

Access Denied: Assessing Physical Risks to Internet Access Networks

Alexander Marder
CAIDA / UC San Diego

Zesen Zhang
UC San Diego

Ricky Mok
CAIDA / UC San Diego

Ramakrishna Padmanabhan
CAIDA / UC San Diego

Bradley Huffaker
CAIDA / UC San Diego

Matthew Luckie
University of Waikato

Alberto Dainotti
Georgia Tech

kc claffy
CAIDA / UC San Diego

Alex C. Snoeren
UC San Diego

Aaron Schulman
UC San Diego

Abstract

Regional access networks play an essential role in connecting both wireline and mobile users to the Internet. Today’s access networks support 5G cellular phones, cloud services, hospital and financial services, and remote work essential to the modern economy. Yet long-standing economic and architectural constraints produce points of limited redundancy that leave these networks exposed to targeted physical attacks resulting in widespread outages. This risk was dramatically shown in December 2020, when a bomb destroyed part of AT&T’s regional access network in Nashville, Tennessee disabling 911 emergency dispatch, air traffic control, hospital networks, and credit card processing, among other services.

We combine new techniques for analyzing access-network infrastructure deployments with measurements of large-scale outages to demonstrate the feasibility and quantify potential impacts of targeted attacks. Our study yields insights into physical attack surfaces and resiliency limits of regional access networks. We analyze potential approaches to mitigate the risks we identify and discuss drawbacks identified by network operators. We hope that our empirical evaluation will inform risk assessments and operational practices, as well as motivate further analyses of this critical infrastructure.

1 Introduction

Regional access networks are an essential component of the Internet infrastructure: they connect end users to the rest of the Internet. In order to balance reliability and performance against the enormous cost of providing last-mile connectivity to vast populations of geographically distributed users, access networks aggregate customer traffic into layers of central offices that are connected with varying degrees of redundancy. Unlike backbone networks, access networks often lack sufficient redundancy to withstand single-facility failures and a recent study showed that third parties can infer these points of limited redundancy [1]. Troublingly, physical attacks against regional access network infrastructure are becoming increasingly common [2–5].

Today’s regional access networks are far more critical than when they were first deployed, with ballooning real-world impacts of network outages. No longer just conduits of landline telephone and cable TV, modern access networks support 4/5G cellular phones, cloud services, hospital and financial services, and the remote work essential to the modern economy. Perhaps the most dramatic illustration of these inter-dependencies occurred in December 2020 when a bomb disconnected an AT&T network facility in Nashville, Tennessee [6]. This single event took the entirety of AT&T’s wireline and wireless network in the Nashville area offline for several days. It also disconnected 911 emergency services [7], grounded flights by taking air traffic control offline [8], prevented hospitals from reaching remote records and health-care providers [9], and even halted credit card processing [10]. We believe these circumstances demand a clear-eyed assessment of the threats to regional access network infrastructure and a reconsideration of the operational trade-offs occurring today.

In this paper, we evaluate the ways in which regional Internet access networks are at risk of physical attack in an effort to better inform the cost-benefit analysis of existing and future deployments. We perform a large-scale measurement campaign to study the impact of infrastructure failures on real-world access networks. Specifically, we continuously monitor users of the primary access networks in several regions of the United States for a year. When we observe large correlated outages, we identify the portion of the access network topology that likely failed using a technique we introduce in this work. To our knowledge, this is the first public study to assess the potential impacts of physical attacks on the regional access network infrastructure in the U.S.

Furthermore, we show how operational practices may facilitate targeted attacks. For example, regulations often require providers to record locations of their diesel fuel storage and battery backup power systems in local hazardous-materials registries. We demonstrate that an attacker often can identify the physical infrastructure serving a particular region based upon a set of design patterns: access networks typically have well-segregated coverage areas. As a result, an attacker can

infer the infrastructure providing service to a particular target area by, e.g., wardriving nearby public WiFi hotspots.

We hope that our work will spur further analyses of this critical infrastructure. This paper makes the following contributions:

- **We identify concrete threats to operational regional access networks.** Through conversations with operators at the largest U.S. access networks and by analyzing recent results on mapping access network topology [1], we describe how the redundant power and packet-transport infrastructure currently in place to withstand natural events is insufficient for intentional attacks.

- **We study the root cause and impact of large access network outages.** We combine inferred network infrastructure maps with continuous reachability measurements to millions of access network customers to detect outages and identify the failed infrastructure. We investigate outages of different magnitudes in detail, including the Nashville bombing. These outages indicate that the scale of an attack’s impact can be expected to range from thousands to hundreds-of-thousands of users, and the duration to span hours to days.

- **We show that targeted attacks can be launched without insider information.** By combining public hazardous-material datasets with targeted use of the ubiquitous traceroute tool, we show that an attacker can learn the location of infrastructure whose failure will disconnect specific areas. We demonstrate feasibility in three different networks.

- **We explore potential ways to mitigate risks.** Access networks must balance infrastructure security with manageability and cost, and we explore trade-offs associated with mitigating physical threats to the infrastructure.

Ethical considerations. The Menlo Report [11, 12] explicitly addresses stakeholders such as network/platform owners in the context of revealing information about critical infrastructure that may provide advantages to adversarial actors. These principles, and feedback from network operators, guide our approach to anonymization and disclosure of details about networks. We anonymize details when we explore the attack surface of different networks (§6 and §7), but do not anonymize networks or locations in case studies (§5) when those details appear in the public press. All three operators we consulted were eager to understand what could be gleaned about their infrastructure by a capable independent third party and how they could raise the bar for attacks.

2 Background: Access Network Topology

Internet Service Providers (ISPs) design access networks with significant redundancy to withstand common failures that occur through random chance, like trees falling on overhead fiber or mains power outages. This redundancy provides some protection against physical attacks as well: networks can con-

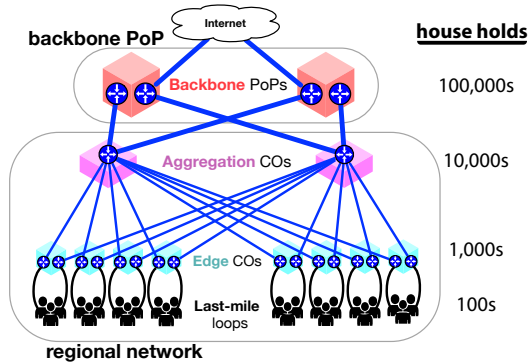


Figure 1: Access network hierarchy: EdgeCO routers aggregate customers and AggCO routers aggregate EdgeCOs.

tinue to function as normal after incurring a fiber or power cut. However, physical attacks that damage the backup systems as well can lead to widespread outage, as we will show. To understand this risk, we describe the general architecture of Internet access networks (§2.1) and discuss where access networks deploy topological redundancy (§2.2).

2.1 Key Topological Elements

Access networks consist of dense deployments of fiber optic cables—and often also powered equipment—in nearly every neighborhood in the geographic regions where they provide service (e.g., a metropolitan area). To provide Internet access, each access network connects back to a small number of Internet backbone routers in one or more Internet Points-of-Presence (PoPs). Providers design their networks to achieve this connectivity efficiently by aggregating traffic through a hierarchy of facilities known as *carrier offices (COs)*: these buildings aggregate traffic with last-mile link technologies and switches, and pass traffic up or down the hierarchy with inter-CO routers.

The general network topology of a typical access network is shown in Fig. 1. An *Edge CO* (or EdgeCO) aggregates traffic from hundreds to thousands of customers over last-mile links; e.g., cable, DSL, and fiber. Similarly, an *Aggregation CO* (or AggCO) aggregates traffic from dozens of EdgeCOs providing service for hundreds of thousands of users—often across metropolitan areas or entire states. *Backbone Points of Presence* (Backbone PoPs) aggregate traffic from one or more AggCOs and provide Internet transit services over a backbone network operated by the ISP or another provider.

2.2 Redundant Infrastructure

Based on the topologies of major U.S. access networks revealed in recent work [1] and conversations with network operators, we explore differences in how ISPs deploy redundancy at different layers of regional access networks.

