Detecting Internet Outages with Active Probing

USC/ISI Technical Report ISI-TR-672, May 2011 *

Lin Quan
USC/Information Sciences Institute
4676 Admiralty Way, Suite 1001
Marina del Rey, CA 90292
linquan@isi.edu

John Heidemann
USC/Information Sciences Institute
4676 Admiralty Way, Suite 1001
Marina del Rey, CA 90292
johnh@isi.edu

ABSTRACT

With businesses, governments, and individuals increasingly dependent on the Internet, understanding its reliability is more important than ever. Network outages vary in scope and cause, from the intentional shutdown of the Egyptian Internet in February 2011, to outages caused by the effects of March 2011 earthquakes on undersea cables entering Japan, to the thousands of small, daily outages caused by localized accidents or human error. In this paper we present a new method to detect network outages by probing entire blocks. Using 24 datasets, each a 2-week study of 22,000 /24 address blocks randomly sampled from the Internet, we develop new algorithms to identify and visualize outages and to cluster those outages into network-level events. We validate our approach by comparing our data-plane results against control-plane observations from BGP routing and news reports, examining both major and randomly selected events. We confirm our results are stable from two different locations and over more than one and half years of observations. We show that our approach of probing all addresses in a /24 block is significantly more accurate than prior approaches that use a single representative for all routed blocks, cutting the number of mistake outage observations from 44% to under 1%. We use our approach to study several large outages such as those mentioned above. We also develop a general estimate for how much of the Internet is regularly down, finding about 0.3% of the Internet is likely to be unreachable at any time. By providing a baseline estimate of Internet outages, our work lays the groundwork to evaluate ISP reliability.

1. INTRODUCTION

End-to-end reachability is a fundamental service of the Internet. Network outages—lack of data reachability—break protocols based on point-to-point communication and often harm the user's experience of Internet applications. Replication and content delivery networks strive to cover up outages, but routing changes still can cause user-visible problems as the network reconfigures [18, 19, 28].

In spite of decades of research on network reliability, in routing and other layers, Internet outages are still pervasive, ranging from minutes to hours and days. Outages have various causes, including system, link and router outages [22, 30], stemming from natural disasters [22, 27], human error [21], or political causes [5–7, 29]. The goal of our research is to systematically find *outages in blocks* of adjacent network addresses, to identify *correlated outage events in the network*, and to understand their statistical characteristics and, where possible root causes.

Our approach uses data from Internet surveys, where ICMP echo response requests are sent to each address in about 22,000 address blocks, every 11 minutes for two weeks [12]. These blocks comprise about 2% of responsive blocks of the Internet. This survey methodology and its limitations have been discussed previously [12] (as we review in Section 3.1); but its *analysis to study outages is new.*

Our analysis begins by distilling responses from addresses into block-level outage reports (Section 3.2). An outage is defined as a sharp drop followed by a sharp increase of block responsiveness, compared to typical responsiveness of the that block. Unlike control-plane studies [18, 23], we detect outages that are not seen in the routing system (Section 4.3), expanding the result observed by Bush et al. [2]. Unlike previous data-plane studies using active probing [13,16,17,20,25], our block-level measurements are more data intensive, but considerably more accurate. In Section 4.8 we show that block-wide probing reduces the number of false conclusions (declaring an outage when some of the block still responded) by 44%; Section 2 covers related work more generally.

The second aspect of our new analysis is to discover network-wide events by correlating block-level outages. We use a simple clustering algorithm to visualize outages in two dimensions, time and space (Section 3.3).

*Lin Quan and John Heidemann are partially supported by the US DHS, contract number NBCHC080035, and Heidemann by the NSF, grant number CNS-0626696. The conclusions of this work are those of the authors and do not necessarily reflect the views of DHS or NSF.

This visualization is useful to get a general understanding of network behavior and correlate outages to countries. We then develop a more general clustering algorithm to find network-wide events from the start- and end-times of block-level outages (Section 3.4). Prior studies of routing employ similar clustering [4, 10] (details in Section 2.4); we further use such clustering with data-plane probing.

We validate our methodology by comparing both selected and randomly chosen events to public BGP archives and news sources (Section 4). We find full or partial control-plane evidence for 38% of our events, but many smaller events seem visible only in the data plane, as suggested by Bush et al. [2]. We report the network impact of major event such as the Jan. 2011 Egyptian outage, and the effects of the March 2011 Japanese earthquake, as well as equally large but less newsworthy events. We also confirm the stability of our results over time and with multiple probe locations, and compare our use of full block probing to use of single-address probing.

Finally, in Section 5, we examine the data we have collected to begin to characterize Internet stability as a whole, and outage duration and size.

This paper makes several new contributions. First, we develop new clustering algorithms to visualize outages and find correlated network events from individual outage observations. Second, we validate our ability to find network network events, and discuss the relative sizes of outages corresponding to several newsworthy events. Finally, we define measures of Internet stability, providing data about stability of typical address blocks, and about the size and duration of outages.

2. RELATED WORK

Previous works have studied network stability with control- and data-plane observations, and with studies of user data and event originators. We review each group next.

2.1 Control-plane Studies

Several prior efforts use control-plane data to study Internet outages, mining data such as routing update messages and syslogs to locate failures.

Markopoulou *et al.* use IS-IS update messages to classify failures in Sprint's IP backbone. They classify outages into maintenance-, router-, optical layer-related problems, and report the percentages of each category [23]. Like them, we use control-plane data (BGP archives, news reports), but we use it only to verify outage events, and use data-plane probes to discover outages.

Teixeira *et al.* show the inherent limitations of BGP data and proposed the addition of an omni server for each Autonomous System (AS), which maintains an AS-level forwarding table [26]. They also analyze how to correlate omni servers to diagnose routing changes. Omni server has its own costs and non-trivial change to current Internet architecture, so it's not practically used. Our measurements employ well studied and supported ping probe techniques, and draw conclusions based on a series of large datasets.

Labovitz *et al.* use injected artificial routing failures to understand the impact of failures on end-to-end routing performance [18]. We choose to collect and analyze everyday outages and measure their statistical characteristics.

Control-plane studies of reachability only predict or reflect the real Internet, thus they have inherent limitations as discussed by Bush *et al.* [2]. These results prompt our use of data-plane measurements to identify outages, and use control-plane data only for the purpose of validation.

2.2 Data-plane Studies

By directly measuring reachability, data-plane studies can be more accurate than control-plane measurements.

An MIT study measures Internet path failures by probing between 31 locations, and correlate these failures with BGP messages [9]. They find that most failures are short (less than 15 minutes) and discuss the relationship between path failures and BGP messages. Our work extends theirs in several ways: we use a much larger target population (about 2% of the responsive Internet), and probe much more intensively (each address in each block every 11 minutes, rather than a few addresses for each block), but from only two sites. As with their work, we validate our work using control plane data.

The PlanetSeer system uses active probing between all the 350 PlanetLab nodes, to discover potential network path anomalies, which are originally reported by passive monitors at 120 nodes [33]. Their work can only find outages between the PlanetLab nodes. In the contrary, while we probe only from two U.S. cities, our targets are a random 2% of the responsive Internet address blocks, so we find outages anywhere in the world.

Kompella *et al.* develop spatial correlation algorithms to localize faults within a tier-1 ISP [16] and use active probing to detect black holes or failures [17]. They detect failures by $O(n^2)$ probes from every node to every other node. We probe from 2 sites to thousands of blocks in the Internet address space, so we are not limited to a single ISP.

Very close to our work, the Hubble system uses continuous probes to individual addresses to identify Internet outages [13]. We replace their method using probes to single targets in each address block, with probes to all addresses in each block. While our approach is much more network-intensive, we show in Section 4.8 that we

greatly reduce the number of false conclusions about network outages. We also describe new algorithms for clustering outages for visualization and into network-wide events.

Researchers from IIJ have studied the reachability of Internet address spaces through traceroutes to and from test and anchor prefixes (in- and out-probes), to find bogus bogon filters. They also further analyze common biases and limitations of reachability experiments [1, 2]. Although providing useful insights, their focus is not primarily on outages. We focus on the outages, or instability in reachability, of end-to-end paths to the Internet *edge*, and what we can learn from correlated reachability issues.

2.3 Log-based Analysis

Above the network layer, other systems have looked at system- and user-level logs to determine outages. For example, UCSD researchers have done careful studies of "low-quality" data sources (including router configurations, e-mail and syslog messages), to discover characteristics and reasons of failures in the CENIC network [30]. Such log analysis requires collaboration with the monitored networks, and so they study only their regional network; we instead use active probing that can be done independent of the target and study a random sample of Internet /24 blocks.

BGP misconfiguration can also be a source of outages. Mahajan *et al.* study routing table messages and email network operators for evidences of BGP misconfiguration. They also use active probing to determine the impact of misconfiguration on connectivity [21]. They report that 0.2% to 1% of prefixes suffer from misconfiguration each day. We confirm their results on the overall Internet reachability, finding about 0.3% of the Internet blocks are expected to suffer from outages on a daily basis. Our approach with active probing allows detection of all types of outages (not just BGP-triggered ones), and finds outages not visible to the control plane (as suggested by Bush [2]).

