Layer 1-Informed Internet Topology Measurement

Ramakrishnan Durairajan
University of Wisconsin
rkrish@cs.wisc.edu

Joel Sommers
Colgate University
jsommers@colgate.edu

Paul Barford
University of Wisconsin
pb@cs.wisc.edu

ABSTRACT

Understanding the Internet's topological structure continues to be fraught with challenges. In this paper, we investigate the hypothesis that physical maps of service provider infrastructure can be used to effectively guide topology discovery based on network layer TTL-limited measurement. The goal of our work is to focus layer 3-based probing on broadly identifying *Internet infrastructure that has a fixed geographic location* such as POPs, IXPs and other kinds of hosting facilities. We begin by comparing more than 1.5 years of TTL-limited probe data from the Ark [25] project with maps of service provider infrastructure from the Internet Atlas [15] project. We find that there are substantially more nodes and links identified in the service provider map data versus the probe data. Next, we describe a new method for probe-based measurement of physical infrastructure called *POPsicle* that is based on careful selection of probe source-destination pairs. We demonstrate the capability of our method through an extensive measurement study using existing "looking glass" vantage points distributed throughout the Internet and show that it reveals 2.4 times more physical node locations versus standard probing methods. To demonstrate the deployability of POPsicle we also conduct tests at an IXP. Our results again show that POPsicle can identify more physical node locations compared with standard layer 3 probes, and through this deployment approach it can be used to measure thousands of networks world wide.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Network topology; C.2.3 [**Network Operations**]: Public networks

General Terms

Algorithms, Design, Measurement

Keywords

Physical Internet, POPsicle probing heuristic

IMC'14, November 5–7, 2014, Vancouver, BC, Canada.

http://dx.doi.org/10.1145/2663716.2663737.

1. INTRODUCTION

Studies that aim to map the Internet's topological structure have been motivated for many years by a number of compelling applications including the possibilities of improving performance, security and robustness (e.g., [19]). While these motivations remain as compelling as ever, the ability to accurately and comprehensively map the Internet has, for the most part, remained beyond our grasp.

The primary challenges to thoroughly mapping the Internet stem from its enormous size, distributed ownership, and constantly changing characteristics. Faced with these challenges, the most widely used approach to Internet mapping has been based on gathering data from network-layer measurements using TTL-limited probes[1]. Great progress has been made on solving some of the specific problems related to using these network layer measurements for understanding Internet topological characteristics. However, the fact remains that layer 3 data are inherently tied to the management policies and operational objectives of service providers, which may be at odds with comprehensive and accurate mapping of the Internet.

So, just what do we mean by an "Internet map"? At the lowest level, the Internet is composed of physical conduits that contain bundles of optical fiber, and that terminate at buildings that house routing and switching equipment. We refer to the collection of these data and their geographic locations as "physical maps" of the Internet. Several recent projects have begun to assemble repositories of physical Internet maps [15, 29]. These maps are valuable because they reflect a ground truth perspective of service provider infrastructure. These are in contrast with maps that have been generated based on layer 3 probes (*e.g.,* [45]), which we refer to as "network-layer maps". Ideally, network-layer maps reflect a timely representation of network topology as well as the dynamic aspects of management and configurations.

In this paper we investigate the hypothesis that physical maps can be used to guide and reinforce the process of collecting layer 3 probe data toward the goal of expanding the scope of physical infrastructure captured in network-layer maps. This conjecture leads directly to two key research questions that are the focus of our work: *(i)* how do physical layer maps compare and contrast with network-layer maps? and *(ii)* how can probe methods used by projects like Ark [25] be improved to reveal a larger portion of physical infrastructure? We contend that some of the challenges

[1]Maps can also be created using BGP updates or application-layer data, however those are not the focus of this paper.

inherent in generating maps from layer 3 probes can be overcome by using the constructive approach of first identifying key infrastructure (POPs, etc.) and then identifying nodes (identified by disambiguating IP addresses or using DNS names) that reside in those locations.

Our study begins by considering physical map data from the Internet Atlas (or Atlas) project and network-layer map data from the Ark project. We focus specifically on infrastructure in North America. The Atlas repository includes data from 78 Internet service providers with over 2600 nodes and over 3580 links. Nodes in the Atlas data refer to hosting centers or points of presence (POPs), with links referring to physical connections between those locations. We use Ark measurements collected from September 2011 to March 2013 (approximately the same period over which the Atlas repository has been assembled). We resolve the IP addresses from this corpus to DNS names and then use location hints to associate these with physical locations (*e.g.*, cities), which becomes the basis for our comparisons.

Several characteristics are immediately evident in the data. Most prominent is the fact that among the 50 networks that are the focus of our comparison study, we observe many more nodes and links in the physical maps. There can be a number of explanations for this observation, including *(i)* the limitations of exploiting DNS naming conventions, *(ii)* the use of tunneling protocols (*e.g.*, MPLS) or the lack of layer 3 services which can render nodes invisible to probes, *(iii)* the limited perspective of the network mapping infrastructure and *(iv)* the fact that layer 3 routing configurations may simply obviate the ability to observe all networks, nodes and links. This supposition is supported by the observation that all Ark probes are confined to a minority subset of networks, with the majority of probes traversing an even smaller subset of networks. Despite this, there are still some nodes/locations/links that appear in the network-layer map but are not indicated in the physical map. This can be explained by physical maps that are out of date or are either intentionally or erroneously incomplete.

The differences between the physical and network-layer maps suggests opportunities for *reinforcement* between the data sets. First, networks observed in Ark that do not appear in Atlas offer clues for searching for new maps that would expand the repository. Second, nodes or links in Atlas that do not appear in Ark can become targets for additional probing that could expand the scope of resulting network-layer maps, thereby making them more useful in target applications. We focus specifically on the possibility of identifying new nodes in layer 3 measurements through targeted probing in the second component of our study.

We define the *targeting problem* as identifying source-destination pairs for layer 3 probes that reveal nodes indicated in the physical maps[2]. Probing sources (or Vantage Points—VPs) are publicly available infrastructure such as looking glass and traceroute servers and PlanetLab nodes from which probes can be sent. Destinations are simply IP addresses that may respond to probes. We began our targeting analysis by identifying a subset of 596 POPs from the physical maps across 25 networks as our target set. We then conducted extensive probe-based measurements using 266 unique sources and 742 destination addresses in the target networks using two core ideas: *(i)* source-destination pairs should be proximal to the target geographically and in address space, and *(ii)* verification of measurements using multiple sources is required. We verify the identification of infrastructure using location hints in DNS names and using records available in PeeringDB [5]. Our analysis shows that probing between sources and destinations that are both *within the same autonomous system as the target(s)* reveals the most physical infrastructure.

The results of our targeting experiments motivate a new heuristic algorithm for probe targeting that we call *POPsicle*. We show that POPsicle finds 2.4 times as many nodes as are identified by Ark. We compare the number of POPs found by POPsicle with POPs found using Rocketfuel [45] and in all cases POPsicle performs better. We also found that IXPs play a critical role in the way probes traverse a given network. Specifically, sources that are co-located with IXPs have the advantage of appearing—from a layer 3 perspective—as being internal to any/all of the networks that are connected at that location. Thus, a single source that is co-located within an IXP may enhance the identification of infrastructure across all networks that connect to the IXP. This has the effect of significantly broadening the scope of the infrastructure that can be identified using our approach. To validate this idea, we deployed POPsicle at the Equinix IXP in Chicago, USA, and measured the number of POPs for 10 ISPs and found that POPsicle reveals almost all POPs compared to Atlas and extra POPs (in certain cases) compared to Ark. We also find through a case study of Cogent network that POPsicle identifies over 90% of the nodes identified in Atlas or by the reverse DNS technique of [21], compared with about 65% of the POPs identified through Ark, and only 25% identified in the most recently available Rocketfuel data.

To summarize, the key contributions of our work are as follows. First, we perform a first-of-its-kind comparison of large repositories of physical and network maps and find that physical maps typically reveal a much larger number of nodes (*e.g.,* POPs and hosting infrastructure). Next, we consider the targeting problem and find that using sources and destinations within the same autonomous system for probing reveals the most physical infrastructure. We develop a layer 1-informed heuristic algorithm for probe source-destination selection called POPsicle that identifies 2.4 times as many nodes as standard probing methods. Finally, we identify the fact that sources co-located as IXPs can be used to amplify POPsicle-based probing broadly throughout the Internet resulting in layer 3 maps that can be more effectively applied to problems of interest. To that end, we deployed our method at a real IXP and found that our method finds almost all POPs compared to Atlas and additional POPs compared to Ark for the ISPs studied.