Some networks deploy redundant last-mile connections using fiber rings, letting them survive a single fiber cut to the ring. Well-provisioned networks may even terminate the ring at two different EdgeCOs to provide CO-level redundancy, although deploying and maintaining multiple last-mile connections is expensive. In most networks it is only economical to deploy a single last-mile link to each customer. Without redundancy, a single cut to a last-mile cable bundle will disconnect all customers downstream from the EdgeCO on that fiber strand. Additionally, depending on the last-mile technology used in the network (DOCSIS cable, DSL, etc.), an attacker may be able to disconnect multiple users by cutting a single link in a neighborhood (e.g., DOCSIS feeder coax).

EdgeCOs aggregate thousands of last-mile links that terminate at specialized devices inside the CO; e.g., CMTS in cable networks or DSLAM in DSL networks. Often, adding redundant last-mile links to different EdgeCOs is cost prohibitive, so customers connect to a single EdgeCO. As a result, an EdgeCO outage will disconnect all downstream last-mile customers. A group of EdgeCOs connect to one or more AggCOs through a fiber ring. When a group of EdgeCOs connects to two or more AggCOs, each AggCO interconnects with each EdgeCO in one direction around the ring, allowing the EdgeCO to survive a single AggCO outage.

Smaller regional networks contain a single AggCO layer with one or two AggCOs. If there is only one AggCO, then an attacker can disconnect the entire region by attacking that one CO; if there are multiple, the network can survive one going down. Larger regions often employ multiple AggCO layers, where some AggCOs might only aggregate traffic from other AggCOs. Some providers split their aggregation layers into two or more subregions and use separate fiber rings with one or two AggCOs, so a failure of one ring will not take down all of the region's EdgeCOs.

At the top of the aggregation hierarchy, one or more AggCOs, which serve as entry points into the regional access network, connect to one or two Backbone PoPs, and occasionally interconnect with large transit ISPs as well. If a region only has one Backbone PoP and that PoP is taken offline, all customers in that region will be disconnected from the Internet. In regions that have more than one AggCO and Backbone PoP, each AggCO usually connects to a different Backbone PoP. This configuration allows the the entire region to fail over to the other Backbone PoP if one Backbone PoP fails.

3 Threat Model

This section describes the physical attacks we consider on regional access networks, where the attacker's objective is to cause widespread connectivity outages. We first discuss how an attacker—without insider knowledge—can damage physical plant, such as fiber and power (§3.1). Then we discuss why existing redundancy insufficiently addresses the threat of intentional attack (§3.2).

3.1 Attacker Capabilities

In this work, we show how an attacker without insider knowledge can cause large-scale outages. We demonstrate that motivated attackers can combine network measurement tools with public information to identify minimum cuts in the access network dependency graph and target specific users.

Attackers can damage underground and overhead fiber.

Access networks are built out of fiber optic cables containing bundles of fiber optic strands that are deployed aerially along telephone poles or underground in cable vaults. In both cases, the fiber runs unprotected over large distances, and attackers can cut them using widely available wire cutters. Attackers can visually identify a provider's cables because they often use fiber ID tags on aerial lines, and marker poles and labeled cable vaults on underground lines. An attacker can reach aerial fiber by climbing telephone poles or damaging the poles themselves [13] and cut underground fiber with digging equipment or by accessing the cable vault. An individual attacker can also cut multiple fiber bundles in different locations before the ISP can repair the fiber. Simply detecting the location of damaged fiber can take minutes to hours [14], in part because the provider must dispatch repair crews to the fault location(s).

Recent examples demonstrate the risks for fiber deployments. For instance, between 2009 and 2016 there were more than a dozen incidents of vandals cutting fiber optic cables in California [4]. Two of the attacks disrupted AT&T's access network for hours and led them to offer a \$250,000 reward for information about the culprits [2, 3].

Attackers can disrupt mains power and backup fuel.

Access networks require power inside facilities and out in the field to maintain network operations. An attacker can cut the mains power serving this infrastructure, forcing the network to rely on backup power, and that backup power may run out; e.g., due to lack of fuel. Also, an attacker can damage the mains and backup power simultaneously, which is what occurred in the Nashville bombing [10].

3.2 Threats to Fiber and Power Redundancy

ISPs design COs and last-mile links with redundancy to continue operating in the face of a single fiber cut or loss of power. Across ISPs, the conventional approach is duplicating nearly every piece of infrastructure related to power and network transport, such that if one component fails, the redundant component can seamlessly take over.

Fiber Rings. ISPs physically deploy fibers in a ring topology to aggregate traffic from multiple COs to the CO in the next hierarchy level because rings are resilient to a single fiber cut at any location on the ring: traffic can route in the




Figure 2: An attacker can easily cut fiber rings when both sides of the ring run in parallel.




Figure 3: An attacker must disable either a CO’s redundant power or redundant fiber to induce a failure.

remaining direction around the ring to maintain connectivity until the fiber cut is repaired [15–18]. If an attacker cuts a fiber ring in two places, however, it will disconnect all COs and customers downstream of both cuts. Some fiber rings are especially susceptible to this attack because economic or geographic constraints might force an ISP to run both sides of the ring close together (Fig. 2).

Backup Power. COs are typically provisioned with backup power that seamlessly maintains operation during a power outage for approximately 24 hours until the mains power returns. Although mains and backup power are largely independent, they generally meet at a central power control system. This presents an opportunity for an attacker to induce an outage that takes both systems offline.

A physical attack can proceed in three phases: (1) The attacker selects the customers they want to take offline, or the ISP they want to damage. (2) The attacker finds the fiber or power nearest to those customers or ISP’s facilities by looking for markings on cables and vaults. (3) The attacker cuts fiber and/or disables power. Fig. 3 shows how an attacker can disconnect part of the access network by either

cutting the fiber ring in two places or disabling all power input into powered network equipment.

4 Experiment Methodology

Having established that access network COs remain vulnerable to intentional attack, we empirically measure the expected impact of a CO outage. Specifically, after a successful attack, we examine how many users would likely experience an outage, and for how long. We design a measurement study of the three largest residential ISPs in the U.S.—Comcast, Spectrum, and AT&T—with two goals: (1) estimate the number of customers connected to COs across the U.S., and (2) leverage weather, accidents, and vandalism to empirically learn the scale and duration of CO outages.

Our analysis of CO outages proceeds in three stages. First, we create maps of each regional access network that capture the CO-level topology (§4.1). Second, we infer the customer IP address space connected to each CO (§4.2). Third, we continually send probes to customers of the access networks to observe when a CO experiences an outage and to measure the outage duration (§4.3). When possible, we add context from news stories to confirm that an attacker could intentionally recreate the failures we observe.

4.1 Mapping Regional Access Topologies

Our experiment touches 22 of the regional access networks that Comcast, Spectrum, and AT&T deploy across 14 U.S. states. We conduct large-scale measurements to create CO-level maps of these regional access networks.

At the core of our technique, we use the traceroute tool to reveal router IP addresses between a measurement vantage point (VP) and an arbitrary destination. Traceroute induces a single response from each router along the path containing the IP address assigned to an interface on the router. To increase the likelihood that our path measurements reveal all active paths through the regional networks, we use measurement VPs distributed across the U.S. Our VPs conduct traceroutes to customers connected to the networks, revealing the IP topology of each regional access network. We use the same techniques as Zhang *et al.* [1] to infer CO interconnections and aggregation hierarchies in each access network from the IP topologies.

4.2 Mapping Customers to COs

The techniques from Zhang *et al.* [1] reveal CO interconnections, and substantial prior work observed last-mile outages [19–22], but no prior work has tied those outages to network facilities. To support tying outages to COs in Comcast and Spectrum, we also create mappings from customer address space to COs; i.e., the IP address ranges used by customers attached to a given CO. In the access networks,




Figure 4: The routers and CMTSes inside EdgeCOs appear in traceroute paths.

each CO assigns addresses to customer devices from a pool of addresses allocated to that CO. That pool consists of hundreds-to-thousands of IPv4 /26 subnets, and we infer the pool of residential /26s for each CO in Comcast and Spectrum.