2.4 Origins of Routing Instability

Routing information distributed as part of BGP is an attractive source of data for outage estimation since BGP naturally centralizes otherwise distributed information. Since outages can occur anywhere on the AS path provided by BGP, reaching an accurate determination of the *originator* of an outage has been the subject of several previous studies.

Chang *et al.* cluster BGP path changes into events, both temporally and topologically [4]. They also provide useful insights on how to infer where network events happen. We observe that many large Internet outages happen across different edge ASes. We develop conceptually similar clustering methods, but based on data-plane observations rather than BGP control-plane information.

Feldmann *et al.* present a method to identify ASes responsible for Internet routing instabilities, using time, views and prefixes [10]. They report that most routing instabilities are caused by a single AS or a session between two ASes. (Chang *et al.* make similar conclusions [4]). They also propose useful insights on cautions of identifying instability originators. Our work uses similar ideas to validate outages with BGP routing data.

3. METHODOLOGY

This section describes our process for outage detection with active probing: raw data collection by Internet surveys, outage identification at individual blocks, visualizing outages, and correlation of outages across different blocks into routing events.

3.1 Input: Active Probing of Address Blocks

Our work begins with Internet surveys that actively probe the Internet address space [12]. We briefly review this existing methodology, then discuss how we regularize it for analysis here.

Reviewing Address-Space Surveys: Existing *Internet surveys* use ICMP pings to probe about 2% of the allocated and responsive Internet address space (about 22,000 /24 blocks), at 11 minute intervals for about two weeks [12]. Responses are classified into three broad categories: positive (*echo reply*), negative (for example, *destination unreachable*), and non-response; we discuss how these are treated in the next section.

The choice of an 11-minute probing interval limits the precision of our estimates of outage times. Our choice of this probing interval is primarily to reduce the burden on the target networks to one probe every 2.5 s, In addition, prior studies of dynamic addresses showed typical use durations are 75 or 81 minutes [3, 12, 15], so 11 minutes can capture the median user of a dynamically-assigned address. Finally, 11 minute is relatively prime to most periodic human events.

Three-quarters of survey blocks are chosen randomly from responsive blocks, while one quarter are selected based on block-level statistics as described previously [12]. In this paper we consider all blocks in each survey. Since selection is not strictly random, it may introduce some bias to our results, but prior work has not shown significant skew, and we are the process of evaluating this question for our new results. In addition, our methodology applies only to blocks where 10% of addresses respond (Section 3.2). Typically 45–50% of probed blocks pass this test (Table 2). Our results therefore do not consider sparely populated blocks, but do reflect the a diverse set of Internet users of public addresses where firewalls admit ICMP, including home

users, server farms, universities, and some businesses.

Normalizing survey data: Since probes are spread out in time and responses return with varying delays, in this paper we simplify survey input by mapping probes into *rounds*, where each round is 11 minutes long. We impose rounds on the raw probe results by taking time since each survey's beginning divided into 11-minute bins. We identify rounds with an index i, and there are N_r total rounds in a survey (thus $(i \in [1\,..\,N_r])$.

We correct for two kinds of errors in mapping observations to rounds: sometimes a round is missing an observation, and sometimes we see duplicate responses in that round. Our collection software is not perfectly synchronized to 11 minute rounds, but takes on average 11 minutes and 3 seconds. (We intentionally chose to correct for minor drift rather than guarantee perfect synchronization over days of continuous operation.) Because this interval is not exactly 11 minutes, about one round in 220 has no observation. We detect such holes and fill them with interpolation from the previous observation.

In addition, we sometimes get multiple observations per round for a single target IP address. About 3% of our observations have duplicate responses, usually a timeout (non-response) followed by a negative response (an error code). These duplicates are rare, and somewhat non-uniformly distributed (for example, about 6% of blocks have over 100 addresses with at least one duplicate response during the whole survey time). When we get duplicate responses, we keep the most recent observation, thus the negative response usually overrides the timeout.

Finally, we observe that the process of associating the IP address of an ICMP reply with its request is not perfect. Multi-homed machines sometimes reply with an address of an interface other than the one which was targeted, the phenomena of IP address aliasing in topology discovery (as described in early work [11] and recent surveys [14]). Since we know all the addresses of each probe we send, we filter replies and discard uninterpretable responses.

3.2 Probes to Outages

Given two weeks of responses from all IP addresses in a block now organized into rounds, we can identify potential outages when we see sharp drops and increases of the overall responsiveness of the block.

Our probing of entire blocks distinguishes our methodology from prior work which typically uses only a single representative per block. Although it sends 256 times more traffic, we show in Section 4.8 that use of a single representative per block results in significant numbers of observation errors. Probing can suffer errors due to lost probes or replies, or a transient failure of computer behind address (perhaps due to a machine reboot), or small power or hardware failures. Single probing can be vulnerable to some of these errors and interpret a small error as an outage for the whole block. Our approach of considering the entire block, with 256 addresses over 11 minutes, is naturally robust to individual or very brief errors and potential false conclusions.

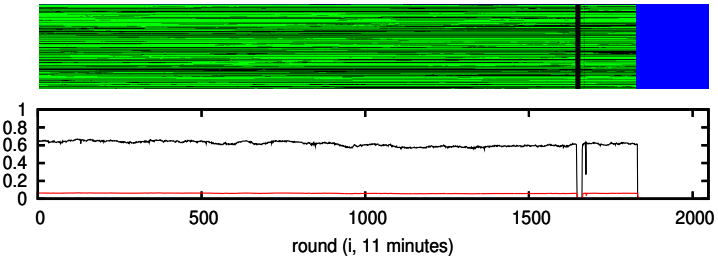

Figure 1: Probe responses (top) and outage evaluation (bottom) for one /24 block with an outage at round 1640. For probes, green shows positive response; black, no response; blue , not probed (rounds beyond 1825). The evaluation (bottom) shows coverage $C(i)$ as the top black line, and the outage threshold $(1-\rho)\bar{C}(i)$ on the bottom. Dataset: Survey S_{30w}.

Before defining coverage, we first define what we observe about individual addresses. Let $r_j(i)$ be 1 if there is a reply for the address j (the last octet of the address) in the block at time i, and 0 if there is no reply, or the error replies network or host unreachable.

$$r_j(i) = \begin{cases} 1, & responsive \\ 0, & otherwise \end{cases}$$

Figure 1 shows a graphical representation of $r_j(i)$, where each green dot indicates a positive response, while black dots are non-responsive (the blue area on the right is after the survey ends). For this block, many addresses are responsive or non-responsive for long periods, as shown by long, horizontal green or black lines. However, there is a certain amount of churn as machines come and go.

The *coverage* of this block, at round i, is defined as:

$$C(i) = \frac{1}{N_s} \sum_{j=1}^{N_s} r_j(i)$$

(Where $N_s = 256$ is the number of IP addresses in the block.) $C(i)$ is a timeseries $(i \in [1\,..\,N_r])$, which represents the overall responsiveness of a block across the whole survey period.

A severe drop and later increase in $C(i)$ indicates an outage for the block. As an example, the black band in Figure 1 shows an outage from round 1640 to 1654. We plot $C(i)$ of this block in Figure 1, when $C(i)$ (black line) drops severely we know an outage starts.

Algorithm 1 formalizes our definition of "a severe

Algorithm 1 Outage detection for a block

Input: $C(i)$: timeseries of coverage, N_r: number of rounds
Output: L: list of outage (start, end) time tuples
$\Omega(i)$: binary timeseries of block down/up information.
Parameters: w: number of rounds to look back, ρ: drop/increase percent to decide outage start/end

```
L = φ
Ω(i) = 0, i ∈ [1..N_r]
for all i ∈ [w + 1..N_r] do
  Ĉ = (1/w) Σ_{j=i−w}^{i−1} C(j) // running average
  if C(i) < (1 − ρ)Ĉ then
    // severe drop ⇒ outage start
    last_outage_start ← i
  else if Ĉ < (1 − ρ)C(i) then
    // severe increase ⇒ outage end
    L = L ∪ {(last_outage_start, i)}
    for all j ∈ [last_outage_start..i] do
      Ω(i) = 1
    end for
  end if
end for
return L, Ω(i)
```

drop": we keep a running average of coverage over window w (by default, 2 rounds or 22 minutes) and watch for changes more than a threshold value ρ (by default, 0.9).

The result of this algorithm is a list of outages and a binary-valued timeseries $\Omega(\cdot)$. This timeseries $\Omega(i), i \in [1..N_r]$, indicats when (i) the block is down $(\Omega(i) = 1)$ or up (0):

$$\Omega(i) = \begin{cases} 1, & blockdown \\ 0, & otherwise \end{cases}$$

A typical two-week dataset is 70GB (compressed), with about ten billion records. We do most processing using Hadoop on a 120-core compute cluster, using a three-step map/reduce job. While we have not tried to optimize our code, we can turn observations into clustered events in about 80 minutes.

