

Selecting Representative IP Addresses for Internet Topology Studies*

Xun Fan
USC/Information Sciences Institute
4676 Admiralty Way, Suite 1001
Marina del Ray, CA, 90292
xunfan@isi.edu

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

ABSTRACT

An *Internet hitlist* is a set of addresses that cover and can *represent* the the Internet as a whole. Hitlists have long been used in studies of Internet topology, reachability, and performance, serving as the destinations of traceroute or performance probes. Most early topology studies used manually generated lists of prominent addresses, but evolution and growth of the Internet make human maintenance untenable. Random selection scales to today’s address space, but most random addresses fail to respond. In this paper we present what we believe is the first automatic generation of hitlists informed censuses of Internet addresses. We formalize the desirable characteristics of a hitlist: *responsiveness*, each representative responds to pings; *completeness*, they cover all the allocated IPv4 address space; and *stability*, list evolution is minimized when possible. We quantify the accuracy of our automatic hitlists, showing that only one-third of the Internet allows informed selection of representatives. Of informed representatives, 50–60% are likely to respond three months later, and we show that causes for non-responses are likely due to dynamic addressing (so no stable representative exists) or firewalls. In spite of these limitations, we show that the use of informed hitlists can add 1.7 million edge links (a 5% growth) to traceroute-based Internet topology studies Our hitlists are available free-of-charge and are in use by several other research projects.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Network topology*; C.2.5 [Computer-Communication Networks]: Local and Wide-Area Networks—*Internet*; C.2.6 [Computer-Communication Networks]: Internetworking

*This work is partially supported by US DHS contract number NBCHC080035, and John Heidemann by 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.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IMC '10, November 1–3, 2010, Melbourne, Australia.
Copyright 2010 ACM 978-1-4503-0057-5/10/11 ...\$10.00.

General Terms

Measurement

Keywords

Internet topology, Topology Representatives, IP Hitlist

1. INTRODUCTION

Smooth operation of the Internet is important to the global economy, so it is essential that Internet users, providers, and policy makers understand its performance and robustness. Although on the surface, individuals care only about their personal performance, a full diagnosis of “why is my web connection slow?” must consider not just the user’s “first mile” connection, but dozens of servers that affect performance [9]. Web content providers invest great effort in optimizing page load times to sub-second values [23] and in building distributed content distribution networks that manage traffic (for example, [14]). Policy makers debate questions about universal access [33], a nation’s relative availability for broadband access [24], and the robustness of what is recognized as critical infrastructure.

To answer these questions, network researchers, operations, and industry have developed a number of tools to map the Internet [17, 15, 26, 20, 22, 8], evaluate performance [32, 20, 23], consider questions about routing and reachability [31, 3], or the performance of replica placement (examples include [32, 10]), and evaluate topology robustness [1]. With the Internet’s lack of centralization and multiple overlapping global “backbones”, active probing plays an essential role in this process, with traceroute and ping and their variants providing the main source of router-level reachability. While one may add AS-level views [4], the Internet’s router-level topology is the focus of this paper. Different router-level studies either target specific networks [26] or the whole Internet. Here we are most interested in observing the whole Internet—more than three billion allocated IPv4 addresses.

Studies of the entire Internet typically employ a *hitlist*—a list of IP address that can *represent* the billions of allocated addresses. The defining characteristic of a hitlist is *completeness*, where a representative is chosen for every autonomous system or, in our case, for every allocated *block* of addresses defined by a /24 prefix, the smallest unit typically present in a default-free routing table. Representatives provide a 256-fold (or more) reduction in scanning size, allowing Internet-wide studies to take place in hours instead of months.

Although completeness is necessary to study the whole Internet, an ideally hitlist is also responsive and stable. A *responsive* representative replies to ICMP messages, allowing traceroute to confirm a path to the edge of the network, and ping to measure round-trip time to an edge host. To support longitudinal studies, the hitlist should be *stable*, with representative identities not changing frequently or arbitrarily.

Although hitlists are easy to define and have been used in topology studies for many years (we review related work in Section 2), they are surprisingly hard to create and maintain. Early hitlists were built manually from well known sites [20], but the size of the Internet and rate of churn in even well-known servers made manual maintenance untenable as it quickly became incomplete. More recent studies have typically used randomly chosen representatives. While randomness has some advantages (it can be statistically unbiased), it sacrifices secondary goals of stability and responsiveness.

The contribution of this paper is to provide a new, automated method of hitlist generation that provides complete coverage while maximizing stability and responsiveness.¹ Our hitlists are constructed (Section 3) by mining data from IP address censuses, complete, ping-based enumerations of the allocated IPv4 address space taken every two to three months [18].

The second contribution of our work is to evaluate our hitlists (Section 4). Our hitlists are 100% complete as of when they are constructed, although when we have no history (in about two-thirds of the blocks) we select representatives at random. We define the *accuracy* of our hitlists as how many representatives are responsive three months after the hitlist is taken. We find that two-thirds of the allocated address space never responds to ICMP probes and so never has responsive representatives. Of the remaining, responsive Internet, our hitlists select representatives that are responsive about 55% of the time. To our knowledge we are the first to study hitlist responsiveness and accuracy.

The final contribution of our work is what hitlists reveal about that nature of the Internet itself. We were surprised that, in spite of such complete input data, the responsiveness of our predicted representatives is not higher. We believe this upper bound on responsiveness characterizes the portion of the Internet that has an inherently high rate of address churn. One corollary of this limit to representative responsiveness is that no manual system could ever have been successful due to natural turnover of addresses in parts of the network. We also characterize the distribution of addresses in each block and show that it strongly reflects address allocation patterns (Section 5).

We make our hitlists available free-of-charge, and they are already being used by several research projects. In Section 6 we discuss the security and policy issues involved in sharing this data.

2. RELATED WORK

Hitlists are used in active probing for studies of topology [17, 15, 19, 26, 30, 22, 25, 8], performance [32, 20, 23], and reachability [31, 4]. and for other purposes [32, 1, 10]. Each of these studies uses some hitlist (sometimes

¹We would like to thank Randy Bush for suggesting the idea that our address censuses data could support hitlist generation.

called a seed or probe list) generated manually, randomly, or automated from several sources. We review each hitlist generation method next.

Early topology work used manually generated lists. Skitter is a well-known measurement tool developed at CAIDA, to study the IPv4-interface-level of the core Internet topology [19]. (With IP alias resolution, this can provide a router-level map.) It uses traceroutes from multiple locations to a hitlist of destinations. Their target address list was manually built from many sources, including tcpdump from the UCSD-CERF link, hostnames from search engines, and intermediate addresses seen from their own traces records. In 2000 their hitlist included about 313,000 destinations, and by 2004 it had grown to 971,080. While their hitlist was of high quality, they found it very labor-intensive to maintain, and responsiveness degraded over time as destinations changed. (They report that their initial, web-server based list shows a 2-3% reduction each month in reachable destinations [19].) The cost of manual list maintenance was one reason that prompted them to change to random probing with Archipelago. More recently, Maennel has maintained a manual list, derived from the Skitter list, but augmented with guided scanning to cover each AS and provide 306,708 representatives. They require reachable addresses study routing reachability [4]. We used a version of their list to seed our initial stable list, but our techniques provide much greater coverage at lower cost. Unlike all of these manual hitlists, our goal is to fully automate hitlist generation to allowing more complete and timely coverage.

Random representative selection allows low-cost generation of hitlists to much larger numbers of networks. Mercator developed *informed random probing* to adaptively adjust its probe list based on prior results [17]. By adaptively growing the hitlist, Mercator strives to quickly and efficiently discover a topology while minimizing hitlist size. Archipelago (Ark) is a measurement platform designed to support traceroute and other measurements [8], effectively a next-generation Skitter. Ark's hitlist covers all routed /24 blocks, choosing a random last-octet within each /24 block. The random hitlists in Mercator and Ark are essential to cover the millions of /24 networks in today's Internet, but Mercator's adaptive algorithm means completeness is uncertain (although efficiency, not completeness, was their goal), and random probing in both Mercator and Ark may sacrifice responsiveness of the destination address. Our hitlist provides complete coverage while maximizing responsiveness (both defined precisely in Section 3.1). In Section 4.5 we evaluate the degree to which informed hitlist generation may improve the number of links discovered in a topology study.

