Trinocular: Understanding Internet Reliability Through Adaptive Probing *

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

ABSTRACT

Natural and human factors cause Internet outages—from big events like Hurricane Sandy in 2012 and the Egyptian Internet shutdown in Jan. 2011 to small outages every day that go unpublicized. We describe *Trinocular*, an outage detection system that uses active probing to understand reliability of edge networks. Trinocular is *principled*: deriving a simple model of the Internet that captures the information pertinent to outages, and populating that model through long-term data, and learning current network state through ICMP probes. It is *parsimonious*, using Bayesian inference to determine how many probes are needed. On average, each Trinocular instance sends fewer than 20 probes per hour to each /24 network block under study, increasing Internet "background radiation" by less than 0.7%. Trinocular is also *predictable* and *precise*: we provide known precision in outage timing and duration. Probing in *rounds* of 11 minutes, we detect 100% of outages one round or longer, and estimate outage duration within one-half round. Since we require little traffic, a single machine can track 3.4M /24 IPv4 blocks, all of the Internet currently suitable for analysis. We show that our approach is *significantly more accurate* than the best current methods, with about one-third fewer false conclusions, and about 30% greater coverage at constant accuracy. We validate our approach using controlled experiments, use Trinocular to analyze two days of Internet outages observed from three sites, and re-analyze three years of existing data to develop trends for the Internet.

Categories and Subject Descriptors

C.2.3 [**Computer-Communication Networks**]: Network Operations—*Network Monitoring*; C.2.5 [**Computer-Communication Networks**]: Local and Wide-Area Networks—*Internet*; C.4 [**Performance of Systems**]: Reliability, availability, and serviceability

Keywords: Internet reliability; network outages; Bayesian inference; adaptive probing

*This research is sponsored by the Department of Homeland Security (DHS) Science and Technology Directorate, HSARPA, Cyber Security Division, BAA 11-01-RIKA and Air Force Research Laboratory, Information Directorate under agreement number FA8750-12-2-0344, and contract number D08PC75599. The U.S. Government is authorized to make reprints for Governmental purposes notwithstanding any copyright.The views contained herein are those of the authors and do not necessarily represent those of DHS or the U.S. Government.

SIGCOMM'13, August 12–16, 2013, Hong Kong, China.

ACM 978-1-4503-2056-6/13/08 ...$15.00.

1. INTRODUCTION

Although rare, network outages are a serious concern since users depend on connectivity, and operators strive for multiple "nines" of reliability. Replicated services and content delivery networks may conceal outages, but not eliminate them, and the size of the Internet means outages are always occurring somewhere. Outages are triggered by natural disasters [22,33], political upheavals [31], and human error [21].

Prior work has generally focused on outages from the perspective of routing. Groups today directly monitor routing [5], track routable prefixes with control- and data-plane methods [18,20], and study traffic to unoccupied addresses [8]. While these approaches are useful to detect and sometimes mitigate *large* outages related to routing, most of the Internet uses default routing [3], and we show that most outages are *smaller* than routable prefixes. While some systems target probing to detect specific kinds of smaller outages [29], to our knowledge, no service today actively tracks outages in all Internet *edge* networks.

The contribution of this paper is to address this gap, providing unbiased, accurate measurements of Internet reliability to all analyzable edge networks. First, we describe Trinocular[1], an adaptive probing system to detect outages in edge networks. Our system is *principled*, deriving a simple model of the Internet that captures the information pertinent to outages, parameterizing the model with long-term observations, and learning current network state with probing driven by Bayesian inference.

Second, using experiments, analysis, and simulation, we validate that these principles result in a system that is *predictable* and *precise*: we detect 100% of outages longer than our periodic probing interval, with known precision in timing and duration. It is also *parsimonious*, requiring minimal probing traffic. On average, each Trinocular instance increases traffic to covered networks by no more than 0.7% of the Internet's "background radiation". This low rate allows a single computer to monitor the entire analyzable Internet, and multiple concurrent instances to identify outage scope.

Finally, we use Trinocular to observe two days of Internet outages from three sites. We also adapt our model to re-analyze existing data, developing three years of trends from measurements of samples of the Internet. This reanalysis includes observations of outages due to Hurricane Sandy in 2012, the Japanese Earthquake in March 2012, and the Egyptian Revolution in January 2012.

2. PROBLEM STATEMENT

Our goal is to provide principled, predictable, precise, and parsimonious record of network outages at the Internet edge.

By *principled*, we mean we build a simple model of network blocks and track their status through learning and ac-

[1]We call our system *Trinocular* after the three states a block make take: up, down, or uncertain.

tive probes (§4). Our simple model is, of course, incomplete and unsuitable to model *all* aspects of the Internet, but we show it is well suited to track outages. We use multi-year network observations to inform our model, establishing the expected behavior of each block (a /24 network prefix). We use Bayesian inference to provide a strong theoretical basis to learn the status of each block, and to decide how many probes to send to improve our belief when it is uncertain. We use periodic probes at fixed-interval, multi-minute *rounds* to detect network outages with a known degree of precision. We use adaptive probing at timescales of seconds to quickly resolve inconsistent information and distinguish transient or non-network behavior (such as packet loss or edge system failure) from outages at the target network. Our default measurements employ three years of quarterly observations at long timescales, rounds of 11 minutes at medium timescales (following [13, 29]), and 3 second intervals for adaptive probes, although these values can be adapted to trade precision for traffic.

By *predictable*, we mean our conclusions about analyzable network blocks provide guaranteed *precision* and positive statements about block status (§5). Our periodic probing bounds the precision of detecting block transitions, and we show that error in estimates of outage duration is uniformly distributed by one half round (±330 s). As with all active probing mechanisms, our approach cannot determine the status of networks that decline to participate, such as those that use firewalls that block probes, nor networks that are too sparse for our techniques. We find 3.4M /24 blocks to be *analyzable* by our method, and we identify non-analyzable blocks. This coverage is 30% greater than current approaches, if one holds accuracy constant.

By *parsimonious*, we mean that we use a minimum number of probes required to establish our belief in edge network state. Long-term history informs our model, and Bayesian reasoning justifies each probe we make, avoiding unnecessary probes. Minimizing probing traffic is critical for a service that operates across the entire Internet. While money can solve the problem of outgoing network capacity at the prober, recipients of probing traffic are very sensitive. Even modest traffic can draw complaints (for which we maintain an opt-out list). We evaluate the impact of our traffic on target networks by comparing it to the amount of background radiation that all public networks observe [34]. We show that at steady state, each Trinocular instance increases background traffic by less than 0.7%, allowing us to run multiple instances to understand outage scope.

Finally, our target is *all edge networks*. We are interested in edge networks because prior work has shown that many networks employ default routing [3], and outages occur inside ISPs [29]. We show that probing all /24s detects many more outages than considering only ASes or routed prefixes (§6). We combine data from three sites to study outage scope, separating outages adjacent to the prober from *partial* and *global* outages affecting some or all of the Internet.

These four characteristics distinguish our work from prior work, which often employs ad hoc mechanisms, does not provide guarantees about outage precision, requires excessive probing, or monitors routable prefixes instead of considering smaller outages in edge networks. They also allow us to provide unique view of Internet reliability, both as a whole, and of specific events (§7).

3. RELATED WORK

We next review prior outage studies by data source.

3.1 Control-plane Studies

Several prior efforts use control-plane data to study Internet outages. Markopoulou et al. use IS-IS update messages to categorize failure types in Sprint's network [23]. Unlike their work, our system uses only data-plane information.

Omni runs servers in each Autonomous System (AS) and uses the forwarding tables and traceroutes to diagnose routing changes [30]. Their approach benefits from non-public routing information, but deployment is challenging. Our work uses centrally-collected measurement and analysis and is easier to deploy since it does not require peering.

Huang et al. combine data from multiple BGP feeds to detect "faint" outages [15]. We also use data from multiple vantage points, but to distinguish between global and local outages; our mechanism can detect small and short events.

BGP misconfiguration is one cause of outages. Mahajan et al. study routing messages and contact network operators about problems [21]. They also use active probing to determine the effects of misconfiguration on connectivity. They report that 0.2% to 1% of prefixes have problems each day. We confirm their results on Internet reachability, finding about 0.6% of the Internet blocks are out, on average.

In general, studies using or triggered by control-plane information are indirect and provide incomplete coverage of all outage types, as discussed by Bush et al. [3]. We further verify this result experimentally (§6.2).

3.2 Data-plane Studies

Several efforts use data-plane probes to detect outages and are close to our work. First, NetDiagnoser [9] and Cunha et al. [6] explore binary tomography to identify routing problems. Their work identifies efficient ways to localize problems with minimal traffic. We also focus on minimizing traffic, but our goal is continuous monitoring of all edge networks, not diagnosing problems in specific ASes.