Because this algorithm detects changes in $C(\cdot)$, it only works for blocks where a moderate number of addresses respond. We typically require 10% of addresses in a block to respond, on average, over the entire survey $(\bar{C} = (1/N_r)\sum_i C(i) \geq 0.1)$. Our selection of 10% is somewhat arbitrary; if we define $\alpha = 0.1$ as this threshold, it is limited by the precision of only $N_s = 256$ addresses in a /24 combined with the need to detect changes in $\lfloor \alpha(1-\rho)N_s \rfloor$, requiring $\alpha > 0.05$ in practice. In Section 4.4 we review the α parameter of our approach, showing that $\alpha = 0.1$ is reasonable choice. Table 1 shows how many blocks are analyzable for Survey S_{30w}.

category	blocks	percentage	
all IPv4 addresses	16,777,216		
non-allocated	1,709,312		
special (multicast, private, etc.)	2,293,760		
allocated, public, unicast	12,774,144	100%	
non-responsive	11,644,391	91%	
responsive	1,129,753	9%	100%
probed	*22,381*		*2%*
too sparse, $\bar{C} < \alpha$	*11,752*		*1%*
analyzable, $\bar{C} \geq \alpha$	**10,629**		**1%**

Table 1: Subsetting for blocks that are *probed* and **analyzable** ($\bar{C} \geq 0.1$), for Survey S_{30w}. Measurements are in numbers of /24 blocks.

3.3 Visualizing Outages

With the above algorithm to find block-level outages, we next develop a simple clustering algorithm to group block-level outages in two dimensions: time and space. We use this algorithm for visualization only; in the next section we show a second clustering algorithm that relaxes the two dimensional constraint.

Our 2-D clustering algorithm (Algorithm 2) orders blocks based on a simple exclusive-or-based distance metric. For blocks m and n, with binary-valued outage timeseries $\Omega_m(i)$ and $\Omega_n(i)$, we define distance:

$$d_v(m,n) = \sum_{i=1}^{N_r} \Omega_m(i) \oplus \Omega_n(i)$$

Outages for blocks m and n are identical if $d_v(m,n)$ is zero.

Figure 2 shows the result of visualization clustering for Survey S_{38c}. The x-axis is time (from 2011-01-27 to 2011-02-10), each row shows the Ω_j uptime for a different /24 block j. We do two steps of filtering (as shown in Table 1): first we exclude the 11,572 sparse blocks where $\bar{C}(i) < \alpha$; then for the reamining 10,629 blocks, we plot only the 500 blocks with most outages. Color is keyed to the country to whom each block is allocated.

We discuss the details of this survey in Section 4, but there are two clusters of blocks that have near-identical outage end times. The cluster labeled (a) covers 19 /24s that are down for the first third of the survey; it corresponds to the Feb. 2011 Egyptian Internet shutdown. The cluster labeled (b) covers 21 /24 blocks for a slightly longer duration; it is an outage in Australia concurrent with a cyclone.

3.4 Outages to Correlated Events

We next wish to associate outages for specific blocks into network events; we use these events later in Section 4 to relate the outages we see to ground truth out-

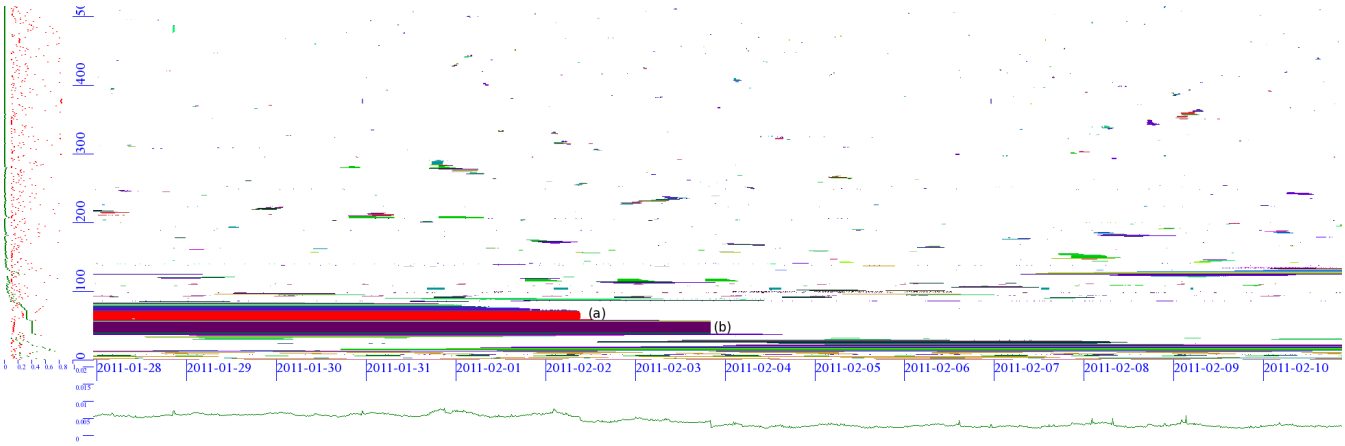

Figure 2: The 500 largest outages of S_{38c}, x axis: time, y axis: address space (blocks). Colors represent countries.

Algorithm 2 Two-dimensional clustering of blocks for visualization.

Input: A: the set of all blocks in a survey, with outage information

Output: B: list of survey blocks, ordered by distance

```
start with block m ∈ A with smallest Σ_{i=1}^{N_r} Ω_m(i) (rounds down)
A = A \ {m}
B.append(m)
while A ≠ φ do
  for all n, s.t. d_v(m, n) = 0 do
    A = A \ {n}
    B.append(n)
  end for
  // pick the next most similar block:
  find m' s.t. d_v(m, m') ≤ d_v(m, n) ∀n ∈ A
  A = A \ {m'}
  B.append(m')
  m = m'
end while
return B
```

ages based on routing and news. While visualization is helpful, the two-dimensional constraint of Algorithm 2 over-constrains clustering since each block can only be adjacent to only two others.

We therefore develop a second clustering algorithm that relaxes this constraint to group block-level outages into network-wide events. We identify events based on similar start- and end-times of outages. This approach may fail if there are two unrelated events with similar timing, but we believe that timing alone is often sufficient to correlate larger events in today's Internet.

Given two outages o and p, each having a start round $s(o)$ and end round $e(o)$ (or $s(p)$ and $s(p)$ for p), we measure their distances by the metric d_e:

Algorithm 3 Finding correlated events

Input: O: the set of all outages in a survey

Output: E: the set of network outage events, each containing one or more outages

Parameters: θ: the threshold to decide if two outages belong to same event

```
while O ≠ φ do
  find first occurring outage o ∈ O
  e = {p : ∀p ∈ O, s.t. d_e(o, p) ≤ θ }
  O = O \ e
  E = E ∪ {e}
end while
return E
```

$$d_e(o,p) = |s(o) - s(p)| + |e(o) - e(p)|$$

Outages that occur at exactly the same time have $d_e(o,p) = 0$. Since routing events often require some time to propagate [18], and outages may occur right on a round edge, we consider outages with small distance (less than θ) to be part of the same event. Currently we set $\theta = 2$ rounds (22 minutes). We have also studied much larger $\theta = 10$ (110 minutes), showing similar results, although less strict matching aggregates many more small events as shown in Section 5.1.

Given this distance measure, event clustering follows by grouping all outages that occur at similar times ($d_e(o,p) \leq \theta$) as shown in Algorithm 3.

3.5 Outages to Internet Availability

Outages and network events are what happens in the network, but they are too raw to characterize network stability as a whole. We therefore define several statistical measures of Internet availability.

As shown in Figure 3, some network events like event (a) affect many blocks for a short period (here, about