Mapping from customer IPs to COs would be trivial if DNS names always indicated the CO for the IP address immediately before the customer in a traceroute path, but many of those addresses either lack a CO identifier or lack DNS names entirely. Instead, we leverage technical details of cable access-network infrastructure to infer comprehensive CO-to-address mappings. In particular, the cable modem termination systems (CMTS) housed inside cable-network COs respond to traceroute probes, so one hop prior to the customer is the CMTS, and two hops prior is a router in the same CO [23] (Fig. 4). By sending traceroutes to every residential customer IP address, we construct a directed interface graph with edges between immediately adjacent hops. We cluster each customer IP address with all preceding addresses within distance two, allowing us to infer a CO mapping for the cluster rather than separate mappings for the individual IP addresses. The transitive closure of each cluster includes the customers, CMTS devices, and routers that all map to the same CO. Finally, we use the CO identifiers that Comcast and Spectrum include in many—but not all—hostnames for their router and CMTS IP addresses to map clusters to COs.

When EdgeCOs have multiple CMTS devices, we might observe different routers prior to disjoint sets of CMTSes, creating two different clusters for a single CO. We evaluated this potential problem on Spectrum’s access networks, which have good hostname coverage for EdgeCO router IP addresses. Clustering the IP addresses created 860 clusters where a hostname let us infer the CO identifier. Only 7.2% of the clusters received an identifier that was also assigned to another cluster, indicating a partial CO cluster. Our approach appears to work well for the other 92.8% of the CO clusters.

For AT&T, traceroutes to most residential customers failed to induce responses from routers within the access network. As a result, we only mapped AT&T customers to COs in one regional network and partially mapped customers in another region. We used the same technique as Zhang *et al.* [1] to estimate the customers connected to EdgeCOs by conducting traceroutes from various locations within the access network.

4.3 Detecting CO Outages

To detect CO outages, we continuously test reachability to the residential customers in each regional access network. Testing reachability of customers—rather than routers in the COs—ensures that any event we detect actually disconnected customers; i.e., the redundancy in the network failed to prevent an outage. We detect CO outages when all customers that depend on the CO experience an outage simultaneously.

We test reachability for Comcast and Spectrum by pinging access network customer addresses every ten minutes from three different VPs. We ping a static set of customer addresses consisting of 50% of the customer addresses for each network across 14 different U.S. states. Using this customer sample allows us to comprehensively detect outages at 10-minute granularity while bringing the financial cost of virtual machines and egress traffic from the cloud within our constraints. Three VPs ping each customer in our set in every ten-minute round, and we consider a customer responsive in a round if it responds to any of the three pings.

To detect CO-level outages, we find 10-minute rounds where all customers of a CO failed to respond to all three VPs. First, we compute the median number of responses for each CO and /26 subnet across all 10-minute rounds in each week of data. To reduce the likelihood of misclassifying last-mile failures, dynamic IP address reassignment, or transient customer device unresponsiveness as CO-level failures, we only consider COs with a median of at least 100 responding customers spread across 20 or more /26 subnets. Next, we iterate over each 10-minute round to identify COs without any responding customers, and the number of consecutive 10-minute rounds with no responding customers quantifies the outage duration. Using the CO interconnection maps we can also infer failures higher up in the access network aggregation hierarchy, when all EdgeCOs dependent on a set of AggCOs fail simultaneously.

We cannot detect outages in AT&T with the same granularity, since AT&T customer devices generally did not respond to our pings. Instead, we use traceroutes toward AT&T customers to observe when portions of an access network disappear at the same time; i.e., when previously observable COs disappear from the traceroutes. CAIDA’s Ark [24] measurement platform uses globally distributed VPs to continually send traceroutes to every IPv4 /24 multiple times a day [25]. To detect outages, we look for periods of time where all traceroutes from Ark VPs fail to observe one or more COs. Ark conducts traceroutes less frequently than we conduct our pings for Comcast and Spectrum, so we can only observe AggCO outages that last for several hours in AT&T.

5 Outage Case Studies

We collected outage data for Comcast and Spectrum between August 2020 and December 2021 (Fig. 5) and looked for




Figure 5: Bars indicate the measurement period for different regions in our study of Comcast and Spectrum. Gaps corresponds to configuration errors that prevented data collection.

CO Type	ISP 1	ISP 2	ISP 3	Total
Backbone PoP	0	0	1	1
AggCO	4	1	0	5
EdgeCO	40	24	0	64
Total	44	25	1	70

Table 1: Observed outages in Comcast, Spectrum, and AT&T.

AT&T outages in the December 2020 Ark traceroutes. We observe 70 outages where our reachability tests failed to reach any customer behind a CO (Table 1). Five outages affected all EdgeCOs downstream of a set of AggCOs, indicating problems either at or near the AggCOs. We observed at least one CO outage in 11 out of the 14 states we probed. The outages mostly lasted between 50–200 minutes, with the median outage lasting 1 hour and 10 minutes (Fig. 6), and typically affecting 4,800–34,000 customers. The longest outage lasted nearly 3 days following Hurricane Ida in Louisiana, and the largest outage disconnected an entire access network in California that serves over 2M customers for 50 minutes.

Our approach cannot distinguish scheduled maintenance outages from failures, and ISPs cannot reroute customers during scheduled maintenance that requires disconnecting a CO. Because networks often perform scheduled maintenance between 00:00–05:59 local time [26, 27], we classify an outage as *overnight* if it occurs within that time window. Fig. 6 shows that while overnight outages tend to be short, they can cover many customers.

The remainder of this section discusses specific outages (Table 2) that suggest the potential impact of successful physical attacks against access networks. We withhold CO locations when not revealed in news stories.

5.1 Case Studies: Backbone PoP Outage

The AT&T backbone PoP failure in Nashville, Tennessee caused widespread outages. On December 25, 2020, a van exploded on the street outside the AT&T Nashville backbone PoP. The explosion disconnected the facility from mains power and caused the backup generators to fail [10]. Battery backups maintained operations for several hours but the PoP went offline when they exhausted.




Figure 6: Outage duration and number of affected customers. We classify each outage as either overnight, when scheduled maintenance is common [26, 27], or daytime.

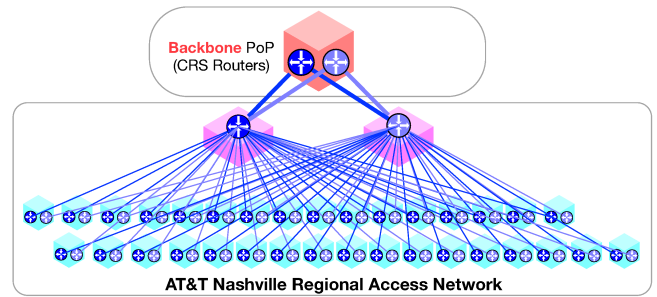


Figure 7: The AT&T Nashville access network relies on a single Backbone PoP. When that facility failed, it disconnected this entire access network from the Internet.

The PoP failure disconnected all AT&T wireline customers in the greater Nashville metropolitan area, but AT&T provides more than residential Internet access over the wireline access network. AT&T wireless also used the access network facilities to reach the AT&T backbone [10]. Worse still, 911 emergency services [7], air traffic control [8], and hospitals [9] all relied on that AT&T access network for communication.

News reports explain how the PoP failed [10], but not why the single PoP failure disconnected hundreds of thousands of AT&T customers in and around Nashville, as well as vital services in the area. To understand why, we generate a topology map of AT&T’s Nashville access network with Ark traceroutes (Fig. 7), and compare that to the observable topology during the 24 hours after the explosion. The maps reveal that all traffic into the Nashville access network passed through two core routers (i.e., CRS routers). During the outage, these two routers—and all routers previously observed downstream of the CRS routers—disappeared from the topol-

Failure Type	ISP	COs	Duration	Customers	Location	Date	Time
Backbone PoP Outage (§5.2)							
Single PoP	AT&T	41	31 hours	229,632	Nashville, TN	2020-12-25	07:10
AggCO Outages (§5.2)							
Multiple AggCOs	Spectrum	44	2 hours	388,608	Maine	2021-04-05	17:20
Degraded Service	AT&T	0	16 hours	0	San Diego, CA	2020-12-20	08:16
EdgeCO Outages (§5.3)							
Multiple EdgeCOs	Spectrum	12	30 minutes	294,400	Los Angeles, CA	2021-02-22	18:00
Single EdgeCO	Comcast	1	40 minutes	3072	Rio Vista, CA	2021-02-25	16:20

Table 2: Our case studies suggest the potential duration and scale of successful attacks against access network COs.

