Bounding Mobile Latency from Network Interconnections

Shefali Tushar Kulkarni
USC/Information Sciences Institute
Los Angeles, United States
shefalit@usc.edu

John Heidemann
USC/Information Sciences Institute
Los Angeles, United States
johnh@isi.edu

Abstract

Mobile latency is determined by how Mobile Network Operators (MNOs) route traffic: from phones, across the MNO's network to *exit points* where it enters the public Internet. Exit points are few, and MNOs are quiet about details, but their role prompts us to ask: What is the latency cost to send traffic through an exit point? We develop a new method to evaluate this question, and demonstrate it on the three largest U.S. MNOs, although in principle it can apply globally. Our new analysis of existing traceroutes shows that AT&T, T-Mobile, and Verizon funnel nationwide traffic through only 9, 5, and 9 exits, and the geographic detour they incur adds up to 62.1 ms compared to an ideal, direct-path. We show that even assuming speed-of-light paths through the exit points, up to 46.9 ms of additional latency remains. To understand if these costs can be overcome we compare actual to theoretically optimal placements, finding that latency of current placements is 18–66% above the population and 32–72% above the market-share weighted optimum with the same number of exits. However, we show this gap is small in absolute terms: perfect placement would recover only a few milliseconds of the tens remaining. Geographic optimization offers little: the inflation users experience comes from routing and other overhead.

CCS Concepts

• **Networks** → **Mobile networks**.

Keywords

Mobile Networks, PLMN, Interconnection, Latency inflation

ACM Reference Format:
Shefali Tushar Kulkarni and John Heidemann. 2026. Bounding Mobile Latency from Network Interconnections. In *Asian Internet Engineering Conference 2026 (AINTEC '26), December 01–03, 2026, Taipei, Taiwan.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3858062.3858394



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.
AINTEC '26, Taipei, Taiwan
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-3086-3/2026/12
https://doi.org/10.1145/3858062.3858394

1 Introduction

Today, phones dominate human access to the Internet. Individuals are continuously connected, and expect fast network access wherever they are. A mobile phone connects to the Internet via the Public Land Mobile Network (PLMN) of the Mobile Network Operator (MNO). The mobile phone uses a 4G or 5G Radio Access Network (RAN), Packet Gateways, Edge, and traffic goes through the PLMN's backhaul to an *exit point* that interconnects to the public Internet (Figure 1).

Large online services and CDNs co-locate servers within PLMN data centers and at IXPs near exit points [11, 21, 38], to reduce the distance traffic travels after leaving the mobile network. Mobile traffic must still traverse the PLMN's backhaul to reach these servers, and a large-scale study finds that bringing data centers closer to their users does not guarantee low end-to-end latency [63]. Our paper is the first to explore the question of *how much latency is added by detours created by the MNO backhaul network and exit points?* We answer this question by locating the current exit points of three U.S. carriers: AT&T, T-Mobile, and Verizon.

We build on two prior measurements of U.S. carriers. Researchers at UCSD mapped the internal topology of the same three largest U.S. MNOs, focusing on the *interior* of the PLMN's network. They highlight the role of Packet Gateways, that bridge the dedicated wireless backhaul to the MNO's internal backbone [62]. We instead focus specifically on the *exit point at the boundary between the MNO and the public Internet.* We use their dataset but go beyond their analysis to explore the number and placement of exit points. Researchers at Northeastern independently collected traceroutes studying the same U.S. carriers in 2024 [16]. They emphasize 5G application and radio performance. We apply our new analysis to their data to validate exit points in 2024.

Our first contribution is we develop a new method that isolates delay required by physical distances from delays added by routing. We define three paths from a handset to a destination: the *observed path*, measured end-to-end; the *direct path* at speed-of-fiber directly between source and destination; and the *exit path* at speed-of-fiber through the operator's own nearest exit point (§3.5). The difference between observed and direct path gives *additional direct-path latency*; the difference between observed and exit-path gives

additional exit-path latency. Computing additional latency requires knowing where the current exit points are, so our first step is a new method to identify exit points from traceroutes (§3.4). We find 9 exit points for AT&T, 5 for T-Mobile, and 9 for Verizon (§4.1).

Second, we apply this method to the three largest U.S. MNOs. (In principle our method can apply to non-U.S. MNOs as well.) Paths through these exits give a median of 62.1 ms (AT&T), 43.9 ms (T-Mobile), and 49.1 ms (Verizon) additional direct-path latency (§4.2). We attribute this residual to hot-potato routing: operators hand traffic off at their nearest exit rather than the one nearest the destination, a known cause of path inflation [54] that has not been quantified for U.S. mobile networks. Granting each MNO the exits it already operates, a median of 23.6 ms, 37.2 ms, and 46.9 ms remains that distance cannot explain (§4.3). Most of the delay is therefore not the geographic detour through exit points but, we argue, routing policy.

Our final contribution examines how well MNOs have placed their exit points: do they account for proximity to users? could they do better? We formulate placement as a facility-location problem using population density and population weighted by MNO market share (§3.6). AT&T, T-Mobile, and Verizon exceed the optimum by 18%, 66%, and 43%, and by 32%, 61%, and 72% under market-share weighting. AT&T and Verizon place most exits near optimal, with the gap coming from a few outliers; T-Mobile's larger gap reflects its use of third-party backbone. Since our optimum ignores cost, peering, and available physical paths, it bounds what any operator could achieve. We show that after a service has five well-placed exits, additional exits have minimal benefits, reducing mean latency, per capita, by under 1% (§4.4).

The impact of our work is to provide operators and regulators a way to evaluate exit-point placement and reason about the trade-offs of adding new exits. Our approach extends to other regions and operators.

Our work analyzes traceroute data from 2021 [62], with validation against an independent collection from 2024 [15], so it poses no ethical concerns about user privacy. While this data is now 5 and 2 years old, we show our results are consistent over time (§5). Our basic data is available to researchers[15, 18, 62], and we will make the data from our analysis available at no cost [32].

2 Related Work

We build on prior studies of PLMNs, mobile latency, geolocation, topology optimization, and facility location.

Mobile network studies: Most closely related to our work are previous studies on the wide-area performance of PLMNs. Zhang *et al.* collected traceroutes from mobile phones physically mailed around the U.S., creating the *ShipTraceroute* dataset, and used it to study the router-level topology of the

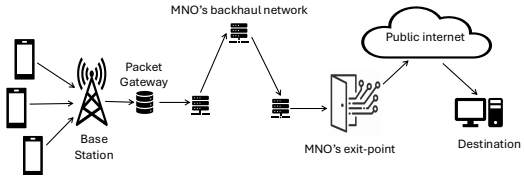


Figure 1: The path from a mobile device, through the PLMN backhaul and exit point to the Internet.

backhaul networks of three U.S. MNOs [62]. Their analysis focuses on the internal network of the PLMN, from the base-station, through the packet gateway (PGW) and the operator's backbone, to the exit point. Our analysis picks up where they leave off, looking at how that exit point affects end-to-end latency (§3.2), and further exploring the impact of the number of exit points. For latency analysis, they report only minimum RTT to the destination; we add theoretical optimal latency to show latency inflation.

