Measuring the Impacts of Power Outages on Internet Hosts in the United States

Scott Anderson(✉), Tucker Bell, Patrick Egan, Nathan Weinshenker, and Paul Barford

University of Wisconsin, Madison, Madison, WI, USA
standerson4@wisc.edu

Abstract. Power outages are a well-known threat to Internet communications systems. While Internet service providers address this threat via backup power systems in datacenters and points-of-presence, office buildings and private homes may not have similar capabilities.

This chapter describes an empirical study that assesses how power outages in the United States impact end-host access to the Internet. To conduct this study, the PowerPing system was created to monitor a power outage reporting website and measure end-host responsiveness in the impacted areas. PowerPing collected power outage and end-host responsiveness data over 14 months from June 2020 through July 2021.

The results reveal that power outages affecting 10% or more customers in U.S. counties occur at a rate of about 50 events/day. The outages typically impact about 3,000 customers and services are restored in just under two hours. The end-host responsiveness characteristics for typical power outage events are also reported. Surprisingly, only a weak correlation exists between power outage impacts and service restoration periods versus end-host responsiveness. This suggests that improving backup power for network devices in office buildings and private homes may enable end-hosts to maintain access to Internet service during typical power outages.

Keywords: Power Outages · Internet End-Hosts · Impacts

1 Introduction

The robust availability of Internet service to end-hosts in office buildings and private homes is essential to day-to-day activities. This was highlighted when people moved from offices and classrooms to their homes during the COVID-19 pandemic. Disruptions to service are not merely irksome or inconvenient; they can have real consequences in terms of lost work time and missed opportunities. The importance of connectivity is directly reflected in service level agreements (SLAs) between Internet service providers (ISPs) and their customers, which typically include specific guarantees on service availability [37]. However, several factors determine the realized availability of service to end-hosts.

Published by Springer Nature Switzerland AG 2024
J. Staggs and S. Shenoi (Eds.): ICCIP 2023, IFIP AICT 686, pp. 62–90, 2024.
https://doi.org/10.1007/978-3-031-49585-4_4

Access to the Internet can be impaired by endogenous and exogenous events that affect single users and groups of users in geographic areas. Endogenous events include misconfigurations and equipment failures. Well-understood best practices exist for minimizing the durations and impacts of such events. Exogenous events include natural disasters, infrastructure failures, accidents and attacks. By definition, these events are outside the direct control of Internet service providers and often require other entities to make repairs before service can be restored. Understanding the causes and effects of exogenous events is essential to improving end-to-end network reliability.

Previous studies on the availability of communications systems in the face of exogenous events have focused on retrospective studies of natural disasters such as hurricanes [11], earthquakes [6] and severe weather events [28,34]. The studies present detailed data about the numbers of end-hosts that lost service and the time required to restore service, and also provide road maps for understanding other types of outage events.

This chapter considers the problem of how power outages impact the availability of Internet service to end-hosts. The focus is on wireline Internet services that are typically delivered to end-hosts via cable, digital subscriber line or fiber, and excludes cellular service. Three principal questions are considered: How do power outages impact Internet service to end-hosts on a day-to-day basis? What are the scopes and durations of typical power outages versus service availability? How can understanding typical power outage events inform new techniques and practices to improve network reliability?

This research differs from previous studies on exogenous events that impact communications systems because power outages are common events that occur daily in the United States [7]. There is also a simple solution to power outages – backup power supplies – assuming that the outage durations are relatively modest.

To conduct the study, a measurement system called PowerPing was developed to monitor the `PowerOutage.us` website that publishes current power outage reports by county in the United States [29]. The data was employed to identify U.S. counties to target with active probe-based measurements of end-host responsiveness. This was accomplished using a ZMap-based probing system that operated in two modes. The first mode conducted background probing of IP addresses geolocated to counties across the United States to establish baselines for responsiveness (i.e., end-hosts that respond to probes). The second mode, conducted when outages were identified, sent probe packets to the IP addresses in the targeted areas until power was restored.

Certain challenges were encountered during the development, configuration and deployment of PowerPing. First, timely data about power outages was required to enable active probing of the affected areas to begin as soon as possible after the outages. The `PowerOutage.us` website was leveraged to obtain outage data at 12-min collection intervals. Second, a database containing IP addresses mapped to U.S. counties as probe targets was required. The database was constructed using Esri's ArcGIS [15] to assign IP addresses to counties based on latitude-longitude coordinates provided by MaxMind [24]. The term "end-hosts" used in this work refers to the IP addresses geolocated by MaxMind that, in many cases, may be home routers instead of computers. Third, PowerPing had to be configured to ensure that ZMap would effectively send and receive probes without biasing results. This was accomplished by deploying PowerPing at three CloudLab sites [8] to evaluate vantage point location bias, probe scaling and consistency. Upon conducting a series of tests, it was discovered that deploying PowerPing in a single location with a maximum probe limit of 60,000 packets/s was adequate to obtain consistent results. Finally, a baseline for responsive IP addresses in each U.S. county was established while minimizing the overall probe load on the network.

PowerPing was deployed to collect data over a 14-month period from June 2020 through July 2021. During this period, there were about 330,000 reported power outages with more than 14,000 outages affecting 10% or more customers in the counties. The power outages varied from impacting fewer than 100 customers to impacting 3.7 million customers in Harris County, Texas on February 16, 2021. The power outage durations varied from less than 24 min to an outage in Linn, Iowa that lasted for ten days starting on August 10, 2020. It was discovered that power outages across the United States follow strong diurnal cycles with the largest numbers of events taking place around midday. This can be explained, in part, by power company reports that outages are typically caused by humans through scheduled maintenance, vehicle accidents and high demand [3,10]. Also, power outages that are relatively significant in their impacts are not uncommon. Outages impacting 10% or more customers in a county occur at a rate of about 50 events per day with service restoration typically completed in under two hours.

Active probing of IP addresses conducted after outage reports reveals a wide range of impacts on service availability. The vast majority of Internet service outages impacted fewer than 1,000 end-hosts in the target areas and the service restoration periods were similar to the power restoration periods of about two hours. In aggregate, across all the power outages at a given time, a strong correlation ($R^2 = 0.99$) exists between the numbers of customers impacted by power outages and the numbers of unresponsive end-hosts. However, at the county level, the correlation is not as strong ($R^2 = 0.66$). Possible explanations for these and other results are discussed later in this chapter.

Ethical considerations related to web scraping and active measurements conducted in this research deserve mention. Low-rate scraping of publicly-available data was conducted with the goal of contributing to the public good; no financial benefits were sought or received. No laws were broken to obtain data [41,44] and ethical principles promulgated by major computing organizations such as the Electronic Frontier Foundation (EFF) [23] and Association for Computing Machinery (ACM) [36] were followed. Active measurements followed established methodologies [14,20,34] and the probing methodology limited the impacts to end-hosts and Internet service providers.

2 Related Work

Several techniques have been developed to measure Internet events and outages. These include active probe-based methods [28,31,34], measuring Border Gateway Protocol (BGP) advertisements [12], measuring changes in Network Time Protocol (NTP) traffic [39], passive techniques such as Chocolatine that leverage Internet background radiation [19], combinations of passive and active measurements such as Disco [35] and analyses of Internet service provider logs [32]. These techniques differ from the work described in this chapter because they mainly focus on network outages without considering their causes.

Of particular relevance to this work are two studies on the impacts of weather events on residential Internet service [28,34]. These studies developed and employed ThunderPing to measure end-host responsiveness in areas affected by severe weather events. The methodology described in this chapter was inspired by ThunderPing, but the objective of understanding end-host responsiveness in areas affected by power outages is different. The distinction is significant because severe weather is just one of several causes of power outages, which also include routine maintenance, human operator error, accidents and overload. Unlike the rare weather events studied using ThunderPing, the following sections demonstrate that power outages are common events, with hundreds of outages occurring every day. Additionally, forecasting is an established science for predicting weather whereas power outages are announced publicly only after they occur. Due to these differences, a completely new code base was developed to study how power outages impact end-host Internet service, helping enhance the understanding of the relationships between the two critical infrastructure sectors.

Tools and techniques for conducting active measurements of Internet hosts have evolved significantly over the years and this research was enabled by the advances. Due to hardware and network limitations in sending and processing active network probes, many early active probing studies focused on small sets of representative IP addresses in their regions of interest [16,20]. In contrast, this research actively probes as many IP addresses as possible in select geographic areas during specific events.

Several tools are available for conducting active Internet surveys, including Nmap and Scamper [22]. However, after evaluating the tools, ZMap was selected for its ability to rapidly and accurately scan large numbers of IP addresses in targeted IP subnets [14]. This study has benefitted from the open release of tools to the research community.

3 Datasets

This section describes the three datasets used in the research that cover current power outages, geographic information on U.S. counties and geographic distributions of IP subnets.

3.1 Power Outages

The two primary sources of data on U.S. power outages are the U.S. government and private power generation and distribution utilities. At the federal government level, the U.S. Energy Information Administration (EIA) publishes data about U.S. energy grid operations, including electricity supply, demand, generation and major disturbances and unusual occurrences [42,43]. However, the EIA data suffers from two major drawbacks that made it inappropriate for this research. First, the data is restricted to very large and/or very long duration outages. Second, there are delays of hours to days before data is published.

