Inferring Regional Access Network Topologies: Methods and Applications

Zesen Zhang
UC San Diego
zez003@eng.ucsd.edu

Alexander Marder
CAIDA, UC San Diego
amarder@caida.org

Ricky Mok
CAIDA, UC San Diego
cskpmok@caida.org

Bradley Huffaker
CAIDA, UC San Diego
bradley@caida.org

Matthew Luckie
University of Waikato
mjl@wand.net.nz

KC Claffy
CAIDA, UC San Diego
kc@caida.org

Aaron Schulman
UC San Diego
schulman@cs.ucsd.edu

ABSTRACT

Using a toolbox of Internet cartography methods, and new ways of applying them, we have undertaken a comprehensive active measurement-driven study of the topology of U.S. regional access ISPs. We used state-of-the-art approaches in various combinations to accommodate the geographic scope, scale, and architectural richness of U.S. regional access ISPs. In addition to vantage points from research platforms, we used public WiFi hotspots and public transit of mobile devices to acquire the visibility needed to thoroughly map access networks across regions. We observed many different approaches to aggregation and redundancy, across links, nodes, buildings, and at different levels of the hierarchy. One result is substantial disparity in latency from some Edge COs to their backbone COs, with implications for end users of cloud services. Our methods and results can inform future analysis of critical infrastructure, including resilience to disasters, persistence of the digital divide, and challenges for the future of 5G and edge computing.

CCS CONCEPTS

• **Networks** → **Public Internet; Logical / virtual topologies; Physical topologies; Network measurement.**

KEYWORDS

Internet topology, Access networks, Mobile networks, Traceroute

ACM Reference Format:

Zesen Zhang, Alexander Marder, Ricky Mok, Bradley Huffaker, Matthew Luckie, KC Claffy, and Aaron Schulman. 2021. Inferring Regional Access Network Topologies: Methods and Applications. In *Internet Measurement Conference (IMC '21)*, November 2–4, 2021, Virtual Event, USA. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3487552.3487812>

Figure 1: An Internet Service Provider (ISP) network can be divided into backbone and regional access.

1 INTRODUCTION

ISP regional access networks are an essential component of an ISP's infrastructure: they bridge millions of users' last-mile access links to the ISP's nearest backbone routers, which may be hundreds of miles away, to reach the Internet. Access networks strategically *aggregate* traffic in order to balance reliability and performance against the cost of providing connectivity over large regions (Fig. 1). Fiber cuts or other hardware failures can lead to large-scale outages spanning neighborhoods, counties, or entire states. Risk of outages motivates ISPs to provide *redundancy* within and across levels of aggregation. But regional access networks are remarkably opaque, which makes it challenging for academics to quantitatively study their role in the continually evolving ecosystem.

We present a measurement-driven exploration of regional access network topologies, through the lens of aggregation and redundancy as the foundations of scalability. Building on advances in Internet measurement methods and tools over the last two decades, we first establish and demonstrate the ability for an independent third-party to infer the topologies of different regional access networks, including aspects of the underlying physical (layer-1) topology, using only active measurements. We then perform measurement campaigns to infer and compare how major U.S. wireline (Comcast, Spectrum, AT&T) and mobile (AT&T, Verizon, T-Mobile) ISPs incorporate aggregation into their regional access network topologies. We show how analyzing these topological differences across providers and even across regions of the same access ISP can yield insights into the propagation of large correlated last-mile link failures [23, 62], sources of edge computing latency [58] and how to minimize it [52], performance limitations of metro-area fiber networks [42], and the evolving Internet ecosystem [21].

This work is licensed under a Creative Commons Attribution International 4.0 License.

IMC '21, November 2–4, 2021, Virtual Event, USA
© 2021 Association for Computing Machinery.
ACM ISBN 978-1-4503-9129-0/21/11...\$15.00
<https://doi.org/10.1145/3487552.3487812>

Our methodology synthesizes the state-of-the-art in three dimensions of Internet router-level topology analysis: combining traceroute paths from distributed vantage points; extracting semantics from DNS hostnames and IPv6 addresses of observed topology; and IP address alias resolution to further refine topology inferences. Our methodology leverages ideas that are well-established in the Internet measurement community, but we used them in different combinations, and with creative refinements, to accommodate the geographic scope, scale, and architectural richness of U.S. regional access ISPs. One of our contributions is effectively a recipe book of how to gain insight into network topology structure under different sets of constraints.

For example, the largest U.S. cable providers today tend to have near-universal reverse DNS on their router IPs, but hostnames are often stale. We devised heuristic methods to filter out stale or misleading hostnames, facilitating a comprehensive mapping of their regional topologies.

Networks without geographically meaningful internal hostnames require another way to infer geographic coverage of routers in the Central Offices (COs), such as probing from many geographically distributed VPs in the regional network. This requirement highlights the most common challenge in inferring internal network topologies, well-known to the research community: many networks provide more (or more accurate) visibility to internal vantage points than to external ones, especially for mobile networks. For wireline networks, we obtained internal vantage points by wardriving public WiFi networks in fast-food chains.

Mobile (cellular) access networks present many challenges: e.g., limited rDNS and blocking external probing. To capture this topology, we used IPv6 address structure and a new approach to gain county-wide internal visibility into mobile networks: cross-country shipping of mobile phones while they actively perform energy-efficient network measurements.

Our measurement methodology contributions are:

- Cost-effective approaches to procuring vantage points; energy-efficient approaches to sustaining mobile vantage points while devices are in transit.
- Analysis of DNS hostnames and IPv6 addresses to infer and geolocate topology, as well as strategically select probe targets to fill in coverage gaps.
- Heuristics that leverage active measurement techniques to filter noise (e.g., stale DNS information) or erroneously inferred hops, and infer missing (e.g., non-responding) IP hops.

Our empirical contributions result from applying our methods to previously unmapped parts of the infrastructure: wireline and mobile regional access networks. We gather enough data to ground the following discoveries:

- Topological redundancy—a metric of resilience—varies widely within and across levels of the hierarchy.
- Layer 3 topology information, including hostnames and IPv6 addresses, can reveal building locations and building-level redundancy within access networks.
- Regional access networks leverage a range of aggregation strategies to accommodate diverse markets, environments, and technologies. One result is substantial disparity in latency from some Edge COs to their Backbone COs.

Figure 2: Routers in EdgeCOs aggregate users while routers in AggCOs aggregate EdgeCOs.

These measurements inform analysis of critical infrastructure, including resilience to disasters, persistence of digital divide, and challenges for edge computing.

2 BACKGROUND

Like any network, regional access networks must balance reliability and performance against the cost of deploying and operating an infrastructure. For these networks, an additional challenge is achieving all these aspirations at scale—in terms of market size and geographic scope—and adapting to evolution of technology and industry structure. The primary architectural mechanism used to achieve this scale is *aggregation*. Regional networks aggregate traffic in *Central Offices (COs)* through a hierarchy of routers (and switches). An *Edge CO* (or EdgeCO) aggregates traffic from many thousands of last-mile links. Similarly, an *Aggregation CO* (or AggCO) houses routers to aggregate traffic from dozens of EdgeCOs, often across metropolitan areas or entire states. *Backbone COs* (or BackboneCO) house equipment to aggregate traffic from AggCOs—and sometimes EdgeCOs, and provide transit services either via the access network provider’s backbone network, or via other ISPs. Fig. 2 illustrates how ISPs use *redundancy* across layers of aggregation to provide resiliency in case of link or node failures.

The topology of this regional network infrastructure frames the performance and resilience of the networks, as traffic must cross the regional access links to reach the nearest Internet PoP. If the nearest PoP is far away, users may experience significant minimum latency [54]. The level of redundancy in any component of the network similarly provides an upper bound on robustness in different parts of the network.

2.1 Evolution of access networks

We provide some historical background for context on the challenges and opportunities for measurement of these networks to study their performance and reliability.

A typical access network is physically constructed of several fiber rings (Fig. 3). These networks generally use three hierarchical rings.

Figure 3: (a) Access networks are physically constructed of a hierarchy of fiber rings. (b) Overlapping Ethernet star topologies are built on top of these rings.

Last-mile links (e.g., Cable/Passive Optical/DSL) are aggregated over *Last-mile Loops* that reach into neighborhoods and terminate at EdgeCOs. EdgeCO traffic is aggregated in *Edge Rings* that terminate at one or more AggCOs. Then AggCO traffic is aggregated on *Core Ring* to BackboneCOs. Early Internet access network architectures used shared SONET on these fiber rings. All traffic passed through every CO in reserved time slots to reach aggregation points at the higher layer, which resulted in suboptimal bandwidth and latency performance. In the 1990s, in response to exploding demand for Internet bandwidth in part due to IP convergence (moving all voice and data services to use IP), access networks gradually replaced these SONET ring topologies with Ethernet-based star topologies. Often repurposing the same physical plant that SONET used, ISPs deployed direct Ethernet links between EdgeCOs, AggCOs, and BackboneCOs using techniques such as Dense Wave Division Multiplexing (DWDM) [80].

With DWDM and the addition of reconfigurable optical switches, operators are now capable of flexibly configuring layer-1 topology on these physical fiber rings. However, there is no standard way to design such a topology, i.e., how many or which EdgeCOs connect to how many or which AggCOs. Designs are driven by demographic (population), geographic (or geological) and regulatory constraints, especially when crossing state boundaries. To improve resilience or accommodate high-traffic regions, some components may use a full mesh (e.g., between the backbone and aggregation layers in Fig. 2) rather than star topology. These design choices have implications for accommodating future services, e.g., low-latency high-bandwidth edge services in COs [52].

Another driving force in network design is reliability and robustness in the face of inevitable failures of components. SONET was especially robust to fiber cuts and router failures, since traffic could travel in either direction around the ring. The Ethernet star topology does not have this feature. To compensate, ISPs add redundant routers and/or links over existing fiber rings, creating “dual ring/star” topologies (Fig. 3).

Each network independently chooses how to implement redundancy: adding routers or links or entire COs, and within or across

Figure 4: Mobile access networks are built of a combination of wireline access networks and mobile-specific networks.

different levels of aggregation. Some regions connect with only one backbone CO; others have only one AggCO. Many COs have redundant routers, and fiber rings have inherent redundancy (i.e., “East” and “West” directions in Fig. 3). Redundant backbone COs can dramatically improve regional reliability given that long-distance backbone also can be prone to failure [21, 25].

2.2 Mobile access networks

Mobile regional access networks leverage wireline access networks to provide mobile Internet access over a large geographic area. Mobile access networks consist of two halves (Fig. 4). Mobile devices communicate wirelessly with base stations, and the base stations aggregate user traffic over existing wireline networks and hidden mobile-specific AggCOs that connect to one or more EdgeCOs; these facilities can also be called Mobile Telephone Switching Offices (MTSO), Mobile Switching Centers (MSC), or mobile data-centers. This overlay network is called the Radio Access Network (RAN). These AggCOs serve as a bridge to connect the RAN to the mobile “Packet Core” which terminates the mobile network at one or more Packet Gateways (PGW) in an EdgeCO. Then these EdgeCOs connect directly to one or more regional BackboneCOs to connect with the rest of the Internet. EdgeCOs in mobile networks are the primary location where low-latency high-bandwidth IP-based edge services can be deployed; Verizon is already deploying edge services at these COs [17].

All mobile Internet traffic traverses both the wireline regional access networks described above and the mobile packet core. Therefore, to understand the aggregation and redundancy of mobile networks we need to also understand the wireline network. Unfortunately, the RAN—and therefore the wireline access network—is not visible from probes sent by mobile devices. However, by observing the topology of wireline providers’ regional networks using their wireline last-mile links, we get insight into the limitations of some of the RANs. Indeed, in this paper we mapped both the wireline and mobile networks of AT&T, and the cable providers we mapped provide backhaul for all major mobile carriers [70].