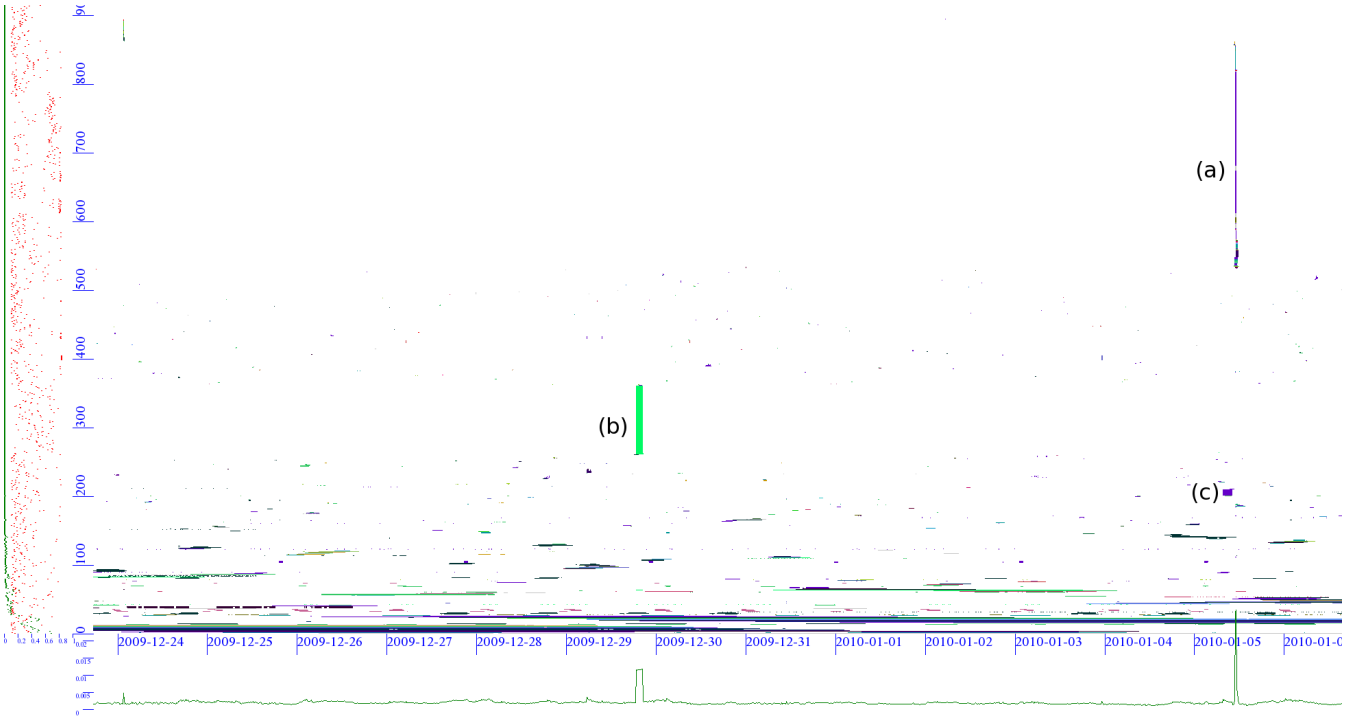

Figure 3: The 900 largest outages of S_{30w}, x axis: time, y axis: address space (blocks). Colors represent countries.

20 minutes), while others like (b) and (c) affect fewer blocks but for longer periods of time (here 2 to 3 hours). We discuss these events in detail in Section 4.2), but they suggest that *marginal distributions* of outages would be useful to characterize this variation.

Given N_b blocks and N_r rounds in a survey, we can compute the time- and space-specific sums:

$$\bar{\Omega}_I(i) = \sum_{b=1}^{N_b} \Omega_b(i), \qquad \bar{\Omega}_B(b) = \sum_{i=1}^{N_r} \Omega_b(i)$$

And we can define the overall outage "area" for a survey as the fraction of time and space that was out over all observations:

$$\bar{\Omega} = \frac{1}{N_b N_r} \sum_{i=1}^{N_r} \sum_{b=1}^{N_b} \Omega_b(i)$$

We use these metrics to evaluate network performance over time in Section 4 and 5.

4. VALIDATING OUR APPROACH

We next validate our approach. We first use several sources to confirm our observations, including BGP announcements, operator discussions, and public news sources; we evaluate both large outages (Section 4.2) and a random sample of all outages (Section 4.3) to avoid any bias due to outage size. We also consider how much of the Internet we cover (Section 4.4).

We discuss the stability of our results as a function of observation location (Section 4.5) and date (Section 4.6). And finally we compare our approach to current approaches in Section 4.8.

4.1 Validation Data Sources and Methodology

In Section 3, we described in detail how we correlate Internet outages across many different blocks into different events. However, we must understand how these observations correspond to *real-world events*. Confirming our observations are valid is not straightforward, since events are months in the past, and network problems are ephemeral and can occur anywhere in the Internet. We primarily use public archives of BGP routing information to verify our data-plane observations with control-plane data. For large events, we also use public news sources to find the root cause.

We next review the datasets we use, how we relate an event to BGP updates in time, and how we relate an event to different Autonomous Systems (ASes) in space.

Datasets: Similar to the datasets of our previous work [3, 12], we consider 24 separate 2-week surveys in this paper, taken from two locations over almost 18 months, as shown in Table 2. Of these datasets, most validation uses S_{30w}, with additional case studies drawn from S_{38w}, S_{38c}, S_{39w} and S_{39c}. We use all datasets to evaluate the stability of our approach in Sections 4.5 and 4.6.

We validate our work with BGP data from Route-

Survey	Start Date	Duration (days)	Blocks (Analyzable)
S_{29w}	2009-11-02	14	22371 (46%)
S_{29c}	2009-11-17	14	22371 (45%)
S_{30w}	2009-12-23	14	22381 (47%)
S_{30c}	2010-01-06	14	22381 (48%)
S_{31w}	2010-02-08	14	22376 (48%)
S_{31c}	2010-02-26	14	22376 (49%)
S_{32w}	2010-03-29	14	22377 (48%)
S_{32c}	2010-04-13	14	22377 (48%)
S_{33w}	2010-05-14	14	22377 (48%)
S_{33c}	2010-06-01	14	22377 (48%)
S_{34w}	2010-07-07	14	22376 (47%)
S_{34c}	2010-07-28	14	22376 (47%)
S_{35w}	2010-08-18	14	22376 (47%)
S_{35c}	2010-09-02	14	22375 (47%)
S_{36w}	2010-10-05	14	22375 (48%)
S_{36c}	2010-10-19	14	22375 (48%)
S_{37w}	2010-11-24	14	22374 (48%)
S_{37c}	2010-12-09	14	22373 (48%)
S_{38w}	2011-01-12	14	22375 (47%)
S_{38c}	2011-01-27	14	22373 (47%)
S_{39w}	2011-02-20	16	22375 (52%)
S_{39c}	2011-03-08	14	22375 (49%)
S_{39w2}	2011-03-22	14	22374 (49%)
S_{40w}	2011-04-06	14	22374 (48%)

Table 2: Internet surveys used in this paper, with dates and durations. Survey numbers are sequential with a letter indicating collection location (w: ISI-west in Marina del Rey, CA; c: Colorado State U. in Ft. Collins, CO). Blocks are analyzable if $\bar{C} \geq 0.1$.

Views [24] and our local BGP taken near our probing sites with BGPmon [32].

Relating events and routing updates in time: To find BGP routing updates relevant to a network event, we search BGP update archives near the start and end of that event. We narrow our search to destination prefixes that become unreachable, and search within 120 minutes of our identified outage time. Our window is fairly broad because our event timing is precise to only ±11 minutes, and we know that routing changes can take minutes to converge.

We then check BGP archives to see if there are relevant matching withdraw and announce messages. We expect to see relevant withdraw messages before event e as the prefix becomes unreachable, and announce messages after e. With both, we claim that e is fully validated, with just one we claim partial validation.

Relating events and routing updates in space: Although the above approach validates routing outages that happen at the destination, we find many outages occur in the middle of the Internet. Narrowing our search to just destination prefixes therefore overly constrains our search.

When our temporal search fails to identify a routing problem, we broaden our search to all ASes on the path, as done by Chang et al. [4] and Fedlmann et al. [10]. We generate an AS-path for the destination prefix by searching in RouteViews BGP snapshots. Finally, we then search for BGP withdraw and announce messages in time windows around the start and end of our network event. Often the first desintation search found an announce message; in that case we look here for withdraw messages for an intermediate AS.

Searching intermediate ASes has two disadvantages. First, the search space is much larger than just considering the destination prefixes. Second, RouteViews BGP snapshots are taken every two hours, so we must widen our window to two hours.

4.2 Network Event Case Studies

We begin by considering three cases where the root cause made global news, then outages near our collection points, and finally three smaller events. These events are medium or larger than typical events we detect. We make no claims that these events are representative of the Internet in general, only that they demonstrate how events found by our tools relate to external observations. In the next section we validate a random sample of events to complement these anecdotes.

Jan. 2011 Internet Outage: Beginning 2011-01-25 the Egyptian people began a series of protests that resulted in the resignation of the Mubarak government by 2011-02-11. In the middle of this period, the government shut down Egypt's external Internet connections.

Our S_{38c} began 2011-01-27 T23:07 +0000, just missing the beginning of the Egyptian network shutdown, and observed the restoration of network service around 2011-02-02 T09:28 +0000. Our survey covered 19 /24 blocks in the Egyptian Internet; they can be seen marked (a) in Figure 2.

We can confirm our observations with widespread news coverage in the popular press [29]. We also confirm the details that we observe with more technical discussions [5,6], and with analysis of BGP data by seeing both withdraws before event and announces after event. We observe outages in a number of Egyptian ASes, including AS8452, AS24835, and AS24863. We see that all Egyptian blocks in our survey go down, and the timing is consistent with BGP messages. We conclude that our approach correctly observed the Egyptian outage.

Feb. 2011 Libyan Outage We also examined the Libyan outages 2011-02-18 to -22 [7]. This period was covered by S_{38c}, but this survey contains only one Libyan block, and coverage for that block was too low (0.014) for us to track outages. Our requirement for blocks with moderate coverage, combined with measuring only a sample of the Internet and Libya's small Internet footprint (1168 /24 address blocks as of March 2011 [31]) shows that we can easily miss smaller outages.