Rather than a random destination, DisCarte's hitlist selects the .1 address in each /24 block of the routed address space. DisCarte [25] adds record route information to traceroute probing to obtain more accurate and complete network topologies. DisCarte requires a responsive destination, and find 376,408 responsive representatives in the .1 address of each routed /24. Our work confirms that the .1 address is responsive twice as often as the address with median responsiveness (Section 5), but we suggest that census-informed representative selection can get much better responsiveness.

Finally, there has been some work in IPv6 topology discovery. The Atlas system uses a manually generated list built from 6bone destinations [30], then expanded it based on discoveries. The full address-space census that is the ba-

sis of our work is only applicable to the IPv4 address space, so combinations of active and passive methods as proposed in Atlas are essential for IPv6. As future work, comparison of both active and passive hitlist generation approaches in the IPv4 space may provide a basis for inferring coverage of passive studies of IPv6.

3. METHODOLOGY

We next describe the requirements of an IP hitlist (Section 3.1), and how we transform census data (reviewed in Section 3.2) using several possible prediction methods (Section 3.3) to get a good quality hitlist. We also provide some details on how our implementation copes with Internet-sized datasets.

3.1 Hitlist Requirements

Our goal is to provide representatives that are responsive, complete and stable.

By *responsive*, we mean each representative is likely to respond to an echo request with an echo reply instead of an ICMP error code. As we describe below, we select representatives that have responded frequently in the past. We do not guarantee that that address responded in the most recent census, but we bias our selection to favor recent results. We consider several prediction functions below in Section 3.3.

By *complete*, we mean we report one representative address for every allocated /24 block. Some groups have used other definitions of completeness, such as one representative per AS, or per routed BGP block. AS- or BGP-complete hitlists will be both sparser and smaller than /24-complete maps, since ASes typically include routes for many prefixes, and routed prefixes often cover blocks larger than /24s. Instead, we select one representative for each /24 block for two reasons. The main reason is so that the hitlist is decoupled from the routing system, to allow independent study of the routing system itself, and since routing tables vary depending on when and where they are taken. Second, it is relatively easy for researchers to derive custom hitlist from ours, perhaps selecting sparser representatives per AS or per routed prefix. We believe our per-/24 representatives thus provide a more general and reusable result than more “cooked” alternatives.

By *stable*, we mean that representatives do not change arbitrarily. We change representatives when a new representative would significantly improve the score for that block, typically because a representative has ceased to be reachable. We promote stability as a goal to simplify longitudinal studies, where frequent changes of representative would make comparisons across time more difficult. Stability also reduces the effects of transient routing outages or packet loss on the long-term hitlist. We implement stability by applying *inertia* as a threshold to changing a previously selected representative. Currently we switch representatives when the switch will improve its score significantly, with inertia of 0.34 (see Section 4.3).

The three goals of responsiveness, completeness, and stability can be in conflict. For example, completeness requires that we select representatives that may be non-responsive. To guarantee representatives for all allocated addresses, we select representatives even for blocks that have no recent responses. We also select representative for blocks that have never responded. In both cases, we annotate these repre-

Censuses	Date	Duration (days)
it11w	2006-03-07	23
it12w	2006-04-13	24
it13w	2006-06-16	31
it14w	2006-09-14	31
it15w	2006-11-08	61
it16w	2007-02-14	50
it17w	2007-05-29	52
it18w	2007-09-14	47
it19w	2007-12-18	48
it20w	2008-02-29	86
it21w	2008-06-17	49
it22w	2008-09-11	35
it23w	2008-11-25	29
it24w	2009-02-03	29
it25w	2009-03-19	29
it26w	2009-05-27	31
it27w	2009-07-27	25
it28w	2009-09-14	30
it29w	2009-11-02	30
it30w	2009-12-23	29
it31w	2010-02-08	30
it32w	2010-03-29	29

Table 1: IPv4 censuses [29] used in this paper.

sentatives with distinguished scores. Stability and inertia can also decrease responsiveness; in Section 4.3 we examine how the inertia threshold reduces churn and how it affects responsiveness.

3.2 Background: Internet censuses

Our main goal with a hitlist is to predict the future: a representative should be responsive in the future. Our tool to make this prediction is data from past responses. Hitlists selection leverages Internet censuses that have been taken regularly since 2003 [18].

Each Internet census is the results of a ping (an ICMP ECHO REQUEST message) sent to every allocated IPv4 address. Censuses are far from perfect: a census must be taken carefully to avoid ICMP rate limiting or transient router errors, and firewalls reduce ping response rates by around 40%. Hitlist, however, prefer hosts that are ICMP responsive, since traceroute consists of iterated, TTL-limited ICMP messages. Firewall-limited censuses are therefore ideal for hitlist generation.

It takes 2–3 months to carry out a full census (the IPv4 space has more than 3 billion allocated unicast addresses, with 4 machines probing at about 6000 probes/s [18]). For this paper we consider censuses starting in Mar 2006 as shown in Table 1, since censuses before this date used a slightly different collection methodology. The results of this paper use all 22 censuses taken over the four years preceding analysis, but we expect to update our results as new censuses become available.

A census elicits a number of responses, including ECHO REPLY messages as well as a variety of errors. Each census is quite large, and more than 3 billion records per census, 22 censuses is over 260GB of raw data. We therefore preprocess all censuses into a *history map* convenient for analysis. A history map consists of a bitstring for each IP address

where each 1 indicates a positive response, and a 0 indicates either a non-response or negative response.

In paper, we only consider echo replies (“positive” responses) as indicating a responsive address. We also explored treating both positive and negative responses (destination unreachable and similar error replies) as predicting a responsive address. However, we found that negative replies only rarely are helpful in predicting future responsiveness. We look at it28w and found 64% of positive responses were also responsive in the next census (165M in it28w, 107M of which respond in it29w). By contrast, of the 50M negative replies in it28w, only 2.7% (1.4M) respond positively in it29w. We therefore believe that negative responses are of little value in predicting future responsiveness.

We next show how this history map can predict future response rates.

3.3 Prediction Method

Of our hitlist goals of responsiveness, completeness, and stability, completeness and stability are under our control, but responsiveness requires predicting the future. Our guidance in this task is the prior history of each address. We next review several *prediction functions* that strive to select the best representative for each /24 block, where best is most likely to respond in the future.

Prediction functions take the prior history of address a as input and weights that history in different ways. Each bit of the history is presented by $r_i(a)$, the response (1 for positive, otherwise 0) of address a to the i th probe, numbered from 0 (oldest) to $N_h - 1$, the most recent observation. We consider several different weights $w(i)$ to get the scores $s(a)$ in the form:

$$s(a) = \sum_{i=0}^{N_h-1} r_i(a)w(i)$$

For each block of addresses, the address with the highest $s(a)$ is selected as the best representative. We may bias this by prior representatives to promote stability. In the case of ties and no prior representative we select any top scoring address in the block at random.

We considered several possible weights $w(i)$. The simplest is $w(i) = 1$, so all responses are *averaged*. To give more recent observations greater influence we consider two biased weights. With *linear* weighting, $w(i) = (i + 1) * 1/N_h$, and for a *power* function, $w(i) = \frac{1}{N_h - i}$. Weighting of each observation for an 8-observation history is shown in Figure 1.

In addition, we can normalize scores by to the maximum possible score (the minimum in all cases is zero), allowing all to fall in the range 0 to 1.

