A Preliminary Analysis of Network Outages During Hurricane Sandy

USC/ISI Technical Report ISI-TR-685b

November 2012, updated February 2013*

John Heidemann Lin Quan Yuri Pradkin
USC/Information Sciences Institute
{johnh, linquan, yuri}@isi.edu

ABSTRACT

This document describes our analysis of Internet outages during the October 2012 Hurricane Sandy. We assess network reliability by pinging a sample of networks and observing those that respond and then stop responding. While there are always occasional network outages, we see that the outage rate in U.S. networks doubled when the hurricane made landfall, then took about four days to recover. We confirm that this increase was due to outages in New York and New Jersey.

1. INTRODUCTION

We are interested in understanding Internet outages. In our work, we have been developing the use of active probing of destination networks to detect network outages [11], and methods to visualize these outages [12]. We probe networks with ICMP echo request messages ("pings"), and interpret networks that cease responding as down. (We summarize our approach in more detail in Section 2.) We know that active probing provides an incomplete view of the network, since it can only see networks that agree to respond (those that are not firewalled), however, we have previously shown that this approach provides a reasonable picture of more than half of the active Internet [5]. We therefore believe that our approach can effectively evaluate network outages.

While other work using active probing focuses on identifying routing problems [7, 8], we are focused on understanding network problems at the edge of the network. Closest to our work is the study of weather outages triggering active probing by Schulman and Spring [13], and Renesys' evaluation of Sandy [2, 9]. Unlike Schulman and Spring, we probe preemptively, not reactively. Unlike Resesys' work, we probe edge networks to directly assess the impact on groups of end users, while they examine BGP routing reports. Other work has interpreted passive observations [3, 4]; we instead focus on active techniques, since we believe active techniques have the potential to provide precise, controlled measurements.

In this brief technical report we summarize our findings for outages related to Hurricane Sandy's effects on the United States at the end of October, 2012 [17]. There are always occasional network outages in a network the size of the Internet. However, we show that the U.S. network outage rate approximately doubled when the hurricane made landfall, and that it took about four days to recover to prior levels. We explore these findings in more detail and provide supporting evidence in Section 3.

Our results are based on small, mostly random sample of Internet prefixes, and so they provide only a partial view of the network. In Section 4 we briefly describe our efforts to further validate of our approach and extend it to cover the whole IPv4 Internet.

2. METHODOLOGY REVIEW

We determine network outages by examining active probing with ICMP echo request messages ("pings") and watching for networks that change status from responsive to mostly or completely non-responsive. We focus on blocks of 256 adjacent IPv4 addresses (/24 address blocks where all addresses are of the form 1.2.3.*). To evaluate outages, we require that, historically, at least 10% of the addresses in the block reply to pings. A full description of our methodology and discussion of these choices can be found in our technical report [11].

Analysis in this paper uses Internet Address Survey

*This material is based on research sponsored by the Department of Homeland Security (DHS) Science and Technology Directorate, Homeland Security Advanced Research Projects Agency, Cyber Security Division BAA 11-01-RIKA and Air Force Research Laboratory, Information Directorate under agreement number FA8750-12-2-0344, and under DHS contract number D08PC75599. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Department of Homeland Security or the U.S. Government.