Second, Hubble finds potential Internet outages by surveying all .1 addresses in each routed prefix and selecting one for regular probes, which trigger traceroutes to confirm and localize a potential outage [17]. We instead regularly probe many selected addresses in each /24 block. Our examination of multiple addresses and /24 blocks detects outages missed by routing and single-address triggering (§6). IPlane captures information about network performance, aggregating information by routable prefixes [20]. We show that it is possible and beneficial to maintain outage information at the granularity of /24 blocks. Our work could be extended with Hubble-like traceroutes to localize outages.

Building on Hubble and iPlane, LIFEGUARD extends this approach to detect and work around local outages caused by routing [18]. Our work's focus on edge networks complements LIFEGUARD's on partial failures in the routing system and the network core. LIFEGUARD detects outages for routable prefixes because that coarser granularity is relevant to re-routing to recover. We instead focus on finer granularity to understand smaller, edge networks, and do not attempt recovery because edge networks are not usually multi-homed.

Schulman and Spring target ICMP probing to study using weather reports [29]. They probe many individual addresses in areas with severe weather from around ten vantage points, and report outages for individual addresses. Like

their work, we are interested in edge networks, but we track blocks, not individual addresses, and we track all that are analyzable, not just those in regions under severe weather. We consider blocks out of concern that tracking single addresses risks confounding outages with human activity (such as suspending a laptop); but a more complete comparison is future work.

In prior work we took censuses of all IPv4 using ICMP, establishing what coverage is possible with active probing [13]. That coverage is an upper bound on our coverage of outage detection. We re-analyze datasets from this work for longitudinal analysis (§7.2), and it inspires our new adaptive probing scheme. In later work, we explored using this data to identify outages [25, 28], and to visualize both outages and BGP changes [26]. This outage work is only preliminary (published as a poster [28] and technical report [25]), and uses methods that require many more probes than Trinocular, and typically underestimate outage duration by 1.5 rounds. We instead use Bayesian analysis to make informed decisions with far less network traffic, and to improve the precision of outage detection to within a half-round.

Finally, Bush et al. study the reachability of Internet address space using traceroute to detect incorrect filtering [2] and to find biases in reachability experiments [3]. We provide additional evidence supporting their observation that default routes are widely used and that control-plane measurements underestimate outages.

3.3 Client-supported Analysis

Client-side observations provide a wider perspective than the centralized methods. Several groups have used meshes of measurement computers [1, 12, 19, 24]. Such experiments can provide strong results for the behavior of the networks between their n vantage points (typically less than 50), and for small n link coverage grows as $O(n^2)$, although edge coverage is only $O(n)$. Without probing outside the mesh, however, these approaches ultimately study only a small fraction of the entire Internet. Other methods of active probing, and our work, aim to provide complete coverage.

In early work, Paxson reports routing failures in about 1.5%–3.3% of trials [24]. A more recent work, the RON system reports 21 "path-hours" of complete or partial outages out of a total of 6825 path-hours, a 0.31% outage rate [1]. Feamster et al. measure Internet path failures from 31 vantage points, correlated to BGP for causes [12]. They find that most failures are short (under 15 minutes) and discuss the relationship between path failures and BGP messages. SCORE is a system that extends measurements to isolate the location of problems [19]. As with most of this work, we validate our findings using control plane data.

Rather than a mesh, PlanetSeer studies traffic from 7–12k end-users to a network of 120 nodes to track path outages [35]. They report that their larger population identifies more anomalies than prior work; we expect our edge coverage of 3.4M blocks will be broader still. In addition, their measurements occur only on connected clients; they miss outages from already disconnected clients.

Choffnes et al. collect information from end systems to detect service-level network events [4]. Our work is different in that we probe to the network edge and do not require extra software or specific operating systems in the edge networks.

Client support in these studies allows better fault diagnosis than our work. Our work complements theirs by providing much larger coverage (3.4M /24 blocks, a large fraction of the Internet edge), rather than "only" meshes of hundreds of nodes, or thousands of end hosts. Our centralized measurement also allows stronger statements about coverage since we do not depend on end hosts that may come or go.

3.4 Passive Data Analysis

Recent work by Dainotti et al. considers Internet outages caused by political censorship [7, 8]. They use a novel approach that combines observations from both control-plane (BGP logs) and data-plane sources (traffic to unoccupied addresses at UCSD network telescope and active probing data from Ark). They focus on using multiple passive data sources, finding their active probes are of limited use because they probe each /24 only every three days. We instead show that a single PC can actively track millions of /24 blocks, providing guaranteed precision for blocks that respond to probes. It is unclear if passive analysis can provide strong statements about precision or coverage, but it does provide important insight into networks that block active probes.

Turner et al. have also mined "low-quality" data sources (router configurations, e-mail and syslogs), to detect failures in the CENIC network [32]. Such log analysis requires collaboration with the monitored networks, thus focuses on a single ISP. In contrast, our active probing is done independent of the target.

4. PRINCIPLED LOW-RATE PROBING

Trinocular carries out principled probing: we define a simple model of the Internet to capture elements essential to outage detection. Trinocular establishes *belief* $B(U)$ that each block is available, and uses Bayesian inference to learn the current status of the network. We drive probing using this model and belief, sending at regular intervals to guarantee freshness, and more quickly when necessary to resolve uncertainty about network state.

4.1 An Outage-Centric Model of the Internet

Trinocular's model of the Internet tracks *block*-level outages, measured with *probes* to *active addresses*, and reasons about them using *belief* changed by Bayesian inference.

We study */24 address blocks* (designated b) as the smallest unit of spatial coverage. Larger blocks, such as prefixes that appear in global routes, may capture outages due to routing changes, but they hide smaller outages. Prior work shows that default routing is widely used [3], and outages occur inside ISPs [29], and we show that outages often occur in sizes smaller than routable prefixes (§6.2).

Trinocular sends only ICMP echo requests as *probes*, each with a 4-byte payload. We chose end-to-end, data-plane probing to detect outages unrelated to routing. We use ICMP because it is innocuous and, compared to other options, less likely to be blocked or interpreted as malicious [13].

In each block, we model which addresses are active, the *ever active addresses*, $E(b)$, a set of up to 256 elements. To interpret the meaning of probe responses, we model the *expected response rate* of $E(b)$ as availability, $A(E(b))$, a value from 0 to 1, never to always responding. These dimensions are independent, so a block where $E(b) = 64$ and $A(E(b)) = 0.5$ has one-quarter of addresses that each respond (on average) half the time. We discard very sparse and very unresponsive blocks as non-analyzable (§4.4).

For blocks when $A(E(b)) < 1$, a negative probe response is ambiguous: it can result from probing temporarily unoc-

probe result	prior U^*	$P(probe\|U^*)$	reason
n	U	$1 - A(E(b))$	inactive addr.
p	U	$A(E(b))$	active addr.
n	$\bar{U}$	$1 - (1-\ell)/\|b\|$	non-response to block
p	$\bar{U}$	$(1-\ell)/\|b\|$	lone router?

Table 1: Bayesian inference from current block state U^* and a new probe.

cupied address, or from the block being down. Our model evaluates the likelihood of these events. We show that this model provides more information per probe than current approaches, allowing lower probe rates (§6.1).

Finally, we judge blocks as either *down* (unreachable), *up* (reachable), or *uncertain*, and denote these states as U, $\bar{U}$, or $U^?$. Belief, $B(U)$ ranges from 0 to 1, with low to high values corresponding to the degree of certainty the block is down or up. Probes influence this belief as described next.

4.2 Changing State: Learning From Probes

Trinocular uses Bayesian inference to weigh each probe's information into our understanding of block status.

Probe responses are either positive, p, or negative or non-responses, n, and they affect belief according to conditional probabilities from Table 1. This table reflects the block size, $|b|$, the combined rate of probe and reply loss, ℓ, and the long-term probability that those addresses reply, $A(E(b))$.

The first two lines of the table represent how belief changes when the block is currently up. They reflect the probability of hitting an active address ($A(E(b))$), or an inactive address (1-$A(E(b))$). In this study we treat $A(E(b))$ as a static parameter and derive this value from analysis of long-term observations, so it reflects both transient address usage and possible loss of probes or replies. Since outages are very rare, they have negligible infuence on $A(E(b))$.

The last two lines characterize what we learn when the block is down. The final line is a positive reply to a block that is down. We consider this case to represent the unusual situation where a single router is up, but all addresses "behind" the router are down. This low-probability event will almost always draw subsequent probes that clarify the block's status. This term uses ℓ, representing the probability of packet loss of the probe or reply. On-line estimation of packet loss is future work; we currently use $\ell = 0.01$, a reasonable but arbitrary value; our results are not sensitive to small changes to this value. The third line is the complement of that probability.