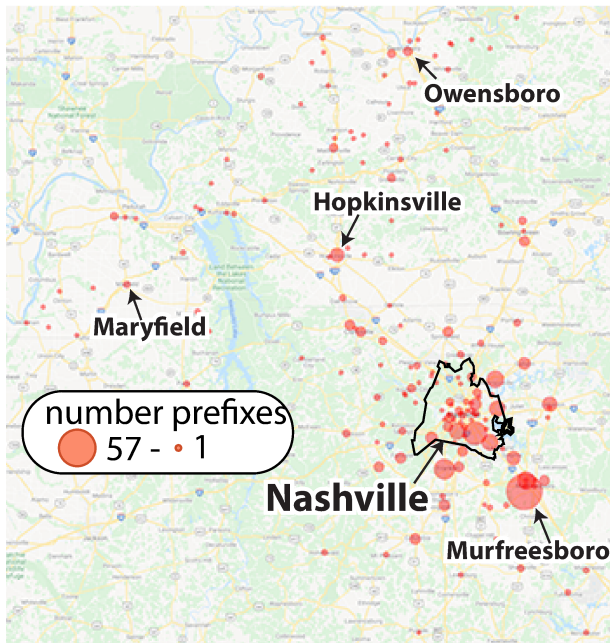


Figure 8: Ark traceroutes reached AT&T customers outside the city limits (black border) via the Nashville CRS routers. /24 prefixes (red dots) geolocated with NetAcuity.

ogy, indicating that the explosion took the CRS routers offline. This explanation is congruent with AT&T outage reports indicating that both CRS routers in Nashville experienced an outage [28], and an operator at AT&T confirmed that a single facility houses both CRS routers.

We confirmed that all AT&T customers throughout the greater metropolitan area relied on those CRS routers (Fig. 8), explaining the geographic scale of the outage. Like Zhang *et al.* [1], we revealed AT&T’s CO-topology in Nashville by conducting traceroute probing from publicly available WiFi access points in April 2021. We connected to three McDonalds’ and eight business WiFi networks available through Instabridge [29] around the city that are AT&T customers, sending traceroutes from each location to destinations outside

AT&T’s network. Every traceroute from the eleven customers passed through one of the CRS routers in the Nashville PoP.

The Ark traceroutes let us retroactively watch as AT&T restored the Nashville access network, and discover that the PoP required only one CRS router. Starting at 16:14 on December 26, the Nashville PoP appeared in paths forwarding traffic to other backbone PoPs. Finally, on the morning of the 27th, we again observed downstream access network COs in the traceroute paths. Consistent with AT&T recovery reports [28], it appears that AT&T initially restored only one CRS router in the PoP along with its fiber connectivity, the minimum needed to restore connectivity to the regional network. We finally observed the second CRS router at 12:00 on December 28th, more than three days after the outage began.

While the bombing likely did not intentionally target the AT&T facility [6], it suggests that intentional attacks could similarly disrupt access network connectivity. AT&T appears to use a single PoP to reach other regional access networks, for instance Zhang *et al.* [1] found one entry PoP housing the two CRS routers in another regional network as well. The outage in Nashville also illustrates the risk of relying on a single access network for many different critical services: a single outage can disrupt nearly all communications in a geographic area. Outages that affect entire regional networks can even disconnect mobile networks [10], so LTE backup might not provide the redundancy that many expect.

5.2 Case Studies: AggCO Outages

Next, we discuss an outage that disconnected all Spectrum customers in the state of Maine, and another that degraded service in AT&T’s San Diego access network.

Two Fiber Cuts Disconnect All Spectrum Customers in Maine. On April 5, 2021, all Spectrum customers in Maine stopped responding to our pings for two hours (Fig. 9). The outage included 1518 /24 subnets, indicating a maximum of 388,608 residential customers. Spectrum disclosed that two separate fiber cuts caused the state-wide outage:




Figure 9: Spectrum customer IP addresses (red dots) were disconnected throughout Maine. Geolocated with NetAcuity.

We've identified two separate fiber breaks in our network, impacting services for Spectrum customers in Maine and New Hampshire... These separate breaks have impacted our redundant path, which normally serves as backup... [30]

This explanation is precisely consistent with our assumptions, since it requires two fiber cuts to disconnect COs.

However, the press release does not indicate why two fiber cuts could disconnect all Spectrum customers in Maine. Spectrum is the largest broadband ISP in the state of Maine, but includes Maine in its larger Northeast regional access network. From our map of Spectrum's Northeast region (Fig. 10), we learn that any IP packet sent to residential customers in Maine must pass through one of two AggCOs in upstate New York. From there, it goes to one of the two AggCOs in Maine. All EdgeCOs in Maine connect to both of the Maine AggCOs, and an EdgeCO needs a connection to only one of the two AggCOs to remain connected to the access network.

The map reveals that fiber cuts between the Maine AggCOs and the entry AggCOs are the only scenario that could disconnect all Maine customers from the Internet but not affect the rest of Spectrum's Northeast regional access network. Without that fiber connectivity, Spectrum customers in Maine could not connect to the rest of the access network or reach Spectrum's backbone. Furthermore, our pings included some Spectrum addresses connected to COs in upstate NY that did not depend on the Maine AggCOs and remained reachable throughout the outage.

Importantly, the outage confirms our hypothesis that the effects of AggCO outages cascade to their downstream EdgeCOs. It also suggests that an attacker might have hours to




Figure 10: Spectrum's Maine sub-region includes two AggCOs leading to every EdgeCO. Two fiber cuts disconnected the AggCOs from the rest of the access network [30].

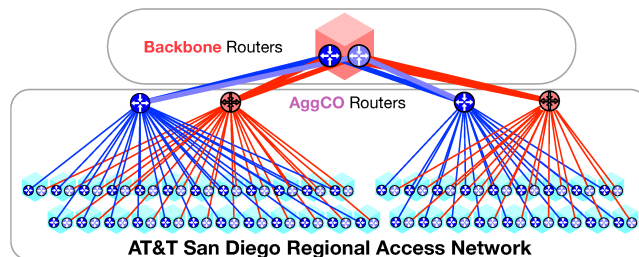


Figure 11: A partial outage appears to disconnect two AggCO routers (red), but customers remained connected.

cut multiple fibers in different locations to cause large-scale outages. The Maine EdgeCOs required only one connection to the upstream AggCOs, but it took at least two hours to bring customers back online, indicating it took Spectrum at least two hours to fix one of the fiber cuts.

Degraded Service After AT&T AggCO Failure in San Diego. We also examined a likely AggCO outage where the 2x redundancy maintained customer connectivity. According to our map of the San Diego AT&T regional access network (Fig. 11), all EdgeCOs connect to two of four AggCOs, which in turn connect to the two CRS routers in the San Diego backbone PoP. On December 18, 2020, two of the four AggCO routers disappeared from the Ark traceroutes for 16 hours (shown in red), leaving only half of the IP-level topology visible. The disappearance of these two AggCO routers suggests that they became disconnected, yet the Ark traceroutes continued to reach customers of the San Diego access network through the remaining AggCO routers.

Although the redundancy maintained AT&T customer connectivity, it appears that the remaining path could not handle




Figure 12: Single EdgeCO outages typically lasted 1–4.5 hours (a) and impacted 3–20K customers (b).

peak traffic demand. Between 19:00 and 22:00 local time, Ark traceroutes inconsistently revealed the San Diego access network CO routers. The most likely explanation is that the increase in traffic during peak Internet usage hours congested the remaining CO interconnections, degrading customer connectivity. This explanation is consistent with DownDetector data [31] showing an increase in customer outage reports starting at the same time. This case shows that even when redundancy prevents a widespread access network outage, an attacker could still cause degraded service.