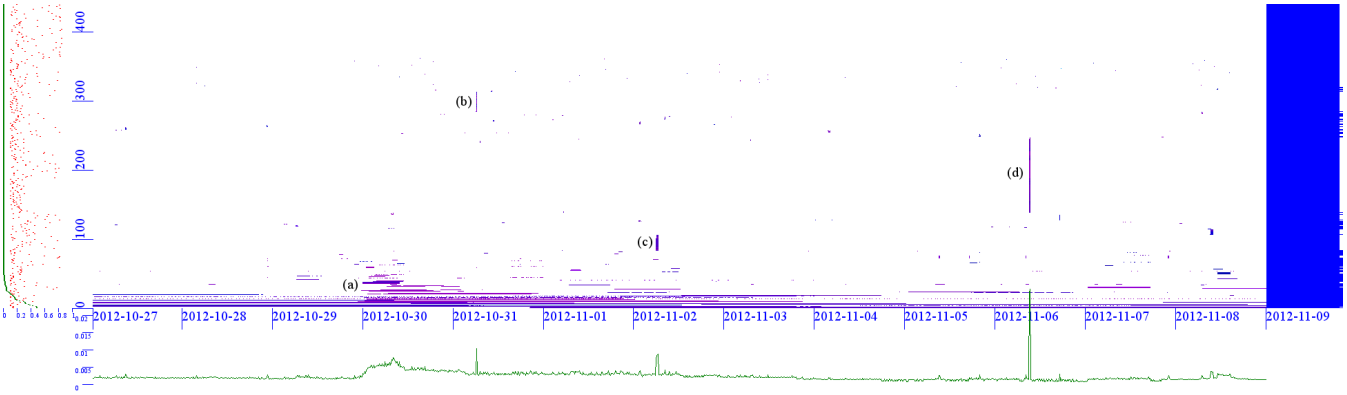

Figure 1: The 400 largest outages of /24 blocks geolocated to the United States. (Dataset: [15]).

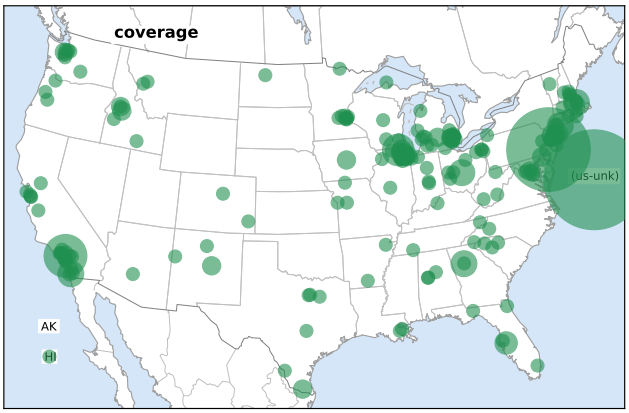

Figure 2: Locations of analyzable blocks in the U.S. (Dataset: [15]).

it50j [15], with probes of 41,582 /24 Internet *blocks* (unique prefixes) starting 2012-10-27 through 2012-11-10. The probing source was in Keio University in Fujisawa, Kanagawa, Japan. (We regular probe from three different sites, with the others in Los Angeles and Colorado, finding similar results.)

These blocks represent a sample of the IPv4 address space. They make up 0.3% of the allocated IPv4 address space, about 0.4% of the routed IPv4 address space. Although about 80% of the IPv4 space is routed, only about 4M /24s have any addresses that respond to our active probing; the reminder is likely not used or is firewalled. Thus the networks we probe are about about 1% of the responsive IPv4 address space, as determined by comparison to our Internet censuses [5]. Only ping-responsive addresses could be used for any outage study that uses active probing. Our specific approach to detecting outages means we can only evaluate a portion of the responsive address space, since we require 10% of addresses (25 addresses) per block to reply over the last three years; about 2.5M /24 blocks meet this criteria.

Most of our analysis focused on the subset of 11,900 /24 blocks that we determine are in the United States. To determine block locations, we map them through Maxmind's GeoLite City database [10]. Maxmind states that this database has 78% correct resolution to city within 40 km, with 16% incorrect solution and 6% unknown cities. As above, many of these blocks are too sparse for our analysis: we find at 4,117 /24 blocks are in the U.S. are responsive enough to analyze. Figure 2 shows the locations of these blocks on the U.S. map, with circle area proportional to the number of networks at each location. To put this number in context, Maxmind identifies about 1.5 billion addresses as located in the U.S., so our 11,900 possible blocks are only 0.2% of American addresses, and of these we can analyze about 0.07% of addresses. With such a small sample, one should be careful about interpreting our results as representative for all U.S. networks.

Visualizations in this paper use clustering algorithms we have previously described [12], where blocks are grouped based on the similarity of their outages and colored by their country code.

3. OBSERVATIONS

We observed a noticeable increase in network outages following Hurricane Sandy. The Hurricane made landfall in the U.S. at about 2012-10-30 t00:00 +0000. When we focus on known U.S. networks, we see about twice the number of network outages for the day following landfall, and above-baseline outages for the four days following landfall.

