

# Primitives for Active Internet Topology Mapping: Toward High-Frequency Characterization

Robert Beverly  
Naval Postgraduate School  
rbeverly@nps.edu

Arthur Berger  
MIT CSAIL / Akamai  
awberger@csail.mit.edu

Geoffrey G. Xie  
Naval Postgraduate School  
xie@nps.edu

## ABSTRACT

Current large-scale topology mapping systems require multiple days to characterize the Internet due to the large amount of probing traffic they incur. The accuracy of maps from existing systems is unknown, yet empirical evidence suggests that additional fine-grained probing exposes hidden links and temporal dynamics. Through longitudinal analysis of data from the Archipelago and iPlane systems, in conjunction with our own active probing, we examine how to shorten Internet topology mapping cycle time. In particular, this work develops discriminatory primitives that maximize topological fidelity while being efficient.

We propose and evaluate adaptive probing techniques that leverage external knowledge (e.g., common subnetting structures) and data from prior cycle(s) to guide the selection of probed destinations and the assignment of destinations to vantage points. Our *Interface Set Cover* (ISC) algorithm generalizes previous dynamic probing work. Crucially, ISC runs across probing cycles to minimize probing while detecting load balancing and reacting to topological changes. To maximize the information gain of each trace, our *Subnet Centric Probing* technique selects destinations more likely to expose their network's internal structure. Finally, the *Vantage Point Spreading* algorithm uses network knowledge to increase path diversity to destination ingress points.

## Categories and Subject Descriptors

C.2.3 [Computer Communication Networks]: Network Operations—*network monitoring*; C.2.1 [Computer Communication Networks]: Network Architecture and Design

## General Terms

Measurement, Experimentation

## Keywords

Internet Topology, Network Topology, Adaptive Probing

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.

## 1. INTRODUCTION

The scale of the Internet makes obtaining representative metrics and characteristics challenging. Compounding this challenge, the Internet is poorly instrumented, lacks measurement and management mechanisms [5], and providers hide information. Researchers therefore must frequently make inferences over limited available data, and may form false conclusions [15].

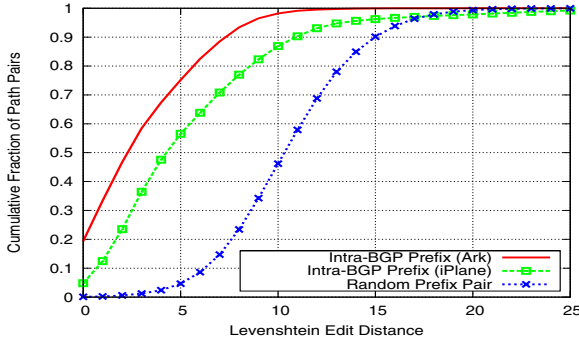
Understanding the complex structure of the Internet is vital for network research including routing, protocol validation, developing new architectures, etc. More importantly, building robust networks, and protecting critical infrastructure, depends on accurate topology mapping.

While dedicated platforms exist to perform topology measurements, e.g. [11, 19], these must balance induced measurement load against model fidelity. Unfortunately, in practice, such balancing results in multiple days worth of measurement to capture even an incomplete portion of the Internet. Employing more vantage points is an effective technique to improve topological recall [23], but does not reduce total load or cycle time.

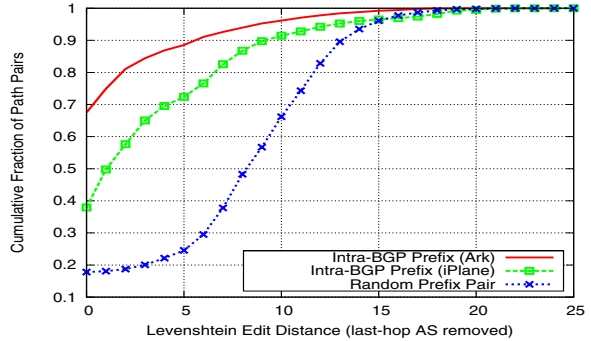
This work proposes primitives toward the eventual goal of performing *high-frequency* active Internet topology measurement. Measurement load hinders the ability to capture small-scale dynamics and transient effects that occur at frequencies higher than the measurement period; effectively creating Nyquist aliasing loss. For example, recent work [20] shows fewer than 50% of Internet paths remain stationary across consecutive days. Our own analysis of set cover techniques [10] finds that the rate of missed interfaces increases in proportion to the time since the covering set was created: implying that “train-then-test” methodologies are insufficient.

Our work therefore focuses on two separable problems via a unified methodology, how to: i) *select destinations* in the network to probe; and ii) *perform the probe*. We examine the hypothesis that by leveraging external network knowledge, e.g. routing, address structure, etc., and adaptive probing, the active traffic load can be significantly reduced without sacrificing the inferred topology fidelity. Our methodology extends prior schemes, e.g. [8] which attempt to reduce measurement overhead, but are artificially parametrized, lossy, and ignore temporal effects across measurement periods. Toward high-frequency Internet topology mapping, we:

1. Quantify unnecessary probing performed by production topology measurement platforms.
2. Develop three algorithms that use network knowledge to intelligently drive adaptive probing.



(a) Unnecessary probing:  $> 60\%$  of intra-BGP traces have  $ED \leq 3$ ; fewer than  $50\%$  of random traces have  $ED \leq 10$ .



(b) Contribution of last-hop AS to path variance:  $\sim 70\%$  of probes to same prefix yield no information gain beyond the leaf AS.

Figure 1: Edit distance (ED) distribution of Ark ( $\simeq 260k$ ) and iPlane ( $\simeq 150k$ ) traces to *different* addresses within the *same* BGP prefix compared to baseline ED between random trace pairs.

## 2. UNNECESSARY PROBING

Several large topology measurement experiments have been deployed, including CAIDA’s Skitter/Archipelago project (Ark) [11, 13], iPlane [19], and DIMES [22]. To better understand the challenges in topology mapping, this section focuses on the existing practice of Ark and iPlane which infer interface-level topologies via traceroute-like [1, 17] probing.

So that measurement is tractable, production systems often follow common assumptions over the Internet’s structure, for instance by probing a target in each subnetwork of size  $2^8$  (herein referred to via common /24 prefix notation). Ark subdivides all routed prefixes (i.e. visible in BGP) into /24’s. A “cycle” of probing is a complete set of measurements to one destination address within each routed /24. The probe target for a given /24 is randomly selected from the  $2^8$  possible addresses.

With approximately half of all IP addresses globally routable, a cycle consists of  $\sim 2^{31-8}/24$ s. Due to this large number of /24’s to be probed in a cycle, approximately 9M, Ark divides the probing work among multiple vantage points (measurement sites). Probing at a /24 granularity requires significant time, and load. With asynchronous, distributed probing to mitigate per-path RTT variance, a full cycle requires multiple days to partially characterize the Internet.

Traces can be distilled into an interface-level representation of the Internet graph. Some traces yield more information than others based on the choice of prior probes. For instance, we expect traces to different addresses within the same BGP prefix to be similar, while probes to very different destination addresses are likely to have a higher information gain.

### 2.1 A Path Pair Distance Metric

To quantify the information gain of intra-BGP traces, we use the Levenshtein, or edit, distance which is a measure of the minimum number of insert, delete or modify operations required to equate two strings.

Let the alphabet of symbols be the unsigned 32-bit integer space,  $\Sigma = \{0, \dots, 2^{32} - 1\}$ . We compute the edit distance ( $ED$ ) between trace pairs using each IP address along the path. An  $ED$  of zero implies that the two paths

are identical, whereas an  $ED$  of one implies that the two traces differ by a single interface addition, subtraction, or replacement. For example, for the following two interface paths,  $ED(t_1, t_2) = 2$ :

$$\begin{aligned} t_1 &= 1.2.6.1, 1.186.254.13, 2.245.179.52, 4.53.34.1 \\ t_2 &= 1.2.6.1, 2.245.179.52, 4.69.15.1 \end{aligned}$$


We use data from a single Ark and single iPlane monitor in a January, 2010 cycle for ED analysis. As a comparative baseline, we also compute ED over an equal number of random trace pairs.

### 2.2 Quantifying Unnecessary Probing

Figure 1(a) shows the cumulative fraction of path pairs in Ark and iPlane as a function of ED. The ED is larger for the randomly selected traceroute path pairs than the pairs from within the same BGP prefix, as determined by a contemporaneous Routeviews [21] BGP routing table. Approximately  $60\%$  of traces to destination in the same BGP prefix have  $ED \leq 3$  while fewer than  $50\%$  of random traces have  $ED \leq 10$ . Thus, as we intuitively expect, there is value to using the BGP structure to drive the probe target selection in order to maximize the information gain.

Next, we wish to quantify the contribution of the last hop autonomous system (AS) to the edit distance of traces to the same BGP prefix, i.e. path difference attributable to subnetting within an AS. For example, Figure 2 depicts the sources of path diversity observed as an “hourglass” with multiple vantage points contributing to diversity into an AS’s ingress points, and the degree of subnetting within the destination AS contributing to the remaining diversity. The “waist” is the set of ingress points for a prefix which may be common to multiple traces or require distributed vantage points in order to be discovered (§4.3 discusses the diminishing return of additional vantage points).

Figure 1(b) is the result of an ED analysis after removing interface hops belonging to the destination AS, as determined by the Routeviews BGP table. We observe that for  $\sim 70\%$  of the probe pairs to the same prefix, there is zero additional information gain beyond the leaf-AS. Therefore, from this off-line analysis of traces from two important



**Figure 2: Topology information gain hourglass: path diversity comes via multiple vantage points and via multiple destinations in a prefix. The hourglass “waist” is the AS ingress point(s).**

topology platforms, we conclude that there exist *significant possible packet savings* by intelligently tuning, e.g. via time-to-live (TTL), the set of hops each trace interrogates. For instance, a basic tracing strategy might start with a TTL suitable to reach the destination and iteratively decrement the TTL until a previously discovered hop, i.e. at the AS ingress, is found.

Moreover, in analyzing pairs of traceroutes to the same destination prefix, but from different vantage points, we find that in  $\sim 30\%$  of the cases, entirely new paths are discovered. Only approximately 10% of the probes from a new vantage point yield less than four previously undiscovered interfaces. Thus, there exists *significant information gain from additional vantage points*.

These potential efficiencies have been recognized, most prominently by the DoubleTree method [8, 7]. Unfortunately, DoubleTree relies on heuristics to tune its probing. In §3 we detail non-parameterized primitives designed to address the low-gain we find here and provide efficiency without sacrificing inference power.

Note that the ED’s for iPlane are higher than for Ark due a non-uniform distribution of traces to prefixes as part of the iPlane logic [18]. Since iPlane provides significantly fewer instances of multiple probes to the same prefix as compared with Ark, we can more readily test our primitives against the latter. We therefore use historic Ark data, as well as our own active probing, for the remainder of this paper.

### 3. ADAPTIVE PROBING METHODOLOGY

This section presents three strategies to illustrate the potential power of adaptive probing in reducing unnecessary probing: 1) *subnet centric probing*; 2) *interface set cover*; and 3) *vantage point spreading*.

#### 3.1 Subnet Centric Probing

A naive strategy of leveraging BGP knowledge is to probe exactly one destination within each advertised prefix. The potential for using BGP routing information was first recognized by Krishnamurty and Wang [14]. While we show that such an approach incurs approximately one-fifth of the normal amount of probing packets sent by Ark, it is too aggressive and misses significant topology information of networks with a rich subnetting structure.

Intuitively, we expect two numerically consecutive IP addresses to be more likely to share paths (and, hence, have a low ED) than two distant addresses. But simply employing address distance is too simplistic and does not capture typical network subnetting structure [4]. For example, the two IP addresses 18.255.255.254 and 19.1.1.1 have a numerical

distance of 2, but they would belong to different networks unless both belonged to a single 18.0.0.0/7 subnetwork.

Instead, we propose to use the knowledge of how networks are subnetted (the preceding example illustrating an example where subnetting is much more probable than no subnetting) to select addresses to probe within each BGP advertised prefix. The motivation is to adapt the number of probes to the degree of subnetting within the prefix to avoid wasted probing. We term this strategy “subnet centric.” The current Ark strategy assumes a fixed subnetting boundary, which may be too granular (wasted probing) or too coarse (missing information). In contrast, we ensure that subsequent destinations in a prefix are as distinct as possible in their most significant bits, i.e., likely part of distinct subnet prefixes. We term this selection of destination the *least common prefix* principle. For example, to choose four probes for prefix 192.168.0.0/16, our algorithm initially picks four addresses from the distinct prefixes: 192.168.0.0/18, 192.168.64.0/18, 192.168.128.0/18, and 192.168.192.0/18.

As probing progresses, we use our pair-wise ED measure to determine whether finer-grained destinations within a prefix are yielding useful additional information. The prefix is continually probed until the ED value of the paths returned by a pair falls within an empirically derived pre-determined threshold  $\tau = 3$ .

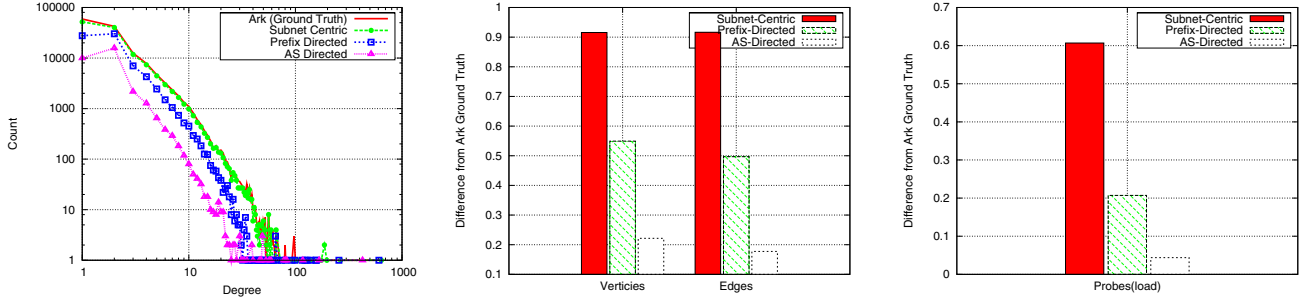
#### 3.2 Interface Set Cover

DoubleTree [8] explores a method to adapt probing in real-time as a measurement cycle progresses. By beginning at a heuristically chosen mid-point and working both back to the vantage point (decreasing TTL) and toward the destination (increasing TTL), DoubleTree achieves packets savings by preempting a trace when a previously discovered interface is observed – the inherent assumption is that subsequent probes along the path will be duplicates of previous traces.

While DoubleTree’s technique partially addresses our findings in §2, it must determine a path’s mid-point and does not cope well with load balancing. More importantly, it treats each cycle independently and is agnostic to information learned in previous probing cycles. Our goal is to leverage this knowledge to reduce the number of trace packets in subsequent cycles.

We hypothesize a greedy *Interface Set Cover* (ISC) scheme that always selects a subset of probing packets based on the interface-level topology of the previous cycle. More specifically, the interface-level topology includes directed edges where the direction of the edge records the direction of a probe. (It is not a problem if an edge is bi-directional, which for the interface-level graph should occur only as outliers.) The ISC scheme iteratively selects paths, and sub-paths, from the directed interface-level topology of the previous round, such that packets would probe interfaces that are not yet accounted for by the paths already selected. The initial “bootstrap” set of destinations may be chosen using our subnet centric probing algorithm. We note that the optimization problem of identifying a *minimum subset* of paths to cover the interfaces discovered is an instance of the well-known NP-complete “Min Set Cover” problem. However, efficient greedy solutions have been shown to be  $\ln n$  approximate of optimal [9].

Formally, let  $\mathbf{P}$  be a set of paths from vantage points to destinations. Each  $\vec{p}_i \in \mathbf{P}$  is a vector of router interfaces



(a) Inferred degree distribution: subnet centric probing well approximates Ark ground truth.

(b) Model fidelity: Subnet centric probing captures  $\geq 90\%$  of the vertices and edges.

(c) Induced load: Subnet centric uses  $\simeq 60\%$  of the ground truth load.

Figure 3: Subnet Centric directed probing performance.

corresponding to the  $i$ 'th path. Let the universe of interfaces be  $\mathbf{I} = \bigcup_{i,j} p_i[j]$ . A sub-path,  $p_i[n : m]$ , is the  $n$  through  $m$ 'th hops of  $\bar{p}_i$ , and includes the case of a full path<sup>1</sup>. The size of sub-path  $|p_i[n : m]|$  is  $m - n + 1$ , i.e. the number of packets to probe hops  $n$  through  $m$ . The ISC problem is to find the set of sub-paths from  $\mathbf{P}$  with minimum total size among all subsets of paths covering  $\mathbf{I}$ . We thus contrast ISC with the full set cover problem that finds a covering set of paths  $\bar{p}_i$  from  $\mathbf{P}$  of minimum size.

We observe that a tension exists between the two conflicting goals of reducing probing traffic and capturing dynamic forwarding paths. Many networks deploy traffic engineering and load balancing. Thus, regardless of whether only full paths or sub-paths are used, we expect that probes will reveal deviations from the prior cycle. When this occurs, we augment ISC with a “change driven” logic: during the interface verification phase, if an interface other than expected is found, ISC begins a DoubleTree-like strategy probing outward in both directions from the unexpected interface. This allows ISC to not only learn of load balancing over multiple cycles, but also adapt to underlying topological changes.

### 3.3 Vantage Point Spreading

Internet-scale mapping involves probing from dozens of different vantage points (VPs). How to divide the probing among VPs presents another opportunity for an adaptive strategy to reduce the probing traffic. However, as our preceding analysis in §2 shows, additional vantage points yield more interface information. Further, in the next section we find that the information gain of additional VPs to the same destination decays slowly. Therefore, due to large value in any additional vantage points, we adopt a simple strategy for assigning destinations to VPs.

Our *vantage point spreading* algorithm simply uses as many distinct VPs as possible for the set of destinations within a given BGP prefix. When combined with *subnet centric probing*, as additional destinations are chosen from determined subnet prefixes, *vantage point spreading* will assign them, if possible, to VPs not yet used for the original BGP prefix, or otherwise distribute them as uniformly as possible when there are more destinations to be probed than VPs.

<sup>1</sup>In practice then, a trace sub-path has the same origin and destination IP addresses, but uses  $TTL=n$  to  $m$ .

## 4. RESULTS

Having defined our three intelligent topology measurement primitives, this section examines their individual performance through a series of Ark experiments.

### 4.1 Subnet Centric Probing

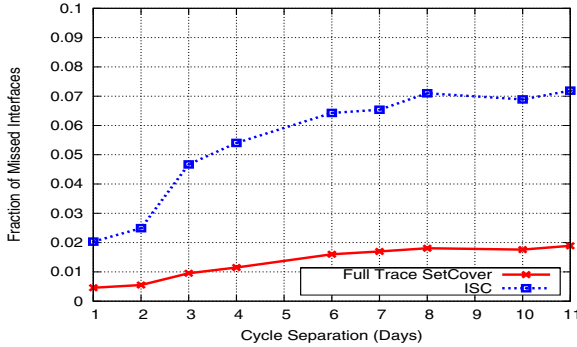
We evaluate the subnet centric probing strategy against a full cycle of Ark probing, which for the sake of this exercise we will consider to be the known ground-truth. Note that this ground-truth is a relative measure, rather than the actual topology which remains unknown. We simulate various strategies by filtering the full Ark probe data from a cycle, i.e. we simulate different resulting topologies by selectively using different available Ark paths. Our performance measures examine the balance between probing load and the topological structure resulting from the probes.

To gain intuition over using external BGP data to drive probe selection, we first follow a naïve strategy similar to [14]. We select a single destination per BGP prefix at random from the available Ark traces; effectively assuming that this single destination is representative of the entire cluster of destinations that fall within its prefix. Similarly, we experiment with an even coarser technique whereby trace destinations are clustered according to their AS and a single destination in the AS is deemed representative of the AS.

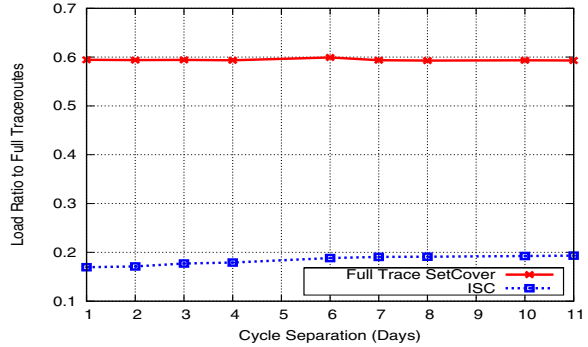
We build the interface-level graph as inferred by the raw Ark data as well as using a single destination per prefix and a single destination per AS. The degree distribution of the resulting inferred graphs is given Figure 3(a). While both naïve strategies capture a structure that appears similar to ground-truth from the full Ark data, there are large numbers of missing interfaces and edges (40-80% of total as shown in Figure 3(b)). However, as shown in Figure 3(c), the prefix clustering method requires approximately one-fifth of the full probing load while the AS clustering method results in even greater probe savings.

Armed with this intuition over the potential probe savings, we ask whether the subnet centric probing algorithm can strike a better balance between consumed probing load and the fidelity of the resulting topology. We observe no qualitative difference<sup>2</sup> in the topology resulting from the

<sup>2</sup>We omit other graph-theoretic measures [24, 16] for brevity, but note that such metrics show similarity to ground truth.



(a) Set cover performance (model interface fidelity relative to traceroute baseline).



(b) Set cover efficiency (induced packet load relative to traceroute baseline).

Figure 4: Comparing full trace topology set covering and ISC techniques over time.

subnet centric approach versus the ground-truth in Figure 3(a). The subnet centric algorithm is able to capture  $\geq 90\%$  of the ground truth vertices and edges while using less than 60% of the ground truth full probing load.

## 4.2 Interface Set Cover

Next, we examine the performance of the Interface Set Cover algorithm, but excluding the “change driven” logic (§3.2). In particular, we are concerned with how the performance of ISC compares with full trace set cover, the degradation of performance over time as the topology changes, and comparative load metrics.

We select 20,000 routed IP destinations at random for these experiments. Each day over a two-week period, we probe the same set of destinations from the same vantage point. The results from the first probing cycle are used to “train” the full set cover and ISC. Figure 4(a) show the fraction of missing interfaces using each set cover technique relative to the interfaces discovered from the full set of traces in that cycle.

We see that after a single day, the full trace set cover misses less than 1% of the interfaces while ISC misses approximately 2%. However, while the full trace set cover uses approximately 60% of the ground-truth probing load, ISC uses less than 20% – a huge savings. Note that for this comparison, we omit consideration of the last hop, the destination. If the destination were included and given that just one vantage point traces to a given destination (as is the case with Ark), then the full trace set cover yields no savings.

The performance of both set cover techniques degrades over time, with ISC degrading faster to 7% interfaces missing relative to ground-truth after 11 cycles. Thus, while set cover techniques can provide a significant savings in probe traffic, they alone do not suffice, as the topology changes over time. Thus, we augment ISC with “change driven” logic. Our expectation, to be tested in future experiments, is that the substantial additional savings in probe traffic with ISC, as compared with full traces, will dominate the amount of additional probing stimulated by the discovered deviations (new and absent interfaces) from the prior cycle.

## 4.3 Vantage Point Influence

To gain intuition over how to assign destinations to vantage points, we first perform a tightly controlled experiment where 2000 randomly selected destinations were each probed from 38 different vantage points. We wish to understand whether adding additional vantage points to probe the same destination increases the discovered topology, and at what point the gain in adding additional vantage points (VPs) diminishes. Figure 5(a) shows the average number of discovered interfaces for each probed destination as a function of the number of vantage points. In addition, the standard deviation error bars shows that the variance in discovered interfaces increases as the number of probing vantage points increases. We find that up to approximately ten vantage points, the number of discovered interfaces is linear, after which the influence of additional vantage points decreases. Yet, the decrease is quite slow – suggesting again that there is significant value in each additional vantage point. This finding contrasts with earlier results [2], suggesting that AS-level peering and interconnections have become richer [6].

Next, we examine vantage point spreading in the context of two other strategies: “random” which models Ark’s current methodology and “single” which uses a single VP to probe all /24’s within a prefix. Figures 5(b) and 5(c) show the number of vertices and edges in the inferred topology using each strategy. As expected, the “single” strategy performs poorly. And while the “random” assignment strategy performs well, we achieve approximately 6% gain in leveraging network knowledge via our VP spreading algorithm.

A reasonable goal for our primitives is substantial savings in probing traffic while attaining as rich or almost as rich interface topology. For the latter, if we consider the criterion of being within 1% the number of discovered interfaces as with full traces, then the above 6% gain in interfaces is well within the scope of concern.

Analytically, for random assignment of /24’s to VPs, and for a prefix with a mask of  $m$  ( $k = 2^{24-m}$  /24’s in the prefix), and for  $N$  vantage points, where  $k \leq N$ , then the probability that all  $k$  /24’s are probed by a unique VP is:

$$P = \prod_{i=0}^{k-1} \frac{N-i}{N}$$



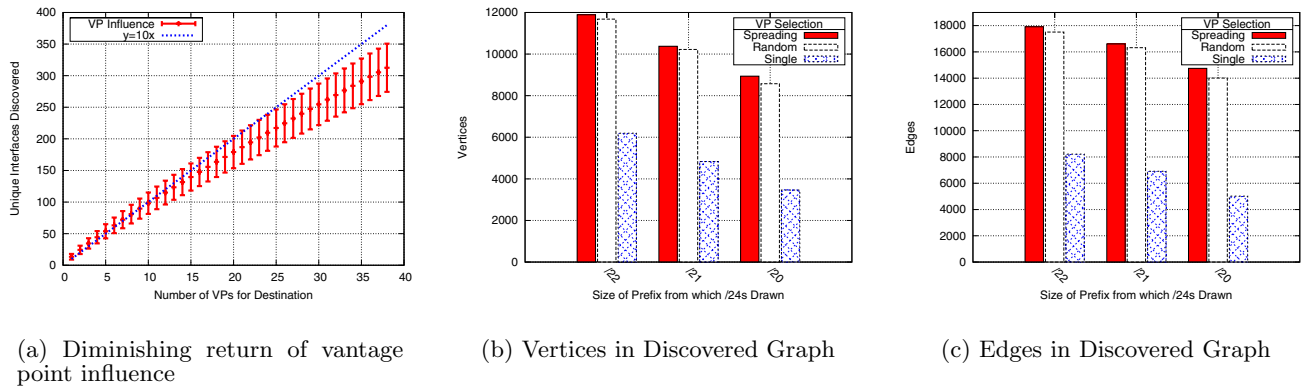


Figure 5: Vantage Point spreading algorithm performance.

Given 23 vantage points, there is only a 25% chance that the 8 /24’s in a /21 would be assigned to distinct vantage points. The chance for the 16 /24’s in a /20 is 0.1%. Empirically, our experiments find, on average, each of the 16 /24’s in a /20 prefix is hit by approximately 12 unique VPs, when performing assignment at random. In summary, vantage point spreading is simple and imposes no additional probing load, and yet the resulting use of additional vantage points attains worthwhile improvement in the estimated topology.

## 5. DISCUSSION

Reducing the number of measurements required to infer network topologies has been explored in the past, notably in DoubleTree [8]. However, our primitives are the first to exploit structural knowledge of the network to reduce measurement cost, while the ISC algorithm is the logical extension of DoubleTree to the multiple-round tracing scenario.

Prior work [12] examines using externally generated and collected synthetic network coordinates to iteratively select probe destinations where the topological distance is most different from the inferred euclidean distance. While their ultimate goal of reducing measurement cost is the same as ours, their problem formulation entails constructing efficient overlay topologies among a known set of nodes by inferring their underlay connectivity. In contrast, we leverage external network knowledge to guide the selection of destinations for topology characterization of an entire AS.

In the big picture, we view the preceding techniques as important building blocks for a new generation of “Internet-scopes” capable of performing one complete round of probing within a day. With the substantial load savings of these primitives, our hope is to utilize the resulting probing budget gain to more completely characterize the Internet – capturing small-scale dynamics and previously hidden structure.

One challenge in combining these primitives into a single system design is that the ISC technique, by nature, has its search space constrained by historical views. To capture the changes in Internet topology, the supplemental “change driven” logic needs to be integrated into ISC and will likely need further refinement.

We also note the complimentary interaction between subnet-centric probing and vantage point spreading. In isolation, VP spreading probes discover the network ingress points while subnet centric probing finds internal network subnet-

ting structure. Used together, however, both goals can be accomplished without exhausting probing budgets. Subnet centric probing is used for stub networks that have a limited number of ingress points whereas vantage point spreading is designed for exploring path diversity of transit networks that have many peering points but not many internal subnets. In other words, we do not need to perform subnet centric probing per vantage point; we can use the same set of probes to accomplish both objectives, by independently choosing their source and destination addresses.

Our abstraction of the narrow waist in Figure 2, and its impact on topology measurement strategy and vantage point selection, is less relevant for core networks. A top-tier network peers with the other top-tiers, in multiple cities, and provides transit for its many downstream networks. Since these interconnections often occur at inter-exchange points, the number of border router interfaces of a top-tier network, though more than for an edge network, is less than the number of its connections to other AS’s. Thus, discovering the topology of a core network, for which additional vantage points is key, has less opportunity for reduction in probing than does edge networks. We intend to quantify the extent of probe reduction possible in measuring core topologies in future work.

Finally, this paper only targets an interface-level graph. An additional alias resolution [3] step, with more probing, is required to reduce an interface-level graph to a router-level graph. We leave the question of how to efficiently perform alias resolution to future work.

## Acknowledgments

The authors would like to thank Young Hyun, k. claffy and CAIDA for measurement infrastructure support, Joel Young, Nick Feamster and Steve Bauer for early feedback, and the anonymous reviewers for their constructive comments. This research was partially supported by the NSF under grants ANI-0520210 and CNS-0721574. Views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of NSF or the U.S. government.

## 6. REFERENCES

- [1] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira. Avoiding traceroute anomalies with paris traceroute. In *Proceedings of the 6th ACM Conference on Internet measurement*, 2006.
- [2] P. Barford, A. Bestavros, J. Byers, and M. Crovella. On the marginal utility of network topology measurements. In *Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement*, 2001.
- [3] A. Bender, R. Sherwood, and N. Spring. Fixing ally's growing pains with velocity modeling. In *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement'* 2008.
- [4] R. Beverly and K. Sollins. An internet protocol address clustering algorithm. In *Proceedings of the 3rd Tackling Computer Systems Problems with Machine Learning Techniques Workshop*, Dec. 2008.
- [5] D. D. Clark. The design philosophy of the DARPA internet protocols. In *Proceedings of ACM SIGCOMM*, pages 102–111, 1988.
- [6] A. Dhamdhere and C. Dovrolis. Ten years in the evolution of the internet ecosystem. In *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*, 2008.
- [7] B. Donnet, T. Friedman, and M. Crovella. Improved algorithms for network topology discovery. In *Proceedings of the 6th Passive and Active Measurement Conference'* volume 3431, 2005.
- [8] B. Donnet, P. Raoult, T. Friedman, and M. Crovella. Efficient algorithms for large-scale topology discovery. In *Proceedings of the 2005 ACM SIGMETRICS*, 2005.
- [9] U. Feige. A threshold of  $\ln n$  for approximating set cover. *Journal of the ACM*, 45(4), July 1998.
- [10] M. Gonen and Y. Shavitt. An  $o(\log n)$ -approximation for the set cover problem with set ownership. *Inf. Process. Lett.*, 109(3), 2009.
- [11] Y. Hyun and k. claffy. Archipelago measurement infrastructure, 2009. <http://www.caida.org/projects/ark/>.
- [12] X. Jin, W.-P. Yiu, S.-H. CHan, and Y. Wang. Network topology inference based on end-to-end measurements. *IEEE Selected Areas in Communications*, 24(12), Dec. 2006.
- [13] k. claffy, Y. Hyun, K. Keys, and M. Fomenkov. Internet mapping: from art to science. In *Proceedings of IEEE Cybersecurity Applications and Technologies Conference for Homeland Security*, Mar. 2009.
- [14] B. Krishnamurthy and J. Wang. On network-aware clustering of web clients. In *ACM SIGCOMM*, pages 97–110, 2000.
- [15] B. Krishnamurthy and W. Willinger. What are our standards for validation of measurement-based networking research? *SIGMETRICS Perform. Eval. Rev.*, 36(2), 2008.
- [16] L. Li, D. Alderson, W. Willinger, and J. Doyle. A first-principles approach to understanding the internet's router-level topology. In *Proceedings of the 2004 ACM SIGCOMM conference*, 2004.
- [17] M. Luckie, Y. Hyun, and B. Huffaker. Traceroute probe method and forward IP path inference. In *Proceedings of the ACM Internet Measurement Conference*, 2008.
- [18] H. Madhyastha. Private communication with iplane author, Mar. 2010.
- [19] 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 USENIX OSDI*, Nov. 2006.
- [20] 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 NSDI*, 2009.
- [21] D. Meyer. University of Oregon RouteViews, 2010. <http://www.routeviews.org>.
- [22] Y. Shavitt and E. Shir. Dimes: let the internet measure itself. *SIGCOMM Comput. Commun. Rev.*, 35(5), 2005.
- [23] Y. Shavitt and U. Weinsberg. Quantifying the importance of vantage points distribution in internet topology measurements. In *Proceedings of IEEE INFOCOM*, Mar. 2009.
- [24] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Network topology generators: degree-based vs. structural. In *Proceedings of the 2002 SIGCOMM Conference*, 2002.