5.3 Case Studies: EdgeCOs Outages

In our study, EdgeCO outages without a corresponding AggCO outage occurred most commonly. Of those, 15 outages disconnected all customers connected to multiple EdgeCOs, and the remaining 49 outages affected a single EdgeCO. The single EdgeCO outages help indicate the expected fallout from an attack against an EdgeCO (Fig. 12); they typically lasted 1–4.5 hours and affected 3–20K customers. We focus specifically on a multi-CO Spectrum outage in Los Angeles and a Comcast EdgeCO outage in Rio Vista.

Nearby Spectrum EdgeCO Outage in Los Angeles Without AggCO Failure. On February 2, 2021 we observed evidence that EdgeCO outages are not always independent. The outage spanned multiple Los Angeles EdgeCOs in Spectrum’s Southern California regional access network (Fig. 13), but the outage did not appear to originate at an AggCO. Starting at 18:00 and lasting 30 minutes, the outage disconnected 8 EdgeCOs from their single upstream AggCO in Los Angeles and degraded service to two other EdgeCOs. News reports confirmed the outage and its duration [32], but Spectrum did not publicly disclose the cause of the outage. This outage shows that even connecting to two EdgeCOs might be insufficient, since an attacker might be able to disconnect nearby COs simultaneously.

EdgeCO Outage Disconnected Customers From 911. An outage in Rio Vista, California highlights that EdgeCO outages can affect customers in ways that customers might




Figure 13: Multiple EdgeCO outage in Los Angeles, California affected up to 294,400 residential customers.

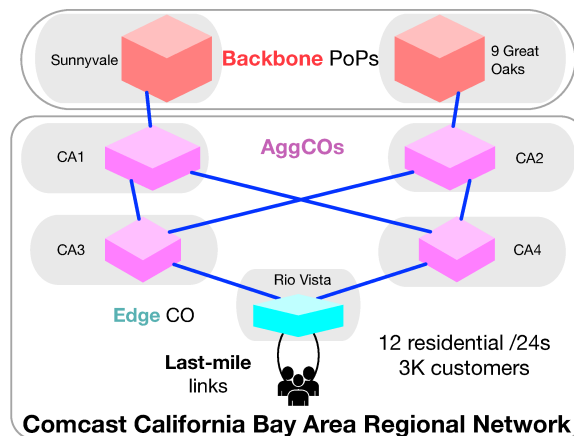


Figure 14: EdgeCO outage in Comcast’s Bay Area regional network, affecting 3K residential customers for 40 minutes.

not expect. At 08:50 on March 23, 2021, we observed a 40-minute Comcast EdgeCO outage in Rio Vista, California that disconnected up to 3,000 customer devices (Fig. 14). The Solano County Sheriff’s office reported the outage [33] to warn that during the outage Comcast-provided phone service could not reach 911 emergency services.

5.4 Security Takeaways

The outages and case studies illustrate three key access network properties that facilitate intentional attacks. (1) Combinations of power failures and fiber cuts frequently disconnect COs, despite their redundant design. If an attacker can disrupt power or fiber connectivity, they will disconnect the CO. (2) Residential customers are typically connected to a single CO, and CO failures disconnect their connected customers. Attackers can target a single EdgeCO to target customers




Figure 15: Three Florida CO locations from hazmat records.

within the local geographic area. (3) Entire regional access networks can fail, as evidenced by the statewide outage in Maine and the complete access network failures in northern California and Nashville. Current access network design in the U.S. allows an attacker to disrupt Internet communication for millions of people by targeting specific COs, without hidden redundancy to maintain connections.

6 Feasibility of Targeted Attacks

Our synthesis of topology mapping with case studies of real outages demonstrates that attackers could disrupt Internet connectivity with physical attacks on COs or last-mile links. However, these case studies do not reveal if it is feasible to perform a targeted attack to disrupt a specific entity or geographic area. In this section, we show that attackers can precisely locate COs and predict the affected geographic area.

Hazardous Materials Records Can Locate COs. Surprisingly, we find that safety regulations increase availability of CO street addresses. To satisfy FCC backup power requirements [34], COs typically use on-site diesel generators and battery cells as redundant power sources. These materials pose fire hazards, so local authorities require the networks to register the capacity and location of storage tanks and other hazardous materials (hazmat) with regulatory bodies. These records are often public [35–37], revealing CO street addresses in a geographic area (Fig 15). We implemented scripts to crawl hazmat records from four different regulators in the US, demonstrating the accessibility of the data.

Wardrive to Predict Service Areas. Customers are not necessarily served by the closest EdgeCO due to regulatory,

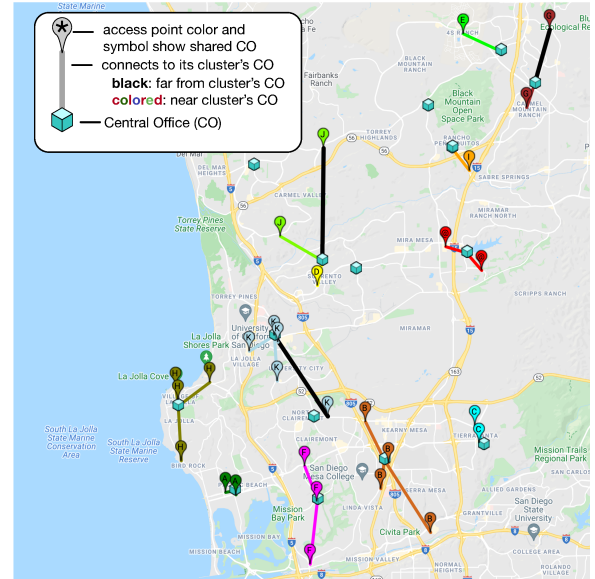


Figure 16: Inferred EdgeCOs for access points (APs) in a San Diego ISP. Marker color identifies APs connected to the same EdgeCO. Black lines indicate that the EdgeCO is not the closest CO.

geographic, and financial constraints, but traceroutes in the target area can reveal the EdgeCO serving an area. Specifically, an attacker can cluster access points to the EdgeCO serving them using a “wardriving” approach to conduct traceroutes via public WiFi access points (APs) in fast-food restaurants and coffee shops, such as McDonald’s and Starbucks. As a proof-of-concept, we conducted traceroutes from 114 public WiFi APs in southwest San Diego County to a server in our lab, and estimated the geographic service areas for each of an ISP’s EdgeCOs (Fig. 16). Notably, 89% of the WiFi APs connected to the geographically nearest EdgeCO.

Match CO Identifiers to Locations. Synthesizing the hazmat records with DNS names can reveal even richer CO topology information. Some access networks include street or neighborhood names as CO identifiers in the DNS hostnames associated with access network router IP address. This allows an attacker to match CO locations in hazmat records to the IP addresses that traceroute reveals.

We matched the CO identifiers in a South Florida access network to the street and city names in public hazmat records (Fig. 17). We validated the mappings with network operators, who asked us to anonymize the network for operational security reasons. This synthesis of physical and topological access network maps reveals the AggCO locations and the interconnections between the AggCO and EdgeCO locations. For example, the map indicates that an attack against AggCOs in Stuart and Pompano Beach could cause widespread outages extending to Palm Beach and Miami.




Figure 17: Map that combines hazmat records, DNS router names, and network topology measurements. Green lines map EdgeCOs (blue) to their corresponding AggCOs (red). Black squares are EdgeCOs that we could not map to DNS names.




Figure 18: COs with large-capacity backup tanks in a Florida access network are located in highly populated areas.

For access networks without useful CO identifiers in their DNS names, an attacker could also use the amount of registered fuel in the facilities to infer the aggregation level of the proximate CO. Compared to EdgeCOs, AggCOs often house equipment with greater power consumption that require more backup fuel. Fig. 18 shows the locations and sizes of backup diesel tanks at COs in a West Florida access network overlaid on top of a population heat map. One facility in the West Tampa neighborhood (dark blue circle) stands out due to its exceptionally large tank size and the number of potential customers nearby.

7 Assessing Outage Potential

After reviewing actual outages, we examine outage potential from intentional attacks based on the access network maps and customers connected to each CO.