3.1 Focusing on U.S. Networks

To support these results, we focus on the 11,900 /24 blocks in the U.S., and specifically the 4,117 we can analyze as described in Section 2.

3.1.1 Outages Over Time

Figure 1 provides a visualization of the 400-U.S. blocks with the largest degree of outages. Each colored point

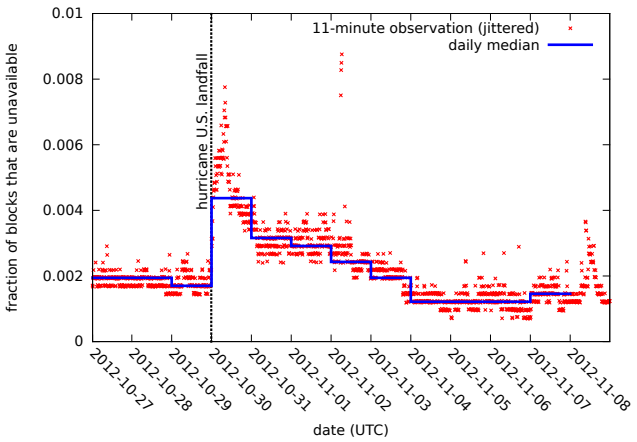

Figure 3: Median daily outages (solid line) for /24 blocks geolocated to the United States, with jittered individual readings (dots). (Dataset: [15]).

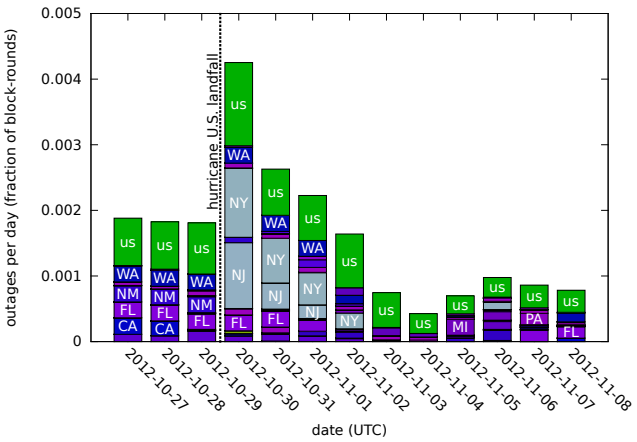

Figure 4: Amount of outages per day, broken down by state, weighted by outage size and duration. States shown where possible, or "US" for unspecified location in U.S. (Dataset: [15]).

represents a network that we classify is down for a specific time. Each row of this figure shows a single /24 network block over nearly two weeks; each column shows all outages in these blocks over an 11-minute observation round. The colors of the points are based on the estimated latitude and longitude of networks, as determined by Maxmind [10], using our world-map color scheme [6].

At label (a), this plot shows a strong cluster of outages at 2012-10-30, corresponding with hurricane landfall. Hurricane-related outages tend to be long, lasting one or more days. We believe these outages correspond to infrastructure damage.

The graph also shows several short-term outages, midday 2012-10-31, 2012-11-02, and 2012-11-06, labeled (b), (c) and (d). We have verified that (b) and (d) both correspond to routing problems. The duration of these outages is only 11–22 minutes, right at the precision of our measurement.

3.1.2 *Amount of Outages*

We know that *some* part of the Internet is always down—in prior work we estimate that about 0.3% of the Internet is down at any given moment [11]. We therefore next want to place these outages into perspective. To do so, Figure 3 plots the exact number of /24 blocks that are down in each round (this value is the marginal distribution of Figure 1). We plot each observation (every 11 minutes) as red points (these are plotted with a small amount of random jitter in the y axis so that consecutive observations are easier to distinguish), and we show 24-hour median values with the dark line.

Figure 3 shows U.S. networks had an outage rate of about 0.2% before landfall. (This rate seems somewhat better than the global average.) This rate jumps to 0.43%, more than double the prior U.S. baseline, for the 24-hours following landfall. The outage level drops slowly over the next four days, first to around 0.3% and finally returning to the baseline on 2012-11-03.