Private power utility companies are the primary source of U.S. outage data. The United States has more than 1,000 power utility companies that collectively serve more than 140 million customers (households). Many of the power utilities maintain online systems that track the occurrences and current status of power outages for their customers [17,26]. The online systems typically present maps of service areas along with pins showing the geographic locations, numbers of customers without power, reasons for the outages and expected resolution times. However, the maps only display current outage data, not historical outage data. Constantly collecting, parsing and storing current data from numerous power utilities to create a dataset of historical data are most challenging.

The `PowerOutage.us` website aggregates data from major U.S. utilities and presents a consolidated national view [30]. More than 680 power utility companies that serve more than 135 million customers across the United States are monitored to provide data about the numbers and percentages of customers without power in most U.S. counties. The data is updated every ten minutes to accommodate updates posted on utility websites. `PowerOutage.us` lists more than 20 companies and government organizations that use its outage data. The website is frequently quoted in news media reports on major power outages [5,25,29,40,45].

This study has leveraged consolidated data from `PowerOutage.us`. However, certain limitations exist compared with the data provided directly by utilities. Power utilities provide accurate and timely information to support their customers whereas `PowerOutage.us` outage data is likely delayed and can be incomplete. Additionally, `PowerOutage.us` does not track about 500 smaller power utilities with a total of 5.5 million customers, so the results of this study would not reflect all outages in the United States. Nevertheless, it is posited that the large data sample is representative of the power outage conditions experienced by most of the U.S. population.

3.2 U.S County Data

This study has sought to measure the impacts of power outages on end-host responsiveness in U.S. counties in the 48 conterminous states. The U.S. Census Bureau identifies the geographic boundaries of 3,108 counties in the conterminous states [13] and Esri ArcGIS [15] was employed to process this data. The Census Bureau also provides county area, population and population density

Table 1. Top ISPs by subnet count in MaxMind data for U.S. counties.

ISP	ASN	Subnets	Network Type
CHARTER	20115	277,471	Cable/Fiber
TWC-MIDWEST	10796	142,494	Cable
TWC-TEXAS	11427	118,322	Cable
BHN	33363	115,763	cable
TWC-PACWEST	20001	92,052	Cable
COMCAST	7922	81,287	Cable
TWC-CAROLINAS	11426	79,053	Cable
TWC-NORTHEAST	11351	67,616	Cable
TWC-NYC	12271	53,692	Cable
ATT-INTERNET4	7018	45,640	Cable/Fiber
UUNET	701	27,327	DSL/Fiber
CENTURYLINK-US-LEGACY-QWEST	209	23,289	DSL/Fiber
ASN-CXA-ALL-CCI-RDC	22773	16,247	Cable
WINDSTREAM	7029	9,600	DSL/Cable/Fiber
FRONTIER-FRTR	5650	8,615	DSL/Fiber

data that was used in the study. The PowerPing tool developed in this study was designed to employ counties as geographical units because they correspond to the smallest geographic resolution considered by `PowerOutage.us`.

3.3 End-Host IP Subnets

An objective of this study was to probe as many IPv4 addresses as possible in target counties during power outages to measure their impacts and durations. The MaxMind database [24] that provides (approximate) geographic locations (latitudes/longitudes) of variable-sized IP subnets worldwide was leveraged for this purpose. ArcGIS was employed to spatially connect the location data of each IPv4 subnet in the MaxMind database with the U.S. Census Bureau county shapefiles to identify subnets in the counties.

The study considered 1,377,238 variable-sized subnets from MaxMind in U.S. counties that were located in power utility service areas tracked by `PowerOutage.us`. The subnets are owned by 9,441 Internet service providers identified by their autonomous system numbers (ASNs); 44 service providers operated more than 1,000 subnets each.

Table 1 shows the top Internet service providers, dominated by large fixed service residential service providers. This study frequently refers to the responsiveness of "Internet hosts" or "end-hosts." Given the representation of Internet service providers listed in the table, the IP addresses used as probe targets in the study would most likely be home routers. Therefore, if they were responsive during power outages, it was assumed that service was available at the corresponding locations.

The MaxMind dataset limitations include inaccuracies in geolocation information, the incompleteness of the identified subnets, the use of subnet address space by Internet service providers in multiple geographic locations and the understanding of baseline end-host responsiveness in subnets. Additionally, Dynamic Host Configuration Protocol (DHCP) churn, i.e., the rate at which hosts change IP addresses, must be considered. North American Internet service providers do not change IP addresses assigned to end-hosts as frequently as providers elsewhere in the world; most U.S. IP addresses are consistently assigned to the same end-hosts for at least several weeks [27]. To account for IP subnet geographic relocation, the IP subnets from MaxMind were updated three times during the course of the study.

4 PowerPing

The PowerPing system developed for the study has two major functions – identifying the numbers of customers without power in 2,987 U.S. counties and conducting active measurements of end-hosts in counties experiencing outages and those not experiencing outages. PowerPing was written in Python 3.6 and is packaged in a GitHub repository for deployment on an Ubuntu 18.04 server in a cloud-based infrastructure.

During the research, PowerPing was deployed on CloudLab nodes [8]. CloudLab is a distributed computing infrastructure deployed from data centers in Utah, Wisconsin and South Carolina that supports experimental research.

4.1 Power Outage Identification

After a power outage occurs, several steps are taken by a power utility and by `PowerOutage.us` to post information online about the outage event. The power utility identifies the occurrence of the outage and posts the location and number of customers affected on its website. `PowerOutage.us` scrapes the power utility website, identifies the new outage and updates its website. The duration between the occurrence of an outage and its posting on `PowerOutage.us` is uncertain. However, the utility and `PowerOutage.us` have incentives to post outage information as soon as possible.

PowerPing scraped the `PowerOutage.us` website to harvest the total number of customers tracked and the number of customers without power in each of the 2,987 U.S. counties. Since power outages are unpredictable, other than scheduled maintenance, data on all counties was collected in 12-min intervals (epochs) to identify changes. The percentages of customers without power were computed during each epoch for three categories of counties – those experiencing outages impacting 10% or more customers, those in which outages were resolved within four hours, and those experiencing outages impacting less than 10% of the customers.

The start of an outage was set to the first epoch when 10% or more customers in a county experienced an outage. An outage was considered to be resolved when less than 2% of the customers in a county were without power. A county with a resolved outage was maintained as a "county of interest" for four hours after resolution, after which the county was removed from the list of counties of interest. The counties of interest list was maintained to accommodate situations where Internet service was unavailable even after power was restored. The county power outage status during each epoch was passed to the active measurement component of PowerPing.

Three issues must be noted with regard to the outage identification component of PowerPing. First, there were inherent delays between the start of a power outage in a county and PowerPing's identification of a power outage in the county. The delays were mostly external to PowerPing – delays in utilities identifying outages and delays in posting outage information on their public-facing websites. However, there also were delays in `PowerOutage.us` posting outage information on its website. Overall, the delays were due to automated processes, except for situations where customers manually informed utilities of outages. These delays are acknowledged, but it was not possible to reduce them any further. In any case, it is posited that the impact is a modest reduction in outage duration measurements. PowerPing was configured to employ a 12-min interval between harvesting outage information. This interval was identified during initial experimentation because it provided a good balance between the load on `PowerOutage.us`, timeliness of outage update reporting and end-host responsiveness probing (described in Sect. 4.2).

The second issue was that `PowerOutage.us` changed its format during the research, which prevented the harvesting of outage information until the code was adapted to process the reported outages. Future changes to `PowerOutage.us` will require additional PowerPing code updates.

The third issue is that only counties with 10% or more customers without power were considered. This convention was adopted for three reasons – it improved system efficiency by limiting the number of active probes sent during an epoch, it reduced the impact of probe traffic on the network and it helped differentiate the impact of an outage on responsiveness versus IP response churn for outages that affected small numbers of customers. However, there is the risk that outages in some of the largest U.S. counties could have been excluded. Nevertheless, the study identified power outages affecting 10% or more customers in five of the ten largest counties as well as in 13 of the 20 largest counties. End-host responsiveness measurements were performed successfully during the power outages in all 13 counties.

4.2 Active Measurement

The active measurement components of PowerPing implement Pre-Processing and IP address probing to assess end-host responsiveness.

Pre-processing. Efficiency was a key PowerPing design requirement due to the frequency of probing and the large numbers of IP addresses in target areas. Certain pre-processing tasks were implemented to address these issues. The tasks included classifying each IP subnet by county, identifying counties with IP subnets tracked by `PowerOutage.us` and specifying optimal system parameters for data collection, storage and processing.

The MaxMind dataset provides the latitudes and longitudes of IP subnets. The ArcGIS system was leveraged to associate each IP subnet with a state and county from the U.S. Census Bureau shapefiles covering all U.S. counties. Of the 3,108 counties in the conterminous United States, 3,093 counties were identified with subnets from MaxMind within their geographic perimeters.