Figure 19: Causing both entry AggCOs to fail would disconnect over a million people in 59% of the regional access networks we study.




Figure 20: An EdgeCO outage would disconnect thousands to tens-of-thousands of customers for the EdgeCOs in our maps.

Customers in Each Regional Network. First, we examine the potential fallout from an attack that disconnects an entire access network, i.e., the entry AggCOs for the network. Nearly all regional access networks in our maps rely exclusively on two entry AggCOs to bridge customers to the Internet, and we can often precisely locate them remotely. If an attacker disables both entry AggCOs, it would disconnect more than 100K customers in all but a single region, and disconnect over 1M wireline customers in 59% of the regions (Fig. 19). The Nashville outage also showed that entire access network outages can disconnect wireless customers that rely on the access network to reach the mobile packet core. The potential to disconnect millions of people, as well as other services that rely on the access networks, makes the regional network itself a compelling target for attack.

Customers Connected to EdgeCOs. EdgeCOs present a softer target for intentional attack than AggCOs; operators indicated they are typically less fortified and might not have continual staff presence. According to our inferred maps, an attack that disables an EdgeCO would disconnect thousands or tens-of-thousands of customers for 92% of the EdgeCOs (Fig. 20), with a median of 12.7K customers. While EdgeCO failures disconnect their wireline residential and business customers, operators told us that wireless customers would often remain connected through nearby cell towers.

	ISP 1	ISP 2	ISP 3
Single AggCO	11.3%	31.3%	100%

Table 3: Percentage of customers that ultimately rely on a single AggCO or backbone PoP. These customers are especially susceptible to natural and intentional disconnections.

Customers Reliant on Single AggCO. Finally, we analyze the fraction of customers that rely on a single AggCO in each network (Table 3), as these customers are especially susceptible to natural outages or intentional attack. In ISP 3, all EdgeCOs connect to multiple AggCOs, but the two regions we investigated rely on a single backbone PoP, so all customers ultimately rely on a single facility. For ISPs 1 and 2, each region relies on multiple backbone PoPs. In ISP 1, some of the smaller regions rely on a single entry AggCO that connects to multiple backbone PoPs, and the customers in these regions lack redundant paths. All of the ISP 2 regions have multiple entry AggCOs, but many subregions connect EdgeCOs to only one AggCO. This topology leaves 31.3% of the customers reliant on a single AggCO, a nearly $3\times$ increase compared to ISP 1.

8 Mitigations and Trade-offs

Our case studies and evaluation of targeted attacks reveal that ISPs are often not prepared for physical attacks on their regional infrastructure. We discussed the threat of intentional physical attack against COs with network operators, who were generally surprised at the level of detail we could reveal. The operators agreed that the threats exist but were unsure how to mitigate them cost effectively. In this section, we review potential mitigations that we discussed with access network operators, along with their perceived drawbacks, to inform future efforts to better secure these critical networks.

Operators consider the possibility of targeted attacks but face inherent tensions between the goals of decreasing the cost and complexity of network deployment, operation, repair, and defending against attacks. Our discussions revealed that the primary concerns for network operators are the cost and complexity of proposed mitigations, as well as retaining their ability to recover from common failure modes. Proposed mitigations that do not account for these concerns are unlikely to gain traction. Below we present the trade-offs operators identified in undertaking five potential mitigations to the attacks we consider.

Hide Locations of Central Offices. The easiest way to cause widespread outages is to find a CO and disconnect either the power or fiber. There are two straightforward ways to precisely locate a CO: searching around a targeted area for the provider’s signage on buildings, or search public databases




Figure 21: A CO’s fiber and power are visible from the street.

to find records of buildings belonging to the targeted provider.

Providers can practice security through obscurity by hiding the location of COs. This is an inexpensive way to hide the infrastructure as the cost will primarily be labor to remove signage from buildings. However, operators thought this could introduce many hidden costs. Operators told us they rely extensively on field technicians and contractors, and removing signs from CO buildings will make it harder for them to easily find the building in case of a problem. Operators also told us that COs are often unstaffed, so signage helps the public report problems to the ISP, such as when a building is on fire.

Similarly, providers can remove CO listings from public databases to prevent an attacker from remotely learning CO street addresses. However, the drawback is that public records of fuel-spill accidents are important for public health and environmental protection, leading governments to mandate them. There may be middle ground where some hazardous materials listings are obfuscated in public records so as to not reveal the purpose of the facility, or its owner. Costs would include paying administrative staff to both obfuscate and reveal the records when deemed necessary.

Hide CO Fiber/Power Lines. Once an attacker finds a CO to target, it is easy to locate fiber and power serving that CO. Fig. 21 shows an example of an EdgeCO that appears to have multiple fibers and power entering its premises on one pole just outside the CO entrance.

Burying the fiber and mains power into a CO—especially some distance away—could prevent an attacker from immediately finding the power and transport after locating a CO. Buried cables are also better protected from the elements than aerial cables. Unfortunately, operators told us that the costs of underground fiber (\$25–75K per mile [38, 39]) and power (~\$500k per mile [40]) are significant. Additionally underground cables are considerably more expensive to repair.

ISPs also label their fiber infrastructure with tags, including underground cable runs, and removing labels from fiber runs would make it harder for attackers to identify fiber belonging to a particular ISP. Operators told us they label the fiber to

prevent accidents and shorten repair time, so removing the labels would likely increase the number and duration of outages due to more common failure modes. This change also introduces the cost of removing labels on splice boxes placed at least every $\sim 1,000$ ft along fiber runs [41].

Increase Last-mile Redundancy. Some access networks do not include redundancy in their shared last-mile links, so a single fiber cut can take thousands of customers offline. Providers could add last-mile redundancy by adding a redundant connection back to the EdgeCO using a ring topology. The primary cost would be the extra network interfaces in the EdgeCOs ($\sim \$24\text{K}$ per 20K customers for CMTS [42]) and re-deploying last-mile fiber in a ring. ISPs could further improve redundancy by connecting customers to two EdgeCOs rather than one. Operators told us that some business customers pay to connect to multiple EdgeCOs, but that doing so for all customers is cost prohibitive.

Another approach is adding backup cellular connectivity to customer premises equipment. Costs include modem equipment and service plans. However, ISPs need to ensure the cellular backup link fails independently. This requires ISPs to provide more transparency about how their wireline access network is used for backhaul in mobile networks.

Make Access Networks Passive. COs depend on both power and fiber for connectivity. Removing the dependency on power would make networks more resilient, and remove an attack vector. Passive network equipment (e.g., optical splitters) are already used in the last mile. It may also be feasible to replace powered routers and CMTSes in EdgeCOs with entirely passive components driven by an AggCO. This technology has not yet been developed, and likely requires longer-term research to develop new passive network technologies. This solution would also incur the cost of upgrading network infrastructure across EdgeCOs.

Hide Access Networks in Measurements. As we demonstrate, an attacker could learn physical topology from wardriving while performing ICMP traceroutes. It is possible to randomize IP address assignment within a given region making it more difficult for an attacker to geolocate infrastructure and users, but operators told us that doing so adds significant network management complexity. ISPs could also disable ICMP responses from their router infrastructure and remove reverse DNS, an inexpensive mitigation. However, this has a key drawback: operators told us that they and their customers rely on traceroute and reverse DNS to troubleshoot and diagnose problems. Also, we demonstrate that it remains possible to find COs with other methods (§6).

9 Related Work

The Internet is designed to be able to route around failures [43], yet large-scale failures are known to occur [20,

22, 44, 45]. Diverse factors cause failures including human error [46, 47], natural phenomena such as earthquakes [48], weather [20], solar activity [49], and equipment failure [50]. Our study focuses on vulnerabilities in access networks, since failures in these networks are challenging to route around.

Attempts to map topological diversity and understand physical network infrastructure vulnerabilities typically focused on backbone networks [51–56] and submarine cable networks [57, 58]. Analytic and probabilistic models were proposed to estimate the risk and survivability of physical attacks [59–61] and natural disasters [62–64]. Our study focuses on the topological diversity of regional access networks; we localized failures to specific EdgeCOs and AggCOs to inform a risk assessment of access network deployments.