3.1.3 *Outage Locations*

To confirm the correlation between the hurricane and these outages, we look at the weighted blocks by state. Figure 4 shows these outages by state.[1] The top "US" portion represents outages that are geolocated in the U.S., but not to a specific state.

This Figure shows that there are *large increases in the amount of outages in New York and New Jersey* (the lighter colored bars in the middle of the graph) after hurricane landfall on 2012-10-30. These problems are generally resolved over the following four days.

While the relative number of outages increases significantly, the absolute amount of outage is still fairly small—at any instant on the day after the hurricane, only about 0.4% of U.S. networks were down (Figure 3). Thus, the overall U.S. Internet is quite reliable, and the hurricane had only a regional effect on the country-wide network. However, the hurricane caused service interruptions for a few specific networks (and therefore specific people), sometimes for extend periods (Figure 1).

We observe that most outages are relatively brief, but the outages from the hurricane (Figure 4 measures outages in *block-times*, the "area" of the outages in Figure 1. The outages in Figure 4 follow those in Figure 3.

Finally, Figure 5 shows a geographic view the U.S. out-

[1] On 2013-02-01 we found an error in our code aggregating outages by location: Figure 4 was using a sample of outages, not all. Before this error was corrected, Figure 4 incorrectly showed lower outage rates than Figure 3. Technical report ISI-TR-685b, issued 2013-02-04, corrects this error.

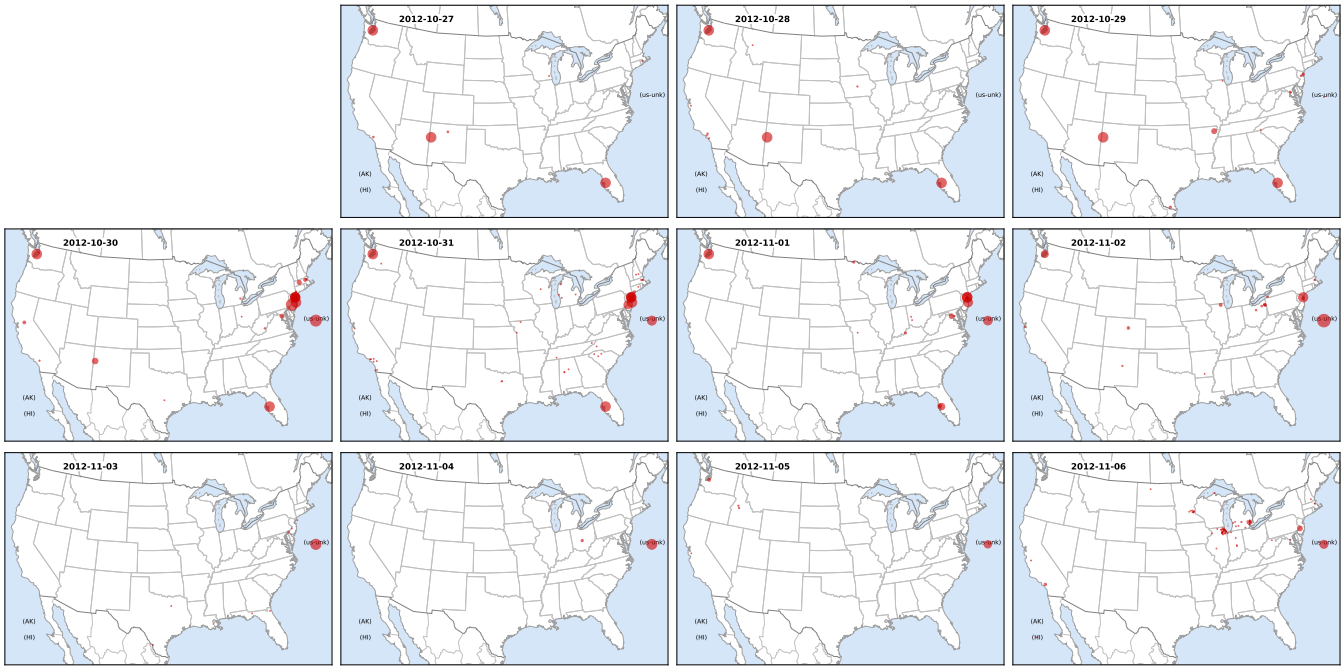

Figure 5: Geographic location of observed U.S. outages, by day. Top row: 3 days before landfall, second: 4 days after landfall, bottom: subsequent days. Circle area represents the block-rounds of outage at each location. (The point in the mid-Atlantic represents U.S. networks with unknown cities.; dataset: [15]).