As an example of the different functions, Figure 2 shows scores for three different weights and different history lengths. For simplicity, we assume $N_h = 8$, shorter than we use in practice (in Section 4.1.2 we vary history duration). We consider three cases, all with 4 of 8 responding, but either responding most recently (Figure 2a), in the middle past (Figure 2b), or alternating response and non-response (Figure 2c). To a first approximation, all three weights are about the same, particularly with intermittent responsiveness in Figure 2c. The differences in decay rates are more obvious when responsiveness is consistent for blocks of time, with power and linear decay faster than average in Figures 2a and 2b. Finally, difference in history duration make a large difference when a block is non-responsive, comparing the left

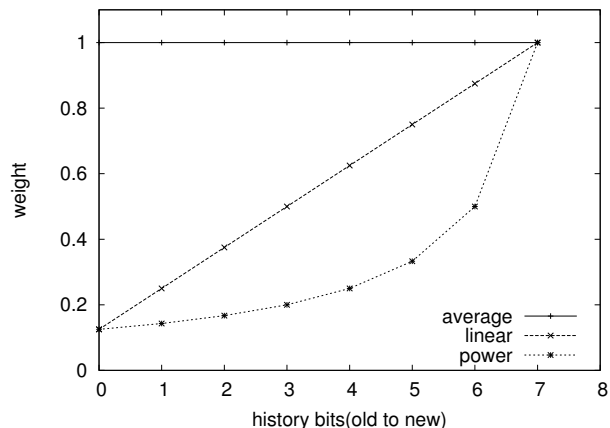


Figure 1: bits weight for different function

and right parts of Figures 2a and 2b, and these effects are even greater when comparing across weights (for example, compare history durations 1–4 of Figures 2a and 2b).

This framework provides flexibility, but requires setting several parameters. We later evaluate which weighting is best (Section 4.1.1), how much history is beneficial (Section 4.1.2), and the underlying reasons addresses are difficult to predict (Section 4.1.3).

3.4 Gone-Dark Blocks

Firewalls are probably the greatest impediment to active probing, since a conservative firewall can suddenly stop traffic an entire block. We will see in Section 4.1.3 that *gone dark* blocks are one cause of poor representative responsiveness. A gone-dark blocks is one that contained responsive addresses for some period of time, but then becomes unresponsive and stays that way, due to firewall or possibly renumbering. While we must select a representative for each allocated block, even if populated only by non-responsive addresses, we would like to indicate our low expectations for gone-dark blocks.

We define a block as *gone dark within history* N_d if, for the most recent N_d observations, no address in the block responded, even though we had some positive response before N_d observations.

We add gone-dark analysis to our hitlist generation by overriding the representative’s score with a designated “gone-dark” value to indicate our skepticism that it will reply. We explored different values of N_d and ultimately select $N_d = N_h = 16$, identifying only those addresses whose responses have aged-out of our history as gone-dark. We use this large value of N_d because this value maximizes the absolute number of responsive representatives, while only decreasing the percentage of responsive, predicted representatives a small amount.

For gone-dark blocks, we still select the representative as the address with the best score. For allocated but never-responsive blocks, we select the .1 address as the representative because that is most likely to be first used (Section 5). In Section 4.1.3 we show the contribution of gone-dark blocks to responsiveness.

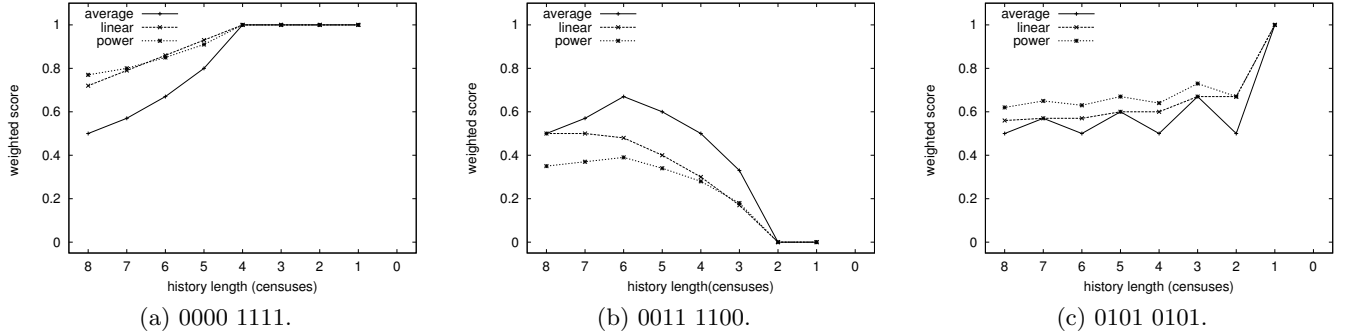


Figure 2: Comparison of three history functions for selected addresses.

3.5 Hitlist Description

To summarize, our hitlist contain three kind of representatives for all allocated /24 blocks: informed and predicted representatives, where we select the best responder; gone-dark representatives, where some address once responded but has not recently; and allocated but never-responsive blocks, where we pick .1 as the representative.

Table 6 lists the hitlists we have publicly released to-date. We identify hitlists by the name of the last census used in their creation, and include the number of censuses in the history. Thus HL28/16 uses 16 hitlists through it28w. When necessary, we add the gone-dark window, so HL28/16-3 uses a window of 3. If no gone-dark window is specified, we disable gone-dark processing.

In addition to these public hitlists, Tables 2 and 3 show unreleased hitlists used to evaluate our methods.

4. EVALUATION

We next evaluate the success of our hitlist: how accurate are its predictions and how complete and stable is it? We first consider how responsiveness is affected by choices in our prediction mechanism. In Section 4.1.3 then look at causes of prediction failure. Finally, we consider completeness and stability.

4.1 Responsiveness

Our primary goal with prediction is *responsiveness*: how accurate is our prediction that the representatives in a hitlist will respond in the future? We can define based on the number responding in the future, N_r , from the number of predicted representatives (including representatives of gone-dark and informed predicted blocks) N_p as:

$$\alpha = \frac{N_r}{N_p}$$

Responsiveness accuracy is affected by our choice of history weighting and length. We consider these next, and then consider structural reasons perfect accuracy is impossible to achieve.

Our general approach to test responsiveness is to generate a hitlist, then evaluate it against ICMP probes in the next census. For example, the first line of Table 2 evaluates HL19/8, generated from the eight censuses from it12w through it19w, tested against it20w. This approach has the advantage of supporting retroactive evaluation of hitlist

hitlist	weighting function		
	average	linear	power
HL19/8	0.50	0.51	0.51
HL21/8	0.53	0.54	0.55
HL23/8	0.53	0.54	0.54
HL25/8	0.53	0.54	0.54
HL27/8	0.54	0.55	0.55

Table 2: Fraction of responsive representatives across 5 different hitlists for three different history weights.

quality under different, controlled conditions. However, it also means each representative is only given one opportunity to be available. For this reason we report exact counts of results, without error estimates such as standard deviation. We evaluate repeatability of our results by considering multiple hitlists at different times.

4.1.1 Comparing History Weights

We first consider how our weighting of prior history affects accuracy. Here we assume a history duration of 8 prior censuses (a reasonable choice as evaluated next in Section 4.1.2), and from that history we predict the results of the next census for the three weights we defined in Section 3.3. Since the network is dynamic, our expectation is that biased weightings will perform best since they favor recent information over older information.

To answer this question, Table 2 compares our three weightings for several predictions. Each line evaluates a different hitlist as generated with three different weights, and evaluated for all predicted representatives (N_p). The most important observation is that *all weights* provide quite similar performance—the worst case responsiveness is only 5% worse than the best. Linear and power functions provide marginally better responsiveness. The examples of the weights in Figure 2 suggests on reason the difference is so small. For many histories, all three weights produce roughly the same relative scores.

4.1.2 Effects of History Duration

A second factor that can affect responsiveness is the duration of history considered in a prediction. Does more history provide more information, or does very old information become irrelevant or even misleading?

To study this question, we considered all history available to us at time of analysis—then we had 18 Internet censuses covering 3.5 years. We consider only the power weighting of history, and look at the responsiveness of our predictions.

Table 3 shows responsiveness of our predictions as a function of history length, for five predictions. We see that very short histories are insufficient: prediction rates are a 1–2% lower when fewer than 8 (about 1.5 years) observations are considered. On the other hand, we see no difference in prediction accuracy for histories from 8 to 16 censuses. (We also looked at history duration with the average function, and found there that long histories became slightly less accurate, although only by 1–2%. This observation argues in favor of a weighting that decays by history, like power.)