2. RELATED WORK

Creating maps of the Internet's topology has been of interest to the research community since the early days of the Internet, and its predecessor, the ARPAnet [37]. Just as maintaining a `hosts.txt` file was feasible in the Internet's infancy, so was the capability of identifying *all* nodes and links in the network [32]. After the privatization and commercialization of the Internet, it became well accepted and understood that the Internet's rapid growth implied that the cataloging efforts of earlier years were no longer possible.

[2] *Efficient targeting* is a related problem that seeks to identify infrastructure with a minimal number of probes. We do not directly consider minimizing probe budget in this study.

Since then, there has been a great deal of effort made to harness layer 3 TTL-limited probes for network mapping since the introduction of the `traceroute` tool [26]. Some efforts (*e.g.*, [40, 45]) have focused on the goal of developing a comprehensive network-layer view of the Internet *i.e.,* unique identification of nodes and links. Other efforts have focused on developing new probing techniques that expand the ability to collect data and thereby improve accuracy and mapping coverage, *e.g.*, [9, 10, 41]. More recent efforts have focused on analyzing and addressing various inaccuracies inherent in probe-based network mapping [43,49]. For example, Roughan, *et al.* and Eriksson, *et al.* develop inferential techniques to quantify the nodes and links that are missed through network-layer mapping [18, 38]. Other researchers have looked closely at the rise of Internet Exchange Points (IXPs) and the effects of IXPs on inaccuracies of network-layer mapping, *e.g.*, [8,10]. Concurrent with the rise of IXPs has come a "flattening" of the Internet's peering structure [13,22,30], which affects the very nature of end-to-end paths through the Internet. Still other researchers have observed that increased use of network virtualization techniques such as MPLS have led to additional inaccuracies in layer 3 mapping, and which are likely to continue to thwart probe-based mapping efforts [14,40,42]. We posit that layer 3 mapping efforts will continue to be important sources of Internet topology information and that complementary efforts to build repositories of physical Internet maps (*e.g.,* [15,29]) will result in representations of Internet topology that are more accurate and applicable to problems of interest than either representation in isolation.

The targeting problem that is a focus of our work is informed by prior studies that analyze the intrinsic importance of measurement infrastructure in Internet topology mapping. Barford *et al.* were among the first to quantify the value of vantage points in discovery of nodes and links in core and edges of the Internet [11]. More recently, Shavitt and Weinsberg consider the problem of bias in measurements based on vantage point distributions and show that a broad distribution of vantage points reduces bias in resulting maps [39]. Our work differs from these studies in that we are focused on using layer 3 probes to identify specific infrastructure targets.

Identifying the geographic location of nodes that have been assigned specific IP addresses (*i.e., IP geolocation*) is a challenging problem that is highly relevant to our study. Some of the earliest work on this problem was done by Paxson, who developed the idea of using DNS hints to identify the geographic locations of nodes that were responding to TTL-limited probes [35]. We use similar methods in our study. Since then, many studies have addressed the problem of IP and POP geolocation using a variety of measurement techniques (*e.g.,* [17, 20, 23, 28, 34, 36, 46, 47]). The fact that POP locations in physical maps are often given at the street address level offers the possibility to improve IP geolocation estimates using standard measurement-based methods. We are also investigating another possibility of leveraging state-of-the-art geolocation techniques (*e.g.,* [24, 46]) to enhance the accuracy of our location extraction approach.

3. DATASETS

In this section we describe the datasets used in our study. One of the key contributions of our work is the comparison of physical topology data from primary sources and network-layer topology data extracted from layer 3 TTL-limited probes, as described below. In the case of physical infrastructure data, we use the latest maps from service providers collected as part of Internet Atlas project [15]. For network-layer topology data, we rely on traceroute data collected as part of CAIDA's Archipelago (Ark) project [25].

3.1 Physical Topology Data

In this study, we rely on the publicly available physical topology data from the Internet Atlas project. Internet Atlas is a visualization and analysis portal for diverse Internet measurement data. The starting point for Internet Atlas is a geographically anchored representation of the physical Internet including *(i)* nodes (*e.g.*, hosting facilities and data centers), *(ii)* conduits/links that connect these nodes, and *(iii)* relevant meta data (*e.g.*, source provenance). The approach for building the repository is to use targeted search to identify maps and other listings of physical Internet infrastructure that are published on the web by ISPs. Though there is no guarantee as to their timeliness or completeness, we use this data as ground truth of service provider infrastructure in this study. Since these data come from primary sources, they also reflect a ground truth perspective of service provider infrastructure. The current repository contains geocoded physical infrastructure data of over 425 ISPs around the world. From this online repository, we obtain detailed geographic information of 7 Tier-1 networks and 71 non-Tier-1/regional networks with a presence in North America consisting of 2611 POPs and 3588 links.

3.2 Network-layer Topology Data

We seek to improve the state-of-the-art in Internet topology mapping by investigating structural characteristics revealed by layer 3 probes. Our goal is to broaden the understanding of Internet topology by investigating how topological characteristics as revealed by layer 3 traceroute probes compare to and contrast with physical structure derived from service provider maps.

The network-layer probe data that we use are collected as part of the Ark project, and include traceroute measurements from a set of 77 monitoring systems distributed around the globe to all routed /24 prefixes in the IPv4 Internet. We used the traceroute data gathered by the Ark project since it represents a canonical system for large-scale Internet topology measurement. We note, however, that the measurements collected in Ark are subject to a variety of network management policies, including blocking or limiting responses to TTL-limited probes, routing configurations and MPLS tunnels, each of which can limit the scope of the measurement data.

Ark is the canonical example of what we might call a *generalized topology probing system.* POPsicle, on the other hand, has a specific goal, which is to discover unique nodes based on guidance from physical maps. Given the difference in goals, the comparisons of the number of unique nodes identified by Ark vs. POPsicle should be interpreted as a comparison between a generalized and a purpose-built system, *i.e.*, POPsicle can be implemented as an extension to Ark or as the basis for designing an entirely new *coordinated large-scale traceroute-based topology measurement system.*

3.3 DNS Data

The DNS data we use are also collected as part of the IPv4 Routed /24 DNS Names Dataset [4], and provide fully-qualified domain names for IP addresses seen in the Ark traces. In this work, our consideration of network-layer topology data is limited by the scope and placement of Ark monitors. However, that project has taken pains to include a broad spectrum of network types (*e.g.*, research, commercial, and educational networks) as vantage points for their monitoring systems, and it provides a widely-used view of the Internet's topology.

Leveraging *location hints* present in domain names to classify IP addresses into POPs is fraught with challenges as described in Section 4. We believe that the accuracy of our results could be improved further either with better techniques of handling DNS naming hints, *e.g.,* using the techniques of Huffaker *et al.* [24], or by using non-DNS-based techniques to classify IP addresses to their corresponding POPs [44].

3.4 Scope of Comparison Study

In this study, we restrict our analysis of Ark data to a period of 19 months, from September 2011 to March 2013, which is contemporaneous with data collection in Atlas. Our focus is on understanding the composite views of networks offered by both data sets over this period. From each individual traceroute in the source data, we extract all the *internal* network IP addresses and links. That is, after processing each traceroute, there is a corresponding interface list (*e.g.*, IP1, IP2, IP3, IP4, *etc.*) and link list (*e.g.*, IP1-IP2, IP3-IP4, *etc.*). For instance, if the traceroute contains a probe of the form A-B-C-D-E (where A, B, C, D and E are IP addresses), we ignore the end point IP addresses (A and E) and extract only the network IP addresses (B, C and D). The interface list thus contains IP addresses B, C and D, and the link list contains B-C and C-D. We merge all the interface/link lists after removing all the duplicate entries to produce a final list of interface IP addresses and links. We then use the corresponding DNS dataset and join the list of interface IPs to their corresponding DNS entries.

4. DATA ANALYSIS

In this section, we describe the methods we use to analyze the network-layer topology data. We begin with discussing results from processing the network-layer data followed by associating geographic locations to the network-layer data.

4.1 Network-layer Data Analysis

In this section, we describe the two-step mapping algorithm that we use to associate a physical location to the IP address interface list obtained from processing network-layer traceroute data, as described above.

Key Idea. One of the aspects of this algorithm is to translate location-based patterns in DNS names that refer to router interfaces to physical (geographic) locations. The influential topology mapping work of Spring *et al.* [45] used such "hints" in their `undns` tool as part of the Rocketfuel project in order to infer locations of network POPs. Many network service providers employ naming conventions that include geographically relevant information such as airport codes, city names, or other location information. By exploiting these conventions and developing rules to infer geographic locations from them, we can build a network-layer topology map.