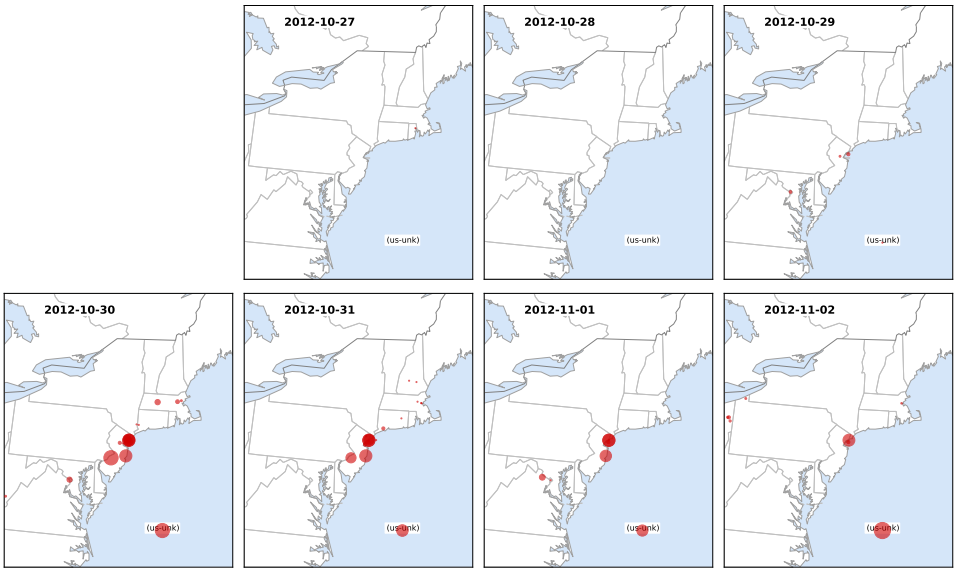

Figure 6: Geographic location of observed outages in the Northeastern U.S., by day. Top row: 3 days before landfall, second: 4 days after landfall. Circle area represents the block-rounds of outage at each location. (The point in the mid-Atlantic represents U.S. networks with unknown cities.; dataset: [15]).

		our analysis		
state	**Renesys**	ever	area	blocks
NY	12%	46% (6)	24% (3.2)	13
NJ	11%	26% (13)	5% (2.7)	50
CT	8%	16% (1)	$\ll$1% (0.008)	6
us-unk	—	7% (4)	2% (1.4)	60

Table 1: Comparing outages with our method (right) with Renesys' report [2].

ages we observed. Circles in this figure are weighted by outage size and duration (as in Figure 4). While these figures show some other outages across the use (particularly on 2012-11-06), they emphasize the localized and long-lasting outages in the New York/New Jersey area at and immediately following landfall. Figure 6 focuses on just the northeastern seaboard.

3.1.4 Comparison To Other Data Sources

To our knowledge, the only other public, quantified measurements of network outages from Hurricane Sandy are from Renesys [2]. Their methodology is based on analysis of BGP data, while ours is based on probes of a sample of edge networks, thus comparison of these two different approaches can help provide confidence in each. The Renesys data has much better coverage than we do (because of our sampling), and provides more precise timing. However, we expect that probing end-systems will reveal outages that that are missed by BGP analysis, since Bush et al. [1] documented that two-thirds of ASes employ default routing. In prior validation of our approach, where we find BGP misses about 60% of outages [11].

Renesys reports extensive outages in the New York area on October 30 (UTC), and gives statistics for outages in the New York area. We compare their reported values to ours in Table 1.