A new probe observation results in a new belief B' based on our old belief B as influenced by this table. After a positive response:

$$B'(\bar{U}) = \frac{P(p|\bar{U})B(\bar{U})}{P(p|\bar{U})B(\bar{U}) + P(p|U)B(U)}$$

After a negative- or non-response:

$$B'(\bar{U}) = \frac{P(n|\bar{U})B(\bar{U})}{P(n|\bar{U})B(\bar{U}) + P(n|U)B(U)}$$

with analogous values for $B'(U)$, and $B(\bar{U}) = 1 - B(U)$.

These equations break down, failing to consider alternatives, if conditional probabilities ($P(probe|U^*)$) go to 0 or 1. We avoid this case by capping $A(E(b))$ to 0.99 for stable blocks, and avoiding very intermittent blocks ($A < 0.1$) as unsuitable for analysis, and we also cap belief to at most 0.99 (and at least 0.01).

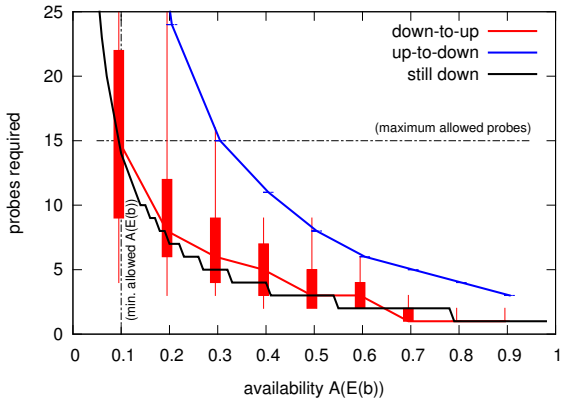

Figure 1: Median number of probes needed to reach a definitive belief after a change in block state. Boxes show quartiles, whiskers 5 and 95%ile; both equal median for outages. Data: analysis and simulation; details: §5.3.

4.3 Gathering Information: When to Probe

Trinocular probes each block with *periodic* probing at medium-timescales coupled with *adaptive* probes sent quickly when we suspect block status may have changed, and *recovery* probes to account for sparse blocks. We probe addresses from $E(b)$ in a pseudorandom order, both to gather information from many addresses and to spread the reply burden.

Periodic probing: We probe each analyzable blocks at a fixed interval so we can bound the precision of our measurements of network outages. Like prior work [13, 29], we use a fixed 11-minute interval for basic probing. The precision in outage measurements follows from this period (see §5.2); we choose it to trade desired precision against traffic. Periodic probing and target rotation are design choices that make Trinocular as lightweight on the target network as possible.

Adaptive probing: We classify a block as down when $B(U) < 0.1$, and up when $B(U) > 0.9$. When a periodic probe causes our belief to become uncertain, or to shift towards uncertainty, we carry out additional, *adaptive, short-timescale* probes to resolve this uncertainty. For adaptive probing, we send new additional probes as soon as each prior probe is resolved until we reach a conclusive belief of the block status. Most probes are resolved by 3 s timeout, so adaptive probes typically occur every 3 s.

Usually a few adaptive probes will quickly resolve uncertainty in our belief; we study this value in §5.3. As address usage becomes sparser, the number of probes to converge grows geometrically (Figure 1). To bound probing, we send at most 15 total probes per round (1 periodic and up to 14 additional adaptive). We cease probing when belief is definitive and not shifting; if we cannot reach definitive belief in 15 probes we mark the block as uncertain. Uncertainty is similar to the "hosed" state in prior work [29]. We speculate that Bayesian analysis could resolve some intermediate states in their work, but detailed comparison is future work.

Recovery probing: There is an asymmetry when blocks transition from down-to-up for intermittently active blocks (low $A(E(b))$). While positive responses are strong evidence the block is up, interpretation of negative responses has increasing ambiguity as A falls. When an intermittent block

comes back up, we still may see several negative responses if probes chance upon temporarily unoccupied addresses.

To account for this asymmetry, we do additional *recovery* probes for blocks that are down. From $A(E(b))$, the probability we get consecutive misses due to k vacant addresses is $(1-A)^k$, resulting in a "false negative" belief that an up block is down. We select k to reach a 20% false-negative rate as a function of A (k is the "still down" line in Figure 1), performing up to $k = 15$ total probes when $A = 0.1$ With recovery probes, false negatives cause outages in sparse blocks that are one third of a round too long, on average.

Traffic: For long-term operation across the Internet, Trinocular must have minimal impact on target networks. Our benchmark is Internet background radiation, the unsolicited traffic every public IP address receives as part of being on the public network. It thus provides a reasonable baseline of unsolicited traffic against which to balance our measurement. A typical unused but routable /8 block receives 22 to 35 billion packets per week [34], so each /24 block sees 2000 to 3300 packets/hour. Our goal is to increase this rate by no more than 1%, on timescales of 10 minutes.

In the best case, we send only 5.4 probes/hour per /24 block in steady state, and if all addresses in a block are active, we probe each address only every other day. This best-case is only a 0.25% increase in traffic. With adaptive and recovery probing, our worst-case probing rate adds 15 probes per 11-minute round, an average probe rate of 82 probes/hour per /24 block, about 5% of the rate of background radiation. Since this worst case will occur only for low-A blocks that change state, we expect typical performance to be very close to best case, not worst case. In §5.3 we show experimentally that median traffic is at 0.4% to 0.7% of our benchmark, our 5% worst case occurs less than 2% of the time.

4.4 Parameterizing the Model: Long-term Observation

We determine parameters $E(b)$ and $A(E(b))$ for each block to weigh the information in each probe.

Initialization: We use long-term, multi-year, Internet censuses to initialize these parameters for each block. Prior work generates regular IP history datasets that provide the information we need [10]. These datasets include the responsiveness of each public, unicast IP address in IPv4 measured 16 times over approximately the last 3 years. We use the full history (16 measurements) to identify $E(b)$. To use recent data, we consider only the 4 most recent censuses to compute $A(E(b))$. We update $E(b)$ every 2-3 months as new history datasets become available, bringing in newly active blocks and retiring gone-dark blocks. Current Internet censuses are specific to IPv4. Our approach applies to IPv6 if $E(b)$ can be determined, but methods to enumerate all or part of IPv6 are an area of active research.

It is very traffic-intensive to track intermittent and sparse blocks with reasonable accuracy (see Figure 1). We therefore discard blocks where addresses respond very infrequently ($A(E(b)) < 0.1$). We also discard blocks that are too sparse, where $E(b) < 15$, so that we are not making decisions based on a very few computers. Because $A(E(b))$ is based on only recent censuses, discard of low $A(E(b))$ blocks removes "gone dark" blocks [10].

Of the 16.8M unicast blocks as of July 2012, we find 14.5M are routed, 8.6M are non-responsive, 0.7M have $E(b) < 15$, 1.5M have $A(E(b)) < 0.1$, leaving 3.4M blocks that are analyzable: 24% of the routed space (and 40% of responsive).

Since most of the Internet is always up, we set belief to indicate all blocks are up on startup.

Evolution: As networks change, model parameters may no longer match the current network. We update our target list and A-estimations every two months as new long-term data is made available. At shorter-timescales, we must handle or adapt when parameter estimates diverge from reality.

Underestimating $E(b)$ misses an opportunity to spread traffic over more addresses. Underestimating $A(E(b))$ gives each probe less weight. In both cases, these errors have a slight affect on performance, but none on correctness.

When $E(b)$ is too large because it includes non-responsive addresses, it is equivalent to overestimating $A(E(b))$. When $A(E(b))$ exceeds the actual A, negative probes are given too much weight and we infer outages incorrectly. Ideally $A(E(b))$ will evolve as a side-effect of probing to avoid false outages when it diverges from the long-term average. Our current system does not track A dynamically (although work is underway), so we detect divergence in post-processing, and identify and discard inaccurate blocks. The result is greater traffic, but few false outages.

Traffic: We do not count long-term observations against Trinocular's traffic budget since it is an ongoing effort, independent of our outage detection. However, even if we take responsibility for all traffic needed to build the history we use, it adds only 0.18 probes per hour per /24 block since collection is spread over 2 months.

4.5 Outage Scope From Multiple Locations

A single site provides only one view of the Internet, and prior work has shown that about two-thirds of outages are partial [17]. We use two approaches to judge outage scope: we detect and eliminate outages where probers are effectively off the network, and we merge views from multiple observers to distinguish between partial and global outages. In §7.1 we report on how frequently these occur in the Internet.