A second study uses data from human driven mobile phones to evaluate coverage, geodiversity, operator diversity, throughput, RTT, and application performance [17]. We use the traceroutes from their later replication of this drive to validate exit-point locations observed in 2024 [16].

Studies characterize commercial mobile performance in the wild, measuring throughput, latency, and device behavior across U.S. cities [19, 40–43], in single cities outside the U.S. [30, 31, 59, 64], at scale through an app [60], from the operator's perspective [48], and under mobility [8, 44, 58]. Zarifis *et al.* diagnose mobile path inflation and conclude it disappeared in the U.S. by 2012 [61]; we find it persists a decade later. At global scale, Varvello *et al.* find that only 20% of 24 mobile networks achieve sub-20 ms latency [56], Ozcan *et al.* show from 9.6 billion RIPE Atlas pings that inter-city delays improved, in North America and Europe [47], and a year-long study finds that 5G has stagnated in delivering consistent latency improvements over LTE [28]. Together these studies provide a rich view of cellular performance and mobility, but they do not examine how core-network topology and exit-point placement constrain wide-area mobile latency. We instead infer the number and placement of MNO exit points, examining if limited exit points keep mobile users from benefiting from fiber and RAN-level improvements.

Latency matters most for interactive applications: Kämäräinen *et al.* measure latency in mobile video and gaming, reporting 20 ms of network delay under optimal conditions and showing the effect of geographic distance [27], while work on Mobile/Edge Cloud places computation near the handset to cut latency for cloud games [63]. Such applications increasingly rely on CDNs to reduce latency [13], but CDN nodes usually sit just outside the MNO's network, and some requests must still reach regional cloud datacenters—in

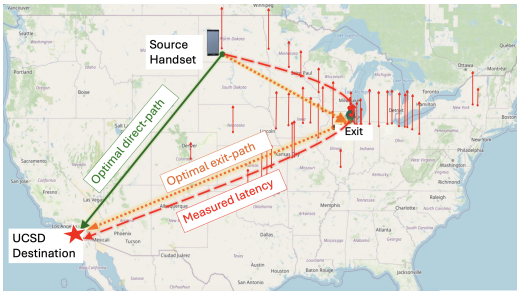


Figure 2: Additional latency from the handset to the Internet: observed (red dashed), optimal direct-path (green solid), optimal exit-path (orange dotted). Vertical red lines give additional direct-path latency through AT&T's Chicago exit.

both cases, exit-point count and placement affect latency. Recent works infer physical network boundaries (submarine cables) from active measurements [50, 57]. We apply similar measurements to a different boundary: private-to-public exit points within the U.S. mobile network, whose locations, unlike cable landing stations, are not public.

Optimal latency and path inflation in wide-area networks: Singla *et al.* compare observed Internet latency against the physical limits of fiber propagation, framing a "speed-of-light" Internet as the ultimate goal [53]. Latency is also the defining advantage in high-speed trading, motivating the use of microwave links instead of fiber to reduce delay [52]. We also measure PLMN latency against speed-of-fiber limits. The exit-point-induced latency we study is a specific instance of *path inflation*, where non-direct routing increases delay. Path inflation has been widely studied in wireless networks, showing that poor routing can significantly inflate latency in IP anycast services [34, 54], and that while such inflation may be acceptable for some applications, including the DNS root even at delays above 100 ms, it remains critical for latency-sensitive services such as CDNs [29]. Consistent with these findings, we observe comparable latency inflation in our exit-point analysis (§4.3). Sundberg *et al.* analyze over billions of RTT samples from a wireless ISP using eBPF, showing internal network latency exhibits large, time-varying tails during busy hours [55]. While their work focuses on last-mile variability, we test the middle-mile network.

Exit-point placement: We formulate optimal exit-point placement as a facility location problem [7, 23, 35], in which facilities (such as warehouses) are placed to minimize travel time. K-means and K-Medians have been recognized as good, approximate solutions to this problem [4, 12, 26], and we build on this method to evaluate optimal exit placement (§3.6). Recent work shows that large-scale placement problems can be solved with provable performance bounds using LP-relaxation heuristics [3], and that optimal placement in wireless mesh networks is NP-hard [24]. Unlike wireless-mesh networks, which are often decentralized and ad hoc, mobile network infrastructure is fixed and carefully planned to give good coverage, We use population-driven and market share measurements to quantify how far real-world MNO exit deployments deviate from the theoretical optimum (§4.4).

3 Methodology

We next discuss a new method that identifies two parts of the mobile latency: latency due to geography, and due to routing. We then apply it to AT&T, T-Mobile, and Verizon.

We define the problem (§3.1), types of latency (§3.2), and data sources (§3.3), then describe our new method for identifying number and location of exit points (§3.4), quantifying additional latency (§3.5), and evaluating optimality (§3.6).

3.1 Problem Statement

In a Public Land Mobile Network (PLMN), traffic from a mobile device reaches the public Internet only after traversing the operator's backhaul network and exiting through a limited number of exit points, as shown in Figure 1. In that backhaul network the PLMN controls the path and exit point. After traffic leaves the MNO AS, external networks determine the path according to their own BGP policies. Because exit-point locations are not public, several important questions remain open: How much additional latency is caused because of geographic detour? How good is exit-point placement, and what placement would optimize latency? How much will adding new exit points improve latency? Being able to reason about these questions can help us evaluate how well a PLMN network architecture serves its customers.

3.2 Latency and Definitions

To understand how PLMN architecture affects user latency we define several components of that latency.

We define three paths between the user and a destination. The *direct path* is the great-circle path from the mobile device to the server, traversed at the speed of light in fiber, roughly two-thirds of its speed in vacuum. In Figure 2, the direct path is the green line from North Dakota to San Diego. The *exit path* follows the same assumptions, but routes from mobile device through the PLMN's nearest exit point to the server—it is the orange path from North Dakota via Chicago to San Diego. The *observed path* is the path taken by the traceroutes, the slightly longer dotted red line.

These paths identify overhead types. *Additional direct-path latency* is the difference between observed and optimal direct-path latency. The optimal direct-path is the minimal physical distance from source to destination, so this difference observes latency due to detouring through an exit-point.

Second, we define *additional exit-path latency* as the difference between observed and optimal exit-path latency. Here

optimal exit-path latency is the physical distance from source to exit-point to destination, so this latency isolates latency from factors *other* than this detour. The optimal paths always assume physical distances only, so they ignore per-router processing delays and extra latency when physical paths are not "as-the-crow-flies", so these latencies are optimistic. These bounds are baselines because the gap between observed and bound is the overhead we set out to measure [53].

Finally, we ask how well the current exit points are placed by solving for the optimal placement of a given number of exits. Real deployments are constrained: fiber follows a limited set of physical paths [9], and data centers have fixed locations. Those constraints can only move placement away from the optimum, so our solution provides an upper-bound on the latency any operator could recover by relocating exits.

3.3 Datasets

Our work draws on several key datasets: ShipTraceroute and the Northeastern dataset for paths from mobile phones to the Internet, OpenCellID to geolocate mobile phones, IP address allocations from ARIN to identify network ownership, IP2Location to geolocate IPs, and finally census data for the U.S. population by county. We describe each below.