We compare their routing impacts to our two different measures. Our "ever" column counts the number of /24 blocks that were ever down on Oct. 30. This value approximates the maximum individual observations in Figure 3. Ever-outages answers the question "how many networks saw any interruptions on this day?". Our "area" column reports the amount of outage weighted by time and space, thus corresponding to the area of the colored regions in Figure 1 and the data in Figure 4. This corresponds to the question "how long were how many people affected by outages".

Renesys does not define exactly what as "12%" outage means, other than to say that they focus on "routing imapcts". We assume that they report how many networks were down for any signifcant time over the day, thus closest to our ever-outages.

As expected, we see more outages than Rensys. This difference reflects the different measurement methods, and the fact that it is feasible to re-route data and thus repair routing problems (as they measure) relatively quickly.

Comparing our ever-outages to area-outages, we see that New Jersey and Connecticut both have much lower area-outages than ever outages. This difference suggests they were able to repair problems fairly quickly. However, the larger area-outage for New York (and also somewhat for New Jersey) suggests outages that take longer to repair. This effect is consistent with reports of flooding and damage to equipment that requires physical replacement and repair.

Methodology differences between these approaches mean we cannot directly compare numbers, but the general magnitudes we observe suggest that both approaches show signifcant outages.

3.2 Context: Outages Across A Sample of the Global Internet

For context, Figure 7 shows the 1400 /24 blocks with the most outages for entire globe as observed in the same dataset. We see several large outages in this picture, including one at 2012-10-30 t00:00 (label: (e)); one mid-day 2012-10-31, split into three parts both labeled (f); and a large one late in 2012-11-04 (g); and a small spike mid-day 2012-11-06 (h). Of these, 2012-10-30 (f) and 2012-11-06 (h) both show up in the U.S. blocks in Figure 1 as (b) and (d). We believe these each of these large events is related to routing problems located near our monitor, and are currently investigating them.

4. FUTURE WORK

While this report summarizes our analysis of Hurricane Sandy, we are working to improve our network outage monitoring system.

We expect to move it from sampled network data to cover all responsive blocks in the IPv4 Internet in the coming months. In preliminary demonstrations we have shown that a single machine can follow the entire analyzable Internet [11]. We also in the process of reducing the amount of traffic required to track outages, while simultaneously increasing the precision of our measurements.

We also expect to explore concurrent measurements from multiple locations, to evaluate the effects of local or regional networking problems as identified in Section 3.2.

Although we are working to track outages across the entire Internet, sampling remains important. Additional work is needed to bound the error in sampled measurements such as we present here.