During each active probe period, up to tens of megabytes of compressed and archived data on ongoing outages and ICMP responses were collected. A standard directory structure, file naming convention and file organization were created for storing and processing the results of each probe period.

End-Host Responsiveness Probing. The IP probing component of PowerPing was informed by prior studies that measured end-host responsiveness [14,20,34]. During each epoch, after U.S. counties were classified according to their power outage status (experiencing outages, recently resolved outages or not experiencing outages), PowerPing identified all the IP subnets in counties with outages, all the IP subnets in counties with outages that were resolved within four hours and all the IP subnets in a select set of counties without outages. Following this, PowerPing sent probes to all the IP addresses in the selected subnets and processed the responses. Finally, it stored the measurement and log data.

All the targeted IP subnets in the three classes of interest were saved in a single "allow list" file for input to ZMap. In accordance with previous research [20], ICMP echo requests were employed as probes. Although ZMap can send probes at a rate of up to 1 Gbps [14], tests of probe rates conducted with network administrators determined that the highest effective rate supported without overwhelming other network traffic was 60 packets/s. When ZMap received a response to a probe, it recorded the responding IP address. Each iteration completed within a variable amount of time, typically five to ten minutes, depending primarily on the numbers of probes sent during an epoch.

Using active probing to identify unresponsive end-hosts required careful consideration. IP address responsiveness is a complex, moving target because end-hosts are naturally cycled on and off the Internet as the devices to which they are attached are moved, and their exact locations are unknown [2]. Therefore, it was difficult to assess how many IP addresses actually existed in a county, how many were typically responsive, how many were responsive prior to an outage, how many were impacted by the outage and how many became responsive after the outage was resolved. To account for these dynamic changes, the end-hosts that responded to all the probes over one-hour each week during a non-outage period were recorded. The corresponding IP addresses were deemed as candidate end-hosts for outages that occurred the same week. If, during an outage period, a response was received from one of the IP addresses, the end-host was considered to be responsive; no response from the IP address led to the end-host being deemed unresponsive.

Another issue was that the probes could be deemed unwanted or even malicious because the packets were sent to IP addresses without the express consent of the administrators. In fact, over more than one year of active probing, only 20 requests to cease the probing of specific IP addresses were received. All the requests were accommodated using ZMap blocklists.

4.3 Deployment

Two important considerations when deploying PowerPing were the selection of measurement vantage points and numbers of probes sent to target IP addresses. Some previous studies have considered these issues [14,20,28,31]. In particular, Wan et al. [46] found that scanning from two vantage points with a single probe increased the network coverage from 95.5% to 98.3%. Additionally, sending two probes instead of one probe increased network coverage from 95.5% to 96.9%.

PowerPing was configured to send one probe from one vantage point to each target IP address during an epoch. This decision could result in false negative responses, but it was made for four reasons. First, since power outages are common events, it is important to limit the impacts of PowerPing probing on the networks. Second, severe power outages that impact wide geographic areas could involve ten million or more end-hosts. Probing such large numbers of end-hosts would push PowerPing up against the 12-min intervals of collection epochs; sending multiple probes would certainly exceed the 12-min collection epochs. Third, there is very little information gain from sending multiple probes instead of a single probe; specifically, network coverage increases from 95.5% for one probe to just 96.9% for two probes. Fourth, Wan et al. [46] observed that vantage points located in the same country as end-hosts have marginally better coverage than vantage points located outside the country and the study described in this chapter only considered end-hosts in the United States.

To verify the design choices, a single server was set up at each of the three CloudLab nodes located at the University of Wisconsin, University of Utah and Clemson University. The servers ran PowerPing to identify power outages and conduct active probing of IP addresses in the impacted U.S. counties. The servers were configured with the same list of IP subnets for each county and were employed simultaneously for one week.

During the testing, differences in the numbers of probe replies received by the servers were observed. Experimentation with different configuration parameters revealed that reducing the ZMap probe rate yielded consistent response rates between the Wisconsin and Utah nodes, but the Clemson node had a consistently lower response rate. However, reducing the ZMap probe rate would increase the time to complete a round of sending probes and processing the responses, limiting the number of IP addresses that could be actively probed during the 12-min epochs.

The difference in active probe network coverage between the CloudLab servers in Wisconsin and Utah was investigated from October 16 through Octo-

Table 2. Network coverage and percentage measurements during power outages.

Network Coverage	Cumulative Percentage (Wisconsin)	Cumulative Percentage (Utah)
99%	90.52%	81.24%
98%	95.66%	93.66%
97%	97.37%	95.64%
96%	97.94%	96.58%
95%	98.24%	97.56%
90%	98.69%	99.74%

ber 25, 2020. During the ten days, each server conducted 10,414 active probe measurements during power outages in 179 counties across 37 states. Consistent with the probing methodology, each server sent a single ICMP probe to each targeted IP address. Three metrics were computed for each county during a measurement period. These included the numbers of IP addresses that responded to each server (R_{wisc} and R_{utah}), total numbers of discrete end-hosts that responded to either server ($R_{total} = R_{wisc} \cup R_{utah}$) and the percentages of end-hosts observed from each server for various network coverage values ($C_{server} = (R_{server}/R_{total}) \times 100$).

Table 2 shows the cumulative percentage measurements taken during outages with indicated network coverage from vantage points at CloudLab sites in Wisconsin and Utah from October 16 through October 25, 2020. In particular, the percentages of end-hosts observed for a network coverage of 97% were $C_{wisc} = 97.37\%$ for CloudLab Wisconsin and $C_{utah} = 95.64\%$ for CloudLab Utah.

Figure 1 shows the total numbers of responses to servers at CloudLab Wisconsin and CloudLab Utah from end-hosts in target counties during power outages from October 16 through October 25, 2020. Specifically, the responses to CloudLab Wisconsin (R_{wisc}) versus the responses to CloudLab Utah (R_{utah}) are plotted for each county for each measurement period to show the consistency across measurements for the two servers. The results demonstrate that less than 4.36% of end-hosts would be expected to be improperly identified as unreachable during more than 97% of measurement periods from a single vantage point. It was posited that this was an acceptable level of uncertainty that would not bias the results significantly because power outages are a common daily occurrence and the study was conducted over a period of 14 months. Furthermore, given the minor differences in response rates, employing multiple vantage points or sending multiple probes to each end-host would be an unnecessary use of Internet resources. As a result, the remaining measurements were conducted using a single server at CloudLab Wisconsin.

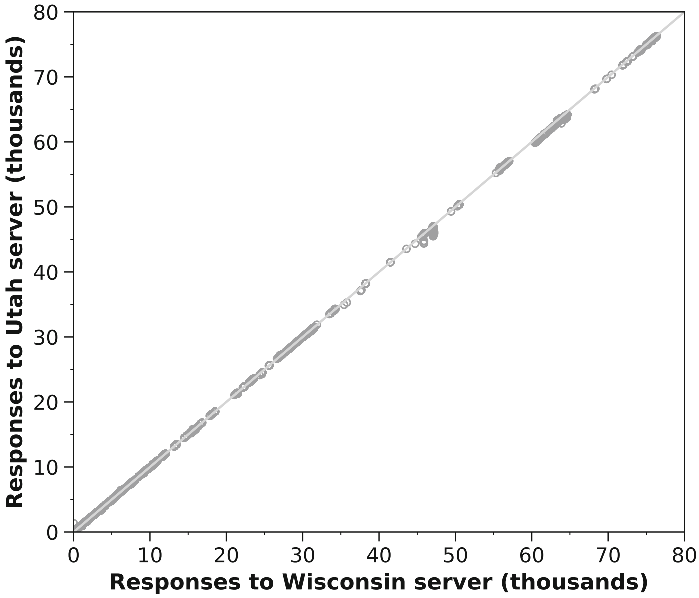

Fig. 1. Total numbers of responses from end-hosts during outages.

4.4 System Design Considerations

While end-hosts require power for operation, there are several reasons why hosts may be reported as responsive during power outages. One reason is the delays between outage occurrences and outage reports on `PowerOutage.us`. Since most power outages are short-lived, they may have already been resolved before they were recognized by PowerPing. Also, the reported numbers of customers affected may not accurately reflect the actual numbers and locations of customers impacted by power outages. For some outages, it was observed that power companies do not update the numbers of customers impacted frequently enough. Instances were routinely observed where the numbers of reported customers with outages did not change, but PowerPing probes had varying response rates. In these instances, the numbers of end-hosts responding to probes may provide more accurate indicators of the numbers of customers without power.

Another reason is that most U.S. counties have multiple power utilities. Although PowerPing determined the number and percentage of customers without power at a given time in a county, it did not distinguish between customers served by different utility providers. Additionally, it was not possible to match individual IP addresses to the utilities that provided power to customers. Active measurements were limited to IP subnets located in counties experiencing power outages, but it was not possible to ascertain that the probed IP addresses belonged to customers impacted by the outages.

Finally, some customers may have used backup power devices such as uninterruptible power supplies for their Internet routers. Internet service providers also maintain redundant power devices and/or backup generators for their network equipment. When customers and Internet service providers utilize backup power during outages, the end-hosts may maintain Internet connectivity during power outages.