Finally, while longer histories may not improve the fraction that respond, it does provide information that allows *more* representatives to be selected. Table 3 shows the absolute number of responders as a function of history duration. Longer history allows 20k more responders with length 16 than with length 8. More history always increases the number responding, although with diminishing returns past a 12 censuses or so.

In practice, the incremental cost of longer history lengths is not large. So we use a history length of 16 censuses in our production lists.

Although 8 censuses provides slightly better results, the fraction responding, only 55%, seems lower than we might expect. We therefore next consider causes of non-responsiveness.

4.1.3 Causes of Failed Responses

We found the observation that our best methods get only 55% responsiveness seems somewhat surprising. Surely such a large amount of history (over three years of full censuses) can be explored somehow to select representatives with greater accuracy. To answer that question, we next explore the causes of why representatives fail to respond. Our conclusion is that *it is unlikely that any prediction can do better than about 70%* because of the use of dynamic address assignment and firewalls.

To support this claim, Table 4 counts prediction failures for HL28/16, tested against it29w (We found roughly similar results in examination of HL31/16 evaluated against it32w.) We see that 44% of representatives are non-responsive (1.8M of the 4M blocks). Two explanations account for the majority of our misses: blocks that use only dynamic address assignment, and “gone-dark” blocks. We consider each of these below.

While dynamic addressing and firewalls are target-specific causes of representative non-responsiveness, measurement error is a possible source of uncertainty. We believe that Internet census-taking methodology reduces these sources of error to random noise for reasons described in prior work [18]. To summarize briefly: we monitor the network hosting the probes for local routing outages. Probes are in pseudorandom order, so routing outages in the middle or near the destination result in lower responsiveness in proportion to outage rates, but randomly distributed. Pseudorandom probing is spread over two months, so the probe rate to any individual /24 is well below typical ICMP rate limits. We considered packet loss and routing outages in the middle or of the network or near probe sources are potential sources of error. For more complete discussion of sources of error

in Internet census-taking, and validation studies, we refer to prior work [18].

Defining stable blocks: Blocks that lack stable addresses makes representative selection inherently difficult. In a block with a stable representative, it will likely remain responsive, but if all addresses in the block are unstable then the probability a representative will respond is equal to the occupancy of that block and independent of prior history. Addresses can lack stability either because the hosts using the addresses are only on intermittently, or because addresses in the block are allocated dynamically to a changing population of computers. Multiple groups have used different techniques to identify dynamically assigned addresses in the Internet [27, 34, 5]. A recent study estimates that about 40% of responsive Internet blocks are dynamic based on Internet address surveys using ICMP probes taken every 11 minutes for two weeks [5]. (We assume here that non-stable blocks are primarily due to dynamic addressing.)

To evaluate the prevalence of stable and non-stable blocks, we would like to identify them from the history that we collect. Prior analysis of surveys used address *availability* and *volatility* to identify dynamic addressing. Availability is the fraction of times the address responds in all probes, while volatility is the fraction of times the address changes between responsive and non-responsive [5]. While appropriate for survey data with 11-minute probes, volatility makes less sense when probes are months apart.

To identify stable blocks with infrequent probes, define a new metric, *truncated availability*, the fraction of time an address responds from its first positive response. More formally, if $r_i(a)$ is the response of address a to the i th probe, the raw and scaled availability, $A^*(a)$ and $A(a)$ (from [5]) and truncated availability, $A^t(a)$ are:

$$\begin{aligned} A^*(a) &= \sum_1^{N_h} r_i \\ A(a) &= A^*(a)/N_h \\ A^t(a) &= A^*(a)/L^*(a) \end{aligned}$$

where $L^*(a)$ is the length of a history, in observations, from the first positive response to the present.

While both volatility and truncated availability are correlated, we found that low volatility and high truncated availability are both good predictors of a stable block. Low A^t values are a good predictor of intermittently used addresses. Continuing the examples in Figure 2, 00001111 has $A^t = 1$, while 01010101 has $A^t = 0.57$.

While A^t is good at differentiating between these solid (00001111) and intermittent (01010101) addresses, it interacts with gone-dark addresses, which will have a string of trailing 0s.

From these, we define a stable representative as $A^t \geq 0.9$. From Table 4 We find that 43% of all representatives are not stable by $A^t < 0.9$, a little higher than other independent observations 40% [5], and of 34–61% [34] (these values are for a random sample of DNS and for Hotmail users, respectively). We do not claim strong validation of this exact percentage because each work is a percentage of a different populations, and different definitions of what is dynamic or not stable. We claim only that our metric is in the right order of magnitude and so provides some insight into sources of non-responsive representatives.

hitlist	predicted representatives (N_p)	Responsive representatives and fraction			
		4	8	12	16
HL19/-	3,091,646 (100%)	1,558,620 (50%)	1,586,303 (51%)	—	—
HL21/-	3,386,540 (100%)	1,813,276 (54%)	1,846,019 (55%)	—	—
HL23/-	3,613,523 (100%)	1,925,322 (53%)	1,948,634 (54%)	1,950,960 (54%)	—
HL25/-	3,794,973 (100%)	2,007,138 (53%)	2,049,607 (54%)	2,059,019 (54%)	—
HL27/-	3,971,208 (100%)	2,135,337 (54%)	2,179,777 (55%)	2,193,062 (55%)	2,200,674 (55%)

Table 3: Responsive representatives with power weighting across 5 different hitlists for different history length.

	HL28/16		HL31/16	
predicted representatives (N_p)	4,055,193	100%	4,307,644	100%
not stable	1,749,471	43%	1,820,806	42%
gone dark	703,987	17%	772,014	18%
responsive (N_r)	2,250,091	56%	2,560,420	59%
non-responsive (N_n)	1,805,102	44%	1,747,224	41%
non-responsive and not stable (only)	590,472	14%	565,654	13%
non-responsive and gone dark only	0	0%	0	0%
non-responsive, not stable and gone-dark	693,832	17%	766,338	18%
non-responsive, just unlucky	520,798	13%	415,232	10%

Table 4: Causes of unsuccessful representatives predicted from HL28/16 and HL31/16, evaluated against responses in it29w and it32w. We don’t apply gone-dark window on prediction here, gone-dark blocks are detected separately with gone-dark window size of 3.

Re-evaluating causes of non-responsive representatives: With these definitions, we return to Table 4. We see that both gone-dark and not-stable blocks contribute three-quarters of our misses. Almost one third are in not-stable only blocks, with almost 40% gone-dark. We therefore claim that three-quarters of our non-responses are due either to new firewalls or selection of representatives in not-stable blocks, neither of which can *ever* have always responsive representatives.

To support the claim that lower A^t values correlate with poorer response, Table 5 breaks out the 1.8M non-responsive representatives by two values of A^t . We see that 29% of non-responses come from stable blocks ($A^t > 0.9$). Representatives with poor truncated availability ($A^t < 0.9$) account for more than two-thirds of non-responses. We conclude there are many unstable blocks, such blocks simply cannot be expected to support stable representatives. Also note, by our definition of gone-dark, dark blocks also qualify as not stable (because $A^t < 0.9$).

To show our choice of threshold for A^t does not alter our conclusion, Figure 3 shows the cumulative distribution of A^t for both non-responsive and responsive representatives. It shows a large difference in responsiveness for any value of A^t .

4.2 Completeness

To evaluate completeness, Figure 4 shows the absolute number of representatives for using 16-deep histories through five different censuses, and Table 6 shows the raw data. We consistently see that about one-third of blocks have some history data allowing an informed selection of representatives (the white region of the graphs, with around 4.2M

non-responsive representatives	
$A^t < 0.9$	1,284,304 (71%)
$A^t \geq 0.9$	520,798 (29%)
total	1,805,102 (100%)

Table 5: Fraction of representatives that are non-responsive, based on A^t (HL28/16 tested against it29w).

blocks). By contrast, about two-thirds of blocks have never responded (the top grey regions)

In addition, this data shows gone-dark selection from Section 3.4. We identify about 0.3–1.5% of allocated blocks as formerly responsive (the black region in the middle of Figure 4).

To guarantee completeness, we select random representatives for never-responsive blocks. However, we can see that we can provide informed choices for only a third of blocks. Finally, we note that IANA only releases new allocation maps quarterly, and routing studies suggest this space becomes routable gradually [3], so we expect our hitlist to be useful for at least three months, about the frequency we update them.