5. CONCLUSIONS

We believe this work shows the relevance of active

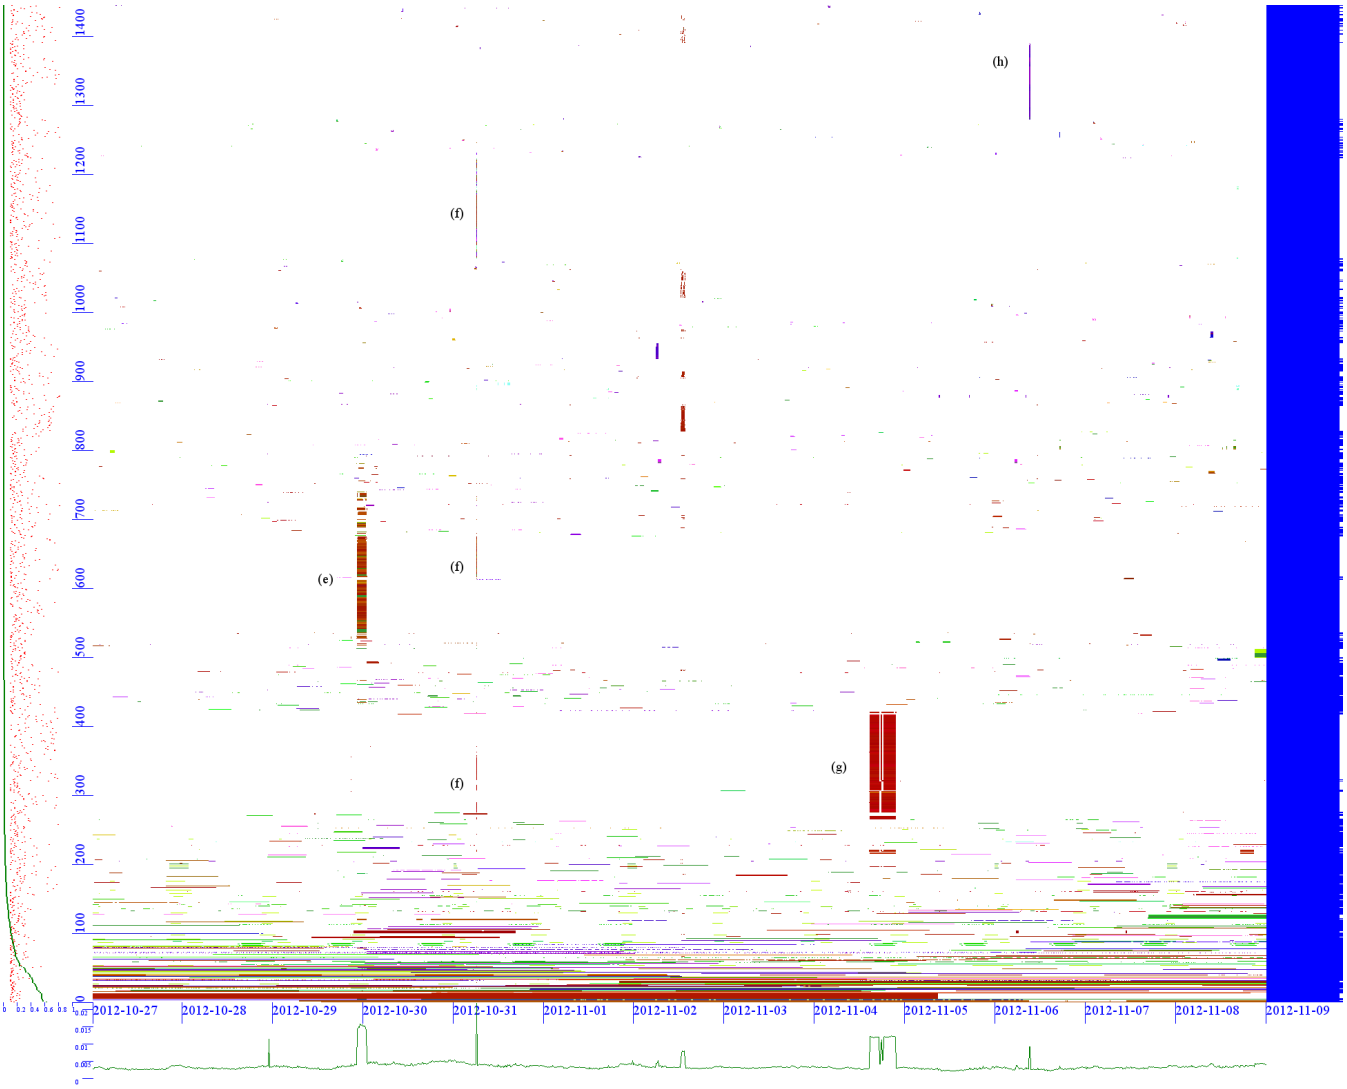

Figure 7: The 1000 largest outages of all /24 blocks. (Dataset: [15]).

probing to understanding the effects of natural disasters on the Internet. Disasters have real impact on humans and property, and while anecdotes of these impacts are easy to assemble, systematic evaluation of physical effects takes time. To the extent that network damage correlates with physical effects, analysis such as ours may provide more rapid, quantitative evaluation of the disaster effects.

Although our system continues to evolve, our analysis of a sample of networks shows the impact of Hurricane Sandy on U.S. networks and their recovery.

Data Availability

The raw [15] and analyzed [16] data from this paper are available at no cost to researchers through the U.S. DHS PREDICT program [14] and by request from the authors.

This work was reviewed by USC's IRB (IIR00000975) and identified as non-human subjects research.

Acknowledgments

We thank Zi Hu for assisting with geolocation analysis in this report. We thank Katsuhiro Horiba, Akira Kato, Midori Kato, Yohi Kuga, and Rod Van Meter, all of WIDE, for hosting our probing infrastructure and providing a BGP feed.

