How to Resolve IP Aliases

Neil Spring, Mira Dontcheva, Maya Rodrig, and David Wetherall
UW-CSE-TR 04–05–04*

May 14, 2004

Abstract

To construct accurate Internet maps, traceroute-based mapping efforts must group interface IP addresses into routers, a task known as alias resolution. In this paper, we introduce two new alias resolution approaches based on inference to handle addresses that cannot be resolved by existing methods based on probe measurements. The first decodes the DNS names assigned by the ISP to recognize the name fragments that identify a router. The second infers aliases from the graph of linked IP addresses and requires no additional measurement traffic. We then experiment with feasible combinations of these techniques and existing ones by resolving aliases during the mapping of PlanetLab, a large wide