3.3.1 Mobile Network Traceroutes. ShipTraceroute [62] is our primary source for known PLMN topology and latency. It consists of hourly traceroutes from three phones, one per U.S. MNO (AT&T, T-Mobile, and Verizon), mailed to locations across 40 states in the continental U.S. in 2021. Targets include servers at UCSD, Google, Netflix, and other peers. We identify exit points from the Google, Netflix, and UCSD traceroutes, and measure additional latency from the UCSD data alone (§3.4). Because GPS is often unavailable, we locate phones by ShipTraceroute's recorded cell tower IDs.

To validate exit points in 2024 (§5), we use traceroutes from a second, independent mobile dataset [16]. It covers same three MNOs with fewer samples.

3.3.2 Mobile Locations. Knowing where each traceroute originates is essential to measuring source-to-destination latency. Since GPS is often unavailable, we map the cell tower IDs in ShipTraceroute to locations using OpenCellID [46], a database of tower locations, Cell IDs, and PLMN metadata. IP geolocation was considered but discarded, as MNOs typically use private IP space, preventing accurate geolocation.

3.3.3 IP to AS Mapping. We detect network transitions along traceroutes using the ISI outage supplement dataset [2], a pre-built map from each /24 prefix to its ASN, aggregated from the five global RIRs. A transition out of the MNO's AS marks an exit point. We use /24 granularity as it is the finest prefix commonly routed and observed in global BGP.

3.3.4 IP2Location. We determine the physical locations of PLMN exit points by combining IP2Location geolocation with IP-to-AS mappings. Geolocation accuracy remains an active research area. We considered MaxMind GeoLite, GeoIP2, and IP2Location-Lite [25]; prior work on Internet infrastructure finds their accuracy comparable, particularly in the United States [14, 36, 45]. We chose IP2Location-Lite for its slightly better city-level coverage. (We reconfirmed this with a mini-study of 200 known, U.S. university locations.) Although city-level geolocation accuracy can err by tens of kilometers, our analysis concerns latency shifts on the order of tens of milliseconds. Since light in fiber travels roughly 200 km per millisecond, city-level accuracy is more than adequate for detecting the effects we observe.

3.3.5 Population Estimation. We use the U.S. Census data from year 2023 to locate the U.S. population by county [22] and state-level MNO market share [1] to approximate subscribers when we evaluate exit-point placement (§3.6).

3.4 Identifying the exit points

Our first objective is to identify exit points: the locations where traffic leaves the PLMN for the public Internet.

To identify these exit points, we extract all traceroutes for an MNO from the ShipTraceroute dataset and determine the network operator of each valid hop using address-to-AS data. An exit point is identified as the first hop where the traceroute transitions from an MNO's IP to a public IP. In most network architectures, the first public IP encountered marks the interface where routing responsibility shifts from the MNO to an external operator. By verifying that the Autonomous System associated with this IP is not part of the MNO, we distinguish between internal and external routing.

We analyze 1,592 successful traceroutes for AT&T, 1,720 for Verizon, and 872 for T-Mobile from the ShipTraceroute dataset. Although the total number of traceroutes is limited, our analysis aggregates measurements at the level of base stations and exit points, and we observe stable exit-point inference and latency estimates with relatively few samples per site. Per MNO, we find 9 exit points for AT&T and Verizon. Prior work reports that T-Mobile uses multiple backbone providers that serve each region at different interconnection points [62]. Consistent with this, we observe that T-Mobile's traceroutes to the UC San Diego server traverse Zayo (AS 6461) and Lumen (AS 3356), which provide Internet transit service to T-Mobile. We therefore provide a second look at the T-Mobile–backbone combination and identify 5 exit points under this model (details in §4.1).

Role of the Destination: Exit points depend on the source and destination locations. We consider traceroutes with a single source: the handset. ShipTraceroute has traceroutes to three destinations: UCSD, Netflix, and Google.

For AT&T and Verizon we infer exit-point locations from traceroutes to all three destinations. For T-Mobile, we use only traceroutes to UCSD, because we find traceroutes to

Google and Netflix from T-Mobile (AS 21928) terminate within T-Mobile's network in different internal locations due to a practice known as *Off-Nets* [20].

Although exit-point determination can use all three destinations, in the next section we discuss that estimation of additional latency requires a single destination.

3.5 Quantifying Additional Latency

Given the identified exit points, we next assess their impact on latency. As described in §3.2, we compare *measured* latency with both the optimal direct-path latency and the optimal latency via the nearest exit point. We determine known latency from traceroutes. Traceroute records the round-trip time (RTT) taken to each hop from the cellphone to a destination, with the last hop showing the end-to-end latency [39].

We approximate the cellphone location for the origin of the traceroute. While GPS data is unavailable, ShipTraceroute dataset provides the Local Area Code (LAC) of the serving cell tower at the time of the traceroute. We map LAC to physical location using OpenCellID. Using OpenCellID metadata, we map each traceroute to its corresponding base station. We use the serving base station's location as a proxy for the handset's location when calculating theoretical latency; the RTT itself is measured end-to-end from the handset.

Our goal is to estimate propagation latency (roughly speed-of-fiber for physical distance) while ignoring queueing delay and other transient effects. We estimate propagation latency as the minimum RTT of all observations from the same device and location, following TCP Vegas' observation that minimum latency reflects physical propagation delay [5].

To keep analysis of additional latency consistent, we consider only traceroutes to one known destination (San Diego), and from known (geolocatable) source locations. Although ShipTraceroute also measures paths to Netflix and Google, we exclude these destinations because computing optimal latency uses the destination's physical location. With anycast and load balancing, the location of Google and Netflix is unknown. The exit points we find are the same for any destination, and only the size of the residual changes. The choice of measurement to a single location on the U.S. west coast will affect the relative size of detours–small detours may be a larger relative change for users on the west coast, but are washed out by the higher absolute propagation latency for east-coast users. Although bias is a limitation of ShipTraceroute and implies our absolute results should be considered preliminary, our methodology of comparing actual and optimal latency generalizes.

For optimal latency, we estimate the round-trip time (to and from the destination) over the shortest physical fiber path, assuming propagation at the speed of light in fiber. The speed-of-fiber, approximately 200 km/ms (about two-thirds the speed of light), determines data propagation through fiber optic media. We calculate the great-circle distance as shortest possible physical distance from cellphone to the fixed destination. As a theoretical optimal, we ignore limitations of physical fiber paths.

To compute the best possible exit-path latency, we identify the exit point from the traceroute (where traffic leaves the PLMN, as per §3.4) and determine its physical location using IP2Location. We then compute the exit-constrained physical distance as the sum of the distances from the handset to the exit point and from the exit point to the destination in San Diego. The estimated latency is obtained by applying the speed-of-fiber to this total distance.

Using the measured latency along with optimal direct and exit-path latency, we compute: *additional direct-path latency* as the difference between known latency and optimal direct-path latency, and *additional exit-path latency* as the difference between known latency and optimal exit-path latency.

