Poster Abstract: Towards Active Measurements of Edge Network Outages[⋆]

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

1 Introduction

End-to-end reachability is a fundamental service of the Internet. We study network *outages* caused by natural disasters [2, 5], and political upheavals [8].

We propose a new approach to outage detection using active probing. Like prior outage detection methods [3, 4], our method uses ICMP echo requests ("pings") to detect outages, but we probe with greater density and finer granularity, showing pings can detect outages without supplemental probing.

The main contribution of our work is to define how to *interpret pings as outages* (§2): defining an outage as a sharp change in block responsiveness relative to recent behavior. We also provide preliminary analysis of outage rate in the Internet edge. Space constrains this poster abstract to only sketches of our approach; details and validation are in our technical report [6]. Our data is available at no charge, see `http://www.isi.edu/ant/traces/internet_outages/`.

2 Methodology

Our method for outage detection begins with active probing, followed by outage identification in individual blocks, and correlation into events.

For this paper, we define a network outage as problems in the network core or near the target that prevent reachability from our vantage point. We watch for and manually remove outages local to the monitors. We know that problems often affect only part of the Internet; evaluation of outages from multiple vantage points to distinguish partial and Internet-wide outages is future work.

2.1 Active Probing of Address Blocks

We collect data with active probing, extending our high-performance probing software used to study the Internet address space [1].

Reviewing Address Probing: We begin with active probing of addresses in some or all analyzable /24 address blocks in the IPv4 address space. Probes are ICMP echo requests (pings) at 11 minute intervals for one to 14 days. Probes are spread over 11 minutes to minimize impact on the target and effects of burst losses. We classify responses as: non-responses, network or host-specific negative replies, other errors, and positive (*echo reply*). We interpret the first two as an inaccessible network, and the later as a reachable network. We survey all addresses in a random sample of 22k or 41k responding /24 blocks.

[⋆] This work is based on research sponsored by the U.S. Dept. of Homeland Security, S&T HSARPA, BAA 11-01-RIKA and Air Force Research Laboratory, Info. Dir., agreements FA8750-12-2-0344, and D08PC75599. The U.S. Gov't is authorized to reproduce and distribute reprints notwithstanding any copyright notation thereon. The views herein are those of the authors and do not necessarily represent the official policies or endorsements of the DHS or U.S. Government.

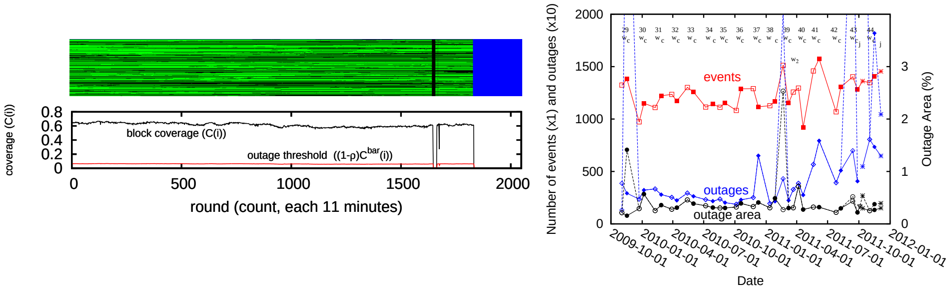

Fig. 1: *Left, top*: responses for one /24 block (green: positive, black: none, blue: not probed). *Left, bottom*: coverage and outage thresholds per round. *Right*: outage events and outage percentage, over 35 2-week surveys.

Outage-specific Steps: For outage analysis, we map probe records into *rounds* with index i. Each round is 11 minutes long, with N_r rounds in a dataset; we account for clock drift and duplicate replies. Our whole-Internet outage system probes 20 addresses in all 2.5M measurable /24 blocks for IPv4 [6].

2.2 Probes to Outages

We identify outages by a sharp drop in overall responsiveness of the block, and recovery by an increase. Let $r_j(i)$ represent the state of each address j in a given block at round i, taking 1 for a reply and 0 if down. Fig. 1 (left) shows a graphical representation of $r_j(i)$: each green dot indicates a positive response, while black dots are non-responsive (the blue area on the right is after the survey ends). In this block many addresses are responsive or non-responsive for long periods, as shown by long, horizontal green or black lines.