While the factors discussed in this section lead to measurement uncertainty, it can be argued that the findings are statistically meaningful because power outages are common events and measurements were collected and analyzed over 14 months. During this time, more than 330,000 outages in 2,495 counties across 48 states were posted on `PowerOutage.us`. Also, by focusing on about 14,000 outages impacting 10% or more customers in counties, nearly all the events with the factors discussed in this section were eliminated. As a result, the negative impact on the findings of this study is expected to be minimal.

5 Results

This section presents the study results that include the characteristics of end-host responsiveness in the absence of power outages (baseline), characteristics of power outages and characteristics of end-host responsiveness during power outages.

5.1 Baseline End-Host Responsiveness

Establishing a baseline of end-host responsiveness in the absence of outages for each U.S. county was essential to the study. The baseline indicates the number of IP addresses as well as the specific IP addresses in each county expected to respond to PowerPing probes. The baseline is employed in the impact and recovery analyses discussed in Sect. 5.3.

The possibility of using existing datasets to identify live end-hosts in subnets was considered. One measurement dataset provides an "IP address space hitlist" upon selecting a single IP address for any /24 subnet to represent all the end-hosts in the subnet [1]. Another dataset provides responsiveness data for hosts running specific services such as HTTP, HTTPS and SSH, but it only collects measurements once a day and does not test responsiveness using ICMP probes [4]. Although these datasets are useful for understanding Internet characteristics at the network subnet and service levels, baseline data was collected during the study due to its focus on individual end-host responsiveness.

Baseline measurements were performed periodically to quantify the responsiveness of end-hosts in each county during non-outage periods. A separate server was set up in the same CloudLab site as the PowerPing server. ZMap was used with the same configuration as the PowerPing server to periodically probe all the IP addresses in each county every ten minutes during a 24-h period. In order to complete probing rounds within ten minutes at a rate of 60,000 packets/s, all the subnets in 100 to 150 counties were selected for probing in each 24-h period. The measurement campaign was conducted for all the counties in the study from August 21 through October 9, 2020.

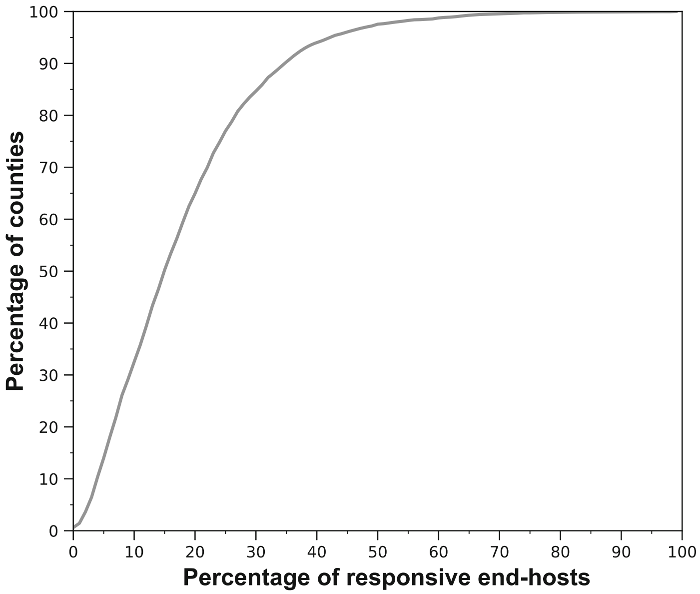

Fig. 2. Cumulative distribution of responsive end-hosts (no outages).

The probing campaign during non-outage periods revealed that the response rate was relatively low in most counties. In an average county, 18.6% of the IP addresses from MaxMind responded to probes. Figure 2 shows the cumulative distribution of the percentages of responsive end-hosts from all the targeted IP addresses in each county during non-outage periods. This is the distribution of responses expected to be received from all the IP addresses in MaxMind. Despite the low response rate, more than 100,000 end-hosts in 183 counties (6% of counties) responded, more than 10,000 end-hosts in 911 counties (29% of counties) responded and more than 1,000 end-hosts in 2,171 counties (70% of counties) responded.

Figure 3 shows a plot of county population from the U.S. Census Bureau versus the number of expected responses from end-hosts for each county during non-outage periods. As expected, the most responses were received from counties with the largest populations: Los Angeles County, California (3.4 million), Cook County, Illinois (1.5 million) and Maricopa County, Arizona (1.2 million). However, the counties with the largest fractions of responses were not associated with the largest metropolitan areas.

Figure 4 (left) shows the distribution of IP addresses in MaxMind by county. Figure 4 (center) shows the numbers of end-host ping responses received. Figure 4 (right) shows the percentages of hosts in MaxMind that responded to target pings.

The numbers of responses received from the counties were consistent over the 24-h measurement periods. ZMap was configured to send one probe to each tar-

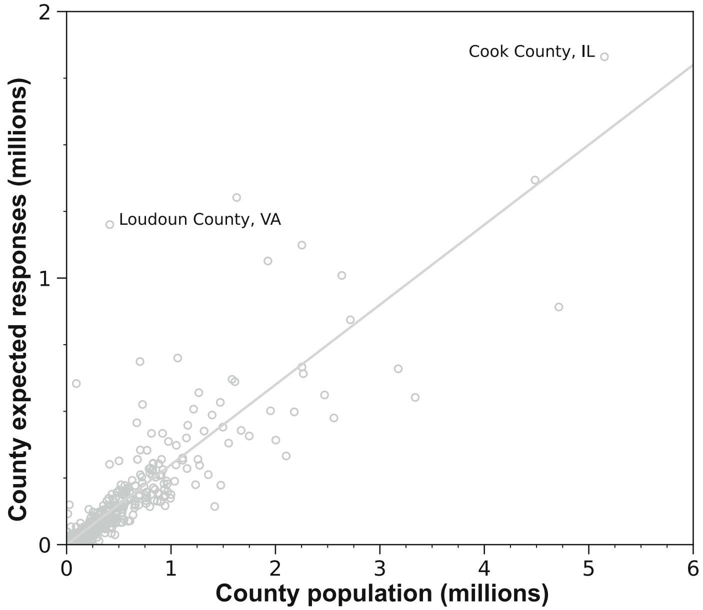

Fig. 3. County population versus expected end-host responses (no outages).

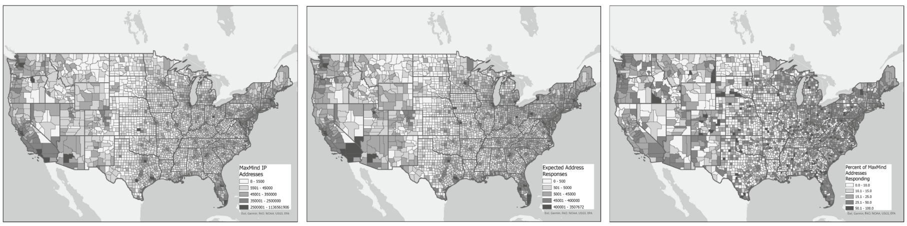

Fig. 4. Geographic distribution of responses by county (no outages).

geted IP address. For this configuration, the ZMap authors measured a 2% single packet loss rate [14]. Figure 5 shows the cumulative distribution of the percentage differences between the maximum and minimum numbers of responses from counties without power outages over the 24-h measurement periods. In 2,766 of 2,987 targeted counties, a difference of 10% or less was measured in the maximum number of responses compared with the minimum number of responses. The differences were less than 2% in 1,391 counties. Diurnal variations in the numbers of responses were not observed.

The baseline of IP address responsiveness was re-evaluated by selecting a uniform random sample of subnets from each county and conducting an additional measurement campaign over a one-month period from March 13 to April 13, 2021. Three to five subnets were selected from the MaxMind dataset for each

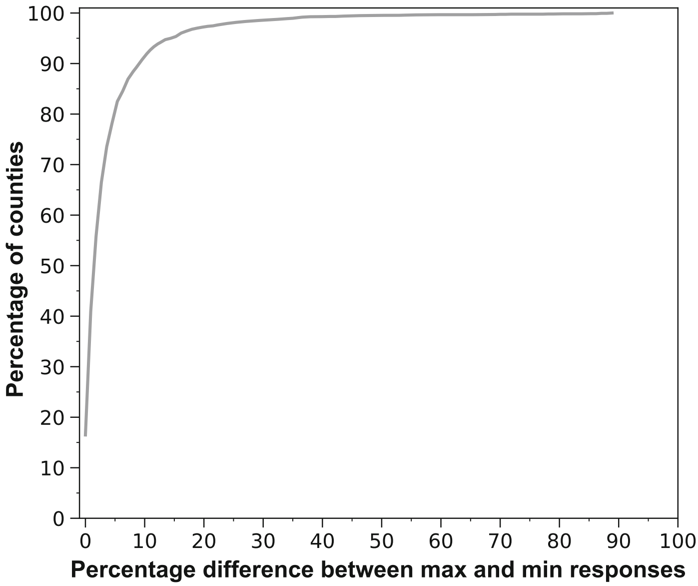

Fig. 5. Distribution of percentage differences in max-min responses (no outages).

county. ZMap was configured to send one ICMP probe to every IP address in the selected subnets every 12 min.