3.6 Determining Optimal Exit Locations

Finally we compare each MNO's known exit-point locations with the optimal possible locations. This estimate complements our known latency measurements and their overhead (§3.5) to find the best-possible latency for the continental U.S. given a limited number of exit points.

Given n exit points, we define their cost as the distance from every person in the continental U.S., through their nearest exit point, to San Diego. The optimal location for a fixed number of n exit points minimizes this objective. Optimizing the location of n sites is the *facility location problem* [10, 23, 37]. While an optimal solution is *NP*-complete, k-means provides a well-known approximate solution [4].

We therefore estimate optimal exit-point location for each MNO by taking their known number of exit points as k, using county-by-county U.S. population data from the Census Bureau to optimize cost using k-means. We optimize exit-point placement using population as a proxy for expected demand, assuming more populous counties contribute more traffic. More granular metrics like per-cell traffic volume or MNO-specific subscriber distributions would be ideal, but are not publicly available. Population data provides a consistent, accessible basis for optimization across MNOs.

Our prior analysis fixes n at each MNO's known count, but in practical systems there is often a temptation to over-provision infrastructure on the assumption that *more is better*. To test that assumption for exit points, we add exits incrementally beyond the optimal set, placing each by k-means over the remaining population. This measures the marginal benefit of exit-point density under idealized conditions (§4.4).

3.7 Limitations

Our analysis is subject to limitations in destination diversity, measurement times, and routing visibility. Although we measure latency only to UCSD, our method generalizes to other

destinations, as discussed in detail in §3.5. Our analysis relies on ShipTraceroute measurements reported in 2021. Comparison with Northeastern data from 2024 provides evidence of exit-point stability at at least few of the locations examined, supporting the continued relevance of our analysis. Further details are discussed in §5. Our traceroute-based analysis cannot directly verify routing policies, distinguish the effects of policies within the MNO from those beyond the exit point, or identify interconnection types. We also considered RouteViews BGP data to assess exit-point and route stability over time, but RouteViews does not host the relevant mobile routes. Future work could address these limitations through operator-provided routing and interconnection data.

4 Results

We next identify exit points (§4.1), how exit location affects latency (§4.2 and §4.3) and optimality (§4.4) in 3 MNOs.

4.1 Exit-Point Locations

We first report the number of observed exit points for each mobile network. Figure 3 presents the number and locations of exit points for each MNO. We find that T-Mobile operates 5 exit points, while AT&T and Verizon each have 9 as of 2021. These exit points are distributed across the continental U.S., a topic further examined in §4.4. Are these exit points sufficient? On one hand, these MNOs remain profitable, and their customers appear satisfied, suggesting that they are adequate. However, Internet-centric companies prioritize broad interconnection; Google, for instance, peers at 25 sites in the continental U.S. as of July 2024 [49] and deploys edge nodes within thousands of ISPs [6]. While the number of exit points is ultimately a business decision for each MNO, it remains critical to understand how those choices affect users. In our next sections, we measure the end-to-end delay caused by the sparse exit-point topology and analyze where additional exit points can deliver most latency reduction.

4.2 Additional Direct-Path Latency

Having identified PLMN exit points, we now examine *additional direct-path latency*, defined in §3.2 as the difference between observed latency and the theoretical direct physical path from the mobile phone to its Internet destination.

Each vertical line in Figure 3 represents a latency observation from a handset to San Diego at different times and across different MNOs. Line height corresponds to additional direct-path latency. The number of observations varies by MNO due to differing traceroute success rates, and we report only one traceroute per unique cellular tower location.

Latency is high even for mobile phones near the destination. The solid lines in Figure 4 show the CDF of additional direct-path latency for each MNO, with median values of 62.1 ms for AT&T, 43.9 ms for T-Mobile, and 49.1 ms for Verizon. For context, the speed-of-fiber latency from Los Angeles to Boston (4200 km) is 21 ms, while measured latencies are around 70 ms, making an additional 40–60 ms significant.

Latency Causes: Figure 2 shows additional direct-path latency for all observations passing through AT&T's Chicago exit point. While Figure 3 shows the additional direct-path latency across all exit points for AT&T, Figure 2 zooms in on a single example: AT&T's Chicago exit point, to make the pattern of increasing latency with distance more visually salient. In this focused view, we can observe that the latency is lowest when the handset is near Chicago, indicating that once traffic reaches the Internet, it follows a relatively direct path. However, latency increases as handsets move farther away, peaking for devices in North Dakota and New York due to traffic being relayed through Chicago, sometimes backtracking before reaching San Diego.

Hot-Potato routing: Handsets in Nebraska and Kansas exit through Chicago rather than Los Angeles (another exit point of AT&T), although Los Angeles lies on the path to San Diego (Figure 2). MNO sends traffic to the exit point nearest the handset, apparently without considering the location of the final destination in the Internet (Figure 3). This early-exit policy is an example of hot-potato routing, long known to inflate network latency [54].

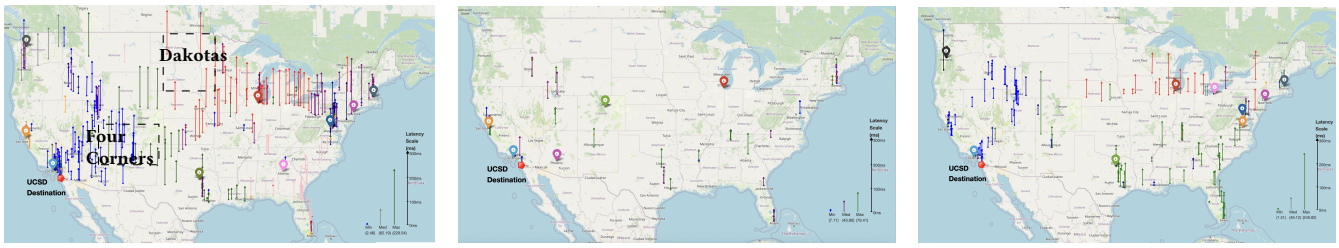
Our latency observations are not affected by radio-specific latency. Mobile phones have sophisticated algorithms to sleep the radio to conserve energy [43, 51]. Moreover, ShipTraceroute initiates each traceroute after an airplane-mode toggle, so only the first probe can incur any radio startup delay. Our observations use final RTT which occurs after the connection is established and the radio is active, so it does not include the radio wake-up delay.

Routes remain stable within regions. The ShipTraceroute authors report no destination-dependent routing or traffic engineering, and no path change as phone moves within a region [62]. We find the same: traceroutes from a given device and location always traverse the same exit point.

4.3 Additional Exit-Path Latency

Since the requirement to pass through an exit point is a large factor in mobile-to-Internet latency, we next look at *additional exit-path latency*, the additional latency that occurs once we accept that we must use a given exit point. Additional exit-path latency represents the cost of routing and indirect paths, but it factors out the physical latency incurred by passing through an exit point.

Figure 5 shows additional exit-path latency for each MNO, with Figure 4 showing corresponding CDFs. The median additional exit-path latency is 23.6 ms (AT&T), 37.2 ms (T-Mobile), and 46.9 ms (Verizon), all less than additional direct-path latency, reflecting the cost of relay through an exit point. We assume the wireless-hop latency is similar across providers, leaving a large 23 ms difference between AT&T and Verizon.