Prober-local outages: Router failures immediately upstream of a prober unenlighteningly suggest that nearly the entire Internet is out. We detect and account for outages that affect more than half the probed blocks.

Partial and global outages: We detect outage scope by merging observations from multiple vantage points. Because each site operates independently and observations tend to occur at multiples of a round, direct comparison of results from different sites will show different timing by up to one round. We correct these differences by taking the *earlier* of two changes that occur, since periodic probes always *delay* detection of a change in block status. We therefore correct disagreements in the merged results only when (a) both sites agree before and after the disagreement, (b) the disagreement lasts less than 1.1 rounds, and (c) the network changes state before and after disagreement. Rules (a) and (b) detect transient disagreement that is likely caused by phase differences. Rule (c) avoids incorrectly changing very short outages local to one vantage point. Merging results thus improves precision. After correction, any remaining disagreement represents partial outages.

4.6 Operational Issues

Our system implementation considers operational issues to insure it cannot harm the Internet.

Probing rate: In addition to per-block limits, we *rate limit all* outgoing probes to 20k probes/s using a simple token bucket. Rate limiting at the prober insures that we do not overwhelm our first-hop router, and it provides a fail-safe mechanism so that, even if all else goes wrong, our prober cannot flood the Internet incessantly. In practice, we have never reached this limit. (This limit is at the prober, spread across all targets. Figure 4 shows that only a tiny fraction of this traffic is seen at each target block.)

We expect our monitor to run indefinitely, so we have implemented a simple *checkpoint/restart* system that saves current belief about the network. This mechanism accommodates service on the probing machine. We restart our probers every 5.5 h as a simple form of garbage collection.

We have run Trinocular for several multi-day periods, and we expect to run Trinocular continuously when adaptive computation of A is added.

Implementation: We use a high-performance ICMP probing engine that can handle thousands of concurrent probes. We use memory-optimized data structures to keep state for each block, leaving CPU cost to match probe replies with the relevant block as the primary bottleneck. We find a single prober can sustain 19k probes/s on one core of our 4-core Opteron. Fortunately, probing parallelizes easily, and with four concurrent probers, a single modest computer can track all outages on the analyzable IPv4 Internet.

5. VALIDATING OUR APPROACH

We validate correctness with controlled experiments, and probe rate by simulation and Internet experiments.

5.1 Correctness of Outage Detection

We first explore the correctness of our approach: if an outage occurs, do we always see it? For a controlled evaluation of this question, we run Trinocular and probe 4 /24 blocks at our university from 3 sites: our site in Los Angeles, and universities 1600 km and 8800 km distant in Colorado and Japan. We control these blocks and configure them in two-hour cycle where the network is up for 30 minutes, goes down at some random time in the next 20 minutes, stays down for a random duration between 0 and 40 minutes, then comes back up. This cycle guarantees Trinocular will reset between controlled outages. We studied these blocks for 122 cycles, yielding 488 observations as dataset $A_{controlled}$, combining data for 4 controlled blocks from datasets A_{1w} (2013-01-19, 4 days), A_{3w} (2013-01-24, 1 day), A_{4w} (2013-01-25, 2 days), A_{7w} (2013-02-12, 2 days)[2].

Figure 2 shows these experiments, with colored areas showing observed outage duration rounded to integer numbers of rounds. We group true outage duration on the x into rounds with dotted black lines. Since periodic probing guarantees we test each network every round, we expect to find all outages that last at least one round or longer. We also see that we miss outages shorter than a round roughly in proportion to outage duration (the white region of durations less than 11 minutes). While these experiments are specific to blocks where addresses always respond ($A(E(b)) = 1$), they generalize to blocks with $A \geq 0.3$ since we later show that we take enough probes to reach a definitive conclusion for these blocks (Figure 1).

These results confirm what we expect based on our sampling schedule: if we probe a block with $A \geq 0.3$, we *always* detect outages longer than one round.

[2] We name datasets like A_{7w} for Trinocular scans of the analyzable Internet (A_{20addr} uses a variant methodology), H_{49w} for Internet histories [11], S_{50j} for Internet Surveys [13]. The subscript includes a sequence number and code for site (w: Los Angeles, c: Colorado, j: Japan).

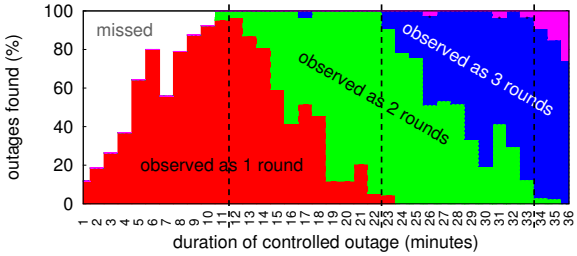

Figure 2: Fraction of detected outages (bar height) and duration in rounds (color), for controlled experiments. Dataset: $A_{controlled}$.

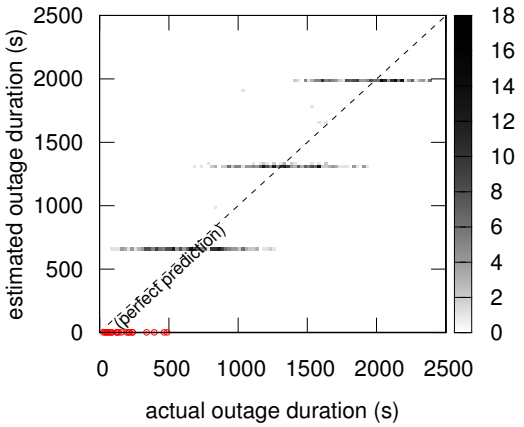

Figure 3: Observed outage duration vs. ground truth. Dataset: $A_{controlled}$ (same as Figure 2).

5.2 Precision of event timing

Figure 2 shows we *do* detect outages. We next evaluate the *precision* of our observed outage durations.

We continue with dataset $A_{controlled}$ in Figure 3, comparing ground truth outage duration against observed outage duration at second-level precision. Our system measures block transition events with second-level precision, but when we examine outage durations, we see they group into horizontal bands around multiples of the round duration, not the diagonal line that represents perfect measurement. We also see that error in each case is uniformly distributed with error plus or minus one-half round. As expected, we miss some outages that are shorter than a round; we show these as red circles at duration 0. Finally, we also see a few observations outside bands, both here and marked with an asterisk in Figure 2. These are cases where checkpoint/restart stretched the time between two periodic probes.

These results are consistent with measurement at a fixed probing interval sampling a random process with a uniform timing. When we compare observed and real event start- and end-times it confirms this result, with each transition late with a uniform distribution between 0 and 1 round.

These experiments use blocks where addresses are always responsive ($A(E(b)) = 1$). We carried out experiments varying A from 0.125 to 1 and can confirm that we see no missed

outages longer than one round and similar precision as long as Trinocular can reach a conclusion ($A > 0.3$). When $0.3 < A < 1$, additional adaptive probes add at most 45 s to detection time (15 probes at 3 s per adaptive probe). For blocks with $A < 0.3$, precision will deteriorate and block status may be left uncertain.

We conclude that periodic probing provides a predictable and guaranteed level of precision, detecting state transitions in just more than a round (705 s, one round plus 15 adaptive probes) for blocks where $A > 0.3$. Greater precision is possible by reducing the round duration, given more traffic.

5.3 Probing rate

Our goal is good precision with low traffic, so we next validate traffic rate. We use simulation to explore the range of expected probing rates, then confirm these based on our Internet observations.

Parameter Exploration: We first use simulation to explore how many probes are needed to detect a state change, measuring the number of probes needed to reach conclusive belief in the new state. Our simulation models a complete block ($|E(b)| = 256$) that transitions from up-to-down or down-to-up. When up, all addresses respond with probability $A(E(b))$. When down, we assume a single address continues to reply positively (the worst case outage for detection).

Figure 1 shows the up-to-down and down-to-up costs. Down-to-up transitions have high variance and therefore have boxes that show quartiles and whiskers 5%ile and 95%ile values. Up-to-down transitions typically require several probes because Trinocular must confirm a negative response is not an empty address or packet loss, but they have no variance in these simulations. Trinocular reaches a definitive belief and a correct result in 15 probes for all blocks with $A > 0.3$.

For down-to-up transitions, 15 probes are sufficient to resolve all blocks in 50% of transitions when $A > 0.15$, and in 95% of transitions when $A > 0.3$. Variance is high because, when A is small, one will probe many unused addresses before finding an active one. This variance motivates recovery probing (the black "still down" line).

Experimentation: To validate these simulations, Figure 4 shows probe rates from A_{7w}, a 48-hour run of Trinocular on 3.4M Internet-wide, analyzable blocks starting 2013-02-12 T14:25 UTC. Here we examine the subset $A_{7w\text{-}5.5h}$ from this data: the first 5.5 hours (30 rounds) from one of the four probers, with 1M blocks; other subsets are similar.