This work builds on prior investigations into cyber attacks on related critical infrastructure: the electric grid. Internet access relies on power, and these prior threat assessments reveal how an attacker can force access networks to rely on backup power sources. Researchers found vulnerabilities in SCADA systems that manage electricity networks [65–68], and real-world attacks that caused electricity outages for hundreds of thousands of endpoints [69, 70]. They also examined how an attacker can coordinate demand attacks over the Internet to cause cascading power grid failures [71–73]. Since these attacks require Internet connectivity to execute, this work provides some insight into how the power redundancy built into access networks may make it possible for an attacker to continue performing an attack even as it causes parts of the access network to lose power.

10 Conclusions

Although successful attacks on access networks require sophistication and planning, their impact on modern society—disconnecting critical infrastructure and economic activity—suggests that motivation for such attacks will increase. Given the increase in interdependence with other critical services, we believe our approach to considering resilience of this infrastructure must evolve. As with other critical ecosystems [73, 74], it would be better not to wait for high-profile attacks before undertaking this effort.

Our empirical approach combined new techniques for analyzing access network infrastructure deployments with measurements of weather-induced and accidental large-scale outages to quantify the potential cascading impact of targeted attacks. We discovered new insights into the physical attack surfaces and resiliency limit of regional access network infrastructure. We also analyzed approaches to mitigating risks we identified, and associated tradeoffs in terms of cost and management complexity. Our results can inform risk assessments and reconsideration of approaches to safeguard this critical infrastructure on which our lives now depend.

Acknowledgements

We thank the anonymous reviewers for their insightful comments. We also are appreciative for the elucidating conversations we had with several US Internet access network operators. These helped to validate our findings about access network vulnerabilities, and provided guidance in understanding the feasibility of mitigating physical attacks. This work was supported in part by National Science Foundation grants CNS-2105393, CNS-1901517, CNS-2120399, CNS-2212241, CNS-1705024, ITE-2226460, OAC-2131987, and OAC-1724853. This work was also supported by DARPA CA HR00112020014. Approved for public release; distribution is unlimited. This work does not necessarily reflect the position or the policy of the U.S. Government, and no official endorsement should be inferred.

References

- [1] Z. Zhang, A. Marder, R. Mok, B. Huffaker, M. Luckie, kc claffy, and A. Schulman, "Inferring regional access network topologies: Methods and applications," *ACM Internet Measurement Conference (IMC)*, 2021.
- [2] S. Gate, "\$250,000 reward in phone cable vandalism," <https://www.sfgate.com/bayarea/article/250-000-reward-in-phone-cable-vandalism-3245341.php>, 2009.
- [3] T. Shin and M. Garske, "AT&T cables vandalized, \$250,000 reward offered for information," <https://www.nbcsandiego.com/news/local/att-cables-cut-vandalized-250000-reward-offered/1924959/>, 2012.
- [4] K. Murphy, "The cyberthreat under the street," <https://www.nytimes.com/2015/11/08/sunday-review/the-cyberthreat-under-the-street.html>, 2015.
- [5] D. Wells, "Vandalism blamed for 13-hour Comcast outage in SLC area," <https://www.fox13now.com/news/local-news/vandalism-blamed-for-13-hour-comcast-outage-in-slc-area>, 2021.
- [6] E. C. Webb, "FBI releases report on Nashville bombing," <https://www.fbi.gov/contact-us/field-offices/memphis/news/press-releases/fbi-releases-report-on-nashville-bombing>, 2021.
- [7] J. Gill and S. Posey, "AT&T outages across Tennessee, Kentucky affecting multiple 911 services," <https://www.wkrn.com/news/local-news/att-outages-across-tennessee-kentucky-affecting-multiple-911-services/>, 2020.
- [8] B. Stelter, K. Jones, and H. Silverman, "AT&T working to restore outages after Nashville explosion," <https://www.cnn.com/2020/12/25/us/nashville-explosion-service-disruptions/index.html>, 2020.
- [9] K. Kruesi, M. Balsamo, and E. Tucker, "FBI at home of possible person of interest in Nashville bomb," <https://apnews.com/article/us-news-nashville-coronavirus-pandemic-tennessee-dc6eb653053967a4187f0ca8276d20a8>, 2020.
- [10] R. Rojas, J. McGee, E. Lee, and S. Cavendish, "When Nashville bombing hit a telecom hub, the ripples reached far beyond," <https://www.nytimes.com/2020/12/29/us/nashville-bombing-telecommunications.html>, 2020.
- [11] Kenneally, Erin and Dittrich, David, "The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research," 2012, <http://ssrn.com/abstract=2445102>.
- [12] Dittrich, David and Kenneally, Erin and Bailey, Michael, "Applying Ethical Principles to Information and Communication Technology Research: A Companion to the Menlo Report," 2013, <http://ssrn.com/abstract=2342036>.
- [13] L. Lollman, "Power pole arson causes major internet outage for cox customers in north phoenix," https://web.archive.org/web/20210119053739/https://www.azfamily.com/news/power-pole-arson-causes-major-internet-outage-for-cox-customers-in-north-phoenix/article_9b699a52-fe73-11ea-b6b6-a74c606e5175.html, 2020.
- [14] D. Takahashi, "Comcast can now pinpoint fiber optic cable breaks within minutes," <https://venturebeat.com/2021/07/22/comcast-can-now-pinpoint-fiber-optic-cable-breaks-within-minutes/>, 2021.
- [15] H.-S. Yang, M. Herzog, M. Maier, and M. Reisslein, "Metro wdm networks: performance comparison of slotted ring and awg star networks," *IEEE Journal on Selected Areas in Communications (J-SAC)*, 2004.
- [16] C. Whitaker, "The Comcast enterprise network story," <https://www.slideshare.net/cwhita002/the-comcast-enterprise-network-story>, 2011.
- [17] M. Gunkel, M. Schneiders, S. Vorbeck, W. Weierhausen, R. Leppla, F. Rumpf, R. Herber, V. Furst, and M. Rodenfels, "Aggregation networks: Cost comparison of WDM ring vs. double star topology," in *ONDM*, 2008.
- [18] Cisco, "Introduction to DWDM technology," https://www.cisco.com/c/dam/global/de_at/assets/docs/dwdm.pdf, 2000.
- [19] R. Padmanabhan, A. Schulman, A. Dainotti, D. Levin, and N. Spring, "How to find correlated internet failures," in *Passive and Active Network Measurement Conference (PAM)*, 2019.