The results obtained during the week of March 14, 2021 were typical of those seen during the additional measurement campaign. During that week, 2,709 counties did not have any power outages that impacted 10% or more customers and 840 active polling iterations were conducted. Although only 5.8% of the end-hosts responding to at least one probe responded to every probe that week, as many as 76.2% of the end-hosts responded to 99% of the probes and 91.3% of end-hosts responded to at least 90% of the probes. Only 7.3% of end-hosts responded to less than 80% of the probes over the entire week. These results indicate consistently high responsiveness levels from IP addresses during non-outage periods.

5.2 Power Outage Characteristics

Power outages are relatively common occurrences and most outages follow distinct cycles. An outage begins with an event that interrupts normal service. Power utilities identify several events that cause outages, the most common being severe weather and motor vehicle accidents. Other events include equipment failures, wildlife interference, high demand, damage from construction work and maintenance [3,10]. An outage is detected by a utility via automated means or customer reports. The utility then deploys the necessary assets to restore power. The outage may be resolved simultaneously for all impacted customers or it may be resolved incrementally for groups of customers.

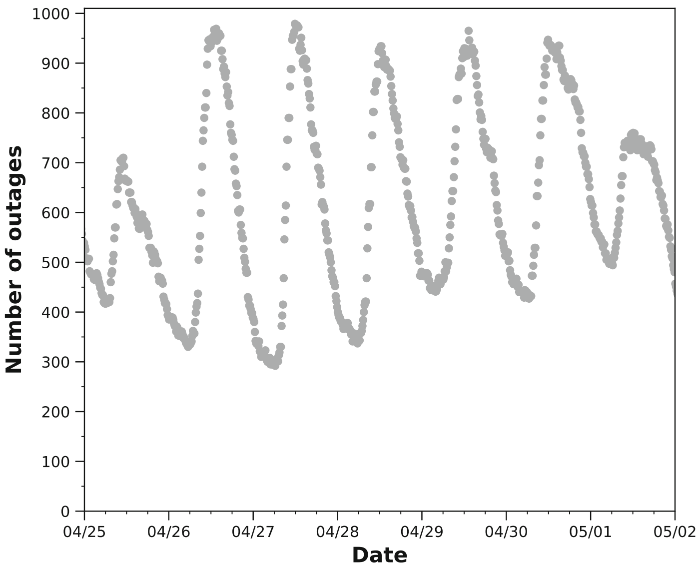

Fig. 6. Power outages detected during measurement epochs over one week.

Most power companies maintain online trackers of known power outages. The online trackers are updated after outages are detected. Additional updates to outages track changing conditions on the ground. Complete resolution of an outage may not be updated on the tracker at the same time there is resolution on the ground.

Figure 6 shows the number of power outages detected in each measurement epoch during the week of April 25, 2021. As shown in the figure, power outages in the United States typically have a strong diurnal pattern. Most outages occur during the early afternoon. A steady increase in the number of reported power outages is seen from early morning until early afternoon. From early afternoon to late evening, a steady decrease is seen in the number of reported outages. The fewest outages occur late at night. This is consistent with previous observations that the majority of power outages are caused by maintenance or operational disturbances, which are more likely to occur during business hours [21]. During the study, fewer power outages were observed on weekends and major holidays. Typically, there were about 50 power outage events per day across the 48 conterminous states that impacted 10% or more customers in a county.

Most outages were short lived – 80% were resolved in under one hour and 90% were resolved in under two hours. A small number of long-duration outages pushed the average outage duration to just under two hours. Figure 7 shows the cumulative distributions of outage durations during each week from September 27 to October 18, 2020. Each outage duration was computed from the time of first report on `PowerOutage.us` to the time the outage was removed from the site.

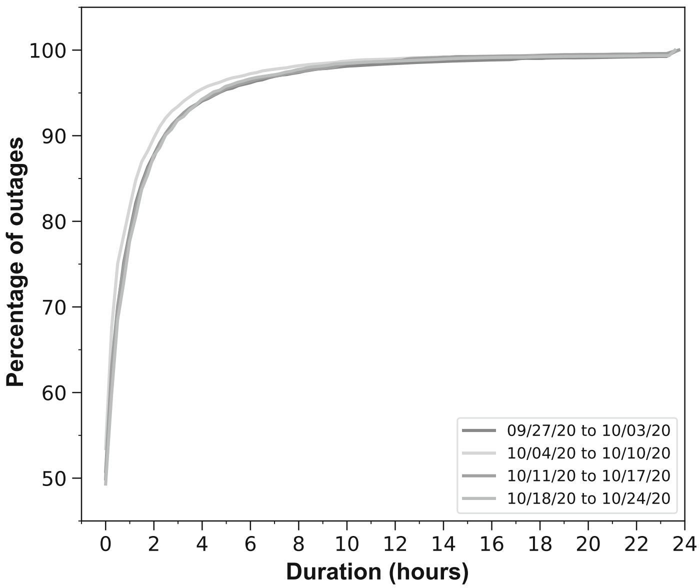

Fig. 7. Cumulative distributions of outage durations over a four-week period.

While the number of power outages follows a consistent diurnal pattern, the study revealed that during most weeks there were strikingly different patterns in the numbers of impacted customers. Figure 8 plots the numbers of customers without power in large counties (top plot), medium counties (middle plot) and small counties (bottom plot) during the week of April 25, 2021. As expected, counties with the largest populations had the most customers without power. The sharp spike in the number of customers without power on the night of April 30, 2021 was due to strong winds and rain that caused power outages along the East Coast [9]. During the study, numerous instances of spikes in the numbers of impacted customers were observed that did not follow diurnal patterns. Also, many counties had small numbers of customers without power (typically fewer than 10) during most probing epochs.

In summary, the study revealed that power outages follow strong diurnal patterns, with most outages occurring on weekday afternoons. Nearly all outages are resolved within an hour. Additionally, the daily numbers of impacted customers have more variations than the daily numbers of outages.

5.3 End-Host Responsiveness During Outages

Two key metrics were identified to assess the impacts of power outages on end-host responsiveness. The first metric is impacts – the percentages of end-host IP addresses (versus the background response rates for counties) that are unresponsive to probes during a power outage. The second is durations – the lengths of

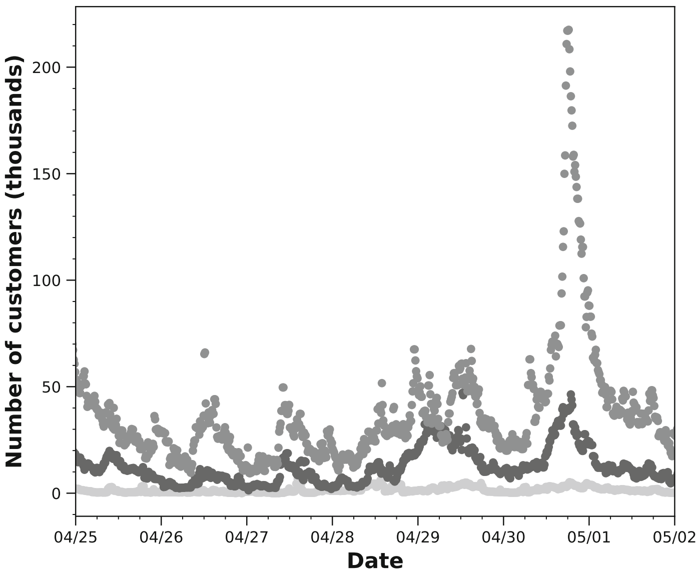

Fig. 8. Customers without power in large, medium and small counties.

time end-host IP addresses in counties are unresponsive during and after power outages.

As far as impacts are concerned, the study revealed that most power outages affect fewer than 1,000 end-hosts. Figure 9 shows the distributions of the numbers of unresponsive end-hosts in counties experiencing power outages for four typical weeks during the study. In most weeks, 80% of the outages affected less than 1,000 end-hosts. The week of October 4, 2020 is a clear outlier. The reason was Hurricane Delta, which struck the Gulf Coast on October 9, 2020, leading to power outages and large numbers of unresponsive end-hosts [33].

A positive correlation was observed between the aggregate numbers of customers without power and aggregate numbers of unresponsive end-hosts during power outages across all counties during each measurement epoch. Specifically, Fig. 10 shows the scatter plot for total customers without power versus total unresponsive end-hosts from February through July 2021 with a correlation $R^2 = 0.99$. However, the correlation results are skewed by the Texas power outages that occurred over four days in February 2021 and impacted up to 4.5 million customers [38].