Feb. 2011 Australian Outage: We also observe a significant Australian outage in S_{38c}. Marked (b)

in Figure 2, by our observations this outage involved about as many blocks as the Egyptian outage. We can partially validate our outage with BPG, but its root cause is somewhat unclear. Tropical Cyclone Yasi made landfall in the Cairns area of Queensland on 2011-02-03. There were news reports about network and power outages around this time, with service outages for Telstra and Optus, two of the largest Australian ISPs [27]. The recovery of the network seems consistent with news reports about telecommunications repairs. However, the start of the outage in our observations is before our survey begins on 2011-01-27, at least five days before the cyclone makes landfall. We observe BGP announce messages for the target blocks around 2011-01-27 T09:00 +0000, before our survey begins, but we cannot locate relevant withdraw messages.

Our observations suggest that this Australian outage was about *as large and long-lasting* as the Egyptian outage, yet the Egyptian Internet outage made global news while the Australian outage was little discussed outside Australia. The Egyptian outage was more newsworthy both because of the political significance, and because it represented nearly all Egyptian traffic. Australia, by comparison, has eight times more allocated IPv4 addresses than Egypt, so though the Australian outage may be as large as the Egyptian one, it does not have the same country-wide impact. We believe this example shows the importance of tools such as ours to *quantify* the size and duration of network outages.

March 2011 Japanese Earthquake: In survey S_{39c}, we observe a Japanese Internet outage, as shown in Figure 4 marked (a). This event is confirmed as an undersea cable outage caused by the Tōhoku Japanese earthquake 2011-03-11 [22].

Unlike most other outages we observe, both the start and the recovery from this outage vary in time. For most blocks, the outage begins at the exact time of the earthquake, but for some it occurs two hours later. Recovery for most blocks occurs within ten hours, but a few remain down for several days.

Local Outages: In addition to outages in the Internet, they also happen near our monitors. (We watch for outages in our data, and by talking with network operations.) Survey S_{39w} shows two such events. In Figure 5, event (b) was planned maintenance in our server room; the blue color indicates absence of data. Event (c) was a second planned power outage that took down a router near our survey machines although probes continued running. Both of these events span all probed blocks, although Figure 5 shows only 500 of the blocks. Finally, event (a) is due to temporary firewalling of our probes by our university due to a miscommunication.

These examples show that our methods have some ability to distinguish local from distant outages. They also revealed an interaction of our probing with Linux iptables. In outage (c), the number of active connections in iptables overflowed. Connection table overflow produces random ICMP network unreachable error replies at the probing host. We were able to confirm and filter these from our data, and have since disabled ICMP connection tracking.

valid.	with.	ann.	count		outage sizes
no	—	—	31	(62%)	1 to 57, median 4
partial	Yes	—	1	(2%)	24
partial	—	Yes	10	(20%)	1 to 27, median 15
yes	Yes	Yes	8	(16%)	1 to 697, median 21
			50	(100%)	

Table 3: Validation of algorithm with counts of missing (—) or found (Yes) withdraw and announce messages, for randomly selected events from Survey S_{40w}. Counts in events; sizes in blocks.

Three Small Events: Finally, we explore three small events in survey S_{30w} as examples of "typical" network outages. These events are shown on Figure 3.

Verizon outage 2010-01-05 T11:03 +0000: In Figure 3, event (a) is a short outage (about 22 minutes) affecting many blocks (about 331 blocks). Many of these destinations belong to AS19262, a Verizon AS. Examination of RouteViews BGP archives confirms this event. Examination of the AS-paths of affected blocks suggests that the outage occurred because of a problem at AS701, another Verizon AS, present in the path of all but 0.6% of destinations. It also confirms the duration, with the BGP withdraw-to-announce time of about 20 minutes.

AT&T/Comcast 2010-01-05 T07:34 +0000: In Figure 3, event (c) is an 165 minute outage affecting 12 blocks. Again, we confirmed this outage in RouteViews BGP archives. The affected destinations were AS7132 (AT&T) and AS7922 (Comcast). Routing archives confirm withdraws and returns of these routes, and AS-paths suggest the root cause was in AS7018 (AT&T WorldNet), likely upstream of the destinations.

Mexico outage 2010-12-29 T18:36 +0000: The event labeled (b) in Figure 3 corresponds to a large number of destinations in AS8151, a Mexican ISP (Uninet S.A. de C.V.). The event is fairly large and long: 105 blocks for 120 minutes. We were unsuccessful in identifying the root cause of this outage in RouteViews data. This survey pre-dates our local BGP feed, and all RouteViews BGP archives are several ASes from our probing site, suggesting the outage may have been visible to us but not seen at the RouteViews monitors, or that some of these blocks may be using default routing as described by Bush et al. [2].

4.3 Validation of Randomly Selected Events

Our outage case studies in the prior section were selected because of their importance and so are biased

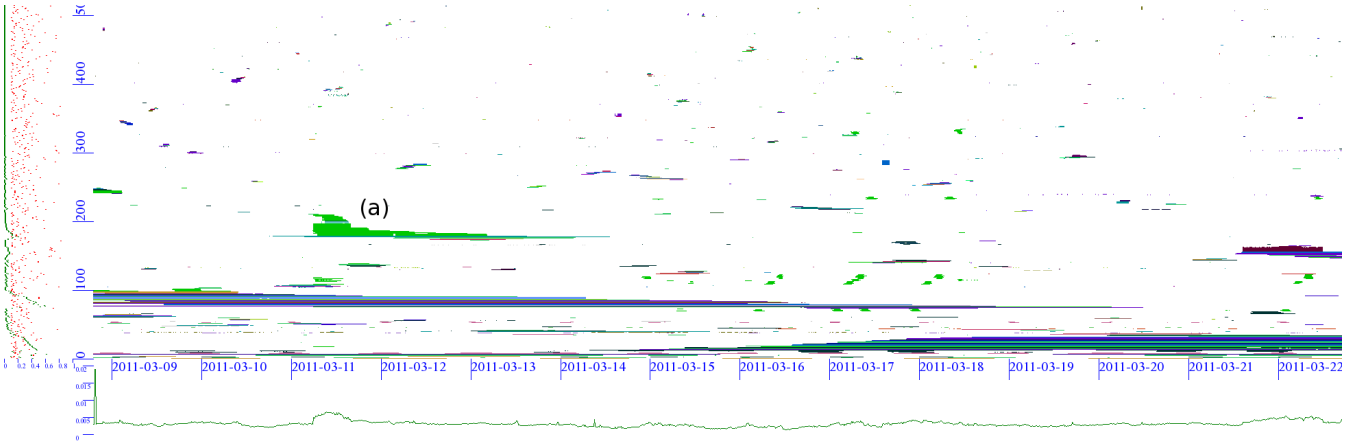

Figure 4: The 500 largest outages in S_{39c}, x axis: time, y axis: address space (blocks). Colors represent countries.

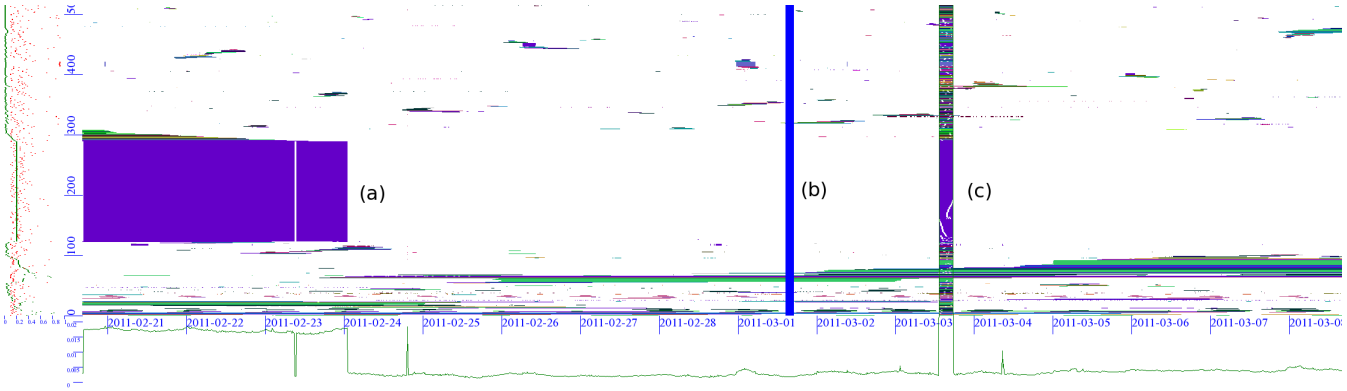

Figure 5: The 500 largest outages in S_{39w}, x axis: time, y axis: address space (blocks). Colors represent countries.