Challenges. Leveraging naming conventions in DNS entries has two important challenges. The first is that these names may be out-of-date or misconfigured, which would lead to invalid geographic inferences. The work of Zhang *et al.* [48] quantified the prevalence such problems and found them to occur infrequently, but to have potentially large impact on topology mapping studies that rely on DNS information. They developed a set of heuristics to avoid such problems, including the detection of POP-level loops within a single provider (which should not occur, assuming that the ISP's intra-domain routing protocols are configured properly). We also use such techniques in our work to avoid problems with exploiting DNS naming conventions. A second challenge with using DNS entries is that there are inherent ambiguities associated with them, *e.g.*, a single string may be used by two different ISPs to refer to two *different* physical locations. To cope with these problems, a set of regular expression patterns can encode different rules to disambiguate location hints from different providers. This approach was also taken in the earlier `undns` tool [45]. Table 1 shows several example patterns and how they are used to resolve ambiguities in naming.

Algorithm. The algorithm for developing a network-layer map from raw traceroute data takes four inputs:

- *regular expression patterns* to extract the location code from DNS entries. The location code is that part of the hostname that contains location data. For example, for `A.B.C.LAX2.D.NET`, the location code is LAX, which is the airport code for Los Angeles, CA, USA;
- *mapping codes* [12] to translate location code obtained from DNS entries to physical location (a latitude/longitude pair);
- the *list of nodes* (along with each corresponding DNS entry) obtained by parsing the traceroute data from Ark as described above;
- and the *list of links* obtained by parsing the traceroute data from Ark, also as described above.

Using these inputs, we associate physical locations to the IP addresses in the interface list using the following steps:

- First, we match the domain names against the regular expression location patterns and extract a location code from every entry.
- Next, we translate the location code to an actual physical location using the mapping codes. The result of this second step is that we have location information associated with every interface IP address that has a DNS entry with location hints embedded in it. We also use Team Cymru's IP-to-ASN mapping service [3] to classify the list of nodes and links into different ISPs based on the Autonomous System (AS) Numbers.

At the end of applying this algorithm we have network-layer maps for different autonomous systems in which the nodes refer to geographic locations of POPs, and links refer to the fact that packets can be forwarded between a pair of POPs. Note that we do not consider intra-POP links, or individual routers in POPs. The result is that we have a network-layer map that can be equitably compared with the physical map available from Internet Atlas.

Table 1: Examples of regular expressions used for extracting location hints from DNS entries.

Regular expression	Explanation
/\.(birmingham)\d*\.(level3)\.net$/i	`birmingham` could refer to a city like Birmingham, UK, but means Birmingham, AL, USA to Level3.
/\.(manchester)\d*\.(level3)\.net$/i	`manchester` could refer to a city like Manchester, NH, USA, but it means Manchester, UK to Level3.
/\.(mad)\.(verizon\-gni)\.net$/i	`mad` could refer to a city like Madrid, Spain, but it means Madison, NJ, USA to Verizon.
/\.(ham)\d*\.(alter)\.net$/i	`ham` could refer to a city like Hamburg, Germany, but it means Hamilton, Canada to alter.net
/\.(cam)\-bar\d*\.(ja)\.net$/i	`cam` could refer to a city like Cambridge WI, USA, but it means Cambridge, MA, USA to bbnplanet, and Cambridge, UK to ja.net

4.2 How are POPs in the same city identified?

To identify POPs located within the same city, we leverage three types of information: *(1)* Personal email communication with network operators and administrators who run the ISPs, *(2)* IP address allocation information from publicly available databases like PeeringDB, and *(3)* naming conventions recorded from ISP websites. In what follows, we give a list of examples for all three cases.

- Tinet (now Intelliquent) has multiple POPs at multiple cities. To identify those POP locations, we contacted one of the network operators [16] from Tinet and identified the naming convention followed by them — the first three letters are city code, and next digit is location code. For instance, `ams10` and `ams20` are two different POPs in Amsterdam.
- Another reliable source of information that is frequently updated and maintained by network operators is PeeringDB. Apart from providing the list of peers at a particular facility (or an IXP), PeeringDB also provides information like address space allocation, network operator contacts, etc. For instance, GTT has multiple POP locations in New York. One of them peers at NYIIX and has `198.32.160.0/24` as its address space, and one another POP peers at Coresite NY with `206.51.45.0/24` as its address space.
- ISPs routinely publish their naming conventions in their websites along with inter-city POP details. For instance, Lumos Networks and Atlantic Metro Communications publicly list all inter-city POP naming conventions [6, 7].

4.3 Associating Geographic Locations with Traceroute Data

We first provide details on results from processing the traceroute data used for building network-layer topologies. Over the 19 months of Ark data considered in our study, we identified 14,593,457 unique interface IP addresses, comprising 31,055 unique ASes. On these traceroute measurements, we applied the algorithm described above to construct network-layer topologies for comparison with the physical networks chosen for our study. Table 2 shows several statistics resulting from applying our algorithm.

As shown in Table 2, there were a number of situations in which we could not reliably use the traceroute data for building network-layer topologies. In particular, over 13M IP addresses did not have an associated DNS name with any (obvious) location information embedded in it[3], which represents 95.16% of all IP addresses observed in our data. Of these, over 6M were unusable because of DNS resolution failures, *e.g.*, `fail.non-authoritative.in-addr.arpa`, which represented 40.31% of all IP addresses observed in our data. While these results certainly limit our ability to compare physical and network-layer topologies for *all* networks, the remaining "usable" trace information represents 4,135 separate autonomous systems, which we argue still represents a significant slice of the Internet.

Table 2: Basic results from processing 19 months of Ark traceroute data using the algorithm described in Section 3

Total traceroutes processed	2,674,959,041
Number of unique interface IP addresses	14,593,457
Number of unique ASes	31,055
Valid DNS entries found	6,936,146
No associated DNS name found	7,657,311
DNS entries with location hints	704,935
Number of ASes with at least one geographically identifiable interface address	4,135

An issue we encountered when applying the algorithm of Section 4.1 was that, in some cases, there were no associated AS numbers indicated by the Team Cymru IP-to-AS mapping service or available in other whois databases. For such networks, we used a manual keyword search (*e.g.*, layer42.net refers to the Layer42 ISP), which was effective for subnets with at least one associated DNS entry.

5. COMPARING LAYER 1 MAPS WITH LAYER 3 PROBE DATA

In this section, we analyze the physical and network-layer topology data. We begin with comparing the two views of Internet topology by considering how each view intersects and differs from one another, and also how the two views of network topology reinforce each other. We focus our discussion on 50 regional and national ISPs with footprints in North America. We focus on these particular networks be-

[3]For example, the DNS naming conventions may not be oriented around physical node location and thus be unusable for our purposes, *e.g.*, entries such as `216-19-195-15.getnet.net` and `173-244-236-242.unassigned.ntelos.net`.

cause there is significant detail within the Internet Atlas data regarding POPs and inter-POP links for these ISPs.

5.1 Comparison of Physical and Network-layer Nodes and Links

We now compare the physical and network-layer topologies obtained from the Atlas data and the Ark data, respectively. Again, the basic entities we compare are *nodes*, which represent city-level points of presence or data centers, and *links*, which represent physical and/or logical connectivity between two city-level POPs.

Table 6 in Appendix A shows the number of nodes and links observed in each topology type, for each of the 50 networks under study. We first see that while all physical networks have non-zero nodes and links, there are some network-layer topologies for which there are zero nodes and/or links observed. There are two reasons for this. First, an interface IP address for a given network may have no clear location information embedded in its associated DNS entry. For example, for 21 out of 50 networks, there were no location hints observable in the related DNS records. This result may be because of non-obvious naming conventions, or simply that there are no name records available. We note that although some ISPs in our list of 50 have been acquired by other companies, the AS number and address blocks assigned to these companies still refer to the original ISP [4].

The second reason we may observe zero nodes and/or links for a given network is that we may simply not have observed *any* interface addresses for a given network in 19 months of traceroute data. This observation was true for 16 out of the 50 networks included in our study. Considering the fact that the Ark project targets *every* routable /24 in the IPv4 Internet, this is a surprising result. Still, there may be a variety of reasons for this observation. First, some ISPs may configure their routers not to respond to hop-limited probes with ICMP time exceeded messages (resulting in "stars" in the traceroute output). Second, some networks may use tunneling protocols such as MPLS, and configure these tunnels to be completely hidden. Third, there may be interfaces controlled by an ISP under study that are configured with IP addresses from a third party, *e.g.*, an IXP. In the end, we were left with 13 networks that had DNS entries for which we could identify a physical location.

To assess how the physical and network-layer views of a network compare, we consider node and link *intersection*, as well as the number of nodes and links *only* observed in one or other other topology. To determine the intersection, we consider a node to intersect each topology if we identify the same POP location in each one. We consider a link to intersect each topology if there are POPs identified in the same two locations in each topology and there is a link identified between them. For example, if we observe nodes in Chicago and Kansas City in both the physical and network-layer topologies for a given ISP, and a link between those two cities, we say the link and two nodes intersect.