4.3 Stability and Inertia

We next consider two aspects of hitlist stability: how much churn is there in the hitlist, with and without a representative inertia, and how much does inertia reduce prediction accuracy.

Recall that inertia is the amount I by which prediction score must improve to change representatives. An inertia $I = 0$ means we always pick the highest rank address in

class	HL28/16	HL29/16	HL30/16	HL31/16	HL32/16
allocated /24 blocks	12,774,056	12,774,056	12,905,128	13,036,541	13,167,613
never responding blocks	8,718,863	8,631,417	86,797,99	8,728,897	8,775,398
predicted blocks	4,055,193	4,142,639	4,225,329	4,307,644	4,392,215
gone-dark blocks	35,341	75,714	109,099	154180	195,216
informed prediction blocks	4,019,852 (100%)	4,066,925 (100%)	4,116,230 (100%)	4,153,464 (100%)	4,196,999
changed representatives	— (5%)	218,419 (8%)	341,765 (7%)	292,079 (7%)	306,588
new representatives	— (4%)	171,428 (2%)	82,690 (2%)	82,315 (2%)	84,571
responsiveness	2,250,091	2,344,539	2,411,662	2,451,351	—

Table 6: Released hitlists to-date, by last census used in prediction (top). The top group of rows show hitlist composition, including churn (changed) and new representatives relative to the prior hitlist. The bottom line, responsiveness, evaluates the hitlist against the census.

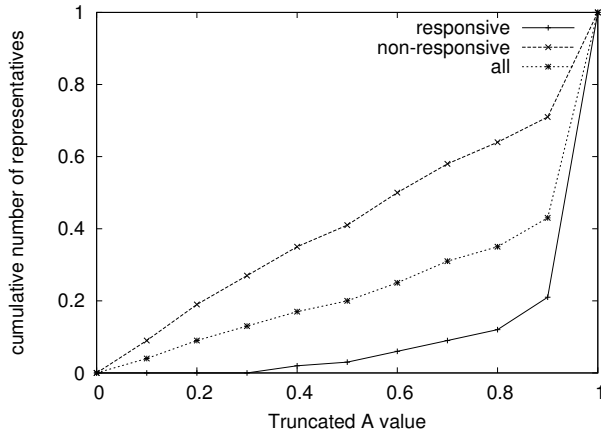


Figure 3: A^t cumulative distribution on responsive part of N_p , non-responsive part of N_p and all N_p

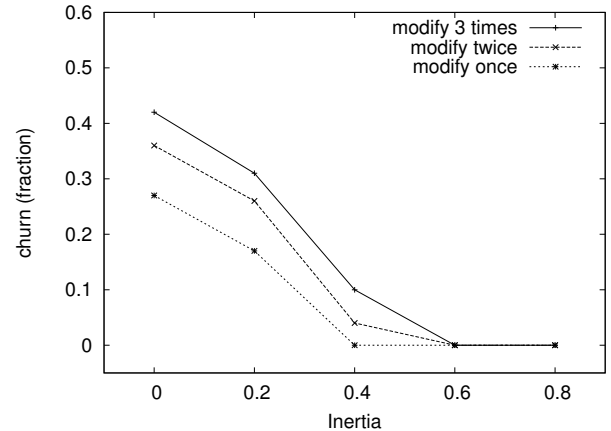


Figure 5: Effects of different inertia on representative churn (HL28/16 modified three times by it29w through it31w; modified twice by it29w, it30w; and once by only it29w).

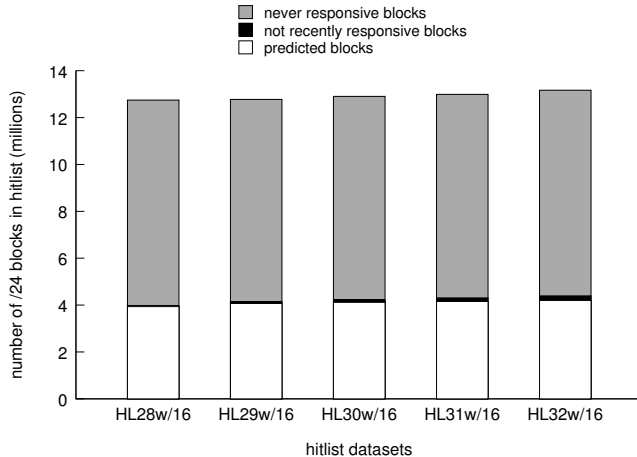


Figure 4: Relative size of hitlist components.

a block as the representative, independent of the representative in a hitlist based on a prior censuses. As inertia approaches 1, we will never switch representatives once chosen. For our production hitlists, we use $I = 0.34$ based on score changes due to weighting (Section 4.1.1) and the following analysis.

Inertia on churn: We first consider how much inertia affects churn. Churn is that rate at which we switch representatives for established blocks. Table 6 shows the amount of churn for four hitlists when using our standard inertia $I = 0.34$. Churn is shown in the “changed representatives” row, so in the HL32/16 column, we see that about 7% of all predictions (306,588 representatives) changed relative to the prior hitlist (HL31/16). This Table 6 shows that the rate of churn is relatively stable over time, with 5–7% of all informed predictions changing each census.

While Table 6 shows churn over time for a fixed inertia, in Figure 5 we vary inertia to observe its effect on churn. To estimate the relationship shown in this figure, we generate HL28/16, then modify it three times with censuses it29w, it30w, and it31w with different levels of inertia. (Here we suspend gone-dark processing to focus only on inertia.) We then evaluate the hitlist against observations from census it32w. We evaluate inertia over several steps for two reasons. First, hitlist staleness is partially a function of time. Second, large values of inertia suppress changes in single or a few censuses.

As expected, Figure 5 shows that higher inertia suppresses churn, because it takes several new negative responses for a representative’s score to change. In fact, weight selection means score can change only by 0.3 from one new census,

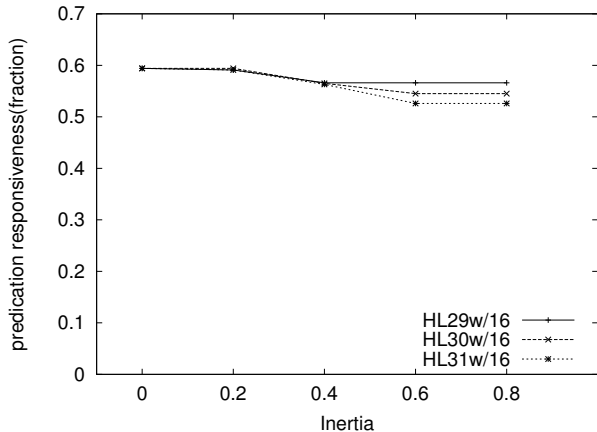


Figure 6: Effects of different inertia on responsiveness (HL28/16; modified once by it29w, then tested against it30w; modified twice by it29w, and it30w, then test against it31w; then modified three times by it29w through it31w and tested against it32w).

and decrease to 0.5 from two new censuses since the weight decrease in our pow weighting, so with three new observations here, an inertia of 0.2 has one observations that might cause change, while $I = 0.4$ has two; $I = 0.6$, three; and $I = 0.8$ requires more than eight observations to change.

Inertia on responsiveness: Inertia is selected to keep hitlists stable, reducing the amount of arbitrary representative turnover in long-running experiments. Such turnover can be eliminated by simply never changing representatives (setting $I = 1$), but prior experience shows that the responsiveness of a static hitlist will degrade over time as servers move, losing as much as 2–3% per month for the early Skitter web-server-based list [19]. We would therefore like to know the trade-off between inertia and representative responsiveness.

Figure 6 shows hitlist responsiveness for different values of inertia after this process. (This analysis was generated with the same multi-step process as Figure 5 described above.) We see that responsiveness degrades slightly for high inertia values, from 59% responsiveness with no inertia, to a low of 53% responsiveness when $I = 0.8$, when there are effectively no changes. We conclude that a moderate inertia has little effect on responsiveness costing at most 6 percentage points, even over eight months.