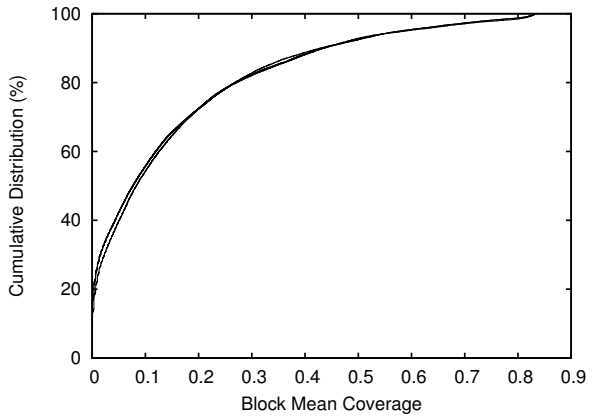

Figure 6: CDF distribution of block mean coverage ($\bar{C}$ or density) in S_{30w}, S_{38c} and S_{38w}.

towards larger events. To provide a more careful study of the validity of our approach, we randomly pick 50 events from a total of 1295 events in Survey S_{40w}. We then attempt to confirm each using BGP information as described in Section 4.1.

Table 3 summarizes our results. We are able to fully or partially confirm 38% of the cases by finding either corresponding BGP withdrawal or announcement messages. Randomly selected events are often small (as confirmed in Section 5.1), and we are more able to verify large events. One possible reason smaller events do not appear in the control plane is that smaller networks more often use default routing. Bush et al. describe how default routing can result in "reachability without visibility", as addresses may be reachable without visibility to the BGP control plane [2]. Our results are consistent with a corollary: "invisible unreachability", as these default-routed addresses can go down without corresponding BGP messages. We are currently working to verify this hypothesis with additional validation.

4.4 Coverage

While the prior case studies establish that we observe real outages and can confirm them in routing and news, our approach only applies to blocks where several addresses respond to our probes. We currently require that, on average, 10% of addresses in a blocks to respond ($\alpha = 0.1$ in Section 3.2).

To understand how this requirement affects the coverage of our approach, Figure 6 shows the CDF of mean coverage per block, for all blocks in Surveys S_{30w}, S_{38c} and S_{38w}. (Distributions for each are nearly identical.)

This graph shows that we do not have sufficient responsiveness in about half of the blocks to draw any conclusions about them using our algorithms. This analysis actually overestimates our coverage compared to a random sample of Internet blocks, because the survey population is drawn from blocks that have had some response from a prior Internet census [12].

While our coverage is limited, *no* strategy dependent on active probing will be successful when probing non-responsive blocks. The approximately 11,000 blocks for which we have sufficient information to track do represent about 1% of the *responsive* blocks in the Internet, a large enough sample to provide stable results as shown in the next two sections. In addition, we know which blocks are unsuitable for analysis, and so make no claims about their status.

To put our coverage into perspective, we compare it to the coverage of the Hubble system [13]. Hubble probes only one address in each /24 block (the address ending in .1). While this address is the most likely to respond [8] among all 256 last octets, it only responds about 0.86% of the times. Thus our coverage should be better than Hubble's. Another alternative would be to track and probe the address most likely to respond, using a hitlist [8]. However that work suggests that hitlist predictions is at best about 50–60% accurate, suggesting that hitlist-based probing should have coverage about equal to ours. Bush et al. use a number of representatives for each target prefix [2], although such a list is likely challenging to maintain. Best coverage might be obtained by combining hitlists and our full probing, to track both sparse but stable blocks and less sparse but dynamic blocks; such a combination is future work.

4.5 Evaluation from Different Locations

Probing location can affect evaluations of network outage. If the first hop ISP supporting the probing site were unreliable, we would underestimate overall network reliability. In Section 4.2 we discussed how we can detect and correct for local outages that would skew our results; here we address this question more generally by comparing results from two different probing sites.

Our probing takes place regularly from two different sites, ISI and CSU, each with several upstream network providers. Probing site is indicated in survey names, with "w" for ISI (ISI-west), and "c" for Colorado State. Network service at ISI is through Los Nettos, a L.A.-based regional network with peerings with Level3 and Verio; and with connections to Internet2 via USC and CalIT2. CSU has connections to Level3, Qwest, Comcast, Internet2 and NLR through Front Range GigaPop, a non-profit Colorado-based regional network.

To evaluate if ISI and CSU differ, Figure 7 indicates ISI surveys with open symbols and CSU with filled symbols, and it calls out survey location at the top of the graph. Surveys generally alternate with CSU following immediately after ISI completes, although Survey S_{39w2} was an extra survey, and Survey S_{40c} was unavailable at time of analysis. Visually, this graph

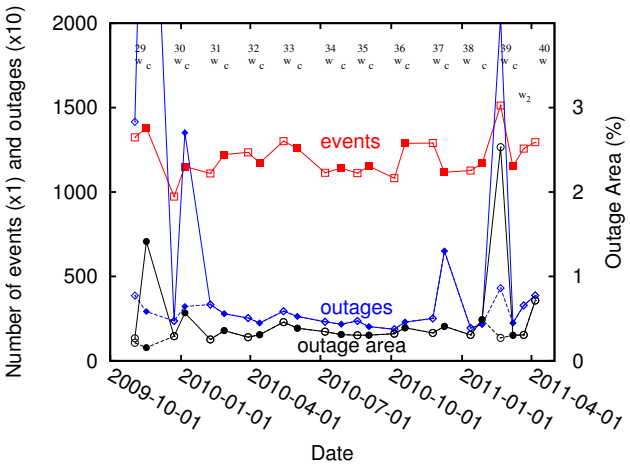

Figure 7: Statistics of Internet events, outages and outage percentage over time. Outages for S_{29c} are 40,627, omitted from the graph for scale. Dotted lines show statistics with local outages removed.

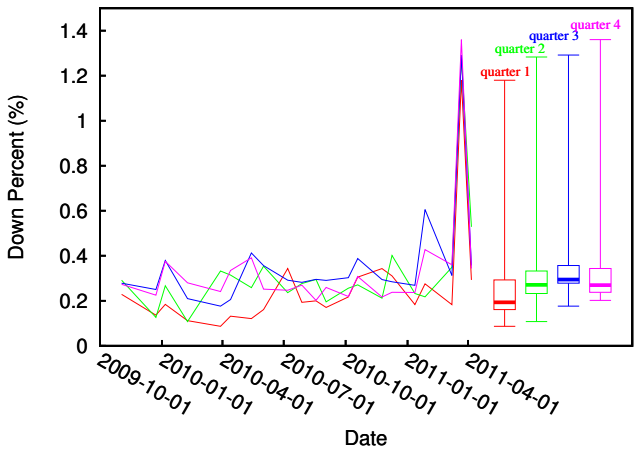

Figure 8: Downtime percentage over time, for 4 different quarters of our dataset.

suggests that ISI and CSU provide similar results.

To strengthen this comparison we carried out Student's t-test to evaluate the hypothesis that our estimates of events, outages, and $\bar{\Omega}$ for the two sites are equal. The test was unable to reject the hypothesis at 95% confidence, suggesting the sites make statistically similar observations.

4.6 Evaluation at Different Dates

In addition to location, we wish to know how consistent the results are across time. Again, the trends in Figure 7 suggest fairly stable results over time, with two exceptions. Surveys S_{29c} and S_{39w} each had extended local outages, for about 41 and 4 hours, respectively. These local outages show up as peaks in on the outage count and $\bar{\Omega}$ estimates; they do not change the event estimate because each outage is mapped to a single network event.

We can detect local outages (Section 4.2). When we remove them from the datasets, as shown in the dotted lines, our estimates of network stability in these surveys are the same as others.

We conclude that the network is fairly stable, with a mean outage level ($\bar{\Omega}$) around 0.34% (standard deviation 0.1%) after local outages are removed (or 0.50% if not removed).

4.7 Evaluation By Quarter Components

Each of our survey dataset uses four components (quarters) as probing target blocks: stable and fixed blocks; stable but randomly selected blocks; randomly chosen blocks, with an odd third octet; and randomly chosen blocks, with an even third octet.

To validate if our results are skewed by selection of blocks, we plot the outage percentage of the four quarters over time (Figure 8). We also plot the outage percentage quartiles of all four quarters in the right part of Figure 8 (raw data in Table 4), showing we are slightly under-estimating the Internet's outages, as Quarter 1 (stable fixed blocks) has less overall outage rates (2.5%), while other three quarters' outage rates are around 3.3%. Our explanation is Quarter 1 are the most stable blocks in our datasets, which are typically broadband or server blocks. The final reported results are not skewed much since the other three quarters dominate.

4.8 Comparison with Other Approaches

Our approach probes all addresses in a block to evaluate outages. The alternative for active probing is to probe a single address, possibly multiple times, as is done in Hubble [13]. Probing an entire block is much more network intensive, but by avoiding the requirement to select a representative address in each block it covers different blocks as discussed in Section 4.4.

In this section we quantify the accuracy of three strategies for active probing: *all*, our approach, probing all addresses; *single*, an approximation of Hubble's approach, probing only the .1 address of each block; *hitlist*, the approach using the best representative address of each block [8]; and *any*, an extreme strategy where we probe all addresses, but consider the block up if any single address responds. We simulate all three methods by replaying a Survey S_{30w} in the results shown here. (We get very similar results when we also evaluate Surveys S_{38w} and S_{38c}.)