The *coverage* of a block at round i is defined as: $C(i) = {N_s}^{-1} \sum_{j=1}^{N_s} r_j(i)$ (where N_s is number targets probed in a block; 256 for experiments and 20 for operation). $C(i)$ is a timeseries of block responsiveness over the observation period. An outage starts when there is a severe drop (90% change or more) of $C(i)$, compared to a running average $\bar{C}$ over the last two rounds. (Exact choice of the threshold is not critical provided it is relatively large [6].) We graph $C(i)$ in Fig. 1 (bottom left), observing that it drops to zero for rounds 1640 to 1654, an outage that shows as a black, vertical band in the top panel. Because we must observe several targets, we exclude blocks that are too sparse. We consider blocks where fewer than 10% of addresses historically respond to be too sparse.

The result of this algorithm is a list of outages, represented as binary-valued timeseries $\Omega(i)$, indicating when the block is down ($\Omega(i) = 1$) or up (0). Outages incorporate data measured over the course of a round. Through controlled experiments we verify that we detect all controlled outages that last 1.9 rounds (about 20 minutes), and typically underestimate duration by about 0.5 rounds.

3 Preliminary Analysis

As an example of our outage detection method, Fig. 2 visualizes outages during the Jan. 2011 Egyptian revolution (Survey S_{38c}). This visualization clusters blocks by similarity (as previously described [7]); here we present this data to

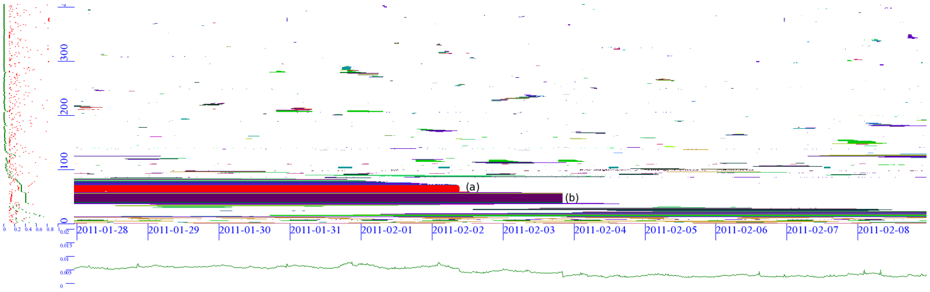

Fig. 2: The 400 largest outages of S_{38c} (see `http://www.isi.edu/ant/outage/38c`).

illustrate outage detection. Fig. 2 shows the 400 blocks with the most outages, with time on the x-axis and each row giving the Ω_j downtime for some /24 block, and colors keyed to country. There are two clusters of blocks that have near-identical outage end times. Cluster (a) covers 19 /24s, corresponding to the Feb. 2011 Egyptian Internet shutdown. Cluster (b) covers 21 /24 blocks for a slightly longer duration, related to flooding in eastern coast of Australia. Our technical report validates these events with external data [6].

This event is one example of the kind of outages we observe. We have been observing from three locations (southern California, Colorado, and Japan) for over two years. Fig. 1 (right) shows data for three years, with different shapes (open, closed, and asterisk) showing different locations. This figure suggests that our results are similar regardless of probing site and date, after we remove outages local to the prober (the dotted lines). Numerically, variation is low: mean outage "area" is 0.33%, standard deviation is only 0.1%. Overall, our data shows the Internet is about 99.7% up, or about 2.5 "nines" of availability.

References

1. J. Heidemann, Y. Pradkin, R. Govindan, C. Papadopoulos, G. Bartlett, and J. Bannister. Census and Survey of the Visible Internet. In *Proc. of ACM IMC*, Oct. 2008.
2. International Business Times. Optus, Telstra see service outages after Cyclone Yasi, 2011. `http://hken.ibtimes.com/articles/108249/20110203/optus-telstra-see-service-outages-after-cyclone-yasi.htm`.
3. E. Katz-Bassett, H. V. Madhyastha, J. P. John, A. Krishnamurthy, D. Wetherall, and T. Anderson. Studying black holes in the internet with Hubble. In *NSDI*, 2008.
4. H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: an information plane for distributed services. In *OSDI*, 2006.
5. O. Malik. In Japan, many undersea cables are damaged. GigaOM blog, `http://gigaom.com/broadband/in-japan-many-under-sea-cables-are-damaged/`, Mar. 14 2011.
6. L. Quan, J. Heidemann, and Y. Pradkin. Detecting internet outages with precise active probing (extended). Technical Report ISI-TR-2012-678, USC/ISI, Feb. 2012.
7. L. Quan, J. Heidemann, and Y. Pradkin. Visualizing sparse internet events: Network outages and route changes. In *ACM Workshop on Internet Visualization*, Nov. 2012.
8. N. Y. Times. Egypt cuts off most internet and cell service. `http://www.nytimes.com/2011/01/29/technology/internet/29cutoff.html`.