On shorter timescales (month-long periods), R^2 correlations ranging from 0.19 (April 2021) to 0.99 (February 2021) were obtained. For comparison, a correlation $R^2 = 0.76$ over the same six-month period was obtained when the week of the Texas power outages was excluded. Figure 11 shows the numbers of customers without power (lighter shade) versus numbers of unresponsive end-hosts (darker shade) in counties with major power outages during the week of April 25, 2021. The graph shows an example of temporal variations across all the

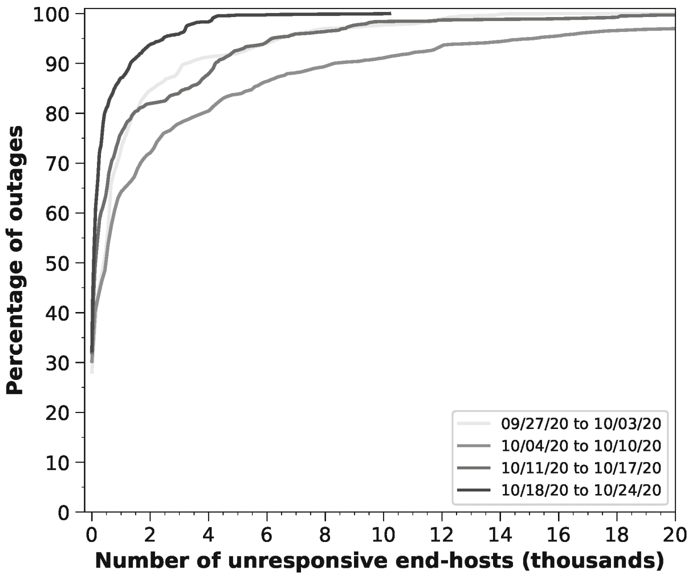

Fig. 9. Cumulative distributions of unresponsive end-hosts in counties.

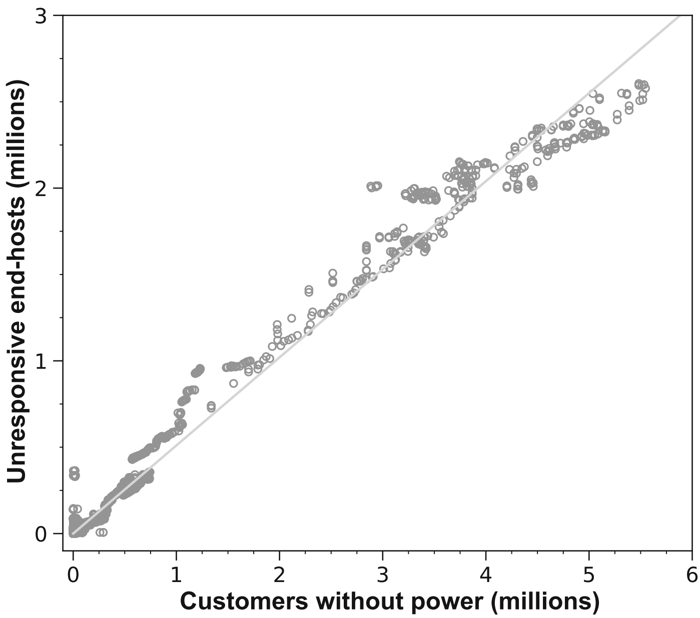

Fig. 10. Customers without power versus unresponsive end-hosts.

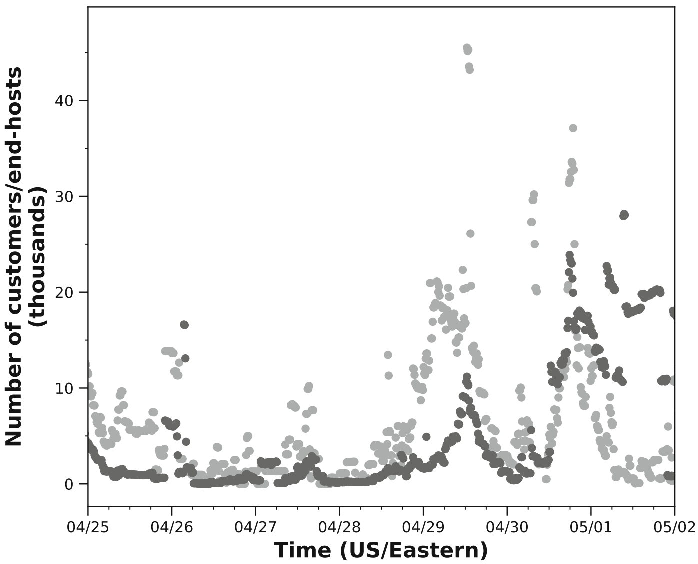

Fig. 11. Customers without power and unresponsive end-hosts (major outages).

major power outages during that week. The number of customers without power corresponds closely with number of unresponsive end-hosts, with the exception of May 1, 2021 due to the Texas power outages.

Unlike at the national aggregate level, it was observed that power outages at the county level often impacted customers without affecting the responsiveness of end-hosts. Across all the power outages from February to July 2021, a correlation $R^2 = 0.66$ was computed for the numbers of customers without power in counties versus the numbers of unresponsive hosts in the corresponding counties. For example, over the week of April 25, 2021, 118 power outages were observed to have increased end-host unresponsiveness during the outages and 98 power outages were observed to have no increase in end-host unresponsiveness.

At the county level, distinct patterns were observed when comparing the percentages of customers without power with the percentages of unresponsive end-hosts. The patterns were placed in four outage classification categories:

- **Category 1:** The percentages of unresponsive end-hosts roughly follow the percentages of customers without power throughout the outages.
- **Category 2:** The percentages of unresponsive end-hosts remain largely unchanged throughout the outages.
- **Category 3:** The percentages of unresponsive end-hosts change smoothly during the collection periods throughout the outages whereas the percentages of customers without power remain constant or undergo frequent large changes.
- **Category 4:** The percentages of unresponsive end-hosts diverge considerably from the percentages of customers without power.

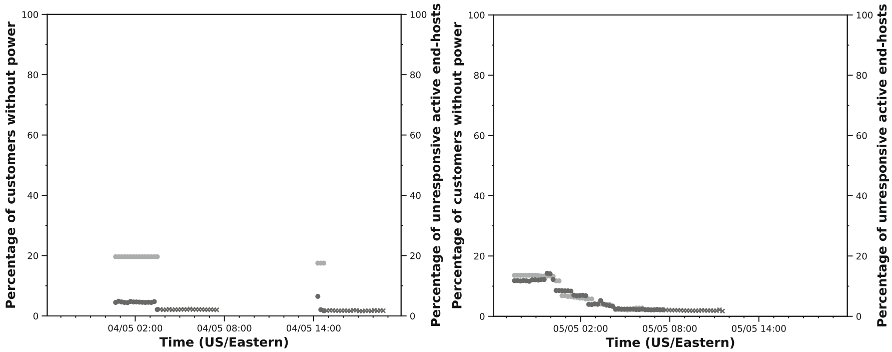

(a) Dawson County, Montana. (b) Barbour County, Alabama.

Fig. 12. Category 1 power outages.

The four categories of power outages and their frequencies of occurrence are discussed in the remainder of this section.

Figure 12 compares the percentages of customers without power against the percentages of unresponsive end-hosts when the unresponsive end-hosts closely track customers without power. Specifically, it presents data for two Category 1 outages in Dawson County, Montana on April 5, 2021 and in Barbour, Alabama on May 5, 2021. As the percentage of customers without power in a county increases, the percentage of unresponsive end-hosts increases, and vice versa. This behavior was observed in geographically-distinct areas for counties of various sizes (by area and population) for outages of varying durations and intensities, as well as for counties with different numbers of subnets and expected numbers of end-hosts that respond to active probing.

Figure 13 shows the behaviors of Category 2, 3 and 4 power outages that do not align with the intuitive behavior of Category 1 power outages. Figures 13(a) and (b) present data for Category 2 outages in Forest County, Wisconsin on March 6, 2021 and in Camp County, Texas on March 15, 2021. In the Forest County outage, the percentage of customers without power decreased smoothly from about 15% to 5% over about one hour, but the percentage of unresponsive end-hosts did not vary during or after the outage. Similar behavior is seen in the Camp County outage, where two different outages impacted almost 40% of the power utility customers. However, no effects on the responsiveness of end-hosts were measured during either outage.

Figures 13(c) and (d) present data for Category 3 outages in McDonald County, Missouri on May 6, 2021 and in Lake County, Michigan on June 18, 2021 where the percentages of customers without power stayed almost constant throughout the outages, but the percentages of unresponsive end-hosts varied. The McDonald County outage shown in Fig. 13(c) lasted about ten hours with a constant 18% of customers without power. Towards the beginning of the outage, approximately 12% of end-hosts were unresponsive; the percentage of unresponsive end-hosts decreased to about 5% approximately two hours into the out-

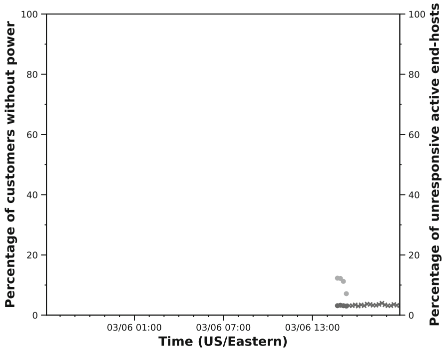

(a) Forest County, Wisconsin.