As one would expect, in most rounds, most blocks finish with just a few probes: about 73% use 4 or fewer per round. This distribution is skewed, with a median of 13.2 probes/hour, but a mean of 19.2 probes/hour, because a few blocks (around 0.18%) reach our probing limit per round. Finally, we report that 0.15% of blocks actually show more than expected traffic (the rightmost peak on the graph). We find that a small number of networks generate multiple replies in response to a single probe, either due to probing a broadcast address or a misconfiguration. We plan to detect and blacklist these blocks.

This experiment shows we meet our goals of generating only minimal traffic, with probing at 0.4% (median) to 0.7% (mean) of background radiation, and bounding traffic to each block.

Probe rate as a function of $A(E(b))$: The above experiment shows most blocks require few probes, and our

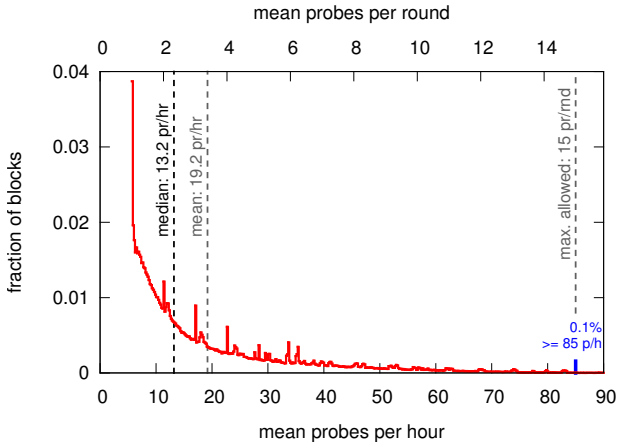

Figure 4: Distribution of probes to each target block. Dataset: $A_{7w\text{-}5.5h}$.

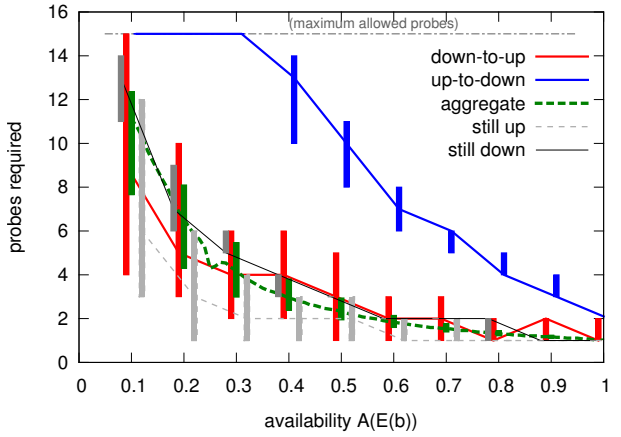

Figure 5: Median number of probes, with quartiles, for aggregate and state transitions. Dataset: $A_{7w\text{-}5.5h}$.

simulations show probe rate at transition depends strongly on address responsiveness. To verify this relationship, Figure 5 breaks down probes required by transition type and each block's $A(E(b))$.

The dotted line and thick quartile bars show aggregate performance across all states. We track blocks with $A > 0.3$ with less than 4 probes per round, with relatively low variance. Intermittent blocks ($A < 0.3$) become costly to track, and would often exceed our threshold (15 probes).

Figure 5 identifies each state transition from Figure 4 separately. We see that the shape of recovery and outages match simulations (Figure 1), although outage detection has larger variance because of imperfect estimation of $A(E(b))$.

Overall this result confirms that Trinocular does a good job of keeping probe rate low, and of adapting the probe rate to meet the requirements of the block.

6. EFFECTS OF DESIGN CHOICES

We next explore two design choices that differ from prior work and contribute to Trinocular accuracy.

6.1 How Many Addresses to Probe

Trinocular sends probes to $E(b)$, the known active addresses in each block. Most prior systems probe a single

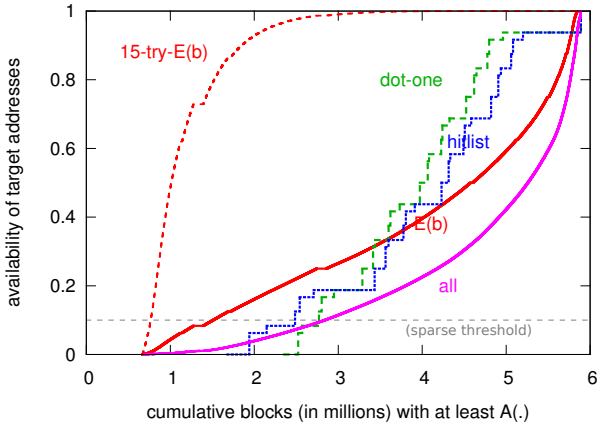

Figure 6: Probability of a positive response (availability) of one address of each block (dot-one, hitlist, ($E(b)$, and all), or any of 15 addresses (15-try-$E(b)$). Dataset: H_{49w}.

address, sometimes a single specific address (such as that ending in .1 [17, 18, 20], or they probe all addresses [28]). We show that alternatives either can cover fewer blocks or gather less information per probe, and may miss outages.

Information Per Probe and Coverage: To evaluate the amount of information each probe may provide, we examine IP history data for each alternative. We begin with history dataset "it49w" [11] (identified here as H_{49w}), summarizing IPv4 censuses from 2006 to 2012. Figure 6 shows the distribution of availability value A for different approaches to selecting probing targets per block: probing .1, a hitlist's single most responsive address [10], all responsive addresses ($E(b)$), and all addresses. This A value correlates to the information a single probe provides about the block, since probing an inactive address does not help determine if the block is up or down.

We first compare probing active to all addresses ($E(b)$ vs. all). The $E(b)$ line has greater availability for all blocks, so $A(E(b)) > A(b)$ and each Trinocular probe provides more information than does probing all.

Dot-one and hitlist do better than $E(b)$ for many of blocks (about 40% of all, from 3.5M to the right), but poorer for about 50% of all blocks (from 0.7M to 3.5M). In many cases (for dot-one, about 2M blocks from 0.7 to 2.5M, about 30% of all blocks), a single address provides *no* coverage where $E(b)$ shows some addresses would respond. Thus, while a single target, may work well for 40% of blocks, particularly when probing includes retries, it provides poor or no coverage for even more blocks—probing $E(b)$ can cover about two-thirds more blocks than .1.

While A characterizes the information provided by a *single* probe, Trinocular sends an adaptive number of probes, allowing low-A blocks to get good coverage. To show an upper bound on Trinocular's ability to find an active address, the curve labeled "15-try-$E(b)$" shows the probabilty that any of 15 probes will respond, suggesting that Trinocular can use multiple probes to provide reasonable results even for blocks with very intermittently responding addresses.

While other systems use secondary methods to improve reliability (perhaps verification with traceroute), or use fewer but larger blocks (§6.2), we show that $E(b)$ *provides about 30% broader coverage than depending on a single address.*

strategy	single	hitlist	Trinoc.	all
samples per /24	1	1	$\|E(b)\|$	256
which addresses	.1	top	ever resp.	all
precision	99.97%	99.98%	100%	(100%)
recall	58.6%	66.6%	96.6%	(100%)

Table 2: Comparing precision and recall of different probing targets. Dataset: S_{50j}.

Effect on Outage Detection: To evaluate the impact of probing choice on outages, we next examine a single dataset with three choices of probe target. We use Internet survey "it50j" [13], a 2-week ICMP probing of all addresses in 40k /24 blocks starting 2012-10-27 (here called S_{50j}). We define any response in probes to all addresses as ground truth since it is the most complete. We define a false outage (fo) as a prediction of down when it's really up in all-probing, with analogous definitions of false availability (fa), true availability (ta), and true outages (to). We then compute precision ($ta/(ta+fa)$), and recall ($ta/(ta+fo)$).

Table 2 compares these design choices. Here we focus on the effect of *number* of targets on precision and recall. While precision is uniformly good (inference of "up" is nearly always correct), recall suffers because there are many false outages. We conclude that probing one target (single and hitlist cases) has lower recall. The problem is that in some blocks the target address is never active, and in others with transient usage it is only sometimes active. Probing multiple addresses handles both of these cases.

Other systems use ICMP as a triggering mechanism for secondary methods that verify outages; for example, traceroutes may recover from a false trigger. However, these systems raise other questions (is the target for traceroute up?), and even when self-correcting, incur additional traffic. We show that *probing $E(b)$ provides 30–40% better recall than probing a single address, even without secondary verification.*

6.2 What Granularity of Blocks

Most previous active probing systems track reachability per routable prefix [17,20] (Hubble operation probes at most 1 target per BGP prefix [16]). However, reachability is *not* correlated with BGP prefixes [3]; we see smaller units.