6. REFERENCES

[1] Randy Bush, Olaf Maennel, Matthew Roughan, and Steve Uhlig. Internet optometry: assessing the broken glasses in internet reachability. In *Proceedings of the ACM Internet Measurement Conference*, pages 242–253. ACM, November 2009.

[2] James Cowie and Doug Madory. Renesys. Presentation at 21st Euro-IX Forum, November 2012. Stockholm, Sweeden.

[3] Alberto Dainotti, Roman Ammann, Emile Aben, and Kimberly C. Claffy. Extracting benefit from harm: Using malware pollution to analyze political and geophysical events. *ACM Computer Communication Review*, 42(1):31–39, January 2012.

[4] Alberto Dainotti, Claudio Squarcella, Emile Aben, Marco Chiesa, Kimberly C. Claffy, Michele Russo, and Antonio Pescapé. Analysis of country-wide Internet outages caused by censorship. In *Proceedings of the ACM Internet Measurement Conference*, pages 1–18, Berlin, Germany, November 2011. ACM.

[5] John Heidemann, Yuri Pradkin, Ramesh Govindan, Christos Papadopoulos, Genevieve Bartlett, and Joseph Bannister. Census and survey of the visible Internet. In *Proceedings of the ACM Internet Measurement Conference*, pages 169–182, Vouliagmeni, Greece, October 2008. ACM.

[6] Zi Hu and John Heidemann. Towards geolocation of millions of IP addresses. In *Proceedings of the ACM Internet Measurement Conference*, page to appear, Boston, MA, USA, 2012. ACM.

[7] Ethan Katz-Bassett, Harsha V. Madhyastha, John P. John, Arvind Krishnamurthy, David Wetherall, and Thomas Anderson. Studying black holes in the internet with Hubble. In *Proceedings of the 5th USENIX Symposium on Network Systems Design and Implementation*, pages 247–262, San Franciso, CA, USA, April 2008. USENIX.

[8] Ethan Katz-Bassett, Colin Scott, David R. Choffnes, Ítalo Cunha, Vytautas Valancius, Nick Feamster, Harsha V. Madhyastha, Tom Anderson, and Arvind Krishnamurthy. LIFEGUARD: Practical repair of persistent route failures. In *Proceedings of the ACM SIGCOMM Conference*, pages 395–406, Helsinki, Finland, August 2012. ACM.

[9] Doug Madory. Hurricane Sandy: Initial impact. Renesys blog `http://www.renesys.com/blog/2012/10/hurricane-sandy-initial-impact.shtml`, October 2012.

[10] Maxmind. Geolite city. Web page `http://dev.maxmind.com/geoip/geolite`, 2012.

[11] Lin Quan, John Heidemann, and Yuri Pradkin. Detecting internet outages with precise active probing (extended). Technical Report ISI-TR-2012-678b, USC/Information Sciences Institute, February 2012. (updated May 2012; this TR superceeds ISI-TR-2011-672.).

[12] Lin Quan, John Heidemann, and Yuri Pradkin. Visualizing sparse internet events: Network outages and route changes. In *Proceedings of the First ACM Workshop on Internet Visualization*, page to appear, Boston, Mass., USA, November 2012. Springer.

[13] Aaron Schulman and Neil Spring. Pingin' in the rain. In *Proceedings of the ACM Internet Measurement Conference*, pages 19–25, Berlin, Germany, November 2011. ACM.

[14] The PREDICT Program. PREDICT: protected repository for the defense of infrastructure against cyber-threats. `http://www.predict.org`, January 2005.

[15] USC/LANDER project. Internet address survey dataset, predict id `usc-lander/internet_address_survey_reprobing_it50j`. web page `http://www.isi.edu/ant/lander`, November 2012.

[16] USC/LANDER project. Internet address survey

dataset, predict id usc-lander/internet_outage_survey_it50j. web page http://www.isi.edu/ant/lander, November 2012.

[17] Wikipedia. Hurricane sandy. http://en.wikipedia.org/wiki/Hurricane_sandy, 2012. Retrieved 2012-11-24.