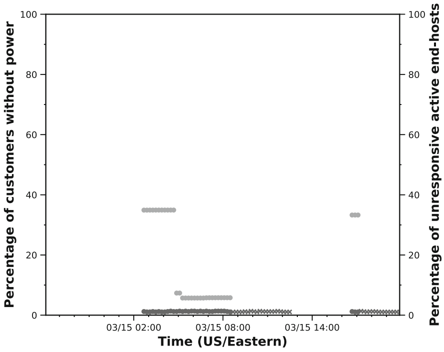

(b) Camp County, Texas.

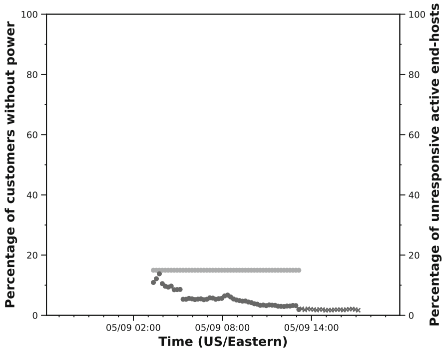

(c) McDonald County, Missouri.

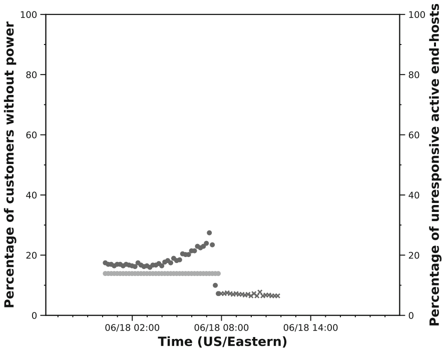

(d) Lake County, Michigan.

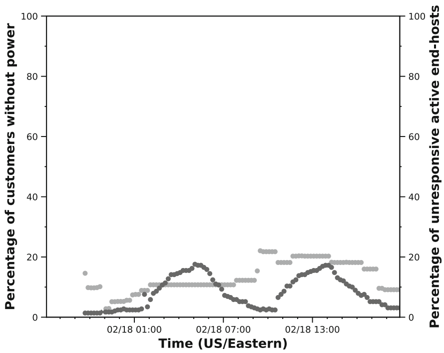

(e) Carroll County, Mississippi.

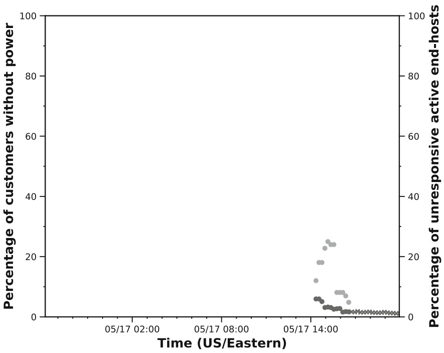

(f) Seminole County, Oklahoma.

Fig. 13. Category 2, 3 and 4 power outages.

age, remained near-constant for three hours and then decreased slowly over the remaining five hours of the outage. After the outage was resolved, a consistent percentage of unresponsive end-hosts remained.

Table 3. Occurrences of the four categories of outages in March 2021.

Category	Outages	Percentage
1	402	38.8%
2	423	40.8%
3	34	3.3%
4	177	17.1%

The Lake County outage in Fig. 13(d) was similar to the McDonald County outage, but it impacted an increasing number of end-hosts from 20% gradually up to nearly 30% at the end of the outage. When the outage was reported as resolved, an immediate drop in the percentage of unresponsive end-hosts occurred. In both the Category 2 situations, it is surmised that active probing was a better predictor of customers with outages than what was reported by the utilities. However, as shown in Table 3, Category 3 outages are the least common of the four categories of outages.

Figures 13(e) and (f) present data for Category 4 outages in Carroll County, Mississippi on February 18, 2021 and in Seminole County, Oklahoma on May 17, 2021 where the percentages of customers without power and the percentages of unresponsive end-hosts varied differently, but the metric that most accurately describes the ground situation could not be established definitively. The Carroll County outage had a slowly-changing percentage of customers without power, ranging from hours with a consistent percentage of customers without power, which increased or decreased in distinct steps from 5% to 20% of customers without power. In contrast, the percentage of unresponsive end-hosts rose and fell in two distinct hills that peaked at 5 am and 2 pm.

The Seminole County Category 4 outage in Fig. 13(f) lasted approximately three hours during which the percentage of customers without power slowly rose to almost 30% and then dropped distinctly to less than 10% of customers without power towards the end of the outage. The percentage of unresponsive hosts conveys a different story. A peak of nearly 10% of unresponsive end-hosts occurred at the beginning of the outage, which dropped steadily over the duration of the outage until nearly all the end-hosts became responsive at outage resolution.

To quantify the frequencies of occurrence of the four categories of outages, measurements were conducted for all the outages in all the counties during the month of March 2021. Table 3 lists the occurrences of the four categories of outages during the month of March 2021. Category 1 and 2 outages were most common whereas Category 3 were rare. Although this could not be confirmed, it appears that Category 2 outages are manifested when backup power at the Internet service providers and customer locations helps maintain connectivity during power outages.

The durations of end-host unresponsiveness were also computed during and after power outages. It was determined that more than 80% of end-hosts became responsive to active probing within one hour of power outage resolution and 90%

recovered within two hours. For several long power outages where power was restored to customers incrementally over hours or days, similar increases in the numbers of responsive end-hosts were observed as power was restored.

In summary, the study revealed that most power outages impact the responsiveness of less than 1,000 end-hosts. In aggregate, the numbers of unresponsive end-hosts are closely correlated with the numbers of customers without power. However, the correlation is not as strong at the county level. Additionally, unresponsive end-hosts typically became responsive within two hours of outage resolution.

6 Maintaining Communications During Outages

The study results reveal that power outages are frequent events and often last less than two hours. A natural question is whether or not it is possible for customers to maintain Internet connectivity during power outages.

In order to maintain Internet connectivity, three types of devices or equipment must have alternate power sources: end-host devices (computers, televisions and smartphones), home network equipment (modems and routers) and Internet service provider network equipment. Disruptions of one or more device/equipment types would result in disruptions of Internet connectivity.

Some customer devices, such as laptops and smartphones, have batteries that provide hours of service during power outages. Other customer devices, such as printers, game consoles, smart speakers and televisions, do not. Customers may install their own battery backups for many of these devices.

At this time, no U.S. Government regulations require customer devices and network equipment to have built-in battery backups. However, situations arise where voice (telephone) service continues during power outages while Internet service is lost. This can occur when customer modems and routers have battery backups. Some models provide battery backup for voice service but not Internet service. Customers may take steps to ensure uninterrupted Internet connectivity by installing batteries internal to devices when the options are available or by plugging modems/routers into uninterruptible power supplies. Less common customer solutions involve the installation of residence-level batteries, power generators or solar panels.

Private communication with Internet service providers via `nanog.org` revealed that it is standard practice to deploy various levels of backup power for their equipment. These include battery backups that provide uninterrupted service for short-term outages and backup power generators at their aggregation centers and points of presence. Additionally, Internet service providers may provide short-term (several hours) battery backups for local nodes in residential neighborhoods.

Interruption of power supply to devices or equipment at any of the levels would interrupt Internet service. The disruptions would be inconvenient (e.g., loss of access to online gaming and streaming video), problematic (e.g., inability to conduct online banking, shopping and business communications) or critical (e.g., disrupting access to emergency communications services, news and weather reports and medical devices that require Internet access) [18]. Given the ubiquity of laptops and other consumer devices with batteries, the study findings suggest that the availability of backup power for network devices is not geographically uniform across the United States and end-host connectivity during power outages could be improved with backup power for network devices at Internet service providers as well as at customer residences.

7 Future Work

This empirical study has clarified the relationships existing between power outages and availability of Internet service to end-hosts in the United States. Several opportunities are available for future research. The PowerPing system may be deployed in other geographic areas to assess regional variations in end-host responsiveness. However, the challenge to an expanded geographic scope is that power outage information is not always reported accurately or in a timely manner.

The study indicates that power utilities may not always update outage status in a timely manner. However, given the correlations existing between power outages and Internet service outages, active measurements of end-host service availability is an alternative to obtain more accurate pictures of the prevalence and extent of power outages. This approach would require ground truth power outage data from a source such as `PowerOutage.us` and a careful probing strategy that minimizes network impact.

Important next steps are conducting similar studies for cellular service interruptions during power outages and to include end-hosts with IPv6 addresses. One challenge is that PowerPing could not be directly adapted to these studies. In fact, different techniques and tools would be required to measure outages involving these technologies.

This study has focused on the complete loss of power, but it is important to consider situations where power utilities reduce electricity supply to customers. One type of situation is brownouts, which occur when electricity demand exceeds generation capacity. This study did not measure periods of power brownouts. With an adequate real-time dataset on brownouts, it would be worthwhile evaluate the impacts of brownouts on Internet service.

8 Conclusions

This chapter describes an empirical study on how power outages impact Internet service availability to end-hosts. The PowerPing system was developed to monitor active power outages in the conterminous United States and probe end-hosts in IPv4 subnets geolocated in counties with power outages. During the 14-month study period, more than 330,000 power outages were monitored, including almost 14,000 outages – approximately 50 outages per day – that impacted 10% or more customers in U.S. counties. Most power outages were observed to last less than two hours. In the aggregate, a strong correlation was determined to exist between power outage impact and duration and end-host responsiveness; however, the correlations were found to be weak at the county level. The findings highlight the diverse impacts on Internet connectivity at the county level. The results suggest that providing improved backup power sources for network devices, especially for modems and routers in customer residences, may be adequate for end-hosts to maintain uninterrupted Internet service during typical power outages.