There can be significant differences in the topology of mobile access networks because each provider can make their own decision about how to aggregate their traffic to BackboneCOs. Indeed, they have significant flexibility because mobile networks are designed to be an overlay on other networks. The primary factors affecting mobile network topology are tradeoffs in performance, economics, and reliability.

3 RELATED WORK

Mapping Wired Networks: In 2002, Spring *et al.* developed Rocketfuel to construct router-level maps of individual networks using focused traceroutes, alias resolution, DNS hostnames, and BGP routing tables [45], and used it to map ten transit networks. Researchers used the Rocketfuel maps to study, interior routing [41], path inflation [64], and the maps sparked lively methodological discussions (e.g., [68, 75]). In 2007, Mao *et al.* developed NetworkMD [43], an approach to infer *failure groups* in the last-mile layer-1 network topologies—i.e., devices such as repeaters whose failure impact downstream modem connectivity through topological dependencies. In 2011, motivated by the incompleteness and methodological limitations of traceroute-based maps at the time, Knight *et al.* constructed the Internet Topology Zoo, parsing information that network operators published on their websites; the majority of their maps are at the PoP level (where a network interconnects with other networks) and half are research and education networks [34]. PoP-level maps are not sufficiently granular to study aggregation structures in access networks. Beginning in 2015, a growing body of research investigated the physical infrastructure behind networks, especially focused on fiber [21, 42], the frequency and impact of fiber cuts [25], and the impact of fiber deployment on end-to-end latency [9]. In this work, we map the router-level aggregation structure of access networks.

Broadband Networks: Substantial work analyzed broadband networks by sending probe packets to user’s gateway [18], deploying home routers or embedded devices with measurement scripts [3, 10, 67], embedding measurements into BitTorrent software [57], and crowdsourcing measurements to end-users [35, 51]. This work illuminated characteristics (e.g., latency, packet loss rate, throughput, and uptime) of the end-to-end [8, 13, 65] and last-mile [7, 24, 66] performance of residential broadband networks without understanding the logical and physical topologies of the access network. This paper leverages the aggregation structure that we discovered to understand the latency observed by end-users.

Mapping Mobile Networks: Previous work studied the geographic coverage of mobile regional networks using the correlation between IP prefixes and location [79] and locating the PoP used by mobile devices in traceroutes [81]. However, these analyses were performed on 3G networks, and do not reveal the underlying access network infrastructure that produces these behaviors. An extensive body of wireless network measurement research has investigated the behavior of network elements unique to wireless—everything from the end-user devices to the mobile-specific middleboxes [2, 14, 22, 36, 49, 71, 76–78]. Connectivity factors can also impede performance of mobile ISPs, e.g., legacy hierarchical routing [20], lack of direct interconnection with content providers [81], peering strategies between mobile virtual network operators (MVNOs) and the underlying network infrastructure [61], and poor selection of DNS servers [55, 82]. In this work we determine the sources of latency limitations in today’s mobile networks, and if those limitations can be overcome by moving services into access network infrastructure. We also reveal new hints in IPv6 addresses of reveal the region, packet gateway, and CO serving a mobile carrier.

4 METHODOLOGY OVERVIEW

Broadly, our topology mapping methods require the ISPs to allow the following measurements of their access networks:

Traceroute: We need the ability to observe routers in each CO with traceroute to uncover access network topology. This is straight forward when the network uses IP routing; if the network uses MPLS to organize routing between COs, then our method requires the ability to observe routers in each CO using traceroute towards exit routers, using the method in [72]. For wired access networks, we prune a small number of links because traceroute can produce false links. Our pruning assumes the access network has a ring/star topology (§2.1). Although we can not validate this assumption for all ISPs we study—ISPs rarely publish information about their internal topology—we found support for the use of ring/star topologies in access networks from Comcast [74], Deutsche Telekom [27], Cisco [30], and Juniper [48].

Alias Resolution: To accurately map IPs to routers, we require the ability to resolve aliases using active probing (e.g., with Mercator [26] and MIDAR [33]).

Reverse DNS or Structured Addressing: To accurately map routers to COs, we either require the operator to ether label some of their router IPs with hostnames in their reverse DNS (rDNS), or they must have clear structure in their router address space that corresponds to the structure of their access network.

How general is our approach? The remainder of this paper (§5, §6, §7) studies six different U.S. access networks – Comcast and Charter (wired), AT&T (wired and mobile), Verizon (mobile), and T-Mobile (mobile). We believe our method can be extended to other access networks outside of the U.S.; for example, China’s top three providers have been shown to provide the measurements we need for topology mapping [69]. Also, Bell Canada and Shaw Communications Canada appear to provide the necessary measurement primitives for our method [38]. However, there are classes of access network where our method will not work. Primarily, these are where traceroute does not observe routers in COs. This is common in countries where the access and retail functions of ISPs are separated; examples of these are New Zealand’s UFB [16] and Australia’s NBN [46]. In these scenarios, the access provider tunnels subscriber traffic to hand-over points where the retail provider is co-located, which might be in entirely different cities, so the aggregation structure in the access network is invisible.

5 CASE STUDY: COMCAST AND CHARTER

Our first case study focuses on two networks amenable to external traceroute-style measurements that also provide CO information in their rDNS. We focus on Comcast and the former Time Warner regional networks acquired by Charter—the largest cable Internet providers in the U.S.—networks that typically include router, building, and network information in their rDNS. Fig. 5a shows a traceroute into Charter’s Southern California region. The rDNS for hop 13 ends with `tbone.rr.com`, indicating a BackboneCO. Each subsequent hop includes an rDNS tag for the regional network `social`, showing the transition from Charter’s backbone into the regional network. Each rDNS name includes a portion of a CLLI code

13	66.109.6.227	bu-ether15. lsan care0yw-bcr00. tbone .rr.com
14	66.109.6.231	agg2. lsan care01r. socal .rr.com
15	72.129.1.1	agg1. sndh caax01r. socal .rr.com
16	72.129.1.141	agg1. sndg caax01h. socal .rr.com
17	76.167.26.170	agg1. sndg caax02m. socal .rr.com

(a) Charter traceroute with CO CLLIs.

14	96.110.41.226	be-1102-cr02. sunny vale.ca. ibone .comcast.net
15	68.86.92.206	ae-72-ar01. beaverton .or. bv erton.comcast.net
16	68.85.243.238	ae-1-rur201. troutdale .or. bv erton.comcast.net
17	162.151.213.86	po-1-1-cbr01. troutdale .or. bv erton.comcast.net

(b) Comcast traceroute with CO locations.

Figure 5: Paths into Charter’s Southern California region and Comcast’s Beaverton, OR region. Each hostname includes a CO identifier and regional network.

geolocating the router. Fig. 5b shows the same pattern in Comcast’s Beaverton, OR region, using CO locations rather than CLLIs.

Similar to many backbone point-of-presence (PoP) rDNS labels, the regional CO tags indicate the location of the carrier offices (COs). In Charter, the CLLI codes uniquely identify a specific building. Comcast sometimes uses the street address for a CO, but more commonly uses neighborhood, or city, names that provide a general geographic location, along with the U.S. state. Inspired by other tools that extract information from rDNS [12, 29, 38, 45], we hand-crafted regular expressions (regexes) to map these CO addresses.

Our methodology for these types of networks—those with rDNS and that can be externally probed—proceeded in two phases: (1) build and annotate CO-level topology graphs for these networks; (2) heuristically refine the graphs to reflect the actual topology.

5.1 Phase 1: Build Router-Topology Graphs

This phase conducts traceroutes to reveal the CO interconnections in each regional network. We conducted our probing from 47 vantage points (VPs) distributed throughout the United States in access, cloud, and transit networks.

First, we tracerouted to an address in every /24 in each regional network to expose at least one router from each EdgeCO. Second, we tracerouted to every address with rDNS matching one of our regexes to find CO interconnections missed in the first step. We identified IP addresses with hostnames matching our regexes in the Rapid7 rDNS dataset [53] which queries for PTR records for every IPv4 address. Directly targeting CO router interfaces observed 5.3x and 2.6x more CO interconnections than the /24 traceroutes for Comcast and Charter, respectively, as some COs responded to the /24 probing using addresses without rDNS. Third, we tracerouted to every intermediate IP address observed in these traceroutes to identify links that are entry and exit routers for an MPLS tunnel [72], allowing us to discard false edges between these COs. This MPLS heuristic proved important in larger Charter regions, where top level AggCOs appeared directly connected to nearly all EdgeCOs, which contradicted information about the Charter topology in Maine that we received from a trusted source.

Finally, we used alias resolution (Mercator [26] and MIDAR [33]) to group IP addresses according to their router. We included all IP addresses with rDNS matching our regexes, as well as all IP addresses routed by each regional network. We annotated each inferred router group with a CO tag, using the most common tag

(a) Graph after removing external edges.

(b) Modified graph that accurately represents the topology.

Figure 6: Initially (a), the regional network graph has extraneous and missing edges. We identify the AggCOs (orange), heuristically refine the graph to reflect the regional network (b), and add the BackboneCO connections (grey).

extracted by our regexes using rDNS names for the router’s interfaces. If a router did not have a most common CO tag among the rDNS for its interface addresses, we removed the CO mapping from any address in the router group with rDNS, to avoid inconclusive and potentially inaccurate mappings. We provide more details for how we mapped IP addresses to COs in Appendix B.1.

5.2 Phase 2: Build CO-Topology Graphs

Using the CO mappings, we extract CO edges from traceroute paths, where immediately adjacent routers in a path map to different COs, and construct initial graphs of the topology for each region. The rest of this phase processes the topology graphs to more accurately reflect regional topologies. This phase (1) removes false inter-region edges; (2) identifies the AggCOs; (3) removes false edges between EdgeCOs; (4) adds missing edges from AggCOs to EdgeCOs; and (5) infers the entry points into each region.

5.2.1 Remove False Inter-Region Edges. Large collections of traceroute paths likely contain some random noise [63], so we discard all edges that appear only in a single traceroute as anomalous. Next, we remove edges that appear to interconnect COs in different regions. While some links cross region boundaries (§5.2.5), many of these links result from outdated rDNS that our alias resolution did not catch. Further details are provided in Appendix B.2.

5.2.2 Identify AggCOs. Visually inspecting the regional graphs, such as the graph in Fig. 6a, showed two features of interest. First, they signaled a hierarchical structure (partially obscured by extraneous intra-region edges), where a few COs appeared responsible

for aggregating connectivity to the rest of the Internet for other COs in the region. We devised a heuristic to distinguish AggCOs from EdgeCOs based on the number of outgoing edges for each CO in the graph. In each region, we infer AggCOs as those with a higher out-degree than the average out-degree for that region plus one standard deviation. This heuristic identified the two orange COs in Fig. 6a, COs 1 and 2, as the AggCOs in this region.

Second, the graphs naturally grouped these AggCOs; when a EdgeCO had more than one incoming edge, the two preceding AggCOs had outgoing edges to nearly identical sets of EdgeCOs. Topologically, this looks like a dual-star topology, with the implication that each AggCO in a subregion directly connects with every EdgeCO in the same subregion. The natural groupings also provide new insights into the physical topology, indicating which geographic regions rely on the same AggCOs, and the level of redundancy to each EdgeCO. Appendix B.3 details how we identified AggCOs, removed false edges, and added missing edges.