4.4 Effects of Probe Frequency

Hitlists are based on periodic censuses, so an important operational question is how the frequency of such censuses affect the quality of the hitlists. Intuitively, more frequent sampling can provide more information on address responsiveness, resulting in better predictions. Our current censuses are taken every two to three months (Table 1). More frequent collection is possible, but should be justified because data collection entails some operational cost.

While Internet censuses cover the whole Internet very few months, we also have access to surveys that probe 1% of the Internet every 11 minutes for two weeks [18]. Here we turn to these much more frequent probes to evaluate the effects of different probe rates. A survey has about 1800

Probe Frequency	Responsive Representatives
Survey, 12 hours	15,484 (68%)
24 hours	15,362 (67%)
48 hours	15,197 (66%)
HL30/16 (3-months)	16,123 (70%)
total blocks analyzed	22,861 (100%)

Table 7: Prediction accuracy (in responsiveness) from different probe frequencies, for a given number of /24 blocks.

observations, but we downsample this information to get observations every 12, 24, or 48 hours, providing 28, 14, and 7 bits of history. We use survey 30 (taken at the same time as it30w in Dec., 2009) to do prediction, and test it against it31w. We use the average weighting function to evaluate these observations, since it seems unnecessary to favor recent observations when all are taken of a short, two-week period. In addition, in Section 4.1.1 we showed that weighting has relatively little effect on responsiveness. We then test the prediction from this history against the next census (it31w) to evaluate prediction accuracy. We compare the representatives found from the survey with those computed in HL30/16.

Table 7 compares responsiveness as a function of probe frequency. First, we see that, of the survey-derived hitlists, more frequent probing provides a slightly better prediction (68% responsive from 12-hour samples vs. 66% for 48 hours). This small improvement is because more frequent probing gives more information on host responsiveness. Second, we see that the census-derived hitlist is a better prediction than any of the survey-derived hitlists, by 2–4%. The main source of this difference is that the survey’s hitlist finds some addresses that don’t appear in the census-derived hitlist. These addresses seems less stable in the next census because the survey-derived hitlist considers only a short period of history, while the census-based hitlist considers many months history and so finds long-term stable addresses, if they exist. In addition, because our survey observations are taken at specific times of day, they may discover addresses that are only up during particular daily hours, while evaluating against the next census tests at random times of day.

4.5 Effects on Other Research

The above sections evaluate hitlists based on our goals: responsiveness, completeness, and stability. But hitlists are a tool to enable other research, so their ultimate benefits come by how they improve the quality of other network performance and topology studies.

Some network performance studies *require* responsiveness in their destinations. These studies include those that evaluate performance [32, 20, 23], consider questions about routing and reachability [31, 3], or the performance of replica placement (examples include [32, 10]). For studies that require end-to-end latency measurements, our representative selection methods optimize reachability within the constraints of sparse measurement. Our work also suggests directions for potential improvements: more frequent measurement could potentially better track reachable addresses

in dynamically assigned blocks. In addition, our approach to stability assists evaluation of long-term performance trends.

Responsiveness is helpful but not essential for many other topology studies (such as [17, 15, 26, 20, 22, 8]). Most topology studies employ traceroutes to study *paths* across the Internet. A traceroute attempts to discover all IP addresses on the path towards a destination, and many such paths are aggregated with alias elimination [17, 26, 21] to produce a router-level map. However, in such studies, the presence or absence of the destination itself affects only the last hop. Topology studies thus *do not require* responsive representatives, but they may benefit from responsiveness. However, many have moved towards the use random or deterministically selected representatives. As one example, topology probing in Skitter [20] began with a manually generated hitlist, but later shifted to random probing in Archipelago [8]. Two reasons for this shift were difficulty in maintaining a responsive hitlist and the recognition that responsive targets are not essential.

Although responsiveness is not essential for topology studies that focus on the core of the Internet, it is important for studies that wish to explore the *edge* of the network. We can get a rough estimate on the number of edge links that are missed by randomly selected representatives: empirically, about 4–7% of the Internet responds to ICMP probes [18], so we expect that 93% of random representatives will not respond (approximating the distribution of responders as uniform, to get a rough estimate). If 55% of our hitlists respond, that will improve edge detection for 38% of blocks. With 1.3 million allocated /24 blocks (as of Nov. 2009), statistics suggest that responsive hitlists will detect about 630,000 *additional* links compared to those found using a random hitlist. Some topology studies examining the core of the Internet find about 33M links, so this increases the size of the discovered Internet topology by 2% [6]. This simple analysis ignores correlation in the data, so it is only approximate.

To confirm this simple analysis, we consider CAIDA’s Internet Topology Data Kit (ITDK-2010-01 [6]). ITDK combines the results of traceroutes from many traceroutes with alias resolution [17, 8] to produce a router-level map of the Internet. ITDK is formed from 42 *cycles* of data, each representing a traceroute to a randomly chosen representative in each /24 prefix.

Unfortunately, we cannot directly evaluate ITDK with different approaches to representative selection because it is the result of several levels of processing (aggregation of multiple cycles of data, and duplicate and alias elimination). In addition, direct comparison would be somewhat misleading since the goal of ITDK is to map the Internet’s core, not edges. Instead, we obtained the raw IPv4 topology data [7] used to generate ITDK. We then compare the results of one cycle with random representatives, to using what we expect would have resulted using our HL29/16 (based on censuses from 2007 through Dec. 2009). We can evaluate the responsiveness of randomly selected representatives by looking at the raw data, while we test hitlist responsiveness by comparing to it30w completed in Jan. 2010. (Here report these results comparing to a single cycle, #739, of the raw IPv4 data. We got very similar results comparing to cycles 740 and 741.)

Table 8 summarizes our observations. First, we see that we only have data to make predictions for just under half

of the blocks (N_p of 4,008,861). For the remaining 51%, we too fall back on random probing. (We actually predict about 100k additional prefixes because we use a more recent routing table compared to the IPv4 raw data, however, we omit these from the table to provide a more fair comparison.)

Second, we see that random probing finds responsive addresses about 9% of the time. This result is somewhat better than we predict (4-7% [18]), a difference probably because our prediction is over all allocated blocks, but here we study only *routable* blocks, and non-routable blocks cannot possibly representatives.

Finally, we find about 1.7 million additional edges if we use our hitlist compared to random probing, about 2.4× more edge links than random probing alone finds. For studies of edge links, this is a significant amount of additional coverage. Compared to the 33M links found in ITDK-2010-01, this use of informed probing would discover about 5% greater coverage. However, care must be taken in making this comparison for several reasons. First, ITDK is composed of 42 cycles, while we compare to only a single cycle. Because each cycle probes a different random destination, total coverage will improve as it eventually finds responsive destinations by chance. (In fact, with enough cycles, random probing may provide better coverage as it converges on probing all possible addresses.) Second, it is important to note that ITDK is designed to study the Internet core, not the edges, and this goal is well met with random probing alone. However, our results show they could get additional coverage with no additional probing effort by changing to representatives selected with a method such as ours.

4.6 Cost of Hitlist Generation

Finally, turning from hitlist quality, we consider the cost of generating hitlists. Hitlists are generated from a number of Internet censuses. We therefore review the cost of taking a census and then look at the cost of processing censuses to produce a hitlist.

Hitlists require Internet censuses as input. New censuses are started every two to three months; during most of that time four machines are actively probing the Internet, while there is a week or two of setup and data curation time at the beginning and end of a census [18]. Each census currently requires 12–15GB of archival storage. Censuses are carried out as part of an ongoing Internet topology research project and are used for multiple purposes, but if carried out exclusively for hitlist generation they would represent an ongoing cost.

Processing hitlist from census data incurs several ongoing costs. First, we maintain a master history file with bitmaps of all censuses to date, indexed by IP address. Each census creates an observation file with about 1–2GB of new positive observations. We merge observations into an ongoing history file of all observations to date. Currently 22 observation files (32GB total) are merged into one 7.5GB history file, and the size of that file grows by about 300MB every new census.