We next compare *block-based* schemes that directly measure each /24 block with *prefix-based* schemes where measurement of a single representative address determines outages for a routable prefix of one or many blocks. Prefix-based schemes require little traffic and get broad coverage. However, their trade-off is that they are imprecise, because the single representative may not detect outages that occur in parts of the prefix that are not directly measured. Block-based schemes, on the other hand, require more traffic and cannot cover blocks where no addresses respond, so they have lower coverage. But because block-based schemes directly measure each block, they provide very good precision.

We first compare how precision and coverage trade-off with block-based and prefix-based measurement schemes, then how this difference in precision affects the accuracy of outage detection.

Methodology: We first must define when block-based or prefix-based methods can cover a prefix. Block-based measurement systems can track outages in blocks that have active addresses that respond. Here we require 20 active addresses with $A > 0.1$ (a slightly stricter requirement than Trinocular). Prefix-based systems expect an active .1 ad-

	in prefixes	in blocks
block-direct	184,996 (44%)	2,438,680 (24%)
prefix-direct	240,178 (57%)	219,294 (2%)
prefix-inferred	—	8,115,581 (81%)
overlap	152,295 (36%)	2,410,952 (24%)
neither	145,268 (35%)	1,908,122 (19%)
total	418,147(100%)	10,051,431(100%)

Table 3: Comparing coverage by granularity. Dataset: A_{20addr}.

dress (the target address). To be generous, we consider all .1 addresses in any /24 of a prefix, not just the first.

However, prefix-based systems only directly measure the target address, and from that infer outages for the rest of the prefix. Prefix-based systems require less probing traffic, but we have shown that Trinocular's probe rate is acceptable.

We evaluate the effects of coverage by re-analyzing an Internet-wide survey taken 2012-08-03 [27], labeled A_{20addr}. As with S_{50j}, this dataset consists of ICMP probes sent to addresses every 11 minutes. But it covers only 20 addresses in each of 2.5M /24 blocks, and only for 24 hours on 2012-08-03. We compare this probe data with default-free BGP routing tables from the same site on the same day.

Precision: We compare precision of coverage in Table 3. In the left column we consider, for each routable prefix, if any of its address blocks are covered by block-based measurements, prefix-based, both, or neither. We see 418k prefixes in the BGP routing table. Of these, prefix measurements directly observe 240k prefixes (prefix-direct, 57%), while block-based measurements include data for only 185k prefixes (44%). Block-based coverage misses some prefixes where all blocks include fewer than 20 addresses; prefix-based coverage misses some prefixes where no .1 addresses respond. Overall, prefix-based probing covers 13% more prefixes, although block-based picks up 8% that prefix-based misses (block-direct minus overlap).

The block-level view (right column) presents a different picture. Prefix-based has much larger coverage (81%) when one considers inferred blocks. This large coverage is due to large prefixes that are sparsely occupied, like MIT's 18/8, where most blocks do not respond but a prefix-based scheme allows 18.0.0.1 to represent reachability to them all. Block-based coverage is also lower because it requires more than one address per block. However, direct measurements in these prefixes are quite few: we observe 10 times more blocks than prefix-direct, but inference allows prefix-based to suggest answers for 3 times more blocks. We next consider how direct and indirect measurements affect accuracy.

Accuracy: We next compare the different granularities of prefix- and block-based measurements affect the accuracy of outages in A_{20addr}.

For prefix-based measurements, we observe the status of one address as representing the entire prefix. It therefore directly observes outages in the block of the representative address, and infers the status of other blocks of the prefix. This approach works perfectly when an outage is *prefix-complete*—all parts of the prefix go up or down together, perhaps due to a common change in routing. It is incorrect when the outage is *prefix-partial*, and can either over- or under-count the size of the outage when some blocks in the prefix are up while others are down.

	sites (% block-time)						
	(1 vantage point)			(2 vantage points)			(3)
status	w	c	j	wc	wj	cj	wcj
all down	0.79	0.92	0.74	0.24	0.22	0.26	0.15
all up	99.21	99.08	99.26	98.53	98.62	98.53	98.01
disagree	—	—	—	1.23	1.16	1.21	1.84

Table 4: Outages observed at three sites over two days. Dataset: A_{7}.

To compare accuracy we simulate a prefix-based scheme by observing outages in the prefix's target. We compare to re-analysis of A_{20addr} with a Trinocular-like scheme, following §7.2, but with all 20 addresses in each block as $E(b)$.

With prefix-based schemes, often a prefix will be declared down, but the data shows that other blocks in the prefix remain up. We find that 25% of all block-rounds inferred to be down are incorrect, so prefix inference often overstates outages. It can also understate outages when small outages do not occur at the direct measured block of the prefix; 37% of block-round outages seen by us are missed by a prefix-based scheme.

A fundamental limitation of prefix-based measurement is that *outages usually do not affect entire prefixes*. To quantify this claim, we examine each routable prefix with *any* block-level outages in A_{20addr}. For each prefix, we evaluate if the outage is prefix-complete or prefix-partial. Any prefix-based measurement scheme will *always* be incorrect for prefix-partial outages, over- or under-reporting depending on the status of the directly measured block. We find that only 22% of all prefix-rounds (that have any outage) are prefix-complete, while 78% are prefix-partial, showing that *most* outages are partial.

7. STUDYING THE INTERNET

We next examine what Trinocular says about the Internet.

7.1 Days in the Life of the Internet

We begin by evaluating Internet-wide outages to evaluate the proportion of local and global outages, and demonstrate Trinocular operation. We collected data tracking outages on 3.4M blocks over two days, starting at 2013-02-12 14:25 UTC, from three universities, labeled w, c, and j in Los Angeles, Colorado, and Japan. This experiment produces three datasets: A_{7w}, A_{7c}, and A_{7j}. For analysis, we then identify blocks where A is inconsistent and remove them, leaving 863k, 865k, and 863k blocks.

When rendered to an image with one pixel per round and block, the data is overwhelming (omitted here for space, but at [27]), forming an image of 270 pixels wide and 3.4M tall. The image confirms widespread diurnal "outages", confirming those we report on later (for example, in Figure 7).

Data from three vantage points lets us begin to evaluate how widespread are the outages we observe. Table 4 shows the level of agreement between the three sites for the period when all three had overlapping coverage. We measure agreement as percentage of block-time, that is, the sum of the duration of outages for each block for a given status, for each combination of the three vantage points.

Comparing columns w, c, and j, we see slight variations between the three sites (from 0.74% to 0.92%). We have seen this magnitude of variation in most measurements, with no strong trend favoring any site. We see a similar variation in these whole Internet measurements (here) as in measure-

ments of a *sample* of the Internet in Figure 10, suggesting that our samples are unbiased.

We can evaluate the degree of local and global outages by comparing each site with the other. The 2-vantage point columns (wc, etc.) show that many outages are local and seen at one site but not another. Overlap of all three vantage points (column wcj) shows only about 0.15% of the Internet is down, suggesting only 16–20% of outages seen by a single site are global. We believe we are converging on reporting only global outages with three independent vantage points, but future work should explore additional vantage points to demonstrate a plateau.

Here we considered two days of the Internet. We are currently running Trinocular actively, and in future work plan to compare long-term observations and compare to other public information about network outages.

7.2 Re-analyzing Internet Survey Data

While §7.1 uses Trinocular to study the global Internet, it provides only a brief snapshot. To get a longer-term perspective we next re-examine existing datasets using the principles behind Trinocular.

We draw on Internet survey data collected over the last three years from Los Angeles, Colorado, and Japan [13]. Surveys start with a random sample of 20k or 40k /24 blocks (about 1–2% of the responsive Internet), then probes all addresses in each block every 11 minutes for two weeks.

Survey data is quite different from Trinocular. All addresses in b are probed, not just $E(b)$, so the survey traffic rate is 100× greater than Trinocular. To adapt Trinocular to this bigger but less-tuned data, we track belief of the state of each block and use all probes as input. Since we probe all addresses, here $E(b) = b$ (Table 1). Since we cannot control probing, we have neither adaptive nor recovery probing, but periodic probing occurs every 2.6 s, slightly more frequent than Trinocular's adaptive probing. This change is both good and bad: frequent periodic probing can improve precision in detection of outage start and end, but many probes are sent uselessly to non-responsive addresses that Trinocular would avoid.

Our reanalysis computes $A(b)$ from the survey itself. This "perfect" value differs from Trinocular operation, where A is computed from possibly outdated IP history. Adapting A from probes is work-in-progress.