Table 6 in Appendix A shows results from the intersection analysis. We also show in the table nodes or links that *only* appear in one or the other topology. We see that, in general, there are more nodes and links observed in the physical topologies than are seen in the network-layer topologies. For the networks for which this observation holds, the number of nodes and links observed is, in some cases, *significantly* larger than those seen using the traceroute data. These results strongly suggest that sole reliance on layer 3 probes to generate physical network maps is likely to result in an incomplete view of Internet topology. On the other hand, the table shows that there are a small number of networks in which we observe *more* nodes and links in the network-layer topology. In particular, we see this for AT&T, Tinet, NTT, Sprint, Layer42, and Hurricane Electric (abbreviated as HE in the table). This observation suggests that while published physical maps *usually* offer an authoritative view of physical infrastructure, the published maps may lag recent deployments which can be observed through layer 3 probing.

More broadly, analysis of the Ark traces shows that there are at least 448 distinct networks in North America (that are not part of Atlas). This number is identified by first searching for all North American location DNS hints and then identifying unique service providers in the DNS names. This compares to the 320 distinct networks in the Atlas repository, which have been identified through extensive search-based methods. An implication for this difference is that measurements from Ark can be used as guidance for identifying service provider networks that could be included in Atlas. For example, many small/regional networks like Adera Networks (CA), Grande Communications (TX) and Atala T (NY) were found in the traces of Ark and such networks could be incorporated into future search-based campaigns.

Of the 448 distinct networks identified through Ark measurements, the vast majority of probes pass through tier-1 and major ISPs, as shown in Figure 1. Thus, while it is likely that the POP-level topology of well-connected ISPs can be largely identified through general probing techniques, smaller ISPs are unlikely to be well-mapped. This observation is supported by prior studies on sampling bias in network topology measurements (*e.g.,* [39]). An implication for this observation is that *targeted probing methods* may be necessary to obtain a more comprehensive topological picture of physical Internet infrastructure.

Lastly, we consider one form of *validation* of physical node locations when we observe the same location for nodes in more than one network. We define the metric N_{Index} as the percentage of nodes identified for a given network that have the same physical location as a node in another network. The right-most column of Table 6 in Appendix A shows the N_{Index} for each network. The intuition for why this metric provides some level of validation of the physical location has to do with common industry practices of using co-location facilities and telecom hotels. While this observation may not hold universally, we believe that co-location practices are generally observed for small regional networks in geographically isolated areas since the costs associated with setting up new facilities is high. For instance, larger national ISPs like Layer42, Napnet, Navigata and Netrail show an N_{Index} of 100 (complete overlap with nodes in other networks), whereas smaller regional cariers such as NetworkUSA, RedBestel and Syringa[5] show an N_{Index} less than 20 (mostly their own locations). The combination of

[4]For example, although BellSouth was acquired by AT&T in 2006, the name BellSouth is still referred to in `whois` databases and appears in recent address block usage reports (`http://www.cidr-report.org/as2.0/`).

[5]NetworkUSA is a regional carrier serving Louisiana, RedBestel operates in the Guadalajara region of Mexico, and Syringa is a regional carrier in Idaho.

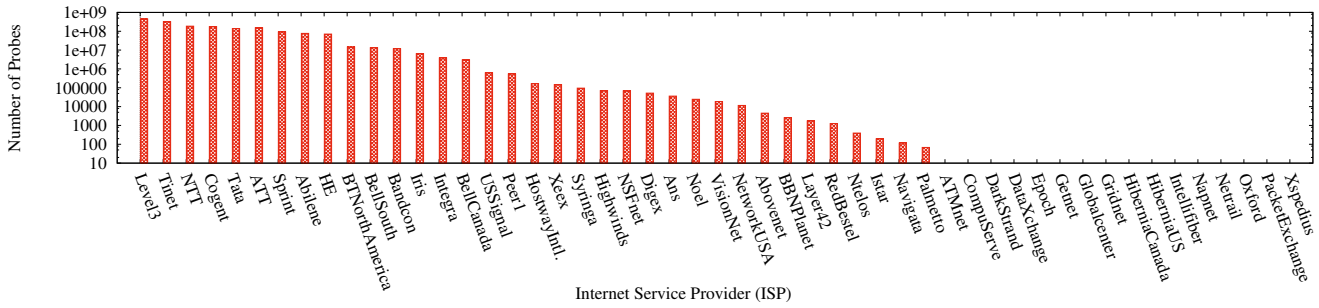

Figure 1: Number of probes sent out by Ark across Internet Service Providers

a high N_{Index} and overlap with traceroute probes provides perhaps the best validation of node locations.

5.2 Case study: Tinet

Tinet represented an interesting special case: the physical topology contained nodes not present in the network-layer topology, and the network-layer topology also contained nodes not present in the physical topology. In particular, there were 65 nodes only present in the physical topology, and 7 nodes that were only observed in the Ark data and network-layer topology. For example, the Tinet physical network map shows four nodes for Amsterdam, Netherlands, one node in San Jose, CA, two nodes in Milan, Italy and two nodes in Washington, DC. However, the network-layer topology revealed additional nodes for these locations. The missing nodes from the physical network may be due to Tinet's network maps not reflecting the most up-to-date deployments. Missing nodes and links in the network-layer view may be due to a variety of reasons, including the inability to gain a broad perspective on Tinet's network from Ark vantage points. What these results indicate is that to gain a *complete* view of a network's topology, multiple data sources must be considered.

5.3 Main findings and implications

The main findings of our comparison of physical and network-layer topologies are as follows.

- We observe many more nodes and links in the physical maps, which may be due to a variety of reasons, but is most critically due to the fact that layer 3 routing configurations simply eliminate the possibility to observe all networks, nodes, and links through end-to-end probing. This likelihood is supported by the fact that all Ark probes are limited to a relatively small subset of networks, with the majority of probes passing through an even smaller set of networks.
- There are still some nodes, locations, and links that appear in the network-layer map but are not observed in physical maps. The likely reason is that the physical maps are out of date or incomplete.
- The observed differences between the physical and network-layer maps suggest opportunities for using one to *reinforce* the other. In particular, networks observed in Ark that do not appear in Atlas offer clues for searching for new maps to expand Atlas. Similarly, nodes or links in Atlas that do not appear in Ark can become targets for additional probing in order to broaden the scope of the resulting network-layer maps.

Indeed, in the next section we focus specifically on how to emit targeted layer 3 probes in order to confirm the existence of nodes identified in physical maps, as well as to identify additional physical nodes.

6. EFFECTS OF VANTAGE POINTS ON NODE IDENTIFICATION

In this section, we examine the effects of source-destination selection on the ability to identify POPs within a service provider using targeted layer 3 probes. Specifically, we examine the differences between using vantage points (probing sources) internal or external to an ISP containing target POP(s), and destinations either internal or external to the ISP. Furthermore, we examine the effects that IXPs may have on probe-based POP identification and how IXP placement may be exploited to aid in node identification by providing a larger set of internal vantage points.

6.1 Effects of vantage point and destination selection

To examine the impact of vantage point and destination IP address selection for identifying all target POPs in an ISP, we leverage publicly available traceroute servers, looking glass servers and Planetlab nodes as VPs[6], and select different combinations of them located within or external to different service providers. In particular, we use three combinations: probing from VPs outside an ISP to destinations inside (denoted VP_{out} to t_{in}), from VPs inside an ISP to destinations outside (denoted VP_{in} to t_{out}) and from VPs inside an ISP and destinations inside (denoted VP_{in} to t_{in}). For each directional modality ($VP_{out} \rightarrow t_{in}$, $VP_{in} \rightarrow t_{out}$, $VP_{in} \rightarrow t_{in}$), we use a greedy approach to identify probe source-destination pairs based on geographic proximity. We choose the VP geographically closest to a target POP, then successively choose from the set of destinations that are also geographically proximal to the target until the target is identified. For instance, a probe from `planet-lab4.wail.wisc.edu` to 184.105.184.158[7] with the aim to identify Hurricane Electric's POP in Los Angeles identified two additional POPs (in Chicago and Denver) in addition to identifying the Los Angeles POP. If we can not identify

[6]We followed principles established in prior work, *e.g.*, [45], to avoid burdening these public servers with excessive load.
[7]`lightower-fiber-networks.gigabitethernet4-10.core1.lax2.he.net`

the POP from a given vantage point, we choose the next closest VP, and so forth (specific details of this method are provided in Section 7).