The *single* case only approximates Hubble; *single* is slightly pessimistic since it tries one probe per round, while to account for packet loss, Hubble probes up to seven times per round if it gets no reply. Our single-

Quarter	Mean	Min	Max	q1	q2	q3
1	0.0025523	0.000866867077143	0.011800197257	0.00161058742452	0.00193414267938	0.00293013873361
2	0.0032409	0.00107654505386	0.0128316850899	0.00232495925303	0.00270794634197	0.00332295497696
3	0.003587	0.00176366728465	0.0129180618287	0.00278342657238	0.00294431314472	0.0035681346342
4	0.0033624	0.00202058108227	0.0136045782656	0.00237974571976	0.00269681834101	0.00343634416036

Table 4: Outage percentage statistics of four quarters.

		all vs. any	single vs. any	single vs. all	hitlist vs. all
true positives	$(A \wedge B)$	99%	56%	56%	79%
true negatives	$(\bar{A} \wedge \bar{B})$	0.30%	0.31%	0.31%	0.02%
false positives	$(A \wedge \bar{B})$	0.01%	0%	0.01%	0.18%
false negatives	$(\bar{A} \wedge B)$	0.03%	44%	44%	20.8%

Table 5: Comparing amount of data considered when estimating outage.

probe estimates are therefore too high by the degree to which packet loss occurs; loss is typically estimated at a few percent. (Note that *all* and *any* are robust to packet loss since they consider multiple probes to make an estimate.)

To compare alternatives, we evaluate methods A and B in pairs, treating A as a trial and B as truth, then count true and false positives (up) and negatives (down) by comparing A against B.

Table 5 compares several combinations of the three amounts of input. We see that our approach (all) is both correct and complete—it misrepresents outages far less than 1% of the time, because it considers all addresses in the block. The penalty of all is that its coverage is lower as it refuses to classify some blocks as discussed in Section 4.4.

Picking a single address, instead, has a fairly high false negative rate (44%). This result means that systems that estimate network outages by tracking a single address are likely to *over-estimate instability* of the network. Using the best representative address (hitlist) compensates but cannot eliminate this error. We believe the source of this error is because for blocks with dynamic address assignment, there is no good single address, as discussed previously [8].

5. EVALUATING INTERNET OUTAGES

Using our new approach to actively measure outages, we next look at the characteristics of outages and events for the Internet as a whole. We look at this data in two ways, first exploring event and outage durations (Section 5.1), then examining network wide stability by exploring marginal distributions ($\bar{\Omega}_B$ and $\bar{\Omega}_I$) across Internet space and time in Section 5.2.

Our observations are fairly stable across both survey location (Section 4.5) and survey time (Section 4.6), at least after detecting and correcting for local outages (such as shown in Figure 7). Here we use five surveys (S_{30w}, S_{38w}, S_{38c}, S_{39w}, S_{39c}) to confirm the consistency of results. We believe our overall results reflect Internet-wide stability within the limits of measurement error, at least as observed from the United States.

5.1 Durations and Sizes of Internet Outages and Events

We first consider the durations and sizes of block-level outages and network-wide events (Figure 9).

Beginning with outages (Figure 9a), we see that half to three-quarters of outages last only a single round. Our current analysis limits precision to one round (11 minutes), but possible future work could examine individual probes to provide more precise timing. All surveys but Survey S_{39w} have the same trend; Survey S_{39w} diverges due to its local outages, but joins the crowd when they are removed (dotted line $S_{39w'}$).

We also see that 80% of outages last less than two hours. While there is no sharp knee in this distribution, we believe this time period is consistent with human timescales where operators detect and resolve problems.

Network events group individual blocks into time-correlated causes, and in Figure 9b event durations, computed as the mean duration of each event's component outages. This figure shows that about one-third of single-round outages cluster into single-round events, since about 40% of events last one round instead of 50–75% of outages. With less strict clustering ($\theta = 10$ instead of $\theta = 2$) this trend grows, with only 20% of events being one round long.

About 60% of events are less than hour long, but there is a fairly long tail of outages to the limits of our observation (2 weeks or 20,000 minutes).

Because local outages correspond to a single event, Survey S_{39w} resembles the other surveys both with and without removal of local outages, and Survey S_{39w} is indistinguishable from $S_{39w'}$.

Finally, Figure 9c shows event sizes. Almost all events are very small: about 62% of events effect only a single block, and 95% are 4 blocks or smaller. Nevertheless, a few large outage events do occur, as discussed in Section 4.2.

5.2 Towards an Internet-wide View of Outages

We next shift our attention to the Internet as a whole.

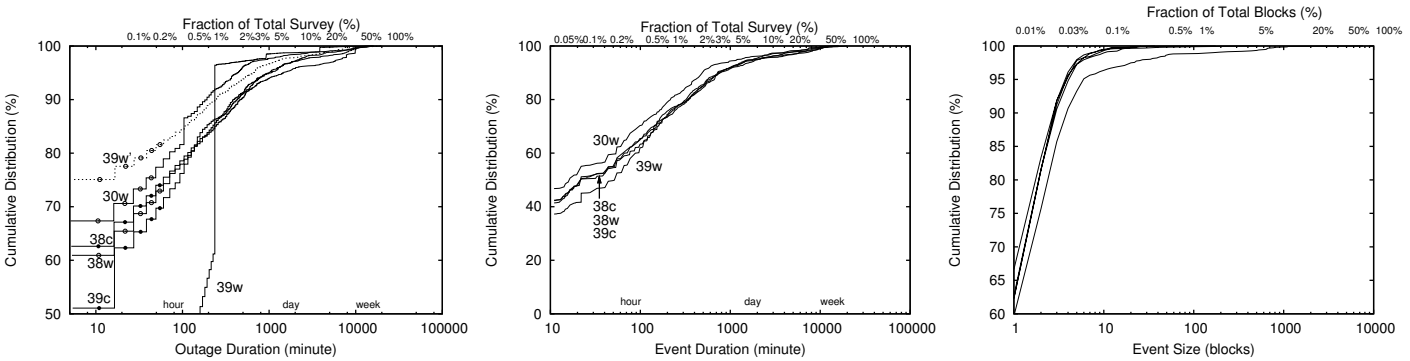

(a) Block outage durations. (b) Network event durations. (c) Network event sizes.

Figure 9: Cumulative distributions of outage and event durations, from Surveys S_{30w}, S_{38c}, S_{38w}, S_{39c}, S_{39w}. The dotted line is Survey S_{39w} with local outages removed. The CDFs of (a) and (c) focus only on portions of the graph.

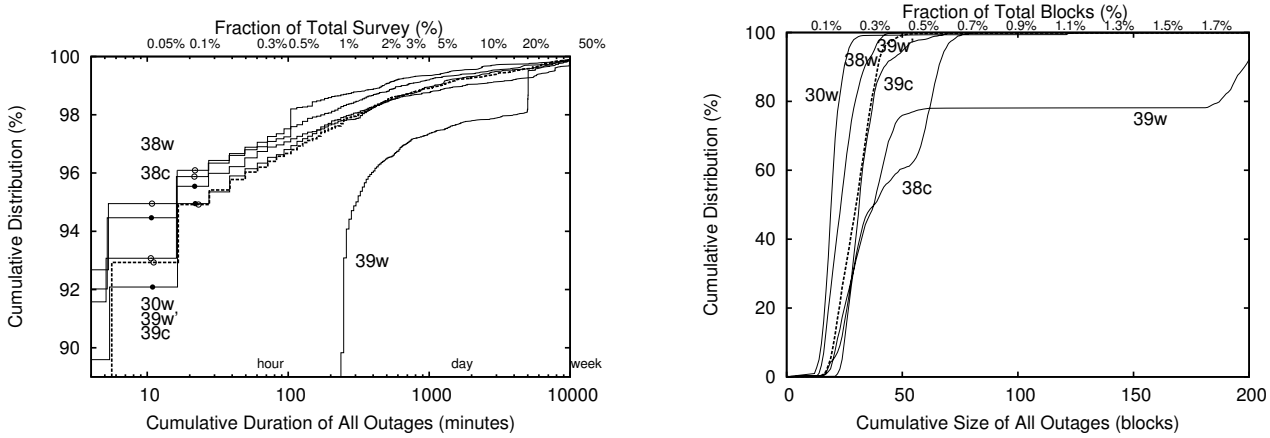

(a) Marginal distribution by rounds, $\bar{\Omega}_I$. (b) Marginal distributions of blocks, $\bar{\Omega}_B$.

Figure 10: Marginal distributions of outage, by round and block, from Surveys S_{30w}, S_{38c}, S_{38w}, S_{39c}, S_{39w}. The dotted line is Survey S_{39w} with local outages removed. The CDF for (a) starts at 80%.

How often is a typical block down, and how much of the Internet is inaccessible at any given time? To understand these questions, Figure 10 shows the marginal distributions of outages by round and block.