(a) AT&T: 9 exits: Seattle, San Francisco, Los Angeles, Dallas, Chicago, Atlanta, Ashburn, New York, Boston; and add. direct path latency.

(b) T-Mobile: 5 exits: Cupertino, Los Angeles, Phoenix, Denver, Chicago; and add. direct path latency.

(c) Verizon: 9 exits: Portland, Los Angeles, Dallas, Chicago, Cleveland, Richmond, Ashburn, NewYork, Boston; and add. direct path latency.

Figure 3: Observed additional direct-path latency (line height, with CDFs in Figure 4) and choice of exit point (line color), corresponding with observed exit points for each PLMN (inverted raindrop markers).

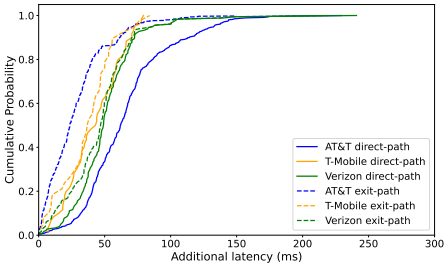


Figure 4: Additional direct-path (solid) and exit-path (dotted) latency per MNO (color).

To better understand how geographic distance influences exit-path performance, we analyze the relationship between round-trip path length and additional exit-path latency. We find negative correlations for AT&T $r = -0.342$, T-Mobile $r = -0.364$, and Verizon $r = -0.181$; with details in Appendix A of our technical report [33]. Given the extra traversal to the gateway, exit-path latency always exceeds the direct path; but our contribution is to quantify the margin of the cost. For T-Mobile and Verizon, this reduction is minimal (4 to 7 ms), as the fiber optic speed allows 5 ms to cover approximately 1000 km. In contrast, AT&T's additional exit-path latency is 38.5 ms lower than its additional direct-path latency, suggesting significant routing inefficiencies, particularly in the Dakotas (North and South Dakota) and the Four Corners region (Regions marked with dashed boxes in Figure 3 and Figure 5 for AT&T.) This discrepancy may stem from greater data availability for AT&T, as T-Mobile and Verizon have limited data for these regions.

4.4 Optimal Exit Location and Number

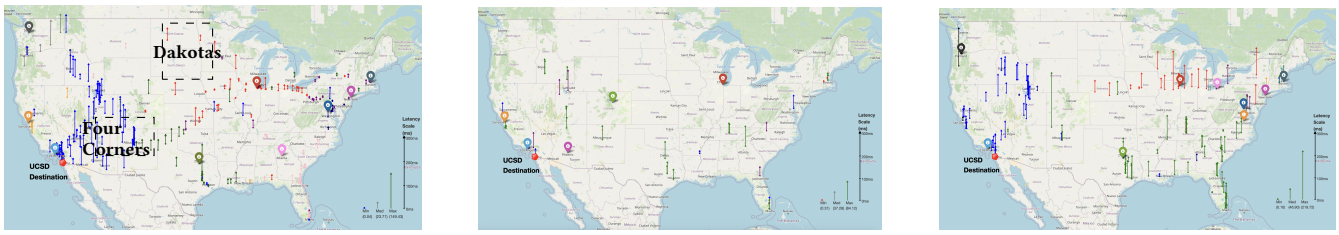
Optimal locations: We evaluate optimal exit locations based on U.S. population density (§3.6). Figure 6 shows computed optimal placements (black teardrops) and their catchments (colored circles) with 9- and 5-exits. We compare these optimal placements with the known exit points for AT&T (Figure 6a), T-Mobile (Figure 6b), and Verizon (Figure 6c).

Optimal exit placements reflect population density, with a higher concentration in eastern and mid-western regions of the U.S. However, the model does not account for existing telecommunications infrastructure. For example, it suggests an exit near the Four Corners region, despite limited local infrastructure, and places an exit in Bakersfield, California, rather than in metropolitan areas such as Los Angeles or San Francisco. Although Bakersfield appears to be an unusual choice, it is strategically located between major population centers and along key transit corridors. This placement could benefit rural users, those in areas with limited Wi-Fi access, and travelers across California by reducing latency.

A comparison of known and optimal placements shows that AT&T's and Verizon's exit points are well-aligned with the population distribution, especially in major metropolitan areas such as Seattle, Los Angeles, Dallas, Chicago, Atlanta, Washington, D.C., and New York. On the other hand, T-Mobile's five exits show larger variance from optimal, with two exit points in California and none on the East Coast.

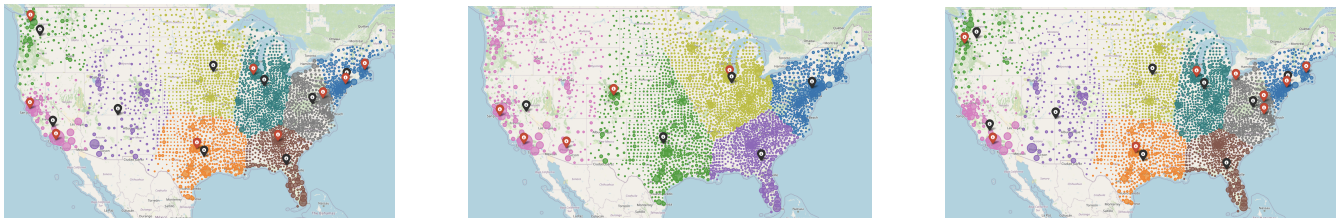
Efficiency of actual placements: To quantify the alignment of known and optimal placements, we compute an *exit-efficiency* ratio, which compares the per-capita distance to the nearest known versus optimal exit. A ratio of 1 indicates perfect alignment, while values above 1 indicate additional latency. AT&T achieves an exit-efficiency ratio of 1.18, while Verizon and T-Mobile have ratios of 1.43 and 1.66, respectively. These values suggest that AT&T's exit points are well-placed, Verizon's are moderately efficient, and T-Mobile's limited coverage contributes to higher latency, particularly in the East and South.

We also analyze exit-point placement with MNO customers. We estimate customers by weighting population data by publicly reported MNO market share [1]. Market-share weighted ratios are 1.32 for AT&T, 1.72 for Verizon, and



(a) Additional exit-path latency of AT&T (b) Additional exit-path latency of T-Mobile (c) Additional exit-path latency of Verizon

Figure 5: Observed additional exit-path latency (line height) and exit-point selection (line color) for each MNO (inverted raindrop markers). Dotted lines in Figure 4 give the CDF.



(a) AT&T 9-exit optimal: Seattle, WA; Bakersfield, CA; 4-Corners, NM; Dallas, TX; Des Moines, IA; Gary, IN; Washington, DC; White Plains, NY; compared to known exits

(b) T-Mobile 5-exit optimal: Reno, NV; Gainesville, TX; Chicago, IL; Macon, GA; New York, NY; compared to known exits

(c) Verizon 9-exit optimal: Seattle, WA; Bakersfield, CA; 4-Corners, NM; Dallas, TX; Des Moines, IA; Gary, IN; Washington, DC; White Plains, NY; compared to known exits

Figure 6: Theoretical optimal exit-point placement (black teardrops) and catchments (colored circles) for 9- and 5-exit-point networks, compared with known exit points (red teardrops), weighted by U.S. population (circle area).