5.2.3 Remove False Edges between EdgeCOs. We knew that regional networks use a ring to connect an AggCO to its EdgeCOs (§2), so the star topology indicated that the fiber ring running from an AggCO to its EdgeCOs bundles separate fiber pairs for each AggCO-to-EdgeCO connection. These fiber pairs create two separate point-to-point connections between an EdgeCO and its AggCOs, bypassing all other EdgeCOs on the ring and eliminating the need to directly connect EdgeCOs to other EdgeCOs. While the graph in Fig. 6a shows that most EdgeCOs only connect to AggCOs, some appear connected to other EdgeCOs, such as edges $9 \rightarrow 12$ and $3 \rightarrow 4$. These edges from EdgeCOs likely result from uncorrected stale rDNS, and we remove them to conform with a fiber ring (Fig. 6b).

5.2.4 Add Missing Edges From AggCOs to EdgeCOs. When an EdgeCO lies along a fiber ring with one of the AggCOs, it will connect directly with the other AggCO on the ring as well. Otherwise, the fiber pairs would bypass the EdgeCO in only one direction. We therefore assume that missing edges, e.g., from AggCO 1 to node 16 in Fig. 6a, likely result from missing rDNS.

Resolving missing edges first requires us to identify the AggCOs on the same fiber rings. Our intuition is that AggCOs on the same ring will directly connect with the same EdgeCOs, so we devised a heuristic that looks for AggCOs that overlap at least 75% of their connections with EdgeCOs, strongly suggesting that they aggregate traffic on behalf of the same EdgeCOs. We then add edges to the graph such that all AggCOs on the same last-mile fiber ring connect to the same set of EdgeCOs. In Fig. 6b, we add the missing edge from AggCO 1 to node 16.

5.2.5 Infer Entry Points Into Each Region. Finally, we add edges back into the graph that cross regional network boundaries, such as BackboneCO entry points and entry points from other nearby regions, but only when overwhelming evidence implies their existence. Returning to the traceroute paths, we extract all triplets of the form $(CO_i, REGION_1) \rightarrow (CO_j, REGION_2) \rightarrow (CO_k, REGION_2)$, where CO_i and CO_j appear in consecutive hops, indicating they directly connect. Given the hierarchical structure of the regional topologies, we only include potential entry points when they appear to lead to EdgeCOs in the region. To avoid misinterpretations caused by stale

Figure 7: The 6 Charter regions include more COs than the 28 Comcast regions.

rDNS, we only include an entry point if we observe it leading to two or more COs in the same region.

Our analysis reveals that all regions in Charter, and all but three of the Comcast regions, connect to at least two BackboneCOs. A Comcast network operator told us that nearly every Comcast region directly connects to two BackboneCOs, so we likely missed three entry points in addition to the 57 backbone entry points we observed across the Comcast regions. In some regions we observe backbone connections and a direct connection to another region; e.g., the Central California region in Comcast appears to connect to two BackboneCOs and the San Francisco regional network. We did not observe direct inter-region connections in Charter.

5.3 Contrasting Comcast and Charter

The key difference between Comcast and Charter is the number of regions they use, impacting the size of the regions and the extent of the aggregation inside each region. We observed only six Charter regions compared to 28 Comcast regions, but the Charter regions tend to cover more geographic area than the Comcast regions; e.g., Charter's Midwest region appears to touch 10 different U.S. states. Thus, a Charter region contains far more COs than a Comcast region (Fig. 7a). Charter also uses more aggregation, and far more AggCOs per region (Fig. 7b), than Comcast, where we define an AggCO as any CO with outgoing edges.

Fig. 8 and Table 1 show the different types of aggregation we observed in Comcast and Charter. The smaller regions often used a single AggCO, small to mid-size regions used two AggCOs for greater redundancy, and the largest regions used multi-layer aggregation where lower aggregation levels might include one or two AggCOs. In the multi-layer topologies, Comcast nearly always connects EdgeCOs to multiple AggCOs, while Charter uses a mix. Charter's choices in aggregation lead to less redundancy to the EdgeCOs than in Comcast; 37.7% of EdgeCOs in Charter connect to only one upstream CO vs 11.4% in Comcast (see Appendix B.4 for important context).

Outside of one Charter region and one Comcast region, the difference in region size does not appear to manifest in greater entry points per regions. In §5.5, we find higher latency to the COs in the Charter regions, likely due to the combination of fewer entries per EdgeCOs and more aggregation lengthening the distance

Figure 8: Three regional access network types.

Aggregation Type	Comcast	Charter
Single AggCO (Fig. 8a)	5	0
Two AggCOs (Fig. 8b)	11	0
Multi-level aggregation (Fig. 8c)	12	6

Table 1: Network types observed in Comcast/Charter.

from the backbone to EdgeCOs. The fewer entries, additional aggregation, and less redundancy to the EdgeCOs that we observe in Charter could also increase the potential points of failure that could disconnect EdgeCOs, and customers connected to those EdgeCOs.

5.4 Validating with Network Operators

We spoke with a network operator at Comcast and an operator at Charter to discuss our topology graphs and interpretations. The two people we spoke to are not experts for each region in their networks, however, both Comcast and Charter use a rough template for all of their regions, with AggCOs connected to EdgeCOs via fiber rings, and use similar technologies in each region.

The Comcast operator had extensive knowledge of one of the largest Comcast regions. We showed the operator our graph of that region, along with a list of the COs that we discovered, and the backbone entries. The operator confirmed that our inferred graph of the region was correct; the graph contained the COs, the second region with its own AggCOs that connect to the first region's AggCOs but not to the backbone, and the correct PoPs connected to the AggCOs. Finally, the operator confirmed that the largest Comcast regions often have two sets of AggCOs—one set connected to the backbone and another set connected to the first set—where each set connects to different EdgeCOs.

The Charter operator was not an expert for any specific Charter regions, but understood their design and general topology from the operator's experience with the Charter backbone. The operator thought we provided a reasonable representation of the regions and the regional topologies, but could not indicate if any COs were

Figure 9: The median RTTs to Massachusetts, Connecticut, Vermont, and New Hampshire in the U.S. from the largest public cloud providers. Connecticut has higher latency being geographically closest to the cloud datacenters.

missing or superfluous. Importantly for our analysis, the operator confirmed that the Charter regions are vast, with layers of fiber rings with their own AggCOs.

We confirmed with both operators that they use fiber rings with star topologies—separate fiber pairs from AggCOs to EdgeCOs—as we inferred in §5.2.1, rather than a ring topology. One network operator informed us that they chose this physical topology because it makes network upgrades simpler. We also asked both operators if the regions contain backup paths that traceroute might not observe. Both operators confirmed that all paths and COs are active, and cited the prohibitive cost of maintaining backup fiber paths or COs as the reason. This implies that traceroute can reveal all of the paths through the regional network, provided the VPs can exhaust the possible entries into the region.

5.5 Impact of Aggregation on Latency

The regional topologies help us better understand the inherent latency limitations imposed by the location of entry points and the aggregation in the regional topologies. To observe RTTs to different EdgeCOs, we conducted 100 pings from a virtual machine (VM) in every U.S. cloud region for Amazon AWS, Microsoft Azure, and Google Cloud to every EdgeCO IP address included in our graphs. Then, we identified the closest location with the lowest minimum RTT to the highest number of EdgeCOs in a region.

Fig. 9 provides the median of the minimum RTTs from the clouds to Comcast EdgeCOs in four states in the Northeast U.S.; in all three clouds the closest location was in Northern Virginia. Massachusetts, Vermont, and New Hampshire all use the same AggCOs in the Boston area, connected to BackboneCOs in New Jersey and New York, and a special purpose PoP in Boston. Surprisingly, although Connecticut is geographically closer to Northern Virginia than the other states, it has worse latency than Massachusetts and New Hampshire. The reason is that the Connecticut regional network does not have its own backbone entries; instead, its AggCOs connect to the backbone through the Massachusetts AggCOs, resulting in a 3.5 ms to 4 ms RTT penalty.

While the aggregation tends to increase latency to EdgeCOs, it presents opportunities to bring cloud applications closer to users without placing edge computing infrastructure in every EdgeCO. Conventional wisdom holds that certain classes of applications, such as augmented or virtual reality, require less than 5 ms of

Figure 10: Although more than 80% of EdgeCOs are more than 5 ms RTT from the nearest cloud VM (a), more than 80% of the EdgeCOs are within 5 ms RTT of their AggCO (b).

Figure 11: EdgeCO in a Comcast regional network.

latency [44], but more than 80% of the Comcast EdgeCOs and 90% of the Charter EdgeCOs have an RTT greater than 5 ms (Fig. 10a) from the nearest cloud location. One approach is to push edge computing to the EdgeCOs, ensuring nearly all users are within the latency constraints, but increasing the cost and complexity of deployment. Another approach could exploit the hierarchy in the regional topologies and place the edge computing infrastructure in the AggCOs. Counting any CO with an outgoing edge as an AggCO, we observe 7.7x as many EdgeCOs as AggCOs across all regions of Comcast and Charter. More than 80% of the EdgeCOs for Comcast and Charter are within 5 ms RTT of the AggCOs, likely bringing the vast majority of regional network customers within the 5 ms requirement (Fig. 10b). Furthermore, the AggCOs are often substantial datacenters, with the security, power, and capacity to include edge computing infrastructure, while EdgeCOs might be houses on residential streets (Fig. 11).

6 CASE STUDY: AT&T

Next, we investigate the topology of AT&T’s wireline regional access networks (Fig. 12). Compared to the cable providers AT&T’s network is relatively opaque. AT&T provides rDNS for their BackboneCO routers but not for other CO routers, and provides rDNS that identify their last-mile IP-DSLAMs and ONTs. However, there is no straight-forward way to identify which IP addresses AT&T assigns to their wired customers or mobile customers, or which addresses they delegate to other networks. While AT&T’s EdgeCO and AggCO router topology can be uncovered by traceroute, this is only

Figure 12: Architecture of AT&T’s access network. Routers in AggCOs and EdgeCOs are unnamed. The backbone router is in the “sd2ca” region.

possible within their respective regional networks. Our methodology for AT&T is similar to §5 at a high-level: we (1) build and annotate router-level topology graphs, and then (2) construct CO topology graphs for each regional network. The low-level methodology details in this section are tailored to the specific challenges presented by AT&T.

6.1 Phase 1: Build Router-Topology Graphs

We bootstrap our discovery of router IPs in the COs by tracerouting from 5 Ark VPs near the region we are mapping to the IP addresses of DSLAMs/ONTs in EdgeCOs. AT&T uses rDNS to label DSLAM/ONT IPs as *.lightspeed.(CLLI).sbcglobal.net, illustrated in Fig. 12. We tracerouted to all 95,821 IPs matching this pattern rDNS in the Rapid7 rDNS dataset. AT&T’s access networks use MPLS tunnels, so these traceroutes only discovered the entry router for the tunnel (the BackboneCO router), and the exit router (an EdgeCO router), missing most routers in AggCOs. Further, the bootstrap traceroutes do not observe all EdgeCO routers, because of MPLS as well as some of DSLAMs/ONTs do not reply and perhaps some do not have rDNS. However, we found that the EdgeCO routers were allocated out of a few prefixes per region. For example, there appear to be 7 /24s used for EdgeCO router IPs in AT&T’s San Diego, CA region (indicated with “sd2ca” in the rDNS entry for the BackboneCO router in Fig. 12). To uncover AggCO routers, we use the same technique as in §5—traceroute to the observed MPLS tunnel exit router [73] in the EdgeCO. Appendix C includes further details about how we infer EdgeCOs.