We parallelize our computation with the Hadoop implementation [11] of map/reduce [12], running over a cluster of about 40 computers with about 120 CPU cores. With this parallelism, join in a new census into an existing history takes about half an hour, and evaluation of a new hitlist takes about another half hour. (our code is written in Perl and not optimized for speed). Our map function groups

/24 routable blocks studied	8,248,027	100%	
no prediction ($\overline{N_p}$)	4,239,166	51%	
prediction (N_p)	4,008,861	49%	[100%]
responsive	2,454,500		[61%]
random probing responsive	730,496	9%	
random probing responsive also in N_p	725,930		[18%]
extra edges found	1,728,570	21%	[43%]

Table 8: Evaluation of random and informed representatives on destinations from one cycle of IPv4 traceroute data.

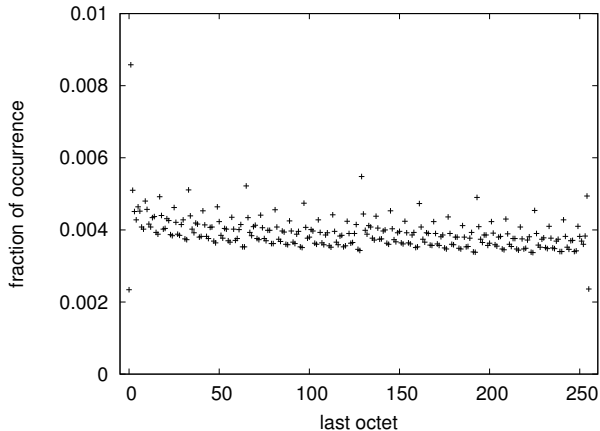


Figure 7: Frequency of responsiveness by last octet of the address (from it29w).

results by block, while the reducer carries out the join or evaluation.

5. OTHER OBSERVATIONS

Given no knowledge about a /24 block, which address is most likely to be responsive?² This question has some bearing on which representative we should select for gone-dark blocks, or for newly-allocated blocks with no census data yet.

Discussions with network operators suggest some network practices are common. Often addresses are allocated sequentially from the start of an block, and network managers often use the first or last address in a block for the routers. Since address blocks are allocated on powers-of-two according to CIDR [16], we expect to see uneven use of the address space. Recent work has confirmed visibility of allocation blocks in census data [5], but not last-octet usage.

To evaluate this question, Figure 7 shows the distribution of responsiveness for the last octet for all in it29w. (We got similar results on it28w and it30w censuses.) Consistent with expectation, the most responsive octet is .1, responding 0.86% of the time, more than twice as often as the median responsiveness (0.38%), and $1.5\times$ more frequent than .129 (0.55%), the next most responsive last octet.

Figure 7 shows a pattern in responsiveness, with responses being most frequent at addresses that are one greater than

²We thank Kim Claffy for suggesting the question of last octet distribution for study.

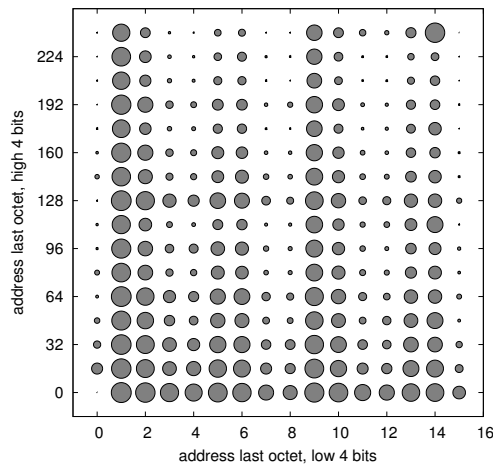


Figure 8: Rank (shown by circle area) of responsiveness by last octet, in a 16×16 grid by address (from it29w).

a power of two. The top ten are ranked .1, .129, .65, .33, .2, .254, .17, .193, .97, .9, and of these only 2 and 254 do not follow this pattern. To show this trend more clearly, Figure 8 shows the rank of each last octet as the area of a circle, with the octets arranged in a sequential grid (so the x axis lists octets sequentially in groups of 16, while each step up the y axis is 16 more than the previous). The vertical lines correspond to more frequent responses, with $x = 1$ showing strong response from .1, .129, .65, etc., and $x = 9$ for .9, etc. The other prominent features are .254 and across $y = 0$ (.1, .2, .3, etc.). While showing ranks exaggerate what can be small absolute differences, these strong patterns show power-of-two allocations affect responsiveness.

6. SHARING HITLISTS

Our goal in generating hitlists is to share them with other research groups carrying out topology studies. We offer them free-of-charge to all, and to date we have provided them to four other projects. Although hitlists are not human subjects, networks are operated by and involve humans. Hitlist use by multiple prompts us to consider their distribution in the context of the Belmont protocols [28], weighing the benefits and potential costs of sharing and designing policies accordingly.

The benefits of sharing hitlists are similar to sharing of other research results. Shared data is a boon to researchers.

A common data source can lower the barrier to entry for future research, and it also makes it easier for researchers to compare their results. (For example, the TREC benchmarks are seen as essential to rapid advances in the field of Information Retrieval [13], although our efforts are far more modest.) As importantly, we expect that the scrutiny of multiple researchers on a common dataset can often identify data or methodological errors that might be otherwise unnoticed. (In Internet topology, the problem of alias resolution is one that is still being refined [2, 21], nearly ten years after the first techniques [17].) For the hitlist creator, a shared result amortizes the operational costs of collection and processing of the input data (Internet censuses) needed to create hitlists. Finally, for the hitlist subjects, network operators in the Internet, a common source allows us to centralize “do-not-probe” blacklists and reduces raw data collection.

Shared hitlists have some costs, however. Most serious is that a hitlist can focus the probing of several researchers on a specific representative address in a network, while independently derived hitlists are more likely to distribute probing load. Second, eventually hitlists will be acquired by malicious users on the Internet. Potential harms are hiding malicious traffic mixed with research traffic, and the slight risk that any list of known active IP addresses may be at risk of additional malicious traffic such as worms or cracking attempts. While a risk, the effort to generate a hitlist is within the reach of a motivated individual, so strong restrictions on hitlists seem unwarranted.

Our current hitlist distribution policies are designed to balance risks with benefits. Although we share hitlists free-of-charge, we provide them subject to a usage agreement. Hitlist users may not redistribute hitlists so we can establish this agreement directly with all users. Tracking hitlist users allows us to estimate load on representatives. We also plan to seed the hitlist with representatives that we monitor to track load. We also hope controlled hitlist distribution delays their acquisition by malicious parties. We expect to review these policies as we gain more experience.

7. FUTURE WORK

Although we are providing our current hitlists for use currently, we see two directions for future work.

First, we would like to better understand the stability of Internet address usage. Our study shows only 50–60% of informed representatives respond three months later, implying a great deal of churn on Internet address usage. We would like to better understand why usage of the remaining addresses cannot be better predicted. More detailed analysis of gone-dark and unstable blocks may provide more information; and a combination of our data with information IANA assignments, RIR registration and block reassignment may explain causes for failure to respond. Alternatively, changes in firewalling may explain address instability.

Second, we expect to get more information about how hitlists are used. Several groups expressed concern that representatives may become overloaded as multiple groups target them with measurement traffic. We are in the process of seed our hitlist with monitors so we can observe hitlist use and compare it to normal “background radiation”. Such an evaluation will help understand hitlist use and characterize general unsolicited Internet traffic. If representative load becomes high, it will suggest operational changes, such as multiple representatives per /24 block.

8. CONCLUSIONS

We have defined the properties that are important to hitlists: representatives that are responsive, stable, and provide complete coverage for the Internet. We have developed a fully automated algorithm that mines data from Internet censuses to select informed representatives for the visible Internet. We employ information that is available for about one-third of the Internet, and when an informed representative is available we see it is 50–60% likely to respond 2–3 months later. We showed that the primary reasons for prediction failure are blocks with dynamic addressing and gone-dark blocks that are probable firewalls.

Our hitlists are available free-of-charge and have already been distributed to four different research groups. Although we do not have external evaluation of how their use changes those studies, our evaluation of one prior study suggests the potential to discover 1.7 million additional edge links.