7.3 Case Studies of Internet Outages

We next examine several cases where Internet outages made global news. We see that systematic measurement of outages can provide information the scope of problems and the speed of recovery. Where possible, we visualize outages by clustering blocks by similarity in outage timing [26], and coloring blocks based on their geolocation.

7.3.1 Political Outages: Egypt and Libya

Two major 2011 outages were caused by political events: most Egyptian routes were withdrawn on 2011-01-27 by the government during what became the 2011 Egyptian revolution, and all Libyan routable prefixes were withdrawn 2011-02-18 during the Libyan revolution. In both cases, we re-examined surveys covering these events (S_{38c} began 2011-01-27, just after Egypt's outage, and ran for 3 weeks to cover Libya). We have strong evidence of the Egyptian outage, with 19 /24 blocks of Egypt's 22k in the survey (visualization omitted due to space). The end of the observed outage is confirmed with news reports and analysis of BGP data.

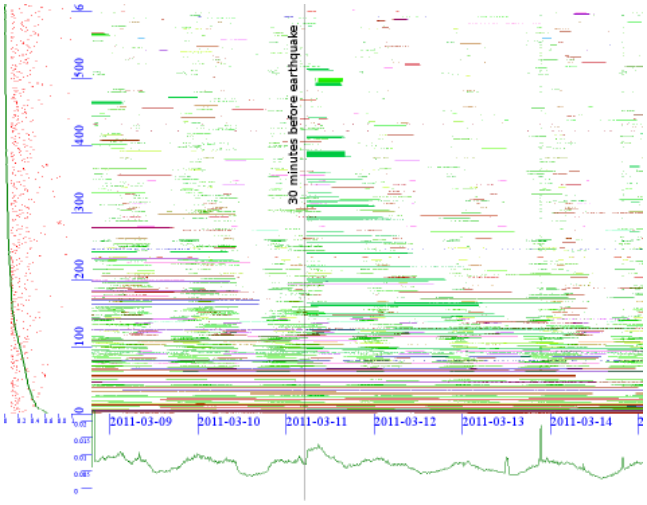

Figure 7: Six days of the 600 largest outages in March 2011 showing results of the Tōhoku earthquake. Dataset: S_{39c}. Colors are keyed to countries.

Libya's Internet footprint is much smaller than Egypt's: only 1168 /24 blocks as of March 2011. Only one of those blocks was in the dataset, and that block is too sparse (only 4 active addresses) to apply Trinocular. However, Trinocular's lightweight probing means that it could have covered the whole analyzable Internet. Had it been active at the time, we would have tracked 36% of Libya's 1168 blocks and likely seen this outage.

7.3.2 March 2011 Japanese Earthquake

In survey S_{39c}, we observe a Japanese Internet outage, in Figure 7 mid-day (UTC) on 2011-03-11. This event is confirmed as an undersea cable outage caused by the Tōhoku Japanese earthquake [22]. We mark a vertical line 30 minutes before the earthquake so as to not obscure transition times; individual blocks do not cluster well because recovery times vary, but the outage is visible as a large uptick in the marginal distribution. Unlike most human-caused outages, both the start and recovery from this outage vary in time. For most blocks, the outage begins at the exact time of the earthquake, as shown by the sudden large jump in marginal distribution less than 6 hours into 2011-03-11, but for some it occurs two hours later. Recovery for most blocks occurs within ten hours, but a few remain down for several days.

This dataset also shows strong evidence of diurnal outages in Asia as the green and white banding seen in the low 300 blocks. These diurnal outages make Trinocular's outage rate slightly higher than our previous approach [26]. We show that these blocks come and go, meeting our definition of outage. Future work may distinguish between cases where networks intentionally go down (such as turning of a laboratory at night) from unexpected outages.

7.3.3 October 2012: Hurricane Sandy

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 UTC. When we focus on known U.S. networks, we see about triple the number of

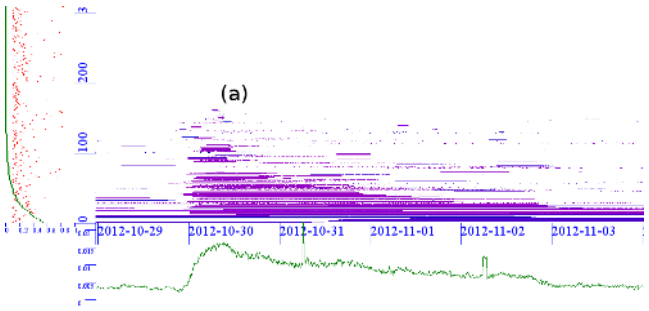

Figure 8: Six days of the 300 largest outages in U.S.-based networks showing Hurricane Sandy. Dataset: S_{50j}.

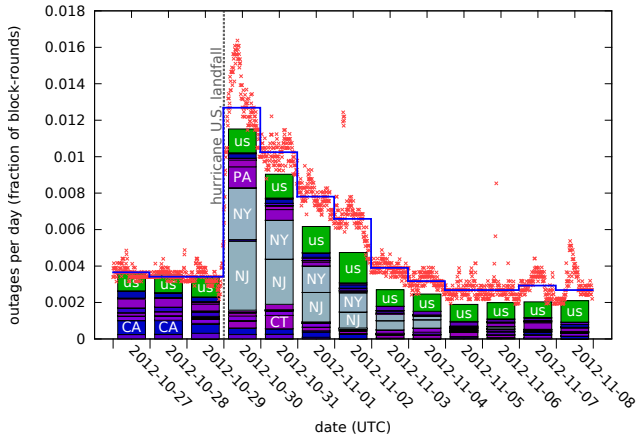

Figure 9: Median number of outages per day, broken down by state, weighted by outage size and duration, with jittered individual readings (dots). Dataset: S_{50j}.

network outages for the day following landfall, and above-baseline outages for the four days following landfall.

Visualizing outages: Figure 8 visualizes the 400 blocks in the U.S. with the largest degree of outages, and label (a) shows a strong cluster of outages at 2013-10-30 (UTC) corresponding with hurricane landfall. Hurricane-related outages tend to be long, lasting one or more days. We believe these outages correspond to residential power outages.

Quantifying outages: We know that *some* part of the Internet is always down, so to place these outages in perspective, Figure 9 plots the exact number of /24 blocks that are down in each round (this value is the marginal distribution of Figure 8). We plot each round as small red points (with small jitter to make consecutive more distinct), and we show 24-hour median values with the dark line.

Figure 9 shows U.S. networks had an outage rate of about 0.36% before landfall. (This rate seems somewhat less than the global average.) This rate jumps to 1.27%, about triple the prior U.S. baseline, for the 24-hours following landfall. The outage level drops over the next four days, and finally returning to the baseline on 2012-11-03.

Locating outages: To confirm the correlation between the hurricane and these outages, we look at the weighted blocks by state. The bars in Figure 9 identify outages by state. 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

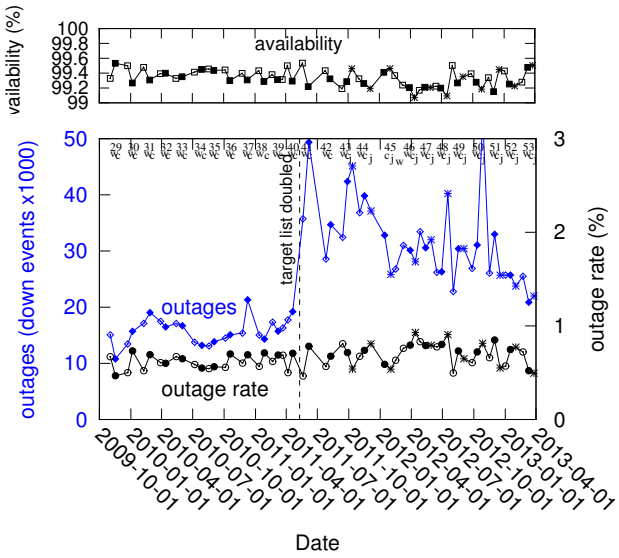

Figure 10: Evaluation of single-site outages in 2-week surveys over three years. Top shows availability, bottom shows Internet events, outages and outage percentage over time. (Dataset varies by time, as shown in the figure.)

colored bars in the middle of the graph) after hurricane landfall on 2012-10-30, about three times the prior baseline. These problems are generally resolved over the following four days. (Because of our more sensitive methodology, we see more outages here than in our prior analysis [14], but our qualitative results are similar.)

While re-analysis of S_{50j} provides insight into Sandy-related problems and recover, survey collection places significant traffic on the targets. Trinocular can cover 3.4M blocks, about 80× more than the 40k in a survey, at about 1% the traffic to each target block.