All the code and data described in this chapter are available to the research community upon request. The views and conclusions in this chapter are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, expressed or implied, of the National Science Foundation or the U.S. Government.

Acknowledgments. This research was supported by the National Science Foundation under Grant nos. CNS 1703592 and CNS 2039146. The authors also wish to thank the operators of CloudLab for their support in this research.

References

1. ANT Lab: IP Address Space Hitlists, University of Southern California Information Sciences Institute, Marina del Rey, California (2021). https://ant.isi.edu/datasets/ip_hitlists/format.html
2. Bano, S., et al.: Scanning the Internet for liveness. ACM SIGCOMM Comput. Commun. Rev. **48**(2), 2–9 (2018)
3. Bowen, C.: 8 common causes of outages, Edison International, Rosemead, California, 27 June 2016. https://energized.edison.com/stories/8-common-causes-of-outages
4. Censys: Censys, Ann Arbor, Michigan (2021). https://search.censys.io/
5. Chappell, B.: Zeta causes 2 million power outages, speeds its way into Virginia, National Public Radio, 29 October 2020
6. Cho, K., Pelsser, C., Bush, R., Won, Y.: The Japan earthquake: the impact on traffic and routing observed by a local ISP. In: Proceedings of the Special Workshop on Internet and Disasters (2011). Article no. 2
7. Chrobak, U.: The U.S. has more power outages than any other developed country. Here's why, Popular Science, 17 August 2020
8. CloudLab: CloudLab, University of Utah, Salt Lake City, Utah (2023). https://cloudlab.us/
9. Constantino, A., Small, M.: Downed wires, trees, outages as strong winds sweep through DC area, WTOP News, 1 May 2021
10. Constellation Energy: 10 common causes of power outages, Houston, Texas, 24 September 2021. https://blog.constellation.com/2020/08/21/10-common-causes-of-power-outages
11. Cowie, J., Popescu, A., Underwood, T.: Impact of Hurricane Katrina on Internet Infrastructure. Renesys, Manchester, New Hampshire (2005)
12. Dainotti, A., et al.: Analysis of country-wide Internet outages caused by censorship. IEEE/ACM Trans. Netw. **22**(6), 1964–1977 (2014)

13. Data.gov: U.S. Census Bureau TIGER Dataset, U.S. General Services Administration, Washington, DC (2021). https://catalog.data.gov/dataset
14. Durumeric, Z., Wustrow, E., Halderman, J.: ZMap: fast internet-wide scanning and its security applications. In: Proceedings of the Twenty-Second USENIX Security Symposium, pp. 605–619 (2013)
15. ESRI, ArcGIS, Redlands, California (2023). www.esri.com/en-us/arcgis/about-arcgis/overview
16. Fan, X., Heidemann, J.: Selecting representative IP addresses for Internet topology studies. In: Proceedings of the Tenth ACM SIGCOMM Conference on Internet Measurement, pp. 411–423 (2010)
17. Georgia Power, GPC Outage Map, Atlanta, Georgia (2021). https://outagemap.georgiapower.com
18. Grandhi, S., Plotnick, L., Hiltz, S.: An Internet-less world? Expected impacts of a complete Internet outage with implications for preparedness and design. In: Proceedings of the ACM on Human-Computer Interaction, vol. 4(GROUP) (2020). Article no. 3
19. Guillot, A., et al.: Chocolatine: outage detection for Internet background radiation. In: Proceedings of the Network Traffic Measurement and Analysis Conference, pp. 1–8 (2019)
20. Heidemann, J., Pradkin, Y., Govindan, R., Papadopoulos, C., Bartlett, G., Bannister, J.: Census and survey of the visible Internet. In: Proceedings of the Eighth ACM SIGCOMM Conference on Internet Measurement, pp. 169–182 (2008)
21. Hines, P., Apt, J., Talukdar, S.: Trends in the history of large blackouts in the United States. In: Proceedings of the IEEE Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century (2008)
22. Luckie, M.: Scamper: a scalable and extensible packet prober for active measurement of the Internet. In: Proceedings of the Tenth ACM SIGCOMM Conference on Internet Measurement, pp. 239–245 (2010)
23. Mackey, A., Opsahl, K.: Van Buren is a victory against overbroad interpretations of the CFAA and protects security researchers. Electronic Frontier Foundation, San Francisco, California, 17 July 2021
24. MaxMind, GeoLite2 Free Geolocation Data, Massachusetts (2020). https://dev.maxmind.com/geoip/geolite2-free-geolocation-data#accessing-geolite2-free-geolocation-data
25. McWhirter, C.: Delta leaves hundreds of thousands without power. Wall Street J., 10 October 2020
26. Pacific Gas and Electric, PGE Emergency Site - Outage Center, San Francisco, California (2023). https://m.pge.com/#outages
27. Padmanabhan, R., Dhamdhere, A., Aben, E., Claffy, K., Spring, N.: Reasons dynamic addresses change. In: Proceedings of the Internet Measurement Conference, pp. 183–198 (2016)
28. Padmanabhan, R., Schulman, A., Levin, D., Spring, N.: Residential links under the weather. In: Proceedings of the ACM Special Interest Group on Data Communications, pp. 145–158 (2019)
29. PowerOutage.us: About PowerOutage.us, Bluefire Studios, South Portland, Maine (2023). https://poweroutage.us/about
30. PowerOutage.us: United States Power Outage Map, Bluefire Studios, South Portland, Maine (2023). https://poweroutage.us
31. Quan, L., Heidemann, J., Pradkin, Y.: Trinocular: understanding Internet reliability through adaptive probing. SIGCOMM Comput. Commun. Rev. **43**(4), 255–266 (2013)

32. Richter, P., Padmanabhan, R., Spring, N., Berger, A., Clark, D.: Advancing the art of Internet edge outage detection. In: Proceedings of the Internet Measurement Conference, pp. 350–363 (2018)
33. Samenow, J., Livingston, I.: Hurricane Delta by the numbers: 101 mph winds and 9.3-foot surge in coastal Louisiana, Washington Post, 12 October 2020
34. Schulman, A., Spring, N.: Pingin' in the rain. In: Proceedings of the ACM SIGCOMM Conference on Internet Measurement, pp. 19–28 (2011)
35. Shah, A., Fontugne, R., Aben, E., Pelsser, C., Bush, R.: Disco: fast, good and cheap outage detection. In: Proceedings of the Network Traffic Measurement and Analysis Conference (2017)
36. Siegel, A., Grosso, A., Rasch, M., Jarvis, R.: Brief for Amicus Curiae United States Technology Policy Committee of the ACM in Support of Neither Party. Association for Computing Machinery, New York (2020). www.acm.org/binaries/content/assets/public-policy/ustpc-amicus-brief-vanburen-v-us.pdf
37. Sommers, J., Barford, P., Duffield, N., Ron, A.: Multiobjective monitoring for SLA compliance. IEEE/ACM Trans. Netw. **18**(2), 652–665 (2010)
38. Subcommittee on Oversight and Investigations, Power Struggle: Examining the 2021 Texas Grid Failure, Virtual Hearing, Committee on Energy and Commerce, U.S. House of Representatives, One Hundred and Seventeenth Congress, Washington, DC, 24 March 2021. www.congress.gov/117/chrg/CHRG-117hhrg46582/CHRG-117hhrg46582.pdf
39. Syamkumar, M., Mani, S., Durairajan, R., Barford, P., Sommers, J.: Wrinkles in time: detecting Internet-wide events via NTP. In: Proceedings of the IFIP Networking Conference and Workshops, pp. 91–99 (2018)
40. Taylor, D., Diaz, J.: Hundreds of thousands of people are without power, New York Times, 30 August 2021
41. U.S. Court of Appeals for the Ninth Circuit, LinkedIn Corporation, Petitioner v. hiQ Labs Inc, Case no. 17–16783, San Francisco, California (2019). www.supremecourt.gov/docket/docketfiles/html/public/19-1116.html
42. U.S. Energy Information Administration: Electric Power Monthly Washington, DC (2023). www.eia.gov/electricity/monthly
43. U.S. Energy Information Administration: Hourly Electric Grid Monitor, Washington, DC (2023). www.eia.gov/electricity/gridmonitor/dashboard/electric_overview/US48/US48
44. U.S. Supreme Court: Van Buren v. United States, Certiorari to the U.S. Circuit of Appeals of the Eleventh Circuit, Washington, DC, 3 June 2021. www.supremecourt.gov/opinions/20pdf/19-783_k53l.pdf
45. Vigdor, N.: More than 3 million homes and businesses have lost power, New York Times, 15 February 2021
46. Wan, G., et al.: On the origin of scanning: the impact of location on Internet-wide scans. In: Proceedings of the ACM Internet Measurement Conference, pp. 662–679 (2020)