Acknowledgments

We thank Kim Claffy and Young Hyun for a detailed reading and comments that improved the paper, and for providing their raw data for comparison in Section 4.5. We also thank the anonymous reviewers for their comments. We are grateful to Yuri Pradkin for developing and operating the Internet census infrastructure that enables this work. We thank Randy Bush, Olaf Maennel, Beichuan Zhang, Ethan Katz-Bassett, and Harsha Madhyastha for their encouragement about this work.

9. REFERENCES

- [1] R. Albert, H. Jeong, and A.-L. Barabási. Error and attack tolerance in complex networks. *Nature*, 406:378–382, July 27 2000.
- [2] A. Bender, R. Sherwood, and N. Spring. Fixing Ally’s growing pains with velocity modeling. In *Proceedings of the 8th ACM Internet Measurement Conference*, pages 337–342, Vouliagmeni, Greece, Oct. 2008. ACM.
- [3] R. Bush, J. Hiebert, O. Maennel, M. Roughan, and S. Uhlig. Testing the reachability of (new) address space. In *Proceedings of the ACM Workshop on Internet Network Management*, pages 236–241, Kyoto, Japan, Aug. 2007. ACM.
- [4] R. Bush, O. Maennel, M. Roughan, and S. Uhlig. Internet optometry: assessing the broken glasses in internet reachability. In *Proceedings of the ACM Internet Measurement Conference*, pages 242–253. ACM, Nov. 2009.
- [5] X. Cai and J. Heidemann. Understanding address usage in the visible internet. Technical Report ISI-TR-2009-656, USC/Information Sciences Institute, Feb. 2009.
- [6] CAIDA. The internet topology data kit—2010-01. <http://www.caida.org/data/active/internet-topology-data-kit/>, Jan. 2010.
- [7] CAIDA. The ipv4 routed /24 topology dataset 2009-12. http://www.caida.org/data/active/ipv4_routed_24_topology_dataset.xml, Jan. 2010.
- [8] K. Claffy, Y. Hyun, K. Keys, M. Fomenkov, and D. Krioukov. Internet mapping: from art to science. In *Proceedings of the IEEE Cybersecurity Applications and Technologies Conference for Homeland Security*

- (*CATCH*), pages 205–211, Alexandria, VA, USA, Mar. 2009. IEEE.
- [9] D. D. Clark, C. Partridge, J. C. Ramming, and J. T. Wroclawski. A knowledge plane for the Internet. In *Proceedings of the ACM SIGCOMM Conference*, pages 3–10, Karlsruhe, Germany, Aug. 2003. ACM.
- [10] E. Cronin, S. Jamin, C. Jin, A. R. Kurc, D. Raz, and Y. Shavitt. Constrained mirror placement on the Internet. *IEEE Journal of Selected Areas in Communication*, 20(7):1369–1383, Sept. 2002.
- [11] D. Cutting. Scalable computing with Hadoop. <http://wiki.apache.org/lucene-hadoop-data/attachments/HadoopPresentation%2Fns/attachments/yahoo-sds.pdf>, May 2006. Lecture note.
- [12] J. Dean and S. Ghemawat. MapReduce: Simplified data processing on large clusters. In *Proceedings of the USENIX Symposium on Operating Systems Design and Implementation*, pages 137–150, San Francisco, California, USA, Dec. 2004. USENIX.
- [13] A. Dekhtyar and J. H. Hayes. Good benchmarks are hard to find: Toward the benchmark for information retrieval applications in software engineering. In *Proceedings of the 22nd International Conference on Software Maintenance*, Philadelphia, Pennsylvania, USA, Sept. 2006. ACM.
- [14] J. Dilley, B. Maggs, J. Parikh, H. Prokop, R. Sitaraman, and B. Weihl. Globally distributed content delivery. *IEEE Internet Computing*, 6(5):50–58, Sept. 2002.
- [15] P. Francis, S. Jamin, C. Jin, Y. Jin, D. Raz, Y. Shavitt, and L. Zhang. IDMaps: A global internet host distance estimation service. *ACM/IEEE Transactions on Networking*, 9(5):525–540, Oct. 2001.
- [16] V. Fuller, T. Li, J. Yu, and K. Varadhan. Classless inter-domain routing (CIDR): an address assignment and aggregation strategy. RFC 1519, Internet Request For Comments, Sept. 1993.
- [17] R. Govindan and H. Tangmunarunkit. Heuristics for Internet map discovery. In *Proceedings of the IEEE Infocom*, pages 1371–1380, Tel Aviv, Israel, Mar. 2000. IEEE.
- [18] J. Heidemann, Y. Pradkin, R. Govindan, C. Papadopoulos, G. Bartlett, and J. Bannister. Census and survey of the visible Internet. In *Proceedings of the ACM Internet Measurement Conference*, pages 169–182, Vouliagmeni, Greece, Oct. 2008. ACM.
- [19] B. Huffaker, M. Fomenkov, D. Moore, and kc claffy. Macroscopic analyses of the infrastructure: measurement and visualization of internet connectivity and performance. <http://www.caida.org/outreach/papers/pam2001/skitter.xml>, Nov. 2001.
- [20] B. Huffaker, M. Fomenkov, D. J. Plummer, D. Moore, and kc claffy. Distance metrics in the internet. In *Proceedings of the IEEE International Telecommunications Symposium*. IEEE, 2002.
- [21] K. Keys. IP alias resolution techniques. Technical report, CAIDA, 2008.
- [22] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: An information plane for distributed services. In *Proceedings of the 7th USENIX Symposium on Operating Systems Design and Implementation*, pages 367–380, Seattle, WA, USA, Nov. 2006. USENIX.
- [23] H. V. Madhyastha, E. Katz-Bassett, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane Nano: Path prediction for peer-to-peer applications. In *Proceedings of the 6th USENIX Symposium on Network Systems Design and Implementation*, Boston, MA, USA, Apr. 2009. USENIX.
- [24] E. Pfanner. Broadband speeds surge in many countries. *New York Times*, page B8, Oct. 1 2009.
- [25] R. Sherwood, A. Bender, and N. Spring. DisCarte: A disjunctive Internet cartographer. In *Proceedings of the ACM SIGCOMM Conference*, pages 303–315, Seattle, Washington, USA, Aug. 2008. ACM.
- [26] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP topologies with Rocketfuel. In *Proceedings of the ACM SIGCOMM Conference*, pages 133–145, Pittsburgh, Pennsylvania, USA, Aug. 2002. ACM.
- [27] M. Sullivan and L. Munoz. Suggested generic DNS naming schemes for large networks and unassigned hosts. Work in progress (Internet draft draft-msullivan-dnsop-generic-naming-schemes-00.txt, Apr. 2006.
- [28] The National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. The Belmont report: Ethical principles and guidelines for the protection of human subjects of research. Technical report, Department of Health, Education, and Welfare, Apr. 1979.
- [29] USC/LANDER Project. Internet IPv4 address space census. PREDICT ID USC-LANDER/internet_address_survey_it11w-20060307. Retrieval information for this and other censuses is at <http://www.isi.edu/ant/traces/>, Mar. 2006.
- [30] D. Waddington, F. Chang, R. Viswanathan, and B. Yao. Topology discovery for public IPv6 networks. *ACM Computer Communication Review*, 33(3):59–68, July 2003.
- [31] F. Wang, Z. M. Mao, J. Wang, L. Gao, and R. Bush. A measurement study on the impact of routing events on end-to-end Internet path performance. In *Proceedings of the ACM SIGCOMM Conference*, Pisa, Italy, Aug. 2006. ACM.
- [32] R. Wolski. Dynamically forecasting network performance using the network weather service. *Journal of Cluster Computing*, 1:119–132, Jan. 1998. Also released as UCSD technical report TR-CS96-494.
- [33] E. Wyatt. Despite ruling, F.C.C. says it will move forward on expanding broadband. *New York Times*, page B3, April 15 2010.
- [34] Y. Xie, F. Yu, K. Achan, E. Gillum, M. Goldszmidt, and T. Wobber. How dynamic are IP addresses? In *Proceedings of the ACM SIGCOMM Conference*, pages 301–312, Kyoto, Japan, Aug. 2007. ACM.