7.4 Longitudinal Re-analysis of Existing Data

Finally, we re-analyze three years of surveys. This data lets us compare the stability of our results over time and across different locations.

Probing location can affect evaluation results. Should the probing site's first hop ISP be unreliable, we would underestimate overall network reliability. We re-analyze surveys collected from three sites (see §7.2), each with several upstream networks. In Figure 10, locations generally alternate, and each location is plotted with a different symbol (W: empty symbols, C: filled, J: asterisks), and survey number and location letter are shown at the graph top. Visually, this graph suggests the results are similar regardless of probing site and for many different random samples of targets. Numerically, variation is low: mean outage rate (area) is 0.64% with standard deviation of only 0.1%. To strengthen this comparison we carried out Student's t-test to evaluate the hypothesis that our estimates of events, outages, and outage rates for our sites are equal. The test was unable to reject the hypothesis at 95% confidence, suggesting the sites are statistically similar.

Besides location, Figure 10 suggests fairly stable results over time. We see more variation after 2011, when the size of the target list doubled to about 40k blocks.

These observations are each from a single vantage point, thus they include both global and local outages. Surveys are

taken for non-overlapping, two week periods because each places a significant burden on the subject networks. Trinocular's much lower traffic rate to targeted blocks (1% that of a survey) allows outage detection to overcome both of these limitations. As demonstrated in §7.1, it can operate concurrently from three sites. We plan to carry out continuous monitoring as Trinocular matures.

8. CONCLUSIONS

Trinocular is a significant advance in the ability to observe outages in the network edge. Our approach is principled, using a simple, outage-centric model of the Internet, populated from long-term observations, that learns the current status of the Internet with probes driven by Bayesian inference. We have shown that it is parsimonious, with each instance increasing the burden on target networks by less than 0.7%. It is also predictable and precise, detecting all outages lasting at least 11 minutes with durations within 330 s. It has been used to study 3.4M blocks for two days, and to re-analyze three years of existing data, providing a new approach and understanding of Internet reliability.

Data Availability and Acknowledgments: The raw and analyzed data from this paper are available at no cost to researchers through the U.S. DHS PREDICT program (`www.predict.org`) and by request from the authors [27]. This work was classified by USC's IRB as non-human subjects research (IIR00000975).

We thank our shepherd, Olaf Maennel, and the anonymous reviewers for comments that made this paper stronger and more readable. We thank John Wroclawski for comments that helped clarify the role of the model, Ethan Katz-Bassett and Harsha Madhyastha for discussion about §6.2, and Ítalo Cunha for a careful reading. 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 BGPmon and at CSU. We also thank Katsuhiro Horiba (WIDE) for providing probing infrastructure and BGP feeds.

9. REFERENCES

[1] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris. Resilient overlay networks. In *Proc. of Symposium on Operating Systems Principles*, pages 131–145, Chateau Lake Louise, Alberta, Canada, Oct. 2001. ACM.

[2] R. Bush, J. Hiebert, O. Maennel, M. Roughan, and S. Uhlig. Testing the reachability of (new) address space. In *Proc. of ACM Workshop on Internet Nework Management*, Aug. 2007.

[3] R. Bush, O. Maennel, M. Roughan, and S. Uhlig. Internet optometry: assessing the broken glasses in Internet reachability. In *Proc. of ACM IMC*, 2009.

[4] D. R. Choffnes, F. E. Bustamante, and Z. Ge. Crowdsourcing service-level network event monitoring. In *SIGCOMM*, 2010.

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

[6] I. Cunha, R. Teixeira, N. Feamster, and C. Diot. Measurement methods for fast and accurate blackhole identification with binary tomography. In *Proc. of 9thACM IMC*, 2009.

[7] A. Dainotti, R. Amman, E. Aben, and K. Claffy. Extracting benefit from harm: using malware pollution to analyze the impact of political and geophysical events on the Internet. *ACM Computer Communication Review*, Jan 2012.

[8] A. Dainotti, C. Squarcella, E. Aben, K. C. Claffy, M. Chiesa, M. Russo, and A. Pescapé. Analysis of country-wide internet outages caused by censorship. In *ACM IMC*, 2011.

[9] A. Dhamdhere, R. Teixeira, C. Dovrolis, and C. Diot. NetDiagnoser: troubleshooting network unreachabilities using end-to-end probes and routing data. In *Proc. of ACM Conference on Emerging Networking Experiments and Technologies*, pages 18:1–18:12. ACM, 2007.

[10] X. Fan and J. Heidemann. Selecting Representative IP Addresses for Internet Topology Studies. In *ACM IMC*, 2010.

[11] X. Fan, J. Heidemann, and R. Govindan. LANDER IP history datasets. `http://www.isi.edu/ant/traces/ipv4_history`, 2011.

[12] N. Feamster, D. G. Andersen, H. Balakrishnan, and F. Kaashoek. Measuring the Effects of Internet Path Faults on Reactive Routing. In *ACM Sigmetrics - Performance*, 2003.

[13] 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.

[14] J. Heidemann, L. Quan, and Y. Pradkin. A preliminary analysis of network outages during Hurricane Sandy. Technical Report ISI-TR-2008-685, USC/Information Sciences Institute, November 2012.

[15] Y. Huang, N. Feamster, A. Lakhina, and J. J. Xu. Diagnosing network disruptions with network-wide analysis. In *Proc. of ACM SIGMETRICS*, pages 61–72, San Diego, California, USA, June 2007. ACM.

[16] E. Katz-Bassett. Private communications, May 2012.

[17] 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 *Proc. of 5th NSDI*, pages 247–262. USENIX, Apr. 2008.

[18] E. Katz-Bassett, C. Scott, D. R. Choffnes, Í. Cunha, V. Valancius, N. Feamster, H. V. Madhyastha, T. Anderson, and A. Krishnamurthy. LIFEGUARD: Practical repair of persistent route failures. In *Proc. of SIGCOMM*, pages 395–406, Helsinki, Finland, Aug. 2012. ACM.

[19] R. R. Kompella, J. Yates, A. Greenberg, and A. C. Snoeren. Detection and Localization of Network Black Holes. In *Proc. of IEEE Infocom*, 2007.

[20] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: An information plane for distributed services. In *Proc. of 7thOSDI*, pages 367–380, Seattle, WA, USA, Nov. 2006. USENIX.

[21] R. Mahajan, D. Wetherall, and T. Anderson. Understanding BGP misconfiguration. In *Proc. of SIGCOMM*, 2002.

[22] 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.

[23] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C. nee Chuah, and C. Diot. Characterization of Failures in an IP Backbone. In *Proc. of IEEE Infocom*, 2004.

[24] V. Paxson. End-to-end routing behavior in the Internet. In *Proc. of SIGCOMM '96*, pages 25–38, Stanford, CA, Aug. 1996. ACM.

[25] L. Quan, J. Heidemann, and Y. Pradkin. Detecting internet outages with precise active probing (extended). Technical Report ISI-TR-2012-678, USC/ISI, February 2012.

[26] L. Quan, J. Heidemann, and Y. Pradkin. Visualizing sparse internet events: Network outages and route changes. In *Proc. of First ACM Workshop on Internet Visualization*, Boston, Mass., USA, Nov. 2012. Springer.

[27] L. Quan, J. Heidemann, and Y. Pradkin. LANDER Internet outage datasets. `http://www.isi.edu/ant/traces/internet_outages`, 2013.

[28] L. Quan, J. Heidemann, and Y. Pradkin. Poster abstract: Towards active measurements of edge network outages. In *Proc. of Passive and Active Measurement Workshop*, pages 276–279, Hong Kong, China, Mar. 2013. Springer.

[29] A. Schulman and N. Spring. Pingin' in the rain. In *Proc. of ACM IMC*, pages 19–25, Berlin, Germany, Nov. 2011. ACM.

[30] R. Teixeira and J. Rexford. A measurement framework for pin-pointing routing changes. In *Proc. of the ACM SIGCOMM workshop on Network troubleshooting*, 2004.

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

[32] D. Turner, K. Levchenko, A. C. Snoeren, and S. Savage. California fault lines: understanding the causes and impact of network failures. In *Proc. of SIGCOMM*, 2010.

[33] Wikipedia. Hurricane Sandy. `http://en.wikipedia.org/wiki/Hurricane_sandy`, 2012. Retrieved 2012-11-24.

[34] E. Wustrow, M. Karir, M. Bailey, F. Jahanian, and G. Huston. Internet background radiation revisited. In *ACM IMC*, 2010.

[35] M. Zhang, C. Zhang, V. Pai, L. Peterson, and R. Wang. PlanetSeer: Internet Path Failure Monitoring and Characterization in Wide-area Services. In *OSDI*, 2004.