Because we can only traceroute to most EdgeCO router IPs from within the same region, we build per-region lists of EdgeCO /24s to probe by associating /24s with the region tag in the BackboneCO router rDNS observed in bootstrap traceroutes. We then traceroute to all IP addresses in these prefixes in the region from a VP within the region. We also performed alias resolution to map individual IP addresses to routers, and then to EdgeCOs and AggCOs. In total, we found 37 AT&T regional networks identified in rDNS, and CAIDA Ark and RIPE Atlas had VPs available in 35 of these regions. However, even in regions where we have many VPs, those VPs are insufficient to reveal the complete topology. Comprehensively revealing the regional network topologies requires finding VPs with different paths. This is particularly important because, as we will discuss in the next phase, mapping router IPs to EdgeCOs requires at least one VP served by each EdgeCO. However, finding topologically diverse AT&T VPs in a region we want to map is

not feasible with existing sources of crowdsourced VPs (Atlas and Ark). Further, AT&T's looking glasses are not suitable VPs because they are located in AT&T's backbone network, which EdgeCO and AggCO routers do not respond to traceroute (ICMP) packets.

To significantly increase the number of VPs inside an AT&T region that we are mapping, we leverage existing network infrastructure that reveals many geographically distributed last-mile links in a region – public WiFi hotspots. Our insight is that many fast food restaurant chains (e.g., McDonald's, Starbucks, and Subway) have many geographically distributed last-mile links, to many EdgeCOs, serving their WiFi hotspots. We call this approach *McTraceroute*. We believe this is the first network topology measurement effort that has made use of geographically distributed WiFi hotspots.

To evaluate how well this technique improved our visibility of a region, we ran traceroutes from all 58 McDonald's in AT&T's San Diego region to all IP addresses in seven /24s that we inferred to contain AT&T's San Diego EdgeCOs and AggCOs. We found 23 McDonald's that used AT&T for their free WiFi services. The diverse location of McDonald's restaurants, whose locations are strategically selected to maximize coverage in an area, provided us opportunities to connect to, and perform measurements from, many EdgeCOs in the regional network.

Next, we investigate how many new paths we observed with each type of VPs in San Diego, to determine if *McTraceroute* significantly increases the number of paths we observe. Considering traceroute paths starting with the second hop, the eight Atlas and two Ark probes in AT&T's San Diego respective regions revealed only half of the IP paths we observed with *McTraceroute*. This indicates that increasing the number of VPs revealed many more paths, despite overlap in the McDonald's EdgeCOs. Note that because the network is opaque, we do not know the true number of paths, however *McTraceroute* provides a significant increase in observed paths over existing VPs.

6.2 Phase 2: Build CO-Topology Graphs

Phase 1 produced the router-level topology shown in Fig. 13a. We inferred two backbone routers (blue), four aggregation routers (red), and 84 EdgeCO routers (black). We inferred the EdgeCO routers as connected redundantly to two aggregation routers each, and all aggregation routers connected to one backbone router. The router-level topology reveals a three-level structure, with two sub-regions that use different aggregation routers.

To infer the CO-level topology, we first map last-mile links to EdgeCO routers. Each last-mile link is served by a single EdgeCO, so if two routers are one hop away from the same last-mile link, we conclude they are both in the same CO. We observed each last-mile link connected to two EdgeCO routers, indicating that each EdgeCO has two routers.

We observed two backbone routers, and both appear fully connected to all aggregation routers. This is unlike the cable networks, where we observed backbone routers connected to one aggregation router. We conclude from this inference that AT&T has only one BackboneCO in this region, and this office contains both core routers. We are less confident about the four aggregation routers, but the highest resilience design would have them operating out of four different COs. Fig. 13b shows the inferred CO-level topology.

(a) Raw router-level topology from *McTraceroute*. We probe from behind the yellow "leaf nodes" (IP-DSLAMs) at the bottom.

(b) Inferred CO-level topology

Figure 13: AT&T San Diego Regional Network

Latency:	3-4ms	4-5ms	5-6ms	6-7ms	9-10ms
EdgeCOs:	5	19	7	2	2

Table 2: Latency from Google Cloud VPs to EdgeCOs in San Diego. Two have >2x the average latency (4.3ms).

6.3 Analysis of AT&T's Topology

AT&T's regional network has a significantly longer history than cable networks, dating back to the early 1900s. Therefore, we expect its structure to reflect design choices constrained by the capabilities of early of telephone networks. AT&T's long distance network, called Long Lines, only reached a single CO in each region. These Long Lines COs now appear to serve as their BackboneCOs. In fact, the BackboneCO we inferred in San Diego still has a Long Lines microwave tower on its roof.

Aggregation. We observed significantly higher EdgeCO density in AT&T's network than in the cable providers we studied. In Charter's San Diego sub-region we observed 16 EdgeCOs, compared to 42 in AT&T's San Diego region. This CO deployment density is consistent with AT&T facing the constraint of local copper telephone service loop lengths. By the time cable networks emerged in the 1990s, Hybrid Fiber Coax allowed for much longer last-mile links from EdgeCOs to customers. We would thus expect, without considering other factors, that AT&T aggregates fewer last-mile links to each EdgeCO than do cable networks. This lower ratio of customers to EdgeCO helps to reduce the scale of outages when an EdgeCO fails (e.g., due to fiber cuts or failed equipment).

To estimate latency differences between EdgeCOs in the San Diego region, we conducted traceroutes from a VM in a Los Angeles Google Cloud datacenter to all of the end user AT&T IP addresses

we could find in the San Diego region. We used Measurement Lab data [39] to extract AT&T customer IP addresses from NDT measurements, and retained the subset of addresses located in San Diego or Imperial County according to NetAcuity [1], a commercial geolocation service. Using only traceroutes that passed through the BackboneCO in San Diego and reached the customer addresses, we inferred that the penultimate traceroute hop corresponds to a device in an EdgeCO.

We could not directly ping these devices, but we could elicit responses by sending an ICMP Echo packet to a customer IP address with the TTL field set to the penultimate probe TTL in the traceroute to that address. To measure latency from Google Cloud in Los Angeles to the EdgeCOs, we conducted 100 probes to each EdgeCO address observed in the traceroutes and used the minimum observed RTT (Table 2). These results show that some EdgeCOs have significantly less latency to the BackboneCO than other EdgeCOs in the region. Two distant EdgeCOs—with connected customers geolocated to Calexico and El Centro, CA—had over twice the average latency of 4.3ms to Google (9–10 msec). This disparity suggests that some AT&T customers will suffer considerably higher latency to cloud services than other users in the region.

Redundancy. AT&T’s network in San Diego has a similar lack of redundancy that we observed in some cable provider regions; namely, the use of only one BackboneCO. In AT&T’s network, these BackboneCOs are fortified for natural disasters, such as Category 5 hurricanes. However, the Christmas 2020 attack on AT&T’s Nashville office, which we assume is the lone BackboneCO in Nashville, took down the entire region, consistent with our inferred topology. Relative to the cable providers, AT&T appears to have more redundancy in their BackboneCO to AggCO paths, with all backbone routers connecting to all Agg routers. Our measurements cannot detect whether these paths take diverse fiber paths.

Validation. Aspects of our inferences match historical documents describing AT&T’s telephone network in San Diego. AT&T’s access network was likely built using these same COs. The first document [5] states that AT&T operates one tandem building in San Diego (CLLI SNDGCA02), consistent with the single BackboneCO that we inferred. The documents also describe 42 subtending COs in San Diego, we believe these match the ~40 EdgeCOs we inferred. A second document [60] shows four “Inter-office” COs in San Diego’s network, we believe this term is AT&T’s term for AggCOs.

7 CASE STUDY: MOBILE CARRIERS

For mapping the regional access networks of all three major mobile carriers, we focused on the portion of the network that bridges the mobile packet core with the rest of the Internet (and edge services). Some mapping challenges are similar those of AT&T’s wireline network: they have no rDNS on routers, and probing requires internal vantage points. However, mobile networks face a significant additional challenge: they have no distributed VPs to provide internal views of the providers’ regional networks. Although, they also present a unique opportunity to observe nationwide network topology: unlike the wireline transparent networks, we can physically move mobile VPs to probe inside different regions. Building on this insight, we introduce a new parcel-based measurement technique,

Figure 14: Improving scamper’s traceroute efficiency

Figure 15: Shipping to 12 destinations covered 40 states

ShipTraceroute, to obtain national coverage of mobile access network regions. Then, we use the large geographically-tagged dataset of traceroutes we collected to infer the topology of the networks.

7.1 Phase 1: Collect router-level topology

We developed ShipTraceroute, a smartphone-based network measurement technique that can send traceroutes from a battery-powered Android device for a prolonged period of time while being shipped inside a truck or railcar. Appendix A.1 describes how shipment of a smartphone running this software complies with U.S. regulations for items shipped in a parcel inside of a truck or train.

7.1.1 Topology collection. We shipped three Samsung Galaxy A71 smartphones (one for Verizon, AT&T, and T-Mobile) to 12 locations in the U.S. The shipment paths traversed 40 states (Fig. 15). During the shipments, the devices attempted to perform a round of traceroutes once per hour. However, signal conditions varied significantly along the routes. Some areas had too weak of a signal inside of the vehicle to perform the traceroutes, particularly in areas where there are no inhabitants. We observed the following success rates for rounds of traceroutes during the journey: 1592/1948 (82%) on AT&T, 1720/2054 (84%) on Verizon, and 872/1153 (75%) on T-Mobile.

The destinations for each round of traceroutes were IPv6 and IPv4 addresses in ASes neighboring the mobile carriers’ networks.¹ The

¹We used Zayo’s AS for T-Mobile because T-Mobile does not have its own IPv4 AS and T-Mobile’s primary backbone provider is Zayo.

reason why we used external destinations (in neighboring ASes) to map mobile access networks, rather than internal addresses like we used in wireline networks, is because mobile networks block traceroutes to internal infrastructure. We used destinations in all neighboring ASes to try and traverse all of the carriers' BackboneCOs in each region (details are in Appendix D). However, quickly we discovered that traceroutes to all of the destinations took the same path inside each of the mobile access networks, allowing us to reduce to a single destination per provider.

We also observed that the path through the mobile network did not change as the phone moved within a region. We found we needed to force the phone to re-register with the core network by putting it into airplane mode before each round of traceroutes to route through all region's EdgeCOs and packet gateways.

Since GPS signals are rarely available inside of shipping vehicles, we logged the device's CELLID each time we started a round of traceroutes. We then converted the CELLID to a geolocation using the OpenCellID public cellular tower geolocation database [4].

7.1.2 Making mobile tracerouting energy efficient. We designed the measurement software on our smartphone to prolong battery life. The goal was to ship the phone by ground transport across the U.S.—a journey that takes about one week—while running measurements each hour, without the battery emptying.

We achieved this without sacrificing measurement fidelity by making two modifications to *scamper* [37]—ShipTraceroute's network probing tool. First, we modified scamper so that it could conduct measurements without without rooting the phone (rooting can disable thermal safeguards). Second, we reduced scamper's energy consumption by modifying its traceroute implementation to send probes to multiple consecutive hops in parallel. This significantly reduces the time that scamper spent waiting for unresponsive hops, and thus reduced the time the phone's radio is fully powered.

We evaluated the energy efficiency of our modified scamper implementation by measuring the energy consumption of a Samsung Galaxy A71 5G performing traceroutes to the 266 IPv4 and IPv6 destinations in AT&T. To measure the device's energy consumption, we fully charged it, and instrumented with a USB-C power monitor on its charging port. This allowed us to measure the energy needed to operate the device. Fig. 14 shows how much we improved energy efficiency: we achieved a 38% reduction in energy from 8.6 mAh with off-the-shelf scamper to 5.3 mAh with ShipTraceroute's scamper. As a result of these improvements, we calculated that our phone can perform hourly traceroutes for ~12 days on one charge, a gain of ~4 days over the off-the-shelf implementation. The other main contributors to power consumption are the energy consumption required to exit airplane mode when we start a measurement (1.4–2.6 mAh), and the trickle of energy consumed when the phone is asleep and in airplane mode between measurements. Although we put the device in airplane mode between traceroute rounds to force it to re-register in the packet core, it also has the additional benefit of reducing energy consumption (14.5 mAh vs. 9 mAh in airplane mode for every 55 minutes asleep).