- [20] R. Padmanabhan, A. Schulman, D. Levin, and N. Spring, "Residential links under the weather," in *ACM SIGCOMM*, 2019.
- [21] A. Schulman and N. Spring, "Pingin' in the rain," in *ACM Internet Measurement Conference (IMC)*, 2011.
- [22] L. Quan, J. Heidemann, and Y. Pradkin, "Trinocular: Understanding Internet Reliability Through Adaptive Probing," in *ACM SIGCOMM*, 2013.
- [23] J. Hu, Z. Zhou, X. Yang, J. Malone, and J. W. Williams, "CableMon: Improving the reliability of cable broadband networks via proactive network maintenance," in *Symposium on Networked Systems Design and Implementation (NSDI)*, 2020.
- [24] CAIDA, "Archipelago," <https://www.caida.org/projects/ark/>.
- [25] CAIDA, "The CAIDA UCSD IPv4 routed /24 topology dataset," https://www.caida.org/catalog/datasets/ipv4_routed_24_topology_dataset/.
- [26] Comcast, "Comcast maintenance notifications for non-service affecting maintenance activities associated with ethernet transport services and ethernet dedicated internet services," <https://business.comcast.com/terms-conditions-ent/maintenance>.
- [27] Spectrum, "Spectrum maintenance update," <https://www.spectrum.net/support/tv/spectrum-tv-maintenance-update>.
- [28] E. Kuhnke, "Nashville," <https://mailman.nanog.org/pipermail/nanog/2020-December/211081.html>, 2020.
- [29] "Instabridge," <https://instabridge.com/en/>, 2020.
- [30] E. D. Murphy and D. Hoey, "Spectrum internet outage hits thousands of customers in Maine and New Hampshire," <https://www.pressherald.com/2021/04/05/spectrum-restores-internet-service-after-outage-affects-thousands-across-maine-and-new-hampshire/>, 2021.
- [31] Ookla, "At&t network outage in dec 18th, 2020," <https://downdetector.com/status/att/news/356027-problems-at-att/>.
- [32] ABC 7, "Spectrum restores service to SoCal customers after brief outage," <https://abc7.com/spectrum-outage-los-angeles-southern-california/10362980/>, 2021.
- [33] Solano County Sheriff's Office, "Comcast outage," <https://www.facebook.com/SolanoSheriff/posts/comcast-has-reported-an-outage-affecting-1016-users-in-fairfield-94533-and-rio-v/2131795846957800/>, 2021.
- [34] FCC, "Katrina panel," <https://www.fcc.gov/katrina-panel>, 2007.
- [35] The State of Florida, Department of Environmental Protection, "Storage tanks and contamination monitoring," https://prodlamp.dep.state.fl.us/www_stcm/reports/DorFacilities, 2021.
- [36] The County of San Diego, "Citizen access portal," <https://publicservices.sandiegocounty.gov/CitizenAccess/Default.aspx>.
- [37] CalEPA, "Calepa regulated site portal," <https://siteportal.calepa.ca.gov/nsite/>.
- [38] Critter Guard, "Pros and cons of underground fiber optic cable," <https://www.critterguard.org/blogs/articles/pros-and-cons-of-underground-fiber-optic-cable>.
- [39] Bloomer Telephone, "Did you know?" <https://bloomer.net/did-you-know/>.
- [40] Connecticut General Assembly, "Undergrounding electric lines," <https://www.cga.ct.gov/2011/rpt/2011-R-0338.htm>.
- [41] C. of Port St Lucie, "Fiber optic network," <https://utility.cityofpsl.com/media/1097/fiber-protocol.pdf>.
- [42] Cisco, "Cisco uBR-MC3GX60V broadband processing engine with full DOCSIS 3.0 support for the cisco uBR10012 universal broadband router," https://www.cisco.com/c/en/us/products/collateral/video/ubr10000-series-universal-broadband-routers/data_sheet_c78-642540.html.
- [43] D. D. Clark, "The design philosophy of the DARPA internet protocols," in *ACM SIGCOMM*, 1988.
- [44] P. Richter, R. Padmanabhan, N. Spring, A. Berger, and D. Clark, "Advancing the art of internet edge outage detection," in *ACM Internet Measurement Conference (IMC)*, 2018.
- [45] G. Aceto, A. Botta, P. Marchetta, V. Persico, and A. Pescapé, "A comprehensive survey on internet outages," *Journal of Network and Computer Applications*, 2018.
- [46] R. Mahajan, D. Wetherall, and T. Anderson, "Understanding BGP misconfiguration," in *ACM SIGCOMM*, 2002.
- [47] S. Janardhan. (2021) More details about the october 4 2021 outage. [Online]. Available: <https://engineering.fb.com/2021/10/05/networking-traffic/outage-details/>

- [48] A. Dainotti, R. Amman, E. Aben, and K. C. Claffy, "Extracting benefit from harm: using malware pollution to analyze the impact of political and geophysical events on the internet," *ACM SIGCOMM Computer Communication Review (CCR)*, 2012.
- [49] S. A. Jyothi, "Solar superstorms: Planning for an internet apocalypse," in *ACM SIGCOMM*, 2021.
- [50] D. Turner, K. Levchenko, A. C. Snoeren, and S. Savage, "California fault lines: Understanding the causes and impact of network failures," in *ACM SIGCOMM*, 2010.
- [51] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson, "Measuring ISP topologies with rocketfuel," in *ACM SIGCOMM*, 2002.
- [52] R. Teixeira, K. Marzullo, S. Savage, and G. M. Voelker, "In search of path diversity in ISP networks," in *ACM Internet Measurement Conference (IMC)*, 2003.
- [53] S. Knight, H. X. Nguyen, N. Falkner, R. Bowden, and M. Roughan, "The internet topology zoo," *IEEE Journal on Selected Areas in Communications (J-SAC)*, 2011.
- [54] R. Durairajan, P. Barford, J. Sommers, and W. Willinger, "Intertubes: A study of the us long-haul fiber-optic infrastructure," in *ACM SIGCOMM*, 2015.
- [55] M. Ghobadi and R. Mahajan, "Optical layer failures in a large backbone," in *ACM Internet Measurement Conference (IMC)*, 2016.
- [56] S. K. Mani, M. N. Hall, R. Durairajan, and P. Barford, "Characteristics of metro fiber deployments in the US," in *Network Traffic Measurement and Analysis Conference (TMA)*, 2020.
- [57] C. Cao, M. Zukerman, W. Wu, J. H. Manton, and B. Moran, "Survivable topology design of submarine networks," *Journal of Lightwave Technology*, 2013.
- [58] Y. Xie and C. Wang, "Vulnerability of submarine cable network of mainland China: Comparison of vulnerability between before and after construction of trans-arctic cable system," *Complexity*, 2021.
- [59] P. K. Agarwal, A. Efrat, S. K. Ganjugunte, D. Hay, S. Sankararaman, and G. Zussman, "Network vulnerability to single, multiple, and probabilistic physical attacks," in *MILCOM*, 2010.
- [60] S. Neumayer and E. Modiano, "Network reliability under random circular cuts," in *IEEE Global Communications Conference (GLOBECOM)*, 2011.
- [61] O. Gold and R. Cohen, "Coping with physical attacks on random network structures," in *IEEE International Conference on Communications (ICC)*, 2014.
- [62] P. N. Tran and H. Saito, "Enhancing physical network robustness against earthquake disasters with additional links," *Journal of Lightwave Technology*, 2016.
- [63] P. Das, M. Rahnamay-Naeini, N. Ghani, , and M. M. Hayat, "On the vulnerability of multi-level communication network under catastrophic events," in *IEEE International Conference on Computing, Networking and Communications*, 2017.
- [64] J. Oostenbrink and F. Kuipers, "The risk of successive disasters: A blow-by-blow network vulnerability analysis," in *IFIP Networking*, 2019.
- [65] C.-W. Ten, C.-C. Liu, and G. Manimaran, "Vulnerability assessment of cybersecurity for SCADA systems," *IEEE Transactions on Power Systems*, 2008.
- [66] S. Sridhar, A. Hahn, and M. Govindarasu, "Cyber-physical system security for the electric power grid," *Proceedings of the IEEE*, 2012.
- [67] C. Barreto, J. Giraldo, A. A. Cardenas, E. Mojica-Nava, and N. Quijano, "Control systems for the power grid and their resiliency to attacks," *IEEE Symposium on Security and Privacy (SP)*, 2014.
- [68] S. S. P. Mittal and H. V. Poor, "BlackIoT: IoT botnet of high wattage devices can disrupt the power grid," in *USENIX Security Symposium*, 2018.
- [69] D. U. Case, "Analysis of the cyber attack on the ukrainian power grid," *Electricity Information Sharing and Analysis Center (E-ISAC)*, 2016.
- [70] A. Cherepanov, "Win32/industroyer: A new threat for industrial control systems," *White paper, ESET*, 2017.
- [71] C. Barreto, A. A. Cárdenas, N. Quijano, and E. Mojica-Nava, "CPS: Market analysis of attacks against demand response in the smart grid," in *Annual Computer Security Applications Conference (ACSAC)*, 2014.
- [72] S. Amini, F. Pasqualetti, and H. Mohsenian-Rad, "Dynamic load altering attacks against power system stability: Attack models and protection schemes," *IEEE Transactions on Smart Grid*, 2018.
- [73] B. Huang, A. A. Cardenas, and R. Baldick, "Not everything is dark and gloomy: Power grid protections against IoT demand attacks," in *USENIX Security Symposium*, 2019.
- [74] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, and T. Kohno, "Comprehensive experimental analyses of automotive attack surfaces," in *USENIX Security Symposium*, 2011.