1.61 for T-Mobile. These again suggest that AT&T operates closer to population than other MNOs. However, Verizon's market-share-weighted exit points are further from optimal than population-weighted and is comparable for T-Mobile.

Exit-point placement is inherently constrained by network topology, capacity limitations, the need to maintain balanced link utilization, and related business constraints, such as high-demand areas and regional prioritization. However, our analysis quantifies the latency cost of current deployment choices, providing a benchmark for potential optimizations. While operators cannot freely place exit points, our findings suggest that targeted adjustments could reduce latency without increasing operational complexity.

Benefits of more: We also extend our analysis to ask: *How significant is the latency improvement achieved by deploying extra exit points?* We expect additional exit points to lower mean latency per capita. Our analysis shows that each new exit yields only modest improvement, reducing latency by 1% or less (details in Appendix B of [33]). Optimizing the locations of existing exits provides a larger reduction in latency. *Better placement* is more beneficial than simply adding more exits.

5 Validation

We validate the stability of inferred exit points by comparing exit-point locations seen in 2021 and 2024. For 2021, we use ShipTraceroute, and for 2024 use Northeastern's 5G driving data [15]. The 2024 data is sparser than 2021, reporting 439 traceroutes for AT&T and 441 for Verizon.

Northeastern's data finds some of exit points we see in 2021: we see 3 of 9 exits (33.3%, in San Francisco, L.A., and Ashburn) for AT&T, and 5 of 9 (55.6%, in L.A., Chicago, Ashburn, N.Y., and Boston) for Verizon. It also finds a *new* exit in Pittsburgh for Verizon. The persistence of these locations across datasets indicates that exit-point placements are stable over years, as expected since an exit requires significant engineering to deploy. While our results confirm many exit points are still active in the same locations (precision is high), sparseness in the Northeastern data prevents assessing completeness. We cannot validate T-Mobile exits because policy changes hide traceroute after 2021.

6 Conclusion

We showed how to identify components of latency in how MNO networks interact with the Internet. We show that exit-point interconnections add latency, but relatively few exit points can provide good latency for the U.S.

Acknowledgments

We thank the authors of ShipTraceroute and the Northeastern (cellular-network-performance-on-wheels) dataset for their timely responses and clarification on the datasets.

This material is based upon work supported by DARPA under Contract No. HR001120C0157 (SABRES), and through NSF under OAC-2530698 (BRIPOD) and 2453092 (TEPPIT). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DARPA or NSF.

References

[1] [n. d.]. Top Mobile Data Providers State-by-State. https://www.webfx.com/blog/internet/top-mobile-data-providers-state-by-state/.

[2] Analysis of Network Traffic (ANT) group. 2023. LANDER:internet outage adaptive a52 supplement-20230401. https://ant.isi.edu/datasets/readmes/internet_outage_adaptive_a52_supplement-20230401.README.txt. [Online; accessed October-2023].

[3] Rasoul Behravesh, David Breitgand, Dean H Lorenz, and Danny Raz. 2024. A practical near optimal deployment of service function chains in edge-to-cloud networks. In *IEEE INFOCOM*. IEEE, 751–760.

[4] P. S. Bradley, K. P. Bennett, and A. Demiriz. 2000. *Constrained K-Means Clustering*. Technical Report MSR-TR-2000-65. Microsoft Research. https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/tr-2000-65.pdf

[5] L.S. Brakmo and L.L. Peterson. 1995. TCP Vegas: end to end congestion avoidance on a global Internet. *IEEE Journal on Selected Areas in Communications* 13, 8 (1995), 1465–1480. doi:10.1109/49.464716

[6] Matt Calder, Xun Fan, Zi Hu, Ethan Katz-Bassett, John Heidemann, and Ramesh Govindan. 2013. Mapping the Expansion of Google's Serving Infrastructure. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Barcelona, Spain, 313–326. doi:10.1145/2504730.2504754

[7] Gérard Cornuéjols, George Nemhauser, and Laurence Wolsey. 1983. *The uncapacitated facility location problem.* Technical Report. Cornell University Operations Research and Industrial Engineering.

[8] Haotian Deng, Chunyi Peng, Ans Fida, Jiayi Meng, and Y Charlie Hu. 2018. Mobility support in cellular networks: A measurement study on its configurations and implications. In *Proc. of the Internet Measurement Conference*.

[9] Ramakrishnan Durairajan, Paul Barford, Joel Sommers, and Walter Willinger. 2015. InterTubes: A Study of the US Long-haul Fiber-optic Infrastructure. In *ACM SIGCOMM*. ACM, London, United Kingdom, 565–578. doi:10.1145/2785956.2787499

[10] H. A. Eiselt and Vladimir Marianov. 2011. Pioneering Developments in Location Analysis: Location Problems—Its Problem Statement, Its Components, and Application. In *Foundations of Location Analysis*, H. A. Eiselt and Vladimir Marianov (Eds.). Springer, 3–6.

[11] Equinix. [n. d.]. Equinix America data centers. https://www.equinix.com/data-centers/americas-colocation.

[12] Feng, Yingying and Zhao, Shasha and Liu, Hui. 2020. Analysis of Network Coverage Optimization Based on Feedback K-Means Clustering and Artificial Fish Swarm Algorithm. *IEEE Access* 8 (2020). doi:10.1109/ACCESS.2020.2970208

[13] Pantelis A Frangoudis, Louiza Yala, and Adlen Ksentini. 2017. CDN-as-a-service provision over a telecom operator's cloud. *IEEE Transactions on Network and Service Management* 14, 3 (2017), 702–716.

[14] Manaf Gharaibeh, Anant Shah, Bradley Huffaker, Han Zhang, Roya Ensafi, and Christos Papadopoulos. 2017. A look at router geolocation in public and commercial databases. In *Proceedings of the 2017 Internet Measurement Conference* (London, United Kingdom) *(IMC '17)*. ACM, New York, NY, USA, 463–469. doi:10.1145/3131365.3131380

[15] Moinak Ghoshal, Omar Basit, Imran Khan, Z. Jonny Kong, Sizhe Wang, Yufei Feng, Phuc Dinh, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2025. Data and scripts from the paper "Replication: Performance of Cellular Networks on the Wheels". GitHub repository. https://github.com/NUWiNS/imc2025-replication-cellular-network-performance-on-wheels-data Measurements collected November 1–8, 2024.

[16] Moinak Ghoshal, Omar Basit, Imran Khan, Z Jonny Kong, Sizhe Wang, Yufei Feng, Phuc Dinh, Y Charlie Hu, and Dimitrios Koutsonikolas. 2025. Replication: performance of cellular networks on the wheels. In *Proceedings of the 2025 ACM Internet Measurement Conference*. 381–396.