7.2 Phase 2: Inferring CO-level topology

Each traceroute collected in phase one revealed a path from the mobile packet gateway (the first hop) until the packet reaches the

BackboneCO. However, it is difficult to infer CO-level topology from these traceroutes because mobile networks have extremely limited rDNS (only Verizon has rDNS).

Fortunately, IPv6 is now widely deployed in cellular networks, and IPv6 addresses' are long enough that providers can encode information in them about where those addresses reside in the topology of their access network. Indeed, we found an early discussion about how to set IPv6 prefixes for LTE infrastructure that described how bits in addresses can be used to indicate what those addresses are used for—infrastructure or users—and what their location is in the network topology [15]. With the large number of geo-tagged samples of IPv6 router addresses in the traceroutes we collected, we looked for patterns in how the bits in the addresses change as the mobile device moves.

7.2.1 AT&T. Fig. 16(a) shows the patterns we observed in AT&T's addresses in their traceroutes. The user address and first hop (packet gateway) /32 prefix are consistent throughout the country, indicating it is the general AT&T mobile user prefix. User addresses also have a more specific /40 prefix that only changes 11 times as we move around the country. This prefix also changes simultaneously with bits 32–47 of the router addresses. We believe this prefix indicates the EdgeCO (and region) that is in use by the device, indeed this prefix can be used to route to the correct BackboneCO router to reach the user. For validation of this result, we discovered an AT&T document from 2014 that also lists 11 mobile datacenters in the U.S. [6]. However, we suspect that as they roll out their 5G network, they will add more EdgeCOs and thus reduce the size of their regions.

The /32 prefix of the rest of hops before leaving AT&T's mobile network are always the same, and different from the user address, so we infer they are the general prefix for AT&T infrastructure (i.e., routers). We observed bits 48–52 of these addresses cycling through several values inside each inferred EdgeCO, and they changed at the same time as bits 32–40 of the user address. Also, these bits changed each time we woke up from airplane mode and re-attached to the cellular network. Therefore, we infer these bits indicate the current packet gateway in the EdgeCO that user is attached to. Table 7 (Appendix D) shows the inferred infrastructure in each region.

7.2.2 Verizon. Fig. 16(b) shows the patterns we observed in Verizon's addresses. All of the first 10 hops are within Verizon's network, but only the first (packet gateway) and the last four hops appear in the traceroutes. The /24 prefix of the user address and first hop stays the same throughout the country, indicating this is Verizon's user address prefix. As the device moves, more specific bits change. The /32 prefix changes 18 times, and the /40 prefix changed 32 times—both were stable within contiguous geographic regions. The /32 prefix changed less frequency in a geographic area than the /40 prefix. One plausible explanation for this behavior is that the /32 prefix identifies the BackboneCO and the /40 identifies the EdgeCO using that BackboneCO. We also observed bits 40–43 in the user address can change when we cycle airplane mode, while other hops stay the same, indicating multiple packet gateways in each EdgeCO.

This explanation is supported by information from the rest of the hops (i.e., infrastructure). The /32 prefix in the user address is likely to represent the BackboneCO because it corresponds with changes in the rDNS of the Verizon backbone hop (i.e., al1ter.net).

Figure 16: Topological hints for mobile networks encoded in IPv6 addresses.

The /40 prefix is likely to be the EdgeCO because when it changes, so do the bits of the addresses in the other hops—the hops to reach the EdgeCO from the BackboneCO—namely, bits 64–75 in all of the infrastructure hops, and bits 32–39 in some of the hops. Table 8 (Appendix D) shows the inferred infrastructure in each region.

Although we were unable to find documentation that validates our inferred topology, we performed several controlled experiments to test our inferences. First, we found Verizon Wireless deploys speedtest servers in their EdgeCOs which contain the names of the EdgeCOs in their rDNS. For example, `cavt.ost.myvzw.com` is the speedtest server in the Vista, California EdgeCO. We performed a controlled drive north from San Diego to Irvine while tracerouting to all of the speedtest servers, and we observed that when the shortest traceroute path switched from the Vista, CA to the Azusa, CA speedtest server, the expected bits in the traceroute hops changed at the same time. Additionally, we performed a long-running stationary experiment verify if the EdgeCO and BackboneCO address bits were stable in a location in San Diego. Indeed, they were generally stable across multiple days, however we did observe a small number of switches to the neighboring EdgeCO connected to the same BackboneCO. This implies the packet core connects to both EdgeCOs and it can switch between them if necessary for load balancing or redundancy.

7.2.3 T-Mobile. Fig. 16(c) shows the patterns we observed in T-Mobile’s addresses. Similar to the other two providers, the user IP prefix /32 stays the same across the entire country. The /40 prefix of the user IP can change each time it leaves airplane mode within a geographic area roughly the size of a city. These /40s are cycled through in a somewhat round-robin fashion, indicating that bits 32–39 likely represent the packet gateway. However, we observed that T-Mobile also cycles through different BackboneCO providers, suggesting that T-Mobile has a different mobile access network topology than the other providers. We infer that T-Mobile has a set of packet gateways in each region, possibly in different EdgeCOs, and with different backbone providers. These packet gateways are likely interconnected by the packet core.

We confirmed with T-Mobile that they have several backbone providers serving each region at different interconnection points within the region. Also a device in one location connects to different packet gateways at different sites (i.e., EdgeCOs), but that they

Figure 17: Inferred Internet topologies of U.S. mobile carriers

prefer the closest site. Their network is designed in this distributed fashion for lower latency and resiliency. Therefore a device can wake up connecting to a different packet gateway than it connected to before it went to sleep.

7.2.4 Summary. We infer topologies among the three providers’ access networks (Fig. 17). AT&T appears to have a single EdgeCO with multiple packet gateways connected to their nearest backbones. Verizon has multiple EdgeCOs sharing the same backbone CO, but the EdgeCOs cover non-overlapping regions. T-Mobile has multiple EdgeCOs in one region, but does not aggregate traffic to a single backbone, rather they aggregate to a variety of backbone providers directly connected to the EdgeCOs. These designs have different tradeoffs. AT&T’s design may be more cost-efficient because equipment and links are centralized to a single EdgeCO per region. However, the lack of diverse CO locations may increase latency. Verizon and T-Mobile appear to have lower latency in part because they have multiple EdgeCOs per region.

7.3 Comparison of US Mobile Access Networks

The topology of mobile access networks has implications for network latency, because user traffic has to traverse to the backbone PoP of the region to reach other Internet hosts. Fig. 18 shows the *minimum latency* we measured from our ShipTraceroute smartphone in different locations to a server located at CAIDA in San Diego. The hexagons indicate where we captured latency: the darker the color, the higher the minimum latency to the server from that

Figure 18: Minimum latency from each location to a single server in San Diego. Colored regions indicate the measurements were handled by the same EdgeCO (inferred from IPv6 addresses). T-Mobile does not aggregate traffic to a single EdgeCO.

location. The colored regions containing multiple hexagons indicate those latency samples were taken from the same mobile access network region according to the IPv6 bit fields for that provider.

AT&T’s regions are much larger than Verizon and T-Mobile, therefore some geographic areas (e.g., Montana and North Dakota) incur high latency to San Diego, due to circuitous paths to the BackboneCO. Verizon’s network generally had lower latency (Fig. 18b), because the larger number of EdgeCOs likely provided shorter average distance to BackboneCOs. As T-Mobile’s distributed topology relies on several backbone providers (Fig. 18c), they had latency similar to Verizon. However, we observed unusually high latency near the border of Florida and Louisiana (Fig. 18c), because during the experiment the device in these regions attached to a distant EdgeCO in South Carolina.

8 FUTURE WORK

Resiliency. The tools and methods we have developed for inferring regional topologies enable a new approach to studying resilience across space and time. The topological differences that we have already observed across different regions have strong implications for resilience against disasters. A promising next direction is to combine these topologies with existing or future data sets on resilience of connectivity.

Edge Computing. Understanding the topology of these regional access networks, and associated performance implications, may be the key to realizing the unachieved potential of the long-hyped edge computing paradigm [28, 59]. In addition to discovering the pyramid structure of the Edge CO and Agg CO topologies, our latency measurements suggest that the AggCO is typically less than 10 msec from both the cloud and customers in the region, which meets the AR/VR latency requirement for edge computing [47]. This result suggests that putting edge computing infrastructure in Agg COs is the most efficient solution. Efforts to offload computation from mobile devices [40] can also leverage an understanding of the effect of distributed EdgeCOs on latency to the cloud.

Scalability of measurement methods. There is opportunity for improving scalability and manageability of our methods. For the AT&T study, we drove to each McDonald’s location in San Diego, connected to their WiFi, and collected traceroutes. This approach is a fun adventure for a graduate student, but operationalizing such a measurement requires crowd sourcing. We could develop an app

that connects to public WiFi spots (while the user waits for their food order), and provides a reward for uploaded results.

We also envision ways to improve the scalability of ShipTraceroute. Besides sending more cellular packets in parallel to save energy, we can arrange for the device to sleep even more between measurements. During a cross-country shipment, a device often stops at a hub for about a day. We could use the device’s accelerometer to pause measurements when the device is at rest.

9 CONCLUSION

We have undertaken a comprehensive measurement study of the topology of U.S. regional access ISPs. Our motivation was to extract insights about architectural choices that ISPs make for how to aggregate traffic, and then empirically assess implications of those insights for the resilience and evolution of the Internet ecosystem. Growing interest in edge computing and 5G co-location, not to mention the pandemic-induced semi-permanent transition to working from home, is placing increasing pressure on these regional networks. We are now entirely dependent on this infrastructure but there has been little attention to independent objective understanding of its resilience and reliability.

This dearth of attention is understandable. While perhaps not the most opaque part of the Internet, these networks are not amenable to straightforward measurement and analysis. Our tools have their limitations, but they allowed us to make surprisingly accurate maps in spite of considerable noise in our input signals, e.g., missing or incorrect DNS or traceroute hops. We were able to identify many different approaches to provisioning redundancy, across links, nodes, buildings, and at different levels of the hierarchy. These measurements can provide a basis for reasoning about sources of performance and reliability impairment in these networks. We believe that sharing our methods, lessons, and results will inform future analysis of critical infrastructure.

ACKNOWLEDGEMENTS

We thank our shepherd Matt Calder and the anonymous reviewers for their insightful comments. This work was supported in part by National Science Foundation grants CNS-2105393, CNS-1901517, and OAC-1724853. Also this work was supported by DARPA grant CA HR00112020014.

