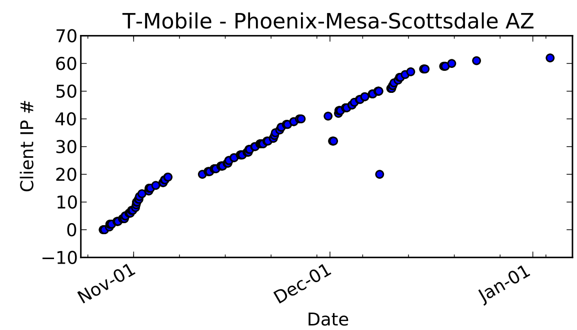

Figure 5: Enumerated visible client IP addresses for a representative client in T-Mobile's network. We observed over 60 unique IP addresses over a 3 month period for this client.

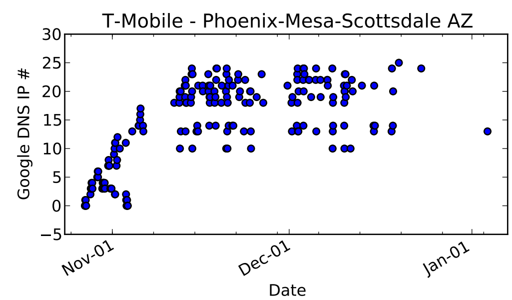

Figure 6: Anycast variability for clients shown for GoogleDNS (8.8.8.8) for a representative client in T-Mobile.

4.4 Continued Work

This work argues for the importance of content server replica selection in next-generation networks. We have shown the impact of replica server selection on the performance on cellular clients, and the impact that the complexity of the cellular core network and its opaqueness have on CDNs. This issue will continue to increase in importance as radio access latency continues to drop in future generations of wireless networks.

Additionally, we highlighted how current CDNs are ill equipped to handle the challenges of content delivery to cellular clients. The traditional heuristics used, including client LDNS and client IP address are not effective in locating clients within networks, and even if they were, CDNs would be unable to measure and

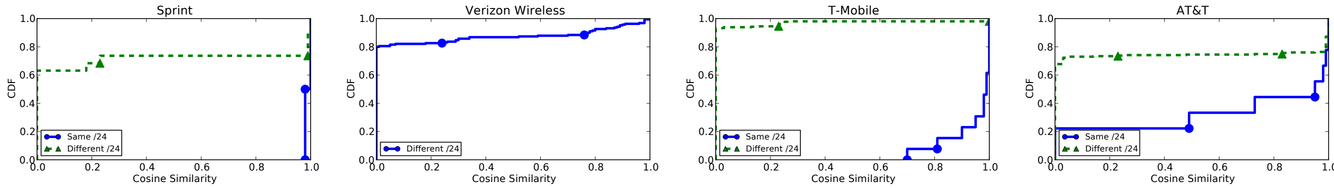

(a) Sprint (b) Verizon (c) T-Mobile (d) AT&T

Figure 4: Cosine similarity of replica servers for buzzfeed.com between resolvers within the same /24 prefix, and those in separate prefixes. Resolvers within the same /24 prefix see very similar sets of replicas (cosine similarity values close to one), and those in separate prefixes see high set independence (values close to zero). Clients changing resolver /24 prefixes are directed towards completely different sets of replica servers.

incorporate changes into their closed-loop systems. What is needed are new means of content delivery specifically designed with cellular clients, and the dynamic, opaque networks, in mind.

5. REFERENCES

[1] 3gpp tr 25.913: Requirements for evolved utra and evolved utran (v9.0.0). Technical report, 2009.

[2] H. A. Alzoubi, M. Rabinovich, and O. Spatscheck. The anatomy of LDNS clusters: Findings and implications for web content delivery. In *Proc. WWW*, 2013.

[3] M. Balakrishnan, I. Mohomed, and V. Ramasubramanian. Where's that phone?: geolocating IP addresses on 3G networks. In *Proc. IMC*, 2009.

[4] M. Casado and M. J. Freedman. Peering through the shroud: The effect of edge opacity on ip-based client identification. In *Proc. USENIX NSDI*, 2007.

[5] D. R. Choffnes and F. E. Bustamante. Taming the torrent: a practical approach to reducing cross-isp traffic in peer-to-peer systems. *ACM SIGCOMM Computer Communication Review*, 38(4):363–374, 2008.

[6] CISCO. Architectural Considerations for Backhaul of 2G/3G and Long Term Evolution Networks. `http://www.cisco.com/c/en/us/solutions/collateral/service-provider/mobile-internet/white_paper_c11-613002.pdf`.

[7] CISCO. CISCO visual networking index: Global mobile data traffic forecast update 2013-2018. Technical report, CISCO Systems Inc., 2014.

[8] J. Damas, M. Graff, and P. Vixie. Extension mechanisms for dns (edns (0)). 2013.

[9] Google. Frequently asked questions - Public DNS – Google Developers. `https://developers.google.com/speed/public-dns/faq`.

[10] C. Huang, A. Wang, J. Li, and K. W. Ross. Measuring and evaluating large-scale CDNs. In *Proc. IMC*, 2008.

[11] G. Intelligence. 4G driving data usage but not all markets reaping the rewards. Technical report.

[12] B. Krishnamurthy, C. Wills, and Y. Zhang. On the use and performance of content distribution networks. In *Proc. IMC*, 2001.

[13] Z. M. Mao, C. D. Cranor, F. Douglis, M. Rabinovich, O. Spatscheck, and J. Wang. A precise and efficient evaluation of the proximity between web clients and their local DNS servers. In *Proc. USENIX ATC*, 2002.

[14] A. Nikravesh, D. R. Choffnes, E. Katz-Bassett, Z. M. Mao, and M. Welsh. Mobile network performance from user devices: A longitudinal, multidimensional analysis. In *Proc. PAM*, 2014.

[15] J. S. Otto, M. A. Sánchez, J. P. Rula, and F. E. Bustamante. Content delivery and the natural evolution of DNS: Remote DNS trends, performance issues and alternative solutions. In *Proc. IMC*, 2012.

[16] J. P. Rula and F. E. Bustamante. Behind the curtain: Cellular dns and content replica selection. In *Proc. IMC*, 2014.

[17] K. Schomp, T. Callahan, M. Rabinovich, and M. Allman. On measuring the client-side dns infrastructure. In *Proc. IMC*, 2013.

[18] A. Shaikh, R. Tewari, and M. Agrawal. On the effectiveness of DNS-based server selection. In *Proc. IEEE INFOCOM*.

[19] J. Sommers and P. Barford. Cell vs. wifi: on the performance of metro area mobile connections. In *Proc. IMC*, 2012.

[20] A.-J. Su, D. R. Choffnes, A. Kuzmanovic, and F. E. Bustamante. Drafting behind akamai (travelocity-based detouring). *ACM SIGCOMM Computer Communication Review*, 36(4):435–446, 2006.

[21] Q. Xu, J. Huang, Z. Wang, F. Qian, and A. G. Z. M. Mao. Cellular data network infrastructure characterization and implication on mobile content placement. In *Proc. ACM SIGMETRICS*, 2011.

[22] K. Zarifis, T. Flach, S. Nori, D. Choffnes, R. Govindan, E. Katz-Bassett, Z. M. Mao, and M. Welsh. Diagnosing Path Inflation of Mobile Client Traffic. In *Proc. PAM*, 2014.