Using a subset of 25 ISP networks that assign DNS names with location hints and that contain 596 target POPs, we analyze the source-destination combinations. Figure 2 shows the fraction of target POPs discovered by these three probing modalities relative to the number of POPs identified in Atlas. The figure shows clearly that the most effective strategy is to send probes from vantage points located *within* an ISP to destinations that are *also within the ISP* (VP_{in} to t_{in}). We further observe that using a VP located within an ISP is more effective than choosing one external to the ISP. We hypothesize that these differences are due to the effects of interdomain versus intradomain routing on probes. In the case of both VP and destination located within an ISP, there is a greater chance for a diversity of paths to be observed due to ECMP, the fact that more information about shortest paths is available, and the greater degree of flexibility that a service provider has in routing packets within its own infrastructure. In the case of either VP or destination being external to the ISP that contains a target POP, interdomain routing protocol effects come into play, such as hot-potato routing and the forced choice of a single best path.

Lastly, we note that in absolute numbers, we observed a total of 188 POPs using VP_{in} to t_{in}, 157 POPs via VP_{in} to t_{out}, and 93 with VP_{out} to t_{in}. For 11 networks we observed zero POPs. Similar to our earlier observations in which we do not see POPs identified in physical maps, this may be due to MPLS deployments, traffic management policies, or routing policies. We intend to further investigate the reason for the invisibility of POPs in future work.

6.2 Using IXPs to expand perspective

Given the result that the most effective probing strategy for identifying physical infrastructure is to choose source-destination pairs that are within an ISP, it is important to recognize that broad deployment of such targeted measurements is inherently limited by the availability of VPs within provider networks. Indeed, the 266 VPs used in this paper are restricted to 248 separate networks, which is substantially less than the total number of networks identified in North America by Ark in Section 5.

Recent work in [10] has highlighted the enormous amount of layer 2 peering that is taking place at IXPs. This leads us to posit that VPs co-located with IXPs might be leveraged to dramatically expand our ability to identify physical infrastructure. Indeed, there is anecdotal evidence that much of the rapid growth in peering at IXPs is being driven by local and regional ISPs and that Tier-1 ISPs have been slower to connect [27]. This offers a tantalizing opportunity since it is generally the smaller networks that are more difficult to map and those networks often do not deploy looking glass servers that are necessary for mapping physical infrastructure.

To consider this possibility, we begin by looking for VPs that are co-located with IXPs in North America. We find that 14 out of 65 IXPs have co-located VPs. Using PeeringDB [5] we find that the total number of unique ISPs that peer at these 14 IXPs is 642. A comparison between these ISPs and those in with VPs used in our study shows that an additional 625 unique networks could be measured from these 14 IXPs alone. This suggests that deployment of VPs in other IXPs could be the starting point for comprehensive mapping of physical Internet infrastructure.

6.3 Main findings and implications

In summary, we consider how to choose sources and destinations for probing in order to identify POPs within a service provider, as well as to discover new POPs. Specifically, we examine whether it is better to use vantage points (probing sources) internal or external to an ISP containing the target POP(s) and destinations either internal or external to the ISP. Our results show that it is best to choose both source and destination to be *within* the ISP that contains the targeted POP(s), which we hypothesize is largely due to intradomain versus interdomain route selection. Further, we observe that co-locating a probing vantage point at an IXP may be particularly useful in that the VP can effectively appear as being internal to all ISPs that peer at the IXP.

7. ENHANCING NODE IDENTIFICATION

In this section, we build on the observations and experiments of Section 6 to describe a new targeted probing algorithm called *POPsicle*. We evaluate POPsicle's effectiveness for reinforcing and confirming information available in physical maps. We deploy POPsicle at an IXP in Chicago, and describe results of experiments carried out at the IXP.

7.1 POPsicle algorithm

POPsicle is designed to send traceroute-like probes toward a target with a known geographic location based on information from a physical map. The objective is to detect the target *at the network layer*. POPsicle is based on the insight that vantage points co-located with IXPs can be used to launch probes in many different networks, and that probe-based detection of target physical infrastructure is most effective when both VP and destination are located within the same service provider network.

Algorithm 1 shows the key steps of POPsicle. The inputs to the algorithm are (1) the name and address prefix(s) of the ISP within which physical targets are to be identified, (2) the specific list of targets (*e.g.*, POPs) to be identified, including their geographic locations according to physical mapping information, and (3) a list of VPs and their known geographic coordinates. The algorithm proceeds by first scanning the target network to identify which hosts are accessible[8]. This step is performed to collect a set of hosts that can be used as probe destinations. The geographic locations of these hosts are then inferred using DNS location hints. Another option at this step would be to use IP geolocation algorithms or tools. However, the accuracy of these techniques is a subject of ongoing research (*e.g.,* [17]) so we do not use them in POPsicle, but they could be easily incorporated.

Next, POPsicle iterates through the list of target nodes to be identified. For each target, we obtain a list of VPs for initiating probes in step 4, ordered by proximity in Euclidean space (using the Haversine formula [1]) to the target. For each VP, we then select a set of destinations that are also ordered by proximity to the target. Destinations, compiled from a variety of sources like Internet Atlas portal and PeeringDB, are IP addresses of infrastructure, like looking glass

[8]We employ the `nmap` tool for this step with the command line `nmap -Pn -sn` *`prefix`*. Even though `nmap` is considered bad, we only did a passive scan without causing any trouble to ISPs

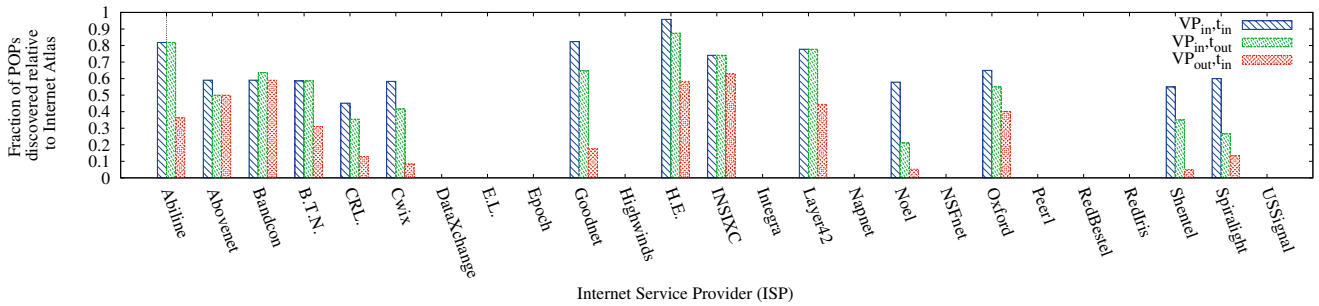

Figure 2: Number of POPs discovered by different probing modalities.

servers, traceroute servers, telecom hotels, and other entities that may simply respond to probes. From this set, we subselect the destinations such that the square of the Euclidean distance between the VP and destination is greater than the sum of the squares of the distance between VP and target and VP and destination. This has the effect of creating a "measurement cone" centered at the VP and directed toward the target node (step 6). These destinations are then iteratively probed using traceroute. For each completed trace we determine whether the target has been found using location hints. If it has, the algorithm completes. If not, we continue until we have exhausted all VPs and their corresponding destination sets. Figure 3 depicts the targeting process of POPsicle.

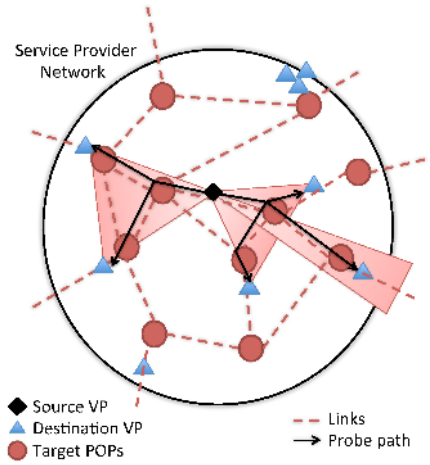

Figure 3: POPsicle targeting process. VPs within the ISP that are geographically closest to the target are selected along with destinations that are geographically closest to the target and "on the other side" of the VP.

POPsicle is based on the notion that target POPs will be part of routes that connect sources and destinations located on either side (from a Euclidean perspective) of the target. We argue that this is likely due to shortest path intra-domain routing. POPsicle is also currently dependent on location hints from DNS for both destination identification and to identify when a target has been discovered. IP geolocation could be used to address the former, while the latter could be addressed through publicly available data (*e.g.,* PeeringDB) by associating IP address ranges with POP locations.