First we consider distribution by rounds in Figure 10a. As expected, we see the vast majority of the blocks in our survey are always up: from 92 to 95% of blocks have no outages over each two week observation. The exception is Survey S_{39w}, where two local outages partitioned the monitors from the Internet for about two hours. When we remove local outages, this survey (as the dotted line) becomes consistent with the others. About 2% of blocks are out once (the step at 11 for one round) and the remaining tail follows the distribution of Figure 9a.

Turning to space, Figure 10b shows marginal distributions of $\bar{\Omega}_B$. Survey S_{39w} is again an outlier due to large local outages, but it resembles the others when local outages are removed.

Considering Figure 10b as a whole, we see that *almost always, some part of the Internet is inaccessible*. Typically 20 to 40 blocks of our survey are unreachable at all times. This result is consistent with our observations from Figure 7 that show 0.2% to 0.4% of the Internet is out, averaged over entire surveys. In addition, we see a set of unusually large outages in Survey S_{38c}, where the 50%ile outage is around 38 blocks, but 80%ile is at 63 blocks. We discuss the root causes for these outages Section 4.2 and Figure 2.

Our analysis of Internet-wide outages is preliminary, but it illustrates the utility of automated methods for detecting and quantifying outages in the data plane.

6. CONCLUSIONS

Researchers have studied Internet outages with control- and data-plane observations for many years. We show that active probing of a random sample of /24 blocks provides a powerful new method to more accurately characterize network outages, and we described algorithms to visualize those outages and cluster them by time into network-wide events. We validate this approach through several case studies, and more rigorously through a random sample of network events. We also show that our approach is stable across time and location, provided one corrects for outages at or very near the observation site. While probing of entire /24 blocks is more network intensive, we find it is quite accurate, particularly compared to use of single representative addresses. Finally, we use our new approach to begin to study Internet-wide reliability and typical outage size and duration.

Acknowledgments

We thank Yuri Pradkin for carrying out survey collection and debugging survey interactions with iptables. We thank Jim Koda (ISI), Brian Yamaguchi (USC), and CSU network operations for providing BGP feeds to assist our evaluation, and Dan Massey, Christos Papadopoulos, Mikhail Strizhov for assisting with BGP-mon and at CSU.

7. REFERENCES

[1] Randy Bush, James Hiebert, Olaf Maennel, Matthew Roughan, and Steve Uhlig. Testing the reachability of (new) address space. In *Proceedings of the ACM Workshop on Internet Nework Management*, pages 236–241, Kyoto, Japan, August 2007. ACM.

[2] 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, New York, NY, USA, 2009. ACM.

[3] Xue Cai and John Heidemann. Understanding Block-level Address Usage in the Visible Internet. In *Proceedings of the ACM SIGCOMM Conference*, pages 99–110, New York, NY, USA, 2010. ACM.

[4] Di-Fa Chang, Ramesh Govindan, and John Heidemann. The Temporal and Topological Characteristics of BGP Path Changes. In *Proceedings of the IEEE International Conference on Network Protocols*, pages 190–199, Atlanta, Georga, USA, November 2003. IEEE.

[5] James Cowie. Egypt leaves the Internet. Renesys Blog `http://www.renesys.com/blog/2011/01/egypt-leaves-the-internet.shtml`, January 2011.

[6] James Cowie. Egypt returns to the Internet. Renesys Blog `http://www.renesys.com/blog/2011/02/egypt-returns-to-the-internet.shtml`, February 2011.

[7] James Cowie. Libyan disconnect. Renesys Blog `http://www.renesys.com/blog/2011/02/libyan-disconnect-1.shtml`, February 2011.

[8] Xun Fan and John Heidemann. Selecting Representative IP Addresses for Internet Topology Studies. In *Proceedings of the ACM Internet Measurement Conference*, pages 411–423, Melbourne, Australia, November 2010. ACM.

[9] Nick Feamster, David G. Andersen, Hari Balakrishnan, and Frans Kaashoek. Measuring the Effects of Internet Path Faults on Reactive Routing. In *ACM Sigmetrics - Performance 2003*, San Diego, CA, June 2003.

[10] Anja Feldmann, Olaf Maennel, Z. Morley Mao, Arthur Berger, and Bruce Maggs. Locating Internet Routing Instabilities. In *Proceedings of the ACM SIGCOMM Conference*, pages 205–218, New York, NY, USA, 2004. ACM.

[11] Ramesh Govindan and Hongsuda Tangmunarunkit. Heuristics for Internet Map

Discovery. In *Proceedings of the IEEE Infocom*, pages 1371–1380, Tel Aviv, Israel, March 2000. IEEE.

[12] 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.

[13] 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 Networked Systems Design and Implementation*, NSDI'08, pages 247–262, Berkeley, CA, USA, 2008. USENIX Association.

[14] Ken Keys. Internet-scale IP alias resolution techniques. *ACM Computer Communication Review*, 40(1):50–55, January 2010.

[15] Manas Khadilkar, Nick Feamster, Matt Sanders, and Russ Clark. Usage-based DHCP Lease Time Optimization. In *Proceedings of the 7thACM Internet Measurement Conference*, pages 71–76, October 2007.

[16] Ramana Rao Kompella, Jennifer Yates, Albert Greenberg, and Alex C. Snoeren. IP fault localization via risk modeling. In *Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation - Volume 2*, NSDI'05, pages 57–70, Berkeley, CA, USA, 2005. USENIX Association.

[17] Ramana Rao Kompella, Jennifer Yates, Albert Greenberg, and Alex C. Snoeren. Detection and Localization of Network Black Holes. In *In Proceedings of IEEE Infocom*, 2007.

[18] Craig Labovitz, Abha Ahuja, Abhijit Bose, and Farnam Jahanian. Delayed Internet routing convergence. In *Proceedings of the ACM SIGCOMM Conference*, pages 175–187, New York, NY, USA, 2000. ACM.

[19] Craig Labovitz, G. Robert Malan, and Farnam Jahanian. Internet routing instability. In *Proceedings of the ACM SIGCOMM Conference*, pages 115–126, New York, NY, USA, 1997. ACM.

[20] Harsha V. Madhyastha, Tomas Isdal, Michael Piatek, Colin Dixon, Thomas Anderson, Arvind Krishnamurthy, and Arun Venkataramani. iPlane: an information plane for distributed services. In *Proceedings of the 7th symposium on Operating systems design and implementation*, OSDI '06, pages 367–380, Berkeley, CA, USA, 2006. USENIX Association.

[21] Ratul Mahajan, David Wetherall, and Tom Anderson. Understanding BGP misconfiguration. In *Proceedings of the ACM SIGCOMM Conference*, pages 3–16, New York, NY, USA, 2002. ACM.

[22] Om 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.

[23] Athina Markopoulou, Gianluca Iannaccone, Supratik Bhattacharyya, Chen nee Chuah, and Christophe Diot. Characterization of Failures in an IP Backbone. In *In IEEE Infocom2004, Hong Kong*, 2004.

[24] University of Oregon. The Route Views Project. `http://www.routeviews.org/`.

[25] Yuval Shavitt and Eran Shir. DIMES: let the internet measure itself. *SIGCOMM Computer Communication Review*, 35:71–74, October 2005.

[26] Renata Teixeira and Jennifer Rexford. A measurement framework for pin-pointing routing changes. In *Proceedings of the ACM SIGCOMM workshop on Network troubleshooting*, NetT '04, pages 313–318, New York, NY, USA, 2004. ACM.

[27] International Business Times. Optus, Telstra see service outages after Cyclone Yasi, Feb. 3 2011. `http://hken.ibtimes.com/articles/108249/20110203/optus-telstra-see-service-outages-after-cyclone-yasi htm`.

[28] Los Angeles Times. Amazon apologizes for temporary server outage. `http://www.latimes.com/business/la-fi-amazon-apology-20110430,0,4604776.story`.

[29] New York Times. Egypt cuts off most internet and cell service. `http://www.nytimes.com/2011/01/29/technology/internet/29cutoff.html`.

[30] Daniel Turner, Kirill Levchenko, Alex C. Snoeren, and Stefan Savage. California fault lines: understanding the causes and impact of network failures. In *Proceedings of the ACM SIGCOMM Conference*, pages 315–326, New York, NY, USA, 2010. ACM.

[31] Webnet77. IpToCountry database, March 2011. `http://software77.net/geo-ip/`.

[32] He Yan, Ricardo Oliveira, Kevin Burnett, Dave Matthews, Lixia Zhang, and Dan Massey. BGPmon: A real-time, scalable, extensible monitoring system. In *Proceedings of the IEEE Cybersecurity Applications and Technologies Conference for Homeland Security (CATCH)*, pages 212–223, Washington, DC, USA, March 2009. IEEE.

[33] Ming Zhang, Chi Zhang, Vivek Pai, Larry

Peterson, and Randy Wang. PlanetSeer: Internet Path Failure Monitoring and Characterization in Wide-area Services. In *Proceedings of the 6th conference on Symposium on Opearting Systems Design & Implementation - Volume 6*, pages 12–12, Berkeley, CA, USA, 2004. USENIX Association.