REFERENCES

- [1] [n.d.]. Netacuity. <https://www.digitalelement.com/solutions/>.
- [2] [n.d.]. Netalyzr. https://play.google.com/store/apps/details?id=edu.berkeley.icsi.netalyzr.android&hl=en_US.
- [3] 2019. RIPE Atlas. <https://atlas.ripe.net/landing/measurements-and-tools/>.
- [4] 2021. OpenCellID. <https://www.opencellid.org/>.
- [5] ATT. 2008. West CA Tandem and subtending end offices. <https://tinyurl.com/7pyywwp7>.
- [6] AT&T. 2014. DRAFT AT&T Iperf Mobile Application User Guide. <https://pdfslide.net/documents/draft-att-iperf-mobile-application-user-guide-aka-iperf-commands-205pdf.html>.
- [7] Vaibhav Bajpai, Steffie Jacob Eravuchira, and Jürgen Schönwölder. 2017. Dissecting Last-mile Latency Characteristics. *ACM Computer Communication Review (CCR)* 47 (2017), 25–34.
- [8] Zachary S. Bischof, John S. Otto, Mario A. Sánchez, John P. Rula, David R. Choffnes, and Fabian E. Bustamante. 2011. Crowdsourcing ISP characterization to the network edge. In *Proc. ACM SIGCOMM*.
- [9] Ilker Nadi Bozkurt, Waqar Aqeel, Debopam Bhattacharjee, Balakrishnan Chandrasekaran, Philip Brighten Godfrey, Gregory Laughlin, Bruce M. Maggs, and Ankit Singla. 2018. Dissecting Latency in the Internet's Fiber Infrastructure.
- [10] CAIDA. [n.d.]. Archipelago. <https://www.caida.org/projects/ark/>.
- [11] CAIDA. [n.d.]. AS-relationships. <https://publicdata.caida.org/datasets/as-relationships/>.
- [12] CAIDA. [n.d.]. DNS Decoded (DDec). <http://ddec.caida.org>.
- [13] Igor Canadi, Paul Barford, and Joel Sommers. 2012. Revisiting broadband performance. In *Proc. ACM Internet Measurement Conference (IMC)*. ACM Press.
- [14] Abhijnan Chakraborty, Vishnu Navda, Venkata N. Padmanabhan, and Ramachandran Ramjee. 2013. Coordinating Cellular Background Transfers using LoadSense. In *Proc. ACM Conference on Mobile Computing and Networking (MobiCom)*.
- [15] Cisco. 2011. LTE Design and Deployment Strategies. https://www.cisco.com/c/dam/global/en_ae/assets/expo2011/saudi-arabia/pdfs/lte-design-and-deployment-strategies-zeljko-savic.pdf.
- [16] Crown Infrastructure Partners. 2021. Ultra fast broadband. <https://www.crowninfrastructure.govt.nz/ufb/what/>.
- [17] Mike Dano. 2019. An Inside Look at Verizon's Edge Computing Capabilities. <https://www.lightreading.com/the-edge/an-inside-look-at-verizons-edge-computing-capabilities>.
- [18] Marcel Dischinger, Andreas Haeberlen, Krishna P. Gummadi, and Stefan Saroiu. 2007. Characterizing residential broadband networks. In *Proc. ACM Internet Measurement Conference (IMC)*.
- [19] Dittrich, David and Kenneally, Erin and Bailey, Michael. 2013. Applying Ethical Principles to Information and Communication Technology Research: A Companion to the Menlo Report. <http://ssrn.com/abstract=2342036>.
- [20] Wei Dong, Zihui Ge, and Seungjoon Lee. 2011. 3G meets the Internet: Understanding the performance of hierarchical routing in 3G networks. In *Proceedings of International Teletraffic Congress*.
- [21] Ramakrishnan Durairajan, Paul Barford, Joel Sommers, and Walter Willinger. 2015. InterTubes: A Study of the US Long-Haul Fiber-Optic Infrastructure. In *Proc. ACM SIGCOMM*.
- [22] Navid Ehsan, Mingyan Liu, and Roderick J. Ragland. 2003. Evaluation of performance enhancing proxies in internet over satellite. *International Journal of Communication Systems* 16 (May 2003).
- [23] FCC. 2020. June 15, 2020 T-Mobile Network Outage Report.
- [24] Romain Fontugne, Anant Shah, and Kenjiro Cho. 2020. Persistent Last-mile Congestion: Not so Uncommon. In *Proc. ACM Internet Measurement Conference (IMC)*.
- [25] Monia Ghobadi and Ratul Mahajan. 2016. Optical Layer Failures in a Large Backbone. In *Proc. ACM Internet Measurement Conference (IMC)*.
- [26] Ramesh Govindan and Hongsuda Tangmunarunkit. 2000. Heuristics for Internet Map Discovery. In *Proc. IEEE Conference on Computer Communications (INFOCOM)*.
- [27] Matthias Gunkel, Malte Schneiders, Sascha Vorbeck, Werner Weiershausen, Ralph Leppla, Frank Rumpf, Ralf Herber, Volker Furst, and Markus Rodenfels. 2008. Aggregation networks: Cost comparison of WDM ring vs. double star topology. In *International Conference on Optical Network Design and Modeling*.
- [28] Wenlu Hu, Ying Gao, Kiryong Ha, Junjue Wang, Brandon Amos, Zhuo Chen, Padmanabhan Pillai, and Mahadev Satyanarayanan. 2016. Quantifying the Impact of Edge Computing on Mobile Applications. In *Proceedings of the 7th ACM SIGOPS Asia-Pacific Workshop on Systems*.
- [29] B. Huffaker, M. Fomenkov, and k. claffy. 2014. DRoP:DNS-based Router Positioning. *ACM Computer Communication Review (CCR)* 44 (2014), 6–13.
- [30] Cisco Systems Inc. 2000. Introduction to DWDM Technology. https://www.cisco.com/c/dam/global/de_at/assets/docs/dwdm.pdf.
- [31] Tive Inc. 2019. Tive Releases Return-By-Mail Supply Chain Tracker. <https://blog.tive.co/tive-releases-return-by-mail-supply-chain-tracker>.
- [32] Kenneally, Erin and Dittrich, David. 2012. The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research. <http://ssrn.com/abstract=2445102>.
- [33] Ken Keys, Young Hyun, Matthew Luckie, and k claffy. 2013. Internet-Scale IPv4 Alias Resolution with MIDAR. *IEEE/ACM Transactions on Networking* 21 (2013).
- [34] Simon Knight, Hung X. Nguyen, Nickolas Falkner, Rhys Bowden, and Matthew Roughan. 2011. The Internet Topology Zoo. *IEEE Journal on Selected Areas in Communication* 29 (2011), 1765–1775.
- [35] Christian Kreibich, Nicholas Weaver, Boris Nechaev, and Vern Paxson. 2010. Netalyzr. In *Proc. ACM Internet Measurement Conference (IMC)*.
- [36] Feng Lu, Hao Du, Ankur Jain, Geoffrey M. Voelker, Alex C. Snoeren, and Andreas Terzis. 2015. CQIC: Revisiting Cross-Layer Congestion Control for Cellular Networks. In *Proc. International Workshop on Mobile Computing Systems and Applications (HotMobile)*.
- [37] Matthew Luckie. 2010. Scamper: a Scalable and Extensible Packet Prober for Active Measurement of the Internet. In *Proc. ACM Internet Measurement Conference (IMC)*.
- [38] Matthew Luckie, Bradley Huffaker, and k claffy. 2019. Learning Regexes to Extract Router Names from Hostnames. In *Proc. ACM Internet Measurement Conference (IMC)*.
- [39] M-Lab. [n.d.]. BigQuery QuickStart. <https://www.measurementlab.net/data/docs/bq/quickstart/>.
- [40] Xiaoqiang Ma, Yuan Zhao, Lei Zhang, Haiyang Wang, and Limei Peng. 2013. When mobile terminals meet the cloud: computation offloading as the bridge. *IEEE Network* (2013), 28–33.
- [41] Ratul Mahajan, Neil Spring, David Wetherall, and Tom Anderson. 2002. Inferring Link Weights Using End-to-End Measurements. In *Proc. ACM Internet Measurement Workshop (IMW)*.
- [42] Sathya Kumar Mani, Matthew Nance Hall, Ramakrishnan Durairajan, and Paul Barford. 2020. Characteristics of Metro Fiber Deployments in the US. In *Proc. IEEE Traffic Measurement and Analysis Conference (TMA)*.
- [43] Yun Mao, Hani Jamjoom, Shu Tao, and Jonathan M. Smith. 2007. NetworkMD: Topology Inference and Failure Diagnosis in the Last Mile. In *Proc. ACM Internet Measurement Conference (IMC)*.
- [44] Nitinder Mohan, Lorenzo Corneo, Aleksandr Zavodovski, Suzan Bayhan, Walter Wong, and Jussi Kangasharju. 2020. Pruning Edge Research with Latency Shears. 182–189.
- [45] N. Spring and R. Mahajan and D. Wetherall. 2002. Measuring ISP topologies with Rocketfuel. In *Proc. ACM SIGCOMM*.
- [46] NBN Co. 2021. NBN Multi Technology Mix. <https://www.nbnco.com.au/learn/network-technology>.
- [47] GSMA Future Network. 2019. Cloud AR/VR Whitepaper. <https://www.gsma.com/futurenetworks/wiki/cloud-ar-vr-whitepaper/>.
- [48] Juniper Networks. 2016. Metro Ethernet Design Guide. https://www.juniper.net/documentation/en_US/release-independent/solutions/information-products/pathway-pages/solutions/metro-ethernet-dg.pdf.
- [49] Ashkan Nikravesh, Hongyi Yao, Shichang Xu, David Choffnes, and Z. Morley Mao. 2015. Mobilyzer: An Open Platform for Controllable Mobile Network Measurements. In *Proc. ACM Conference on Mobile Systems, Applications, and Services (MobiSys)*.
- [50] US Department of Transportation. [n.d.]. 49 CFR §172.185 – Lithium cells and batteries. <https://www.law.cornell.edu/cfr/text/49/173.185>.
- [51] Ookla. [n.d.]. Speedtest. <http://www.speedtest.net>.
- [52] Larry Peterson, Tom Anderson, Sachin Katti, Nick McKeown, Guru Parulkar, Jennifer Rexford, Mahadev Satyanarayanan, Oguz Sunay, and Amin Vahdat. 2019. Democratizing the Network Edge. *ACM Computer Communication Review (CCR)* (2019).
- [53] Rapid7 Labs. 2021. Reverse DNS (RDNS). https://opendata.rapid7.com/sonar.rdns_v2/.
- [54] John P. Rula, Zachary S. Bischof, and Fabian E. Bustamante. 2015. Second Chance: Understanding Diversity in Broadband Access Network Performance. In *Proc. of the ACM SIGCOMM Workshop on Crowdsourcing and Crowdsourcing of Big (Internet) Data*.
- [55] John P. Rula and Fabian E. Bustamante. 2014. Behind the curtain: Cellular DNS and content replica selection. In *Proc. ACM Internet Measurement Conference (IMC)*.
- [56] Samsung. 2021. Keep your Galaxy device at its normal operating temperature. <https://www.samsung.com/us/support/answer/ANS00076952/>.
- [57] Mario A. Sánchez, John S. Otto, Zachary S. Bischof, and David R. Choffnes. 2013. Dasu: Pushing Experiments to the Internet's Edge. In *Proc. Symposium on Networked Systems Design and Implementation (NSDI)*.
- [58] Mahadev Satyanarayanan. 2017. The Emergence of Edge Computing. *Computer* 50 (2017), 30–39.
- [59] Mahadev Satyanarayanan. 2017. The Emergence of Edge Computing. *Computer* 50 (2017), 30–39.
- [60] Southwestern Bell Company (SBC). 2005. Rates and Tariffs Quote Sheet—OC3c Purchased Under the OC-n Point to Point Service Offering. <https://tinyurl.com/uwh3puzj>.
- [61] Paul Schmitt, Morgan Vigil, and Elizabeth Belding. 2016. A study of MVNO data paths and performance. In *Proc. Passive and Active Measurement Conference*.