Algorithm 1: POPsicle algorithm

input: $targetNet$ = target network
input: L^T = list of targets to be identified
input: L^S_{vp} = list of source VPs with known coordinates

```
// Scan target network to find reachable hosts
1  scanResults = scan(targetNet);
2  L^D_vp = inferLocations(scanResults);
3  foreach t in L^T do
       // Choose destination VPs that are closest to
          reachable hosts
4      S^t_vp = geographicallyNearest(t, L^S_vp);
5      foreach vp in S^t_vp do
           // Greedily choose probing destinations within
              a cone extending from vp to t
6          D^t_vp = searchCone(vp, t);
7          foreach dst in D^t_vp do
8              send probe from vp to dst;
9              if t found then
10                 record success for t;
11                 goto step 3;
```

7.2 POPsicle Evaluation

We selected 30 looking glass servers from the Atlas database that satisfied the following criteria: (1) the server is co-located with an IXP in North America, and (2) the ground truth information of the POPs is available either from Internet Atlas or in PeeringDB [5]. The vast majority of providers we selected for analysis are regional providers since we found them to be poorly represented in Ark probing results and thus prime candidates for detailed study. In terms of the number of networks used in this study, the coverage of our technique can appear limited. We had to remove several networks from our study due either to the incompleteness of the physical or network maps, or due to the lack of DNS locations hints.

The selection of these 30 looking glass servers resulted in 13 service provider networks that were the focus for our evaluation. We began by examining Internet2, which we consider a special case since complete ground truth for all the layer1, layer2, and layer3 devices is available [2]. POPsicle-directed probing found 10 out of the 10 POPs in Internet2 that house layer 3 infrastructure.

We initiated probing on the remaining set of 12 ISPs using POPsicle-directed probing to verify and map the POPs for

Table 3: Summary results of network POPs identified with POPsicle, Atlas, Ark, and Rocketfuel for POPsicle deployed at publicly accessible looking glass servers.

	POPsicle	Atlas	Ark	Rocketfuel
Abovenet	13	22	13	13
BellCanada	34	48	30	29
Centauri	7	14	3	—
Cyberverse	2	2	2	—
Data102	2	2	2	—
HopOne	4	4	4	—
HE	23	24	23	8
Inerail	3	25	3	—
Internet2	10	10	10	10
Interserver.net	2	2	1	—
Steadfast.net	3	3	3	—
Towardex	7	8	6	—
XO	42	80	42	39

each of those networks. Table 3 shows the results from all of our probing experiments. Overall, for 8 out of 13 ISPs, we see all or almost all of the POPs identified in physical maps. These 8 ISPs include Cyberverse, data102, HopOne, Hurricane Electric (HE), Inerail, Interserver.net, Steadfast.net, and Towardex. For several ISPs, we also observed additional POP locations which we verified using PeeringDB. We also compare with the most recently available measurements from Rocketfuel, which are from experiments carried out in 2008 [40]. Although the Rocketfuel measurements are not especially recent, we note that it is likely that POP deployments are fairly stable. We observe, for example, that POPsicle and Rocketfuel identify the same number of POPs for 3 out of 5 ISPs. Lastly, we note that Rocketfuel data were unavailable for 8 ISPs.

In the following we discuss various special cases and observations related to results for each ISP:

- For BellCanada, POPsicle identified significantly more POPs than were revealed in Ark data. Additional locations identified were in New York, Palo Alto, Seattle, and Woodbridge. We confirmed these locations with Equinix Palo Alto, NYIIX, and SIX exchange points in PeeringDB. The Woodbridge location could not be confirmed in PeeringDB.
- For Centauri Communications, POPsicle identified four additional POP locations in comparison with Ark, including Palo Alto, San Francisco, San Jose, and Sunnyvale. These locations were all confirmed by SFIX and SFMIX in PeeringDB.
- For cyberverse, data102, Steadfast.net, Inerail, Internet2, Hurricane Electric and XO Communications POPsicle identified the same POPs as were observed using the Ark data.
- For HopOne, POPsicle found one extra POP location in Palo Alto (which is not seen in either Ark or physical topology maps), which was confirmed in PeeringDB. POPsicle did *not* observe a node in Mclean, VA, which was seen in the Ark data.
- For Interserver.net, POPsicle identified one additional POP location in New Jersey which is confirmed by Equinix New York IX.
- For Towardex, POPsicle found an extra POP in Boston which is confirmed in PeeringDB (Boston IX).

In addition to mapping POPs of the 13 ISPs described above, we evaluated POPsicle's effectiveness for mapping and confirming additional *infrastructural* nodes that have known/published physical locations. This test set included data centers, DNS root servers, NTP servers (both stratum 1 and stratum 2), and IXPs. Table 4 shows results of these experiments, as well as summary results of the POP-identification experiments. We can see from the table that POPsicle is able to identify network-layer locations for this larger and much more diverse set of devices. In total, it finds 1.04 times more POPs, 1.54 times more data centers, 9 times more DNS servers, over 11 times more NTP servers, and 1.48 times more IXPs[9] (in North America) compared to nodes found by standard end-to-end layer 3 probing campaigns. Overall, POPsicle reveals and confirms 2.4 times more physical node locations versus standard probe-based topology measurement methods.

7.3 IXP deployment of POPsicle

We observe in Section 6 that a VP co-located with an IXP can provide what appears to be an *internal* probing source for *any* ISP that peers at the IXP as depicted in Figure 4. From such a vantage point, a tool implementing the POPsicle algorithm could be employed to map and identify POPs and other nodes of interest in any one of the adjacent ISPs.

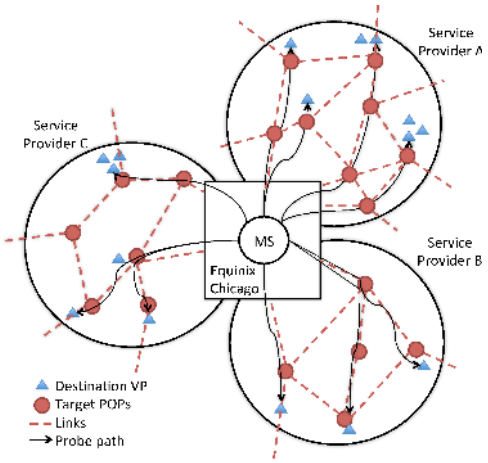

Figure 4: "Multiplexing" an IXP-based measurement server across multiple ISPs using POPsicle.

To substantiate this idea, we deployed a tool implementing POPsicle on a server at the Equinix Chicago Internet Exchange with the help of network operators, and we conducted a week-long measurement study. We chose 10 ISPs that peer at Equinix Chicago for targeted probing. These ISPs were chosen because (1) there was information available in PeeringDB, or we had operator contacts who could verify our inferences, and (2) location hints were available in DNS for IP addresses within the ISP. Unfortunately, the vast majority of ISPs that peer at Equinix Chicago do not have publicly available ground truth information and/or location hints available via DNS thus we could not include

[9]We expect our result to coincide with [10] if we have access to more vantage points.

Table 4: Summary of results from mapping infrastructural nodes.

	POPs (for 13 ISPs)	Datacenters	DNS Servers	NTP Servers	IXPs	Total locations
POPsicle	149	487	9	627	37	1309
Ark	143	315	1	55	25	539
Atlas	244	641	13	827	65	1790
POPsicle compared to Atlas	61.07%	75.98%	69.23%	75.82%	56.92%	73.13%
Ark compared to Atlas	54.60%	49.14%	7.69%	6.65%	38.46%	30.11%
Improvement	1.04x	1.54x	9x	11.40x	1.48x	2.42x

them in this initial study. Also, we note that the ISPs we considered in our POPsicle deployment have, unfortunately, little overlap with the ISPs we consider in Section 5 (and which appear in Table 3) due to our requirement that we have ground truth and location hint information available — only Hurricane Electric and HopOne are in common.

Table 5 shows the results of our IXP-based POPsicle deployment. We observe from the table that POPsicle finds all nodes for 8 out of the 10 ISPs (as compared with the Atlas physical topology data). In the Ark measurements, 6 out of 10 ISPs are fully mapped. For the two ISPs the POPsicle is not able to fully map, a very likely possibility is that the unobserved POPs are invisible to layer3 probes due to configured router policies [31], thus the results we show may be the best that can be achieved through active probing. Overall, our results suggest that POPsicle could be deployed more broadly to accurately map (to the extent possible) ISPs for which we do not have ground truth.

Special Cases. *(1)* The number of POPs found for HE in Table 3 is 23 but in Table 5 the number of POPs found for HE is 24. That is, POPsicle deployed at Equinix Chicago saw an extra node in Calgary, Canada (YYC) which is verified with Datahive IX. One possible implication of this result is that such probe-based measurements are biased towards the vantage points selected. *(2)* For 2 ISPs (PaeTec and Atlantic Metro) some POPs were not visible to our probes, which we intend to investigate further in future work. There is anecdotal evidence that ISPs typically do not expose certain locations to traceroute probes (or any access methods from outside) even when layer 3 services are available at that particular location due to security reasons [31].