[17] Moinak Ghoshal, Imran Khan, Z. Jonny Kong, Phuc Dinh, Jiayi Meng, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2023. Performance of Cellular Networks on the Wheels. In *Proc. of the 2023 ACM on Internet Measurement Conference* (Montreal QC, Canada) *(IMC '23)*. ACM, New York, NY, USA, 678–695. doi:10.1145/3618257.3624814

[18] Moinak Ghoshal, Imran Khan, Z. Jonny Kong, Yufei Feng, Phuc Dinh, Jiayi Meng, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2025. Data from the paper "Performance of Cellular Networks on the Wheels". GitHub Repository https://github.com/NUWiNS/imc2023-cellular-network-performance-on-wheels-data. https://github.com/NUWiNS/imc2023-cellular-network-performance-on-wheels-data

[19] Moinak Ghoshal, Z Jonny Kong, Qiang Xu, Zixiao Lu, Shivang Aggarwal, Imran Khan, Yuanjie Li, Y Charlie Hu, and Dimitrios Koutsonikolas. 2022. An in-depth study of uplink performance of 5G mmWave networks. In *Proc. of the ACM SIGCOMM Workshop on 5G and Beyond Network Measurements, Modeling, Use Cases*.

[20] Petros Gigis, Matt Calder, Lefteris Manassakis, George Nomikos, Vasileios Kotronis, Xenofontas Dimitropoulos, Ethan Katz-Bassett, and Georgios Smaragdakis. 2021. Seven Years in the Life of Hypergiants' Off-Nets. In *Proceedings of the ACM SIGCOMM Conference*. ACM, Virtual, 516–534. doi:10.1145/3452296.3472928

[21] Google. [n. d.]. Data Center locations published by Google. https://www.google.com/about/datacenters/locations/.

[22] US Gov. [n. d.]. Population Density by county. https://www.census.gov/data/datasets/time-series/demo/popest/2020s-counties-total.html.

[23] S. Louis Hakimi. 1964. Optimum Locations of Switching Centers and the Absolute Centers and Medians of a Graph. *OR* 12 (1964).

[24] Bing He, Bin Xie, and Dharma P. Agrawal. 2007. Optimizing the Internet Gateway Deployment in a Wireless Mesh Network. In *2007 IEEE International Conference on Mobile Adhoc and Sensor Systems*. 1–9. doi:10.1109/MOBHOC.2007.4428603

[25] IP2Location. [n. d.]. IP2Location Geolocation Lite. https://lite.ip2location.com/ip2location-lite.

[26] Vitor Jesus and Rui L Aguiar. 2012. Figures of merit for the placement (in) efficiency of interconnected CDNs. In *2012 IEEE Symposium on Computers and Communications (ISCC)*. IEEE, 000277–000282.

[27] Teemu Kämäräinen, Matti Siekkinen, Antti Ylä-Jääski, Wenxiao Zhang, and Pan Hui. 2017. Dissecting the End-to-end Latency of Interactive Mobile Video Applications. In *Proceedings of the 18th International Workshop on Mobile Computing Systems and Applications* (Sonoma, CA, USA) *(HotMobile '17)*. Association for Computing Machinery, New York, NY, USA, 61–66. doi:10.1145/3032970.3032985

[28] Imran Khan, Moinak Ghoshal, Joana Angjo, Sigrid Dimce, Mushahid Hussain, Paniz Parastar, Yenchia Yu, Xueting Deng, Sumit Hawal, Shirui Huang, et al. 2025. How mature is 5G deployment? A cross-sectional, year-long study of 5G uplink performance. *Computer Communications* 237 (2025), 108153.

[29] Thomas Koch, Ethan Katz-Bassett, John Heidemann, Matt Calder, Calvin Ardi, and Ke Li. 2021. Anycast In context: a tale of two systems.

In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference* (Virtual Event, USA) *(SIGCOMM '21)*. Association for Computing Machinery, New York, NY, USA, 398–417. doi:10.1145/3452296.3472891

[30] Konstantinos Kousias, Mohammad Rajiullah, Giuseppe Caso, Ozgu Alay, Anna Brunstrom, Luca De Nardis, Marco Neri, Usman Ali, and Maria-Gabriella Di Benedetto. 2022. Coverage and performance analysis of 5G non-standalone deployments. In *Proceedings of the 16th ACM Workshop on Wireless Network Testbeds, Experimental evaluation & CHaracterization*. 61–68.

[31] Konstantinos Kousias, Mohammad Rajiullah, Giuseppe Caso, Ozgu Alay, Anna Brunstrom, Luca De Nardis, Marco Neri, Usman Ali, and Maria-Gabriella Di Benedetto. 2022. Implications of handover events in commercial 5G non-standalone deployments in Rome. In *Proceedings of the ACM SIGCOMM Workshop on 5G and Beyond Network Measurements, Modeling, and Use Cases*. 22–27.

[32] Shefali Kulkarni and John Heidemann. 2026. Data from the paper "Bounding Mobile Latency from Network Interconnections". Web page https://ant.isi.edu/datasets/mobile_latency. https://ant.isi.edu/datasets/mobile_latency

[33] Shefali Tushar Kulkarni and John Heidemann. 2026. *Bounding Mobile Latency from Network Interconnections (extended).* Technical Report arXiv:2610.19156 [cs.NI]. USC/ISI. https://arxiv.org/abs/2610.19156

[34] Zhihao Li, Dave Levin, Neil Spring, and Bobby Bhattacharjee. 2018. Internet Anycast: Performance, Problems, and Potential. In *Proc. of the ACM SIGCOMM Conference*. ACM, Budapest, Hungary, 59–73. doi:10.1145/3230543.3230547

[35] Baoding Liu. 2009. Facility location problem. *Theory and practice of uncertain programming* (2009), 157–165.

[36] Ioana Livadariu, Thomas Dreibholz, Anas Saeed Al-Selwi, Haakon Bryhni, Olav Lysne, Steinar Bjørnstad, and Ahmed Elmokashfi. 2020. On the Accuracy of Country-Level IP Geolocation. In *Proceedings of the Applied Networking Research Workshop* (Virtual Event, Spain) *(ANRW '20)*. Association for Computing Machinery, New York, NY, USA, 67–73. doi:10.1145/3404868.3406664

[37] Vladimir Marianov and Daniel Serra. 2011. Median Problems in Networks: Introduction. In *Foundations of Location Analysis*, H. A. Eiselt and Vladimir Marianov (Eds.). Springer, 31–43.

[38] Meta. [n. d.]. U.S. datacenter fleet. https://datacenters.atmeta.com/us-locations/.

[39] Scott Mitchell. [n. d.]. How to Read a Traceroute. https://www.inmotionhosting.com/support/server/ssh/read-traceroute/.

[40] Arvind Narayanan, Eman Ramadan, Jason Carpenter, Qingxu Liu, Yu Liu, Feng Qian, and Zhi-Li Zhang. 2020. A first look at commercial 5G performance on smartphones. In *Proceedings of The Web Conference*. 894–905.

[41] Arvind Narayanan, Eman Ramadan, Rishabh Mehta, Xinyue Hu, Qingxu Liu, Rostand AK Fezeu, Udhaya Kumar Dayalan, Saurabh Verma, Peiqi Ji, Tao Li, et al. 2020. Lumos5G: Mapping and predicting commercial mmWave 5G throughput. In *Proceedings of the ACM internet measurement conference*. 176–193.