- [62] Aaron Schulman and Neil Spring. 2011. Pingin' in the rain. In *Proc. ACM Internet Measurement Conference (IMC)*.
- [63] Rob Sherwood, Adam Bender, and Neil Spring. 2008. DisCarte: A Disjunctive Internet Cartographer. In *Proc. ACM SIGCOMM*.
- [64] Neil Spring, Ratul Mahajan, and Thomas Anderson. 2003. The Causes of Path Inflation. In *Proc. ACM SIGCOMM*.
- [65] Srikanth Sundaresan, Walter de Donato, Nick Feamster, Renata Teixeira, Sam Crawford, and Antonio Pescapè. 2012. Measuring home broadband performance. *Commun. ACM* 55 (2012), 100–109.
- [66] Srikanth Sundaresan, Nick Feamster, and Renata Teixeira. 2016. Home network or access link? Locating last-mile downstream throughput bottlenecks. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 9631 (2016), 111–123.
- [67] Srikanth Sundaresan, Renata Teixeira, Georgia Tech, Nick Feamster, Antonio Pescapè, and Sam Crawford. 2011. Broadband Internet Performance : A View From the Gateway. *ACM Computer Communication Review (CCR)* 41 (2011), 134–145.
- [68] Renata Teixeira, Keith Marzullo, Stefan Savage, and Geoffrey M. Voelker. 2003. In Search of Path Diversity in ISP Networks. In *Proc. ACM Internet Measurement Conference (IMC)*.
- [69] Ye Tian, Ratan Dey, Yong Liu, and Keith W Ross. 2012. Topology mapping and geolocating for China's Internet. *IEEE Transactions on Parallel and Distributed Systems* 24 (2012), 1908–1917.
- [70] Hoang Tran. 2012. Case Study: Ethernet Cell Site Backhaul Request for Quotation Process.
- [71] Dalibor Uhlir, Dominik Kovac, and Jiri Hosek. 2015. Multi Service Proxy: Mobile Web Traffic Entitlement Point in 4G Core Network. *International Journal of Advances in Telecommunications, Electrotechnics, Signals and Systems* 4 (2015).
- [72] Yves Vanaubel, Pascal Mérindol, Jean-Jacques Pansiot, and Benoit Donnet. 2015. MPLS Under the Microscope: Revealing Actual Transit Path Diversity. In *Proc. ACM Internet Measurement Conference (IMC)*.
- [73] Yves Vanaubel, Pascal Mérindol, Jean-Jacques Pansiot, and Benoit Donnet. 2017. Through the Wormhole: Tracking Invisible MPLS Tunnels. In *Proc. ACM Internet Measurement Conference (IMC)*, 29–42.
- [74] Chris Whitaker. 2011. The Comcast Enterprise Network Story. <https://www.slideshare.net/cwhita002/the-comcast-enterprise-network-story>.
- [75] Walter Willinger, David Alderson, and John C. Doyle. 2009. Mathematics and the Internet: A Source of Enormous Confusion and Great Potential. *Notices of the American Mathematical Society* 56 (2009), 586–599.
- [76] Daoyuan Wu, Rocky Chang, Weichao Li, Eric Cheng, and Debin Gao. 2017. MopEye: Opportunistic Monitoring of Per-app Mobile Network Performance. In *Proceedings of USENIX ATC*.
- [77] Xiufeng Xie, Xinyu Zhang, Swarun Kumar, and Li Erran Li. 2015. piStream: Physical Layer Informed Adaptive Video Streaming Over LTE. In *Proc. ACM Conference on Mobile Computing and Networking (MobiCom)*.
- [78] Xiufeng Xie, Xinyu Zhang, and Shilin Zhu. 2017. Accelerating mobile web loading using cellular link information. In *Proc. ACM Conference on Mobile Systems, Applications, and Services (MobiSys)*.
- [79] Qiang Xu, Junxian Huang, Zhaoguang Wang, Feng Qian, Alexandre Gerber, and Zhuoqing Morley Mao. 2011. Cellular data network infrastructure characterization and implication on mobile content placement. In *Proc. ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*.
- [80] Hyo-Sik Yang, M. Herzog, M. Maier, and M. Reisslein. 2004. Metro WDM networks: performance comparison of slotted ring and AWG star networks. *IEEE Journal on Selected Areas in Communications* (2004).
- [81] Kyriakos Zarifis, Tobias Flach, Srikanth Nori, David Choffnes, Ramesh Govindan, Ethan Katz-Basett, Z Morley Mao, and Matt Welsh. 2014. Diagnosing path inflation of mobile client traffic. In *Proc. Passive and Active Measurement Conference*.
- [82] Shiwei Zhang, Weichao Li, Daoyuan Wu, Bo Jin, Rocky Chang, Debin Gao, Yi Wang, and Ricky Mok. 2019. An Empirical Study of Mobile Network Behavior and Application Performance in the Wild. In *Proceedings of IEEE IWQoS*.

A ETHICS CONSIDERATIONS

A.1 Transportation of Lithium Batteries

In the US, lithium batteries are considered hazardous materials, and shipment of them must comply with regulations set out by the US Department of Transportation (i.e. US 49 CFR §172.185 [50]). The relatively low capacity of batteries in smartphones, and the fact that they are contained within equipment (i.e., the smartphone), allows for them to be shipped by ground. Rules about shipment of a powered-on devices however, are not clearly specified. We

are aware of a device with a similar operating mode that is widely in use today: shipment tracking devices that use Cellular radios and GPS to report package locations during shipment. One is even available directly from the US Postal Service [31].

Our institution's shipment coordinator—who routinely deals with shipments of hazardous materials such as medical supplies—contacted the US Department of Transportation for clarification on shipping powered-on smartphones. They confirmed that as long as the smartphones do not create a dangerous evolution of heat, or have the risk of catching fire while in transit, shipping powered on devices is permitted. To ensure there were no hazardous conditions our devices could enter while running this software, we thoroughly tested our smartphones in extreme environmental conditions that could be experienced during shipment in trucks/railcars. We operated the at 44° C, and -2° C and for several hours while running our measurement, and the phone continued to operate properly. The smartphones we use also have an automatic thermal shutdown feature as an additional safeguard [56].

A.2 Characterizing Critical Infrastructure

Although this study does not involve experiments with human subjects, there are sensitivities with revealing information about critical infrastructure that may provide advantages to adversarial actors.

Although the Belmont report outlined principles relating to human subjects, the 2012 Menlo Report proposed a framework specifically targeting computer and information technology research [32]. Its companion report provided a set of case studies applying the framework [19]. The Menlo Report is a more appropriate framework for our analysis because it explicitly addresses stakeholders such as network/platform owners and providers but also acknowledges that they may warrant different consideration from that of individuals.

Our considered view is that the benefit of our research exceeds potential risk to infrastructure. We are now entirely dependent on this infrastructure but there has been little attention to independent objective understanding of its resilience and reliability. Given increasing attention to the need for regulatory oversight of the Internet as critical infrastructure, it is important to understand just how much a capable independent third party can accurately infer about various aspects of Internet infrastructure. We need to understand this capabilities in order to know what adversarial actors could likely achieve, as well as to know how benign actors might help to reduce the burdens of government by providing independent confirmation of claims of reliability and/or resilience of critical network infrastructure.

We also have long-standing cordial relationships with engineers at the providers we have studied, who are aware and supportive of our work. Specifically, we discussed our inferences with Comcast, Charter, AT&T, and T-Mobile engineers throughout our study, for the purposes of validation of our findings.

	Comcast	Charter
Initial	204,744	54,079
Alias Resolution		
Changed	2.35%	1.10%
Added	2.76%	0.80%
Removed	0.86%	0.20%
	208,640	54,407
Point-to-Point Subnets		
Changed	0.04%	0.05%
Added	1.27%	0.48%
	211,295	54,670

Table 3: To observing CO interconnections in traceroute, we map IP addresses to COs, and account for outdated and missing information.

B DETAILS ABOUT COMCAST AND CHARTER MAPPING

B.1 Mapping IP Addresses to Hostnames

The traceroute probing yields IP address paths, and we attempt to map each individual addresses to a backbone or regional CO. We use both dig and the Rapid7 rDNS dataset to perform reverse lookups on the addresses, prioritizing the dig names to reduce potentially stale names in Rapid7. Comcast and Charter appear to connect both their backbone and regional routers with point-to-point links, so we also lookup names for all IP addresses in the same /30 subnet as a traceroute IP address. The /30 subnet includes all addresses possibly used in a point-to-point link with that address. Using regular expressions, we extract CO and region identifiers from the names, creating an initial mapping from IP addresses to COs. We perform two steps to improve the CO mappings to account for missing and outdated rDNS names (Table 3): (1) resolve router aliases to map groups of addresses to COs, and (2) add additional constraints using point-to-point subnet addresses.

First, we use Mercator and Midar alias resolution to infer addresses that belong to the same router, since these addresses reside in the same CO. We included all of the traceroute addresses, as well as the additional addresses in their /30 subnets. If more addresses in an inferred router map to one CO than any other CO, we remap all addresses in the group to that CO. We do not apply a minimum threshold for the number of router IP address hostnames containing a CO identifier. In the event of tie, we remove all CO mappings for the addresses to avoid potentially misleading information. The alias resolution modified or added more CO mappings for Comcast (5.1%) than Charter (1.9%).

Next, we use point-to-point subnets to further refine the CO mappings. Interconnected router interfaces must have IP addresses from the same IP subnet, and network operators usually assign these addresses from a point-to-point subnet; e.g., /30 or /31 subnets in IPv4, both of which include two usable interface addresses. Based on the IP addresses in our traceroutes, it appears that Comcast typically uses /30 subnets, while Charter uses /31 subnets to interconnect routers in different COs. Routers typically respond to traceroute

Figure 19: The two paths reveal x followed by two different addresses, y and z . Presuming that y and z belong to /30 subnets, we use the other address in each subnet (y' and z') to correct the CO mapping for x .

with the inbound interface address, so the other address in the point-to-point subnet often belongs to the router at the prior traceroute hop. If that other address has a CO mapping provided by rDNS or alias resolution, we can use that information to refine the mapping for the prior hop.

Fig. 19 illustrates our approach to using point-to-point IP subnets to further refine the CO mappings with two traceroutes through a router in a regional access network, and initial CO mappings for each of the addresses. The initial mappings indicate that IP address x belongs to a router in CO1, but both paths reveal subsequent addresses where the other address in each subnet (y' and z') map to CO2. y' and z' most likely belong to the same router as x , so we use them as possible indications that we initially mapped x incorrectly. Here, more addresses map to CO2 than CO1, so we re-map x to CO2. If x lacked an initial CO mapping, then we would use the mappings for y' and z' to infer a mapping for x .

B.2 Removing CO Adjacencies

Initially, we collect all immediate IP address adjacencies where both addresses have a CO mapping (Table 4). MPLS tunnels can cause false links to appear in traceroute, so we use the approach by Vanaubel *et al.* [72] to reveal MPLS exits and the tunnel IP addresses by conducting follow-up traceroutes to all IP addresses mapped to COs in the original traceroute collection. If a pair of addresses appears adjacent in our initial probing, but are separated by one or more hops in the additional MPLS traceroutes, we remove the pair since it is likely the entry and exit of an MPLS tunnel. In our maps, we only observed MPLS tunnel behavior in one Charter region, although we observed this behavior throughout the region.

Although we attempt to adjust outdated rDNS CO references, outdated CO mappings remain problematic. To combat some of the stale CO mappings, we remove any adjacencies where each address maps to a CO in a different regional network. Prior knowledge of the extensive use of aggregation in each region, and conversations with network operators, indicated that a small number of entries

	Comcast		Charter	
	IP Adjs	CO Adjs	IP Adjs	CO Adjs
Initial	95,671	4777	64,667	3994
Backbone	26.07%	7.39%	11.67%	5.02%
Cross-Region	4.45%	18.78%	1.78%	2.37%
Single	0.06%	1.15%	0.03%	0.43%

Table 4: The unique adjacent IP address adjacencies (IP Adjs) and unique CO adjacencies (CO Adjs) pruned to account for stale rDNS and traceroute path corruptions.

into each region exist, so we remove likely invalid cross-region adjacencies. This removed far more of the unique CO adjacencies for Comcast than for Charter, likely due to more outdated rDNS in Comcast, although the cross-region CO adjacencies accounted for less than 5% of the IP adjacencies in both networks. We also remove adjacencies representing potential entries from the backbone into each region, where one of the IP addresses map to a backbone PoP, and we infer entries into each region in §5.2.5.