Case Study: Cogent. In [21], Ferguson *et al.* present an analysis of Cogent Communication's network based on using reverse DNS records, as well as location-based naming hints. We used the dataset made public by these authors to evaluate, compare and validate POPsicle's probe-based measurement of Cogent's network. We processed the DNS names from Ferguson *et al.*'s dataset using the modified version of location inference technique developed by Chabarek *et al.* [12] and identified 187 POP locations. We then used POPsicle deployed at the Equinix Chicago IXP to target routers within Cogent's network, and it identified 173 POPs. In Appendix A, we see that there are 186 POP locations identified in the Atlas physical topology; it is likely that the additional POP identified in the Ferguson *et al.* dataset is a more recent deployment than was found in Atlas. Also in the table of Appendix A, we see that there are 122 POPs identified through the Ark probes. Lastly, we note that in the most recent Rocketfuel data, there are only 45 POP locations identified. Altogether, these results show that POPsicle's probing technique is very effective for discovering locations of physical infrastructure like POPs, is much better than existing probe-based techniques, and nearly as good as exhaustive use of reverse DNS records.

Table 5: Summary results of network POPs identified with POPsicle deployed at the Equinix Chicago IXP.

ISP Name	POPsicle	Atlas	Ark
BTN	29	29	28
HE	24	24	23
Internet2	10	10	10
PaeTec	54	61	54
Nexicom	9	9	9
HopOne	3	3	3
Indiana Gigapop	2	2	2
MOREnet	4	4	4
Atlantic Metro	9	12	8
Steadfast.net	3	3	3

7.4 Main findings and implications

We describe a new targeted probing technique called POPsicle that is designed to reveal and confirm the presence and location of physical infrastructure such as POPs. To evaluate our method, we used publicly accessible looking glass servers deployed at IXPs, and made a custom deployment of POPsicle at the Equinix Chicago IXP. POPsicle finds 2.4x more physical nodes than Ark probes, and in our custom deployment in Chicago, POPsicle finds nearly all POPs identified in the Atlas physical topologies. In a case study of Cogent's network, POPsicle identified more than 90% of the POPs known through Atlas as well as through the recently described technique based on using reverse DNS records [21]. Moreover, it found many more POPs than Ark probes, or the most recent Rocketfuel measurements.

Overall, our results show that an IXP deployment provides a prime location from which to launch targeted topology discovery probes. Since Rocketfuel maps are commonly used in networking studies that require realistic and representative network topologies, we view this deployment paradigm as having significant potential for generating machine-readable topological information on an on-going basis. We plan to investigate the possibility for additional IXP deployments and a full-fledged system for generating up-to-date network topology data in future work.

The peering model in which IXPs operate is different across different continents. For instance, an IXP in Europe is completely different from an IXP in North America. On one hand the peering model in North America typically involves a commercial colo-operator who also operates the peering equipment. On the other hand, the exchange points in Europe tend to be non-profit, community-based organizations, and the colocation and peering equipment operators

are different [8]. We believe that such a peering model will lead to differences on the results that we observe for networks in our study compared to networks in Europe.

8. CONCLUSIONS AND FUTURE WORK

The high level objective of this paper is to move closer to the goal of having comprehensive and accurate maps of the Internet's topology that can be applied to a wide range of problems. The starting point of our study is to understand how physical and network-layer maps differ. To that end, we compare large repositories of physical and network maps and find that physical maps typically reveal a much larger number of nodes (*e.g.,* POPs and hosting infrastructure). For the selected networks, we find that: *(i)* the physical maps typically show many more nodes/links than the network-layer maps, *(ii)* there is often a high amount of overlap in nodes/links that appear in both data sets, and *(iii)* network-layer maps sometimes include some nodes/links that are not in physical maps due to incomplete or out-of-date published topologies.

These results motivate the development of probing techniques for targeting the identification of nodes with known or suspected physical locations. We develop a layer 1-informed heuristic algorithm for probe source-destination selection called POPsicle that identifies 2.4 times as many nodes as standard probing methods. Finally, we identify the fact that sources co-located as IXPs can be used to amplify POPsicle-based probing since an IXP-based vantage point can be considered to reside within all of the service providers that peer at the IXP. To that end, we deployed POPsicle at a real IXP and found that it finds almost all POPs compared to Atlas, and additional POPs compared with Ark.

In future work, we plan to use POPsicle to more broadly confirm and map network-layer nodes by exploiting additional available IXP-based VPs and by deploying it to new IXPs. We also intend to examine potential efficiency gains in POPsicle's algorithm by more aggressively pruning the search space of destination VPs. Future efforts will include benchmarking versus simple probing methods and more targeted approaches (like iPlane [33]) that will enable us to reason about and quantify the efficiency and effectiveness of the tool in a broader deployment. Lastly, we are considering how to fully automate and integrate POPsicle with the Internet Atlas in order to accurately and quickly assemble multi-layer maps of network service providers.

9. ACKNOWLEDGMENTS

We thank our shepherd, Michael Bailey, and the anonymous reviewers for their invaluable feedback. We also thank Michael Blodgett, Peter Helmenstine (Telx), Sven Engelhardt (GTT), and Adam Eaglestone (Integra) for their feedback and helpful discussions. We also thank Jeff Bartig (Equinix Chicago and Mad-IX), Michael Hare (DoIT) and Bruce LaBuda (DoIT) for helping us run the POPsicle experiments at Equinix Chicago.

This material is based upon work supported by the National Science Foundation under grants CNS-1054985, CNS-0905186, ARL grant W911NF1110227, DHS BAA 11-01 and AFRL grant FA8750-12-2-0328. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the NSF, ARL, DHS or AFRL.

10. REFERENCES

[1] Haversine formula. `http://en.wikipedia.org/wiki/Haversine_formula`.

[2] Internet2 Network Map. `http://www.internet2.edu/media/medialibrary/2013/07/31/Internet2-Network-Infrastructure-Topology.pdf`.

[3] Team Cymru IP-to-ASN service. `http://www.team-cymru.org/Services/ip-to-asn.html`.

[4] The CAIDA UCSD IPv4 Routed /24 DNS Names Dataset - September 2011–March 2013. `http://www.caida.org/data/active/ipv4_dnsnames_dataset.xml`.

[5] The PeeringDB. `https://www.peeringdb.com/`.

[6] Atlantic Metro Communications. `http://www.atlanticmetro.net/resources/maps.php`, Accessed February 2013.

[7] Lumos Networks. `https://www.lumosnetworks.com/sites/default/files/POP-Colocation-List-Feb2013.xlsx`, Accessed February 2013.

[8] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a large European IXP. In *Proceedings of ACM SIGCOMM conference*, 2012.

[9] 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 ACM SIGCOMM Internet measurement conference*, 2006.

[10] B. Augustin, B. Krishnamurthy, and W. Willinger. IXPs: mapped? In *Proceedings of ACM Internet measurement conference*, 2009.

[11] P. Barford, A. Bestavros, J. Byers, and M. Crovella. On the Marginal Utility of Network Topology Measurements. In *Proceedings of ACM Internet Measurement Workshop*, 2001.

[12] J. Chabarek and P. Barford. What's in a Name? Decoding Router Interface Names. In *Proceedings of ACM HotPlanet*, 2013.

[13] A. Dhamdhere and C. Dovrolis. The Internet is flat: modeling the transition from a transit hierarchy to a peering mesh. In *Proceedings of ACM CoNEXT*, 2010.

[14] B. Donnet, M. Luckie, P. Mérindol, and J.-J. Pansiot. Revealing MPLS tunnels obscured from traceroute. *ACM SIGCOMM Computer Communication Review*, 2012.

[15] R. Durairajan, S. Ghosh, X. Tang, P. Barford, and B. Eriksson. Internet Atlas: A Geographic Database of the Internet. In *Proceedings of ACM HotPlanet*, 2013.

[16] S. Engelhardt. Personal communication, 2014.

[17] B. Eriksson, P. Barford, B. Maggs, and R. Nowak. Posit: A Lightweight Approach for IP Geolocation. *ACM SIGMETRICS Performance Evaluation Review*, 2012.

[18] B. Eriksson, P. Barford, J. Sommers, and R. Nowak. Inferring Unseen Components of the Internet Core. *IEEE Journal on Selected Areas in Communications*, 2011.

[19] B. Eriksson, R. Durairajan, and P. Barford. Riskroute: A framework for mitigating network outage threats. In *Proceedings of ACM CoNEXT*, 2013.

[20] D. Feldman, Y. Shavitt, and N. Zilberman. A structural approach for PoP geo-location. *Computer Networks*, 56(3), February 2012.