[42] Arvind Narayanan, Muhammad Iqbal Rochman, Ahmad Hassan, Bariq S Firmansyah, Vanlin Sathya, Monisha Ghosh, Feng Qian, and Zhi-Li Zhang. 2022. A comparative measurement study of commercial 5G mmWave deployments. In *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*. IEEE.

[43] Narayanan, Arvind and Zhang, Xumiao and Zhu, Ruiyang and Hassan, Ahmad and Jin, Shuowei and Zhu, Xiao and Zhang, Xiaoxuan and Rybkin, Denis and Yang, Zhengxuan and Mao, Zhuoqing Morley and others. 2021. A variegated look at 5G in the wild: performance, power, and QoE implications. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. 610–625.

[44] Yunzhe Ni, Feng Qian, Taide Liu, Yihua Cheng, Zhiyao Ma, Jing Wang, Zhongfeng Wang, Gang Huang, Xuanzhe Liu, and Chenren Xu. 2023. {POLYCORN}: Data-driven cross-layer multipath networking for high-speed railway through composable schedulerlets. In *20th USENIX Symposium on NSDI*. 1325–1340.

[45] Abdullah Yasin Nur. 2023. Accuracy and Coverage Analysis of IP Geolocation Databases. In *2023 International Balkan Conference on Communications and Networking (BalkanCom)*. 1–6. doi:10.1109/BalkanCom58402.2023.10167899

[46] OpenCellID. [n. d.]. OpenCellID database. https://www.opencellid.org/.

[47] Selim Ozcan, Ioana Livadariu, Georgios Smaragdakis, and Carsten Griwodz. 2023. Longitudinal Analysis of Inter-City Network Delays. In *2023 7th Network Traffic Measurement and Analysis Conference (TMA)*. IEEE.

[48] Paniz Parastar, Andra Lutu, Özgü Alay, Giuseppe Caso, and Diego Perino. 2023. Spotlight on 5G: Performance, device evolution and challenges from a mobile operator perspective. In *IEEE INFOCOM 2023*. IEEE, 1–10.

[49] PeeringDB. 2024. Google Interconnection Facilities. https://www.peeringdb.com/net/433. Accessed 2024-07-26..

[50] Alagappan Ramanathan and Sangeetha Abdu Jyothi. 2023. Nautilus: A Framework for Cross-Layer Cartography of Submarine Cables and IP Links. *Proc. ACM Meas. Anal. Comput. Syst.* 7, 3, Article 46 (Dec. 2023), 34 pages. doi:10.1145/3626777

[51] Sanae Rosen, Haokun Luo, Qi Alfred Chen, Z Morley Mao, Jie Hui, Aaron Drake, and Kevin Lau. 2014. Discovering fine-grained RRC state dynamics and performance impacts in cellular networks. In *Proceedings of the 20th annual international conference on Mobile computing and networking*. 177–188.

[52] Andriy Shkilko and Konstantin Sokolov. 2020. Every Cloud Has a Silver Lining: Fast Trading, Microwave Connectivity, and Trading Costs. *Journal of Finance* 75, 6 (Dec. 2020), 2899–2927. doi:10.1111/jofi.12969

[53] Ankit Singla, Balakrishnan Chandrasekaran, P. Brighten Godfrey, and Bruce Maggs. 2014. The Internet at the Speed of Light. In *Proceedings of the ACM Workshop on Hot Topics in Networks*. ACM, Los Angeles, CA, USA. doi:10.1145/2670518.2673876

[54] Neil Spring, Ratul Mahajan, and Thomas Anderson. 2003. The causes of path inflation. In *Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications* (Karlsruhe, Germany) *(SIGCOMM '03)*. Association for Computing Machinery, New York, NY, USA, 113–124. doi:10.1145/863955.863970

[55] Simon Sundberg, Anna Brunstrom, Simone Ferlin-Reiter, Toke Høiland-Jørgensen, and Robert Chacón. 2024. Measuring Network Latency from a Wireless ISP: Variations Within and Across Subnets. In *Proceedings of the 2024 ACM on Internet Measurement Conference*. 29–43.

[56] Matteo Varvello and Yasir Zaki. 2023. A worldwide look into mobile access networks through the eyes of amigos. In *2023 7th Network Traffic Measurement and Analysis Conference (TMA)*. IEEE, 1–10.

[57] Caleb J Wang, Ying Zhang, Qianli Dong, Esteban Carisimo, Ramakrishnan Durairajan, and Fabián E Bustamante. [n. d.]. Calypso's Voyage: Charting Traceroute Paths to Submarine Cables. ([n. d.]).

[58] Jing Wang, Yufan Zheng, Yunzhe Ni, Chenren Xu, Feng Qian, Wangyang Li, Wantong Jiang, Yihua Cheng, Zhuo Cheng, Yuanjie Li, et al. 2019. An active-passive measurement study of TCP performance over LTE on high-speed rails. In *The 25th Annual International Conference on Mobile Computing and Networking*.

[59] Dongzhu Xu, Anfu Zhou, Xinyu Zhang, Guixian Wang, Xi Liu, Congkai An, Yiming Shi, Liang Liu, and Huadong Ma. 2020. Understanding operational 5G: A first measurement study on its coverage, performance and energy consumption. In *Proc. of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures,*

and protocols for computer communication. 479–494.

[60] Xinlei Yang, Hao Lin, Zhenhua Li, Feng Qian, Xingyao Li, Zhiming He, Xudong Wu, Xianlong Wang, Yunhao Liu, Zhi Liao, et al. 2022. Mobile Access Bandwidth in Practice: Measurement. *Analysis, and Implications. In SIGCOMM* 2, 2 (2022), 3–1.

[61] Kyriakos Zarifis, Tobias Flach, Srikanth Nori, David Choffnes, Ramesh Govindan, Ethan Katz-Bassett, Z Morley Mao, and Matt Welsh. 2014. Diagnosing path inflation of mobile client traffic. In *International Conference on Passive and Active Network Measurement*. Springer, 23–33.

[62] Zesen Zhang, Alexander Marder, Ricky Mok, Bradley Huffaker, Matthew Luckie, K C Claffy, and Aaron Schulman. 2021. Inferring Regional Access Network Topologies: Methods and Applications. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Virtual Event. doi:10.1145/3487552.3487812

[63] Zhehui Zhang, Shu Shi, Varun Gupta, and Rittwik Jana. 2019. Analysis of Cellular Network Latency for Edge-Based Remote Rendering Streaming Applications. In *Proceedings of the ACM SIGCOMM 2019 Workshop on Networking for Emerging Applications and Technologies* (Beijing, China) *(NEAT'19)*. Association for Computing Machinery, New York, NY, USA, 8–14. doi:10.1145/3341558.3342199

[64] Dariusz Zmysłowski and Jan M Kelner. 2024. Mobile Network Operators' Assessment Based on Drive-Test Campaign in Urban Area for iPerf Scenario. *Applied Sciences* 14, 3 (2024), 1268.