Finally, we remove any CO adjacencies that only appear once in the traceroute paths. Traceroute output occasionally contains anomalous output that results from network path changes during the probing. When COs appear interconnected in only one traceroute path, we conclude that the apparent interconnection might result from anomalous traceroute output, so we remove them. This removed 55 CO adjacencies for Comcast, and 9 for Charter.

B.3 Refining Region Graphs

After removing likely invalid adjacencies, we use the remaining adjacencies to create graphs of each regional network, with a directed edge from one CO to another corresponding to each CO adjacency. Access networks in the US generally use a star topology to connect EdgeCOs, so we attempt to conform our revealed CO topology to a star topology. The revealed topologies still contain noise, primarily in the form of misleading rDNS creating false adjacencies between EdgeCOs, as well as unrevealed CO interconnections. Our goal is to modify the graphs to conform to the likely physical star topology with as few modifications as possible.

First, we infer the cores of the stars; i.e., the AggCOs in each region. We expect that AggCOs should have more outgoing edges than other COs in the region, despite false CO adjacencies and some EdgeCOs actually connected only to another EdgeCO. To separate likely AggCOs from EdgeCOs we consider any CO with more than the mean outgoing edges plus one standard deviation a AggCO.

We then enforce the role of the AggCOs by removing any (x, y) edge from one EdgeCO to another EdgeCO, unless x has multiple outgoing edges to EdgeCOs that do not interconnect with AggCOs. In general, we expect that edges between EdgeCOs typically result from outdated rDNS, but when a CO appears to aggregate connectivity for multiple COs that otherwise lack connectivity, we conclude that the CO might function as a small AggCO. In total, we removed 26.9% of the unique CO edges in Comcast and 10.6% of the Charter CO edges. The higher fraction of removed Comcast edges reflects prior experience with stale Comcast rDNS.

Next, we infer related AggCOs that connect to the same set of EdgeCOs. Networks often connect an EdgeCO to two AggCOs

to increase resiliency to AggCO failure, and we expect that two AggCOs that connect to the same EdgeCO typically connect to the same set of EdgeCOs, since access networks use bundled fiber rings to connect AggCOs to many EdgeCOs. We evaluate each combination of AggCO pairs in the same region, concluding a relationship between the two AggCOs AGG_x and AGG_y if at least 3/4 of the EdgeCOs connected to AGG_x overlap with EdgeCOs connected to AGG_y , and the overlap accounts for at least half of the EdgeCOs connected to AGG_y . The overlap requirements help ensure that we only pair AggCOs with substantial downstream EdgeCO overlap. We also pair two AggCOs if one AggCO has 3/4 overlap with the other AggCO, and neither AggCO would otherwise have a relationship. To reflect the fact that EdgeCOs connect with fiber rings, we add edges to ensure that that all related AggCOs connect to the same EdgeCOs in the regional network graphs. This added 7.8% new edges to Comcast, and 6.1% new edges to Charter.

B.4 Redundant AggCO Connections

We inferred that 11.4% and 37.7% of the EdgeCOs in Comcast and Charter connected to a one other CO, respectively, but we never observed any CO-level redundancy for the Charter regional network in the southeastern US. This region is the only large regional network in Comcast or Charter where we did not observe any CO-level redundancy, suggesting we inferred an incomplete CO topology for the region. Excluding the southeast, 29.0% of the Charter EdgeCOs connect to a single upstream CO. Furthermore, of the EdgeCOs connected to one other CO, 33.7% of the Comcast COs and 42.2% of the Charter COs connect to another EdgeCO (not AggCO). Considering only the EdgeCOs connected to an AggCO, and excluding the Charter southeast region, 10.5% of the Comcast EdgeCOs and 18.4% of the Charter EdgeCOs connect to a single AggCO.

C DETAILS ABOUT AT&T MAPPING

AT&T's regional network routers do not use rDNS names, so we cannot extend the DNS-based geolocation method (§5) to cluster AT&T's IP addresses into physical facilities. Additional visibility challenges arise from operational practices such as MPLS tunneling and ICMP filtering, which can both hide physical router topology from external traceroutes. These challenges make it critical to have a sufficiently large and strategically selected set of targets.

Target selection. To find responsive destinations with known geographic locations, we extracted location hints from rDNS names of the IP-DSLAMs connected to end-user modems (denoted as *lspgw*). From our pilot tests using Ark and RIPE Atlas, we found that AT&T encoded the rDNS names of *lspgws* with the regular expression $([\d-]+-1).lightspeed.([a-z]{6}).sbcglobal.net$, where the first part of the name is the dashed decimal notation of the corresponding IP address and the second part is a CLLI code-like 6-character string that represents the city and the state. For example, *sndgca* and *nsvltn* denoted San Diego, CA, and Nashville, TN, respectively. We denoted each unique combination as a *region*.² To obtain a comprehensive list of *lspgws*, we used Rapid7's rDNS dataset [53], which periodically resolves rDNS names of the entire IPv4 address

²Note that If the geolocation hint is stale we generally find some anomaly in the traceroute that reveals its staleness, e.g., a traceroute with backbone IP addresses in between nodes with the same geolocation hint likely involves a stale geolocation hint.

	Address	rDNS	reply-ttl
1	192.168.1.254		64
2	107.210.168.1	107-210-168-1.lightspeed .sndgca.sbcglobal.net	63
3	71.157.16.42		59
4	108.89.115.1	108-89-115-1.lightspeed .sndgca.sbcglobal.net	61

(a) Intra-region probing traceroute result. From a VP in San Diego, CA probe to a *lightspeed gateway (lspgw)* in the same city. The third hop is the IP of an EdgeCO router.

	Address	rDNS	reply-ttl
1	192.168.1.254		64
2	107.129.92.1	107-129-92-1.lightspeed .sntcca.sbcglobal.net	63
3	71.148.149.186		62
4	71.145.1.52		61
5	12.83.39.213		251
6	12.123.215.237		55
7	71.157.16.42		55
8	108.89.115.1	108-89-115-1.lightspeed .sndgca.sbcglobal.net	54

(b) Inter-region probing traceroute result. From a VP in Santa Cruz, CA to the same *lspgw* in San Diego, CA. The path first traversed COs in Santa Cruz region (hops 3-5), then AT&T's backbone network (hops 6-7), and finally San Diego region (hops 8-9).

Figure 20: Traceroute examples of regional probing of AT&T.

space, to find hostnames (and IPs) that matched the regex. We found 95,821 IPs in 37 regions in the September 2020 dataset.

AT&T blocked traceroute measurements toward most of the *lspgws* from the public Internet, but allows traceroutes from within a region and from nearby regions. We used four CAIDA Ark VPs in and nearby San Diego in AT&T to conduct ICMP paris-traceroutes to *lspgws* IPs. This process partially revealed the topology that connected EdgeCOs and AggCOs in a region.

To observe the rest of the topology, we needed to expose MPLS tunnels between the BackboneCO and the *lspgws* that hide the AggCOs and many EdgeCOs. To expose these tunnels we needed to discover which IP prefixes are assigned to the EdgeCO routers in the region we are mapping. We used both *intra-* (McTraceroute) and *inter-region* (Ark) traceroutes to *lspgws* to discover these prefixes. Fig. 20a and Fig. 20b show samples of intra- and inter-region probing to a *lspgw* in San Diego from a RIPE Atlas VP in San Diego, CA and an Ark VP in Santa Cruz respectively. The San Diego VP reaches *lspgws* in the same region directly without crossing the backbone (Fig. 20a). The traceroute from the Santa Cruz VP traverses AT&T's backbone network, which uses prefix 12.0.0.0/8, to reach other regions (Fig. 20b). We then extract a preliminary list router prefixes from hops between two *lspgws* in intra-region probing (i.e., hop 3 in Fig. 20a) and between the backbone and the destination *lspgws* in inter-region probing (i.e., hop 7 in Fig. 20b).

	Address	rDNS	reply-ttl
1	192.168.1.254		64
2	107.210.168.1	107-210-168-1.lightspeed .sndgca.sbcglobal.net	63
3	71.157.16.114		62
4	75.20.78.58		61
5	75.20.78.55		60
6	71.157.16.42		59

Table 5: Targeted traceroutes to egress interfaces of MPLS tunnels reveals the paths hidden by the MPLS in intra-region probing (hop 4-5).

Central Office type	prefix
Edge CO	71.157.6.0/24
	71.148.118.0/24
	71.148.71.0/24
	71.148.104.0/24
	71.148.70.0/24
Aggregation CO	71.157.16.0/24
	75.20.78.0/24

Table 6: San Diego AT&T CO prefixes

We applied the Direct Path Revelation (DPR) technique [73] to reveal the network paths in MPLS tunnels. We targeted inter- and intra- region traceroute measurements to all of the addresses in the EdgeCO router prefixes we discovered, which correspond to the egress interface of the tunnel (i.e., hop 3 in Fig. 20a and hop 7 in Fig. 20b), which allowed us to discover hidden links in the regional network. Table 5 shows a sample traceroute within the San Diego region that revealed an additional link (hop 4 and 5 in Table 5) that was hidden in traceroutes to *lspgws*. Table 6 shows all the IP prefixes for routers we discovered in AT&T's San Diego region.

D DETAILS ABOUT MOBILE MAPPING

Target Selection. We used the AS relationship dataset [11] to identify each mobile ISP's neighboring ASes. We found 266/406/213 neighboring ASes for AT&T/Verizon/T-Mobile, respectively. We then conducted a pilot test to compile lists of target IPs for each ISP. For each neighboring AS, we found one IPv4 and one IPv6 destination that were responsive to traceroute probes. We used the corresponding target list of the current mobile ISP to perform traceroute measurements.

The ShipTraceroute results showed that the network paths to all the targets shared the same paths within the mobile network until exiting the PGWs. Table 7 and Table 8 show the number of PGWs we inferred using region bits in AT&T and Verizon IPv6 addresses, respectively.

Table 7: Inferred number of AT&T PGW in each region.

Region name	BTH	CNC	VNN	ALN	HST	CHC	AKR	ALP	NYC	ART	GSV
Region bits in IP addresses	2030	2040	2090	2010	20a0	20b0	2000	2020	2050	2070	2080
MTSO number	2	5	5	5	5	5	3	6	4	3	3

Table 8: Inferred number of Verizon PGW in each region.

Backbone Region Name	SEA		SJC		LAX		
Wireless Region Name	RDMEWA	HLBOOR	SNVACA	RCKLCA	LSVKNV	AZUSCA	VISTCA
Region bits in IP addresses	100f:b0	100f:b1	1010:b0	1010:b1	1011:b0	1012:b0	1012:b1
PGW numbers	1	1	2	2	2	2	3
Backbone Region Name	CHI						PHIL
Wireless Region Name	HCHLIL	NWBLWI	SFLDMI	STLSMO	BLTNMN	OMALNE	ESYRNY
Region bits in IP addresses	1008:b0	1008:b1	1009:b1	100a:b0	1014:b1	1014:b1	1002:b1
PGW numbers	2	2	1	1	3	2	1
Backbone Region Name	DEN		DLLSTX			MIA	
Wireless Region Name	AURSCO	WJRDUT	ELSSTX	HSTWTX	BTRHLA	MIAMFL	ORLHFL
Region bits in IP addresses	100e:b0	100e:b1	100c:b2	100d:b0	100d:b1	100b:b0	100b:b1
PGW numbers	2	2	1	2	2	2	2

Backbone Region Name	ATL			IAD		NYC	BOS	
Wireless Region Name	CHRXNC	WHCKTN	ALPSGA	CHNTVA	JHTWPA	WLTPNJ	WSBOMA	BBTPNJ
Region bits in IP addresses	1004:b0	1004:b1	1005:b0	1003:b0	1003:b1	1017:b0	1000:b0	1000:b1
PGW numbers	4	2	2	2	1	2	2	1