[21] A. D. Ferguson, J. Place, and R. Fonseca. Growth analysis of a large isp. In *Proceedings of ACM SIGCOMM Internet Measurement Conference*, 2013.

[22] P. Gill, M. Arlitt, Z. Li, and A. Mahanti. The flattening Internet topology: Natural evolution, unsightly barnacles or contrived collapse? In *Proceedings of PAM*. 2008.

[23] B. Gueye, A. Ziviani, M. Crovella, and S. Fdida. Constraint-Based Geolocation of Internet Hosts. *IEEE/ACM Transactions on Networking*, 2006.

[24] B. Huffaker, M. Fomenkov, and K. Claffy. DRoP:DNS-based Router Positioning. *ACM SIGCOMM Computer Communication Review (CCR)*, Jul 2014.

[25] Y. Hyun, B. Huffaker, D. Andersen, E. Aben, M. Luckie, kc claffy, and C. Shannon. The IPv4 Routed /24 AS Links Dataset: September 2011–March 2013. `http://www.caida.org/data/active/ipv4_routed_topology_aslinks_dataset.xml`.

[26] V. Jacobson and S. Deering. Traceroute, 1989.

[27] A. Kapela. Personal communication, 2014.

[28] E. Katz-Bassett, J. John, A. Krishnamurthy, D. Wetherall, T. Anderson, and Y. Chawathe. Towards IP Geolocation Using Delay and Topology Measurements. In *Proceedings of ACM Internet Measurement Conference*, 2006.

[29] S. Knight, H. X. Nguyen, N. Falkner, R. A. Bowden, and M. Roughan. The Internet Topology Zoo. *IEEE Journal on Selected Areas in Communications*, 2011.

[30] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet inter-domain traffic. In *Proceedings of ACM SIGCOMM Conference*, 2010.

[31] B. LaBuda. Personal communication, 2014.

[32] M. Lottor. RFC 1296: Internet Growth (1981-1991). `http://www.ietf.org/rfc/rfc1296.txt`, January 1992.

[33] H. 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 USENIX Symposium on Operating Systems Design and Implementation(OSDI '06)*, November 2006.

[34] V. Padmanabhan and L. Subramanian. An Investigation of Geographic Mapping Techniques for Internet Hosts. In *Proceedings of ACM SIGCOMM Conference*, 2001.

[35] V. Paxson. *Measurement and Analysis of End-to-end Internet Dynamics*. PhD thesis, University of California at Berkeley, 1997.

[36] A. Rasti, N. Magharei, R. Rejaie, and W. Willinger. Eyeball ASes: from geography to connectivity. In *Proceedings of ACM Internet Measurement Conference*, 2010.

[37] L. Roberts. The Arpanet and computer networks. In *Proceedings of the ACM Conference on the HPW*, 1986.

[38] M. Roughan, S. J. Tuke, and O. Maennel. Bigfoot, sasquatch, the yeti and other missing links: what we don't know about the AS graph. In *Proceedings of ACM Internet measurement conference*, 2008.

[39] Y. Shavitt and U. Weinsberg. Quantifying the Importance of Vantage Points Distribution in Internet Topology Measurements. In *Proceedings of IEEE INFOCOM*, 2009.

[40] R. Sherwood, A. Bender, and N. Spring. Discarte: a disjunctive internet cartographer. In *Proceedings of ACM SIGCOMM conference*, 2008.

[41] R. Sherwood and N. Spring. Touring the Internet in a TCP Sidecar. In *Proceedings of ACM Internet Measurement Conference*, 2006.

[42] J. Sommers, P. Barford, and B. Eriksson. On the prevalence and characteristics of MPLS deployments in the open Internet. In *Proceedings of ACM Internet measurement conference*, 2011.

[43] L. Spinelli, M. Crovella, and B. Eriksson. AliasCluster: A Lightweight Approach to Interface Disambiguation. In *Proceedings of the Global Internet Symposium*, 2013.

[44] L. Spinelli, M. Crovella, and B. Eriksson. AliasCluster: A lightweight approach to interface disambiguation. In *Proceedings of the Global Internet Symposium*, 2013.

[45] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP topologies with Rocketfuel. *ACM SIGCOMM conference*, 2002.

[46] Y. Wang, D. Burgener, M. Flores, A. Kuzmanovic, and C. Huang. Towards Street-Level Client-Independent IP Geolocation. In *Proceedings of USENIX NSDI*, 2011.

[47] B. Wong, I. Stoyanov, and E. Sirer. Octant: A Comprehensive Framework for the Geolocation of Internet Hosts. In *Proceedings of USENIX NSDI*, 2007.

[48] M. Zhang, Y. Ruan, V. Pai, and J. Rexford. How DNS misnaming distorts internet topology mapping. In *Proceedings of USENIX ATC*, 2006.

[49] Y. Zhang, R. Oliveira, Y. Wang, S. Su, B. Zhang, J. Bi, H. Zhang, and L. Zhang. A framework to quantify the pitfalls of using traceroute in AS-level topology measurement. *IEEE Journal on Selected Areas in Communications*, 2011.

APPENDIX

A. COMPARISON RESULTS

Table 6: Summary comparison of nodes and links observed in physical and network-layer topologies for networks with a footprint in North America.

ISP	Physical		Network-layer		Nodes			Links			N_{Index}
	Nodes	Links	Nodes	Links	Intersection	Only in P	Only in N	Intersection	Only in P	Only in N	
AT&T	25	57	39	72	25	0	14	51	6	21	100
Cogent	186	245	122	172	122	64	0	171	74	1	63
NTT	47	216	65	229	47	0	18	189	27	40	57
Tinet	122	132	64	79	57	65	7	79	53	0	37
Sprint	63	102	67	108	63	0	4	98	4	10	54
Level3	240	336	129	237	129	111	0	237	99	0	63
Tata	69	111	0	0	0	69	0	0	111	0	40
Abiline	11	14	8	13	8	3	0	13	1	0	100
Ans	18	25	0	0	0	18	0	0	25	0	94
ATMnet	21	22	0	0	0	21	0	0	22	0	100
Bandcon	22	28	14	22	14	8	0	22	6	0	100
BBNPlanet	27	28	0	0	0	27	0	0	28	0	100
BellCanada	48	65	22	0	22	26	0	0	65	0	56
BellSouth	50	66	0	0	0	50	0	0	66	0	76
BTNorthAmerica	33	76	0	0	0	33	0	0	76	0	85
CompuServe	11	17	0	0	0	11	0	0	17	0	100
DarkStrand	28	31	0	0	0	28	0	0	31	0	96
DataXchange	6	11	0	0	0	6	0	0	11	0	100
Digex	31	38	0	0	0	31	0	0	38	0	97
Epoch	6	7	0	0	0	6	0	0	7	0	100
Getnet	7	8	0	0	0	7	0	0	8	0	100
Globalcenter	9	36	0	0	0	9	0	0	36	0	89
Gridnet	9	20	0	0	0	9	0	0	20	0	100
HiberniaCanada	10	14	0	0	0	10	0	0	14	0	60
HiberniaUS	20	29	0	0	0	20	0	0	29	0	100
Highwinds	18	53	0	0	0	18	0	0	53	0	80
HostwayIntl.	16	21	0	0	0	16	0	0	21	0	94
HE	24	37	23	41	23	1	0	34	3	7	100
Integra	27	36	0	0	0	27	0	0	36	0	74
Intellifiber	70	97	0	0	0	70	0	0	97	0	77
Iris	51	64	0	0	0	51	0	0	64	0	27
Istar	19	23	0	0	0	19	0	0	23	0	84
Layer42	9	12	10	6	9	0	1	4	8	2	100
Napnet	6	7	0	0	0	6	0	0	7	0	100
Navigata	13	17	0	0	0	13	0	0	17	0	100
Netrail	7	10	0	0	0	7	0	0	10	0	100
NetworkUSA	35	39	0	0	0	35	0	0	39	0	34
Noel	19	25	2	0	2	17	0	0	25	0	16
NSFnet	13	15	0	0	0	13	0	0	15	0	92
Ntelos	48	61	0	0	0	48	0	0	61	0	48
Oxford	20	26	0	0	0	20	0	0	26	0	50
PacketExchange	21	27	0	0	0	21	0	0	27	0	100
Palmetto	45	70	0	0	0	45	0	0	70	0	49
Peer1	16	20	0	0	0	16	0	0	20	0	100
RedBestel	82	101	0	0	0	82	0	0	101	0	9
Syringa	66	74	0	0	0	66	0	0	74	0	9
USSignal	61	79	0	0	0	61	0	0	79	0	46
VisionNet	22	23	0	0	0	22	0	0	23	0	23
Xeex	24	34	4	3	4	20	0	3	31	0	96
Xspedius	34	49	0	0	0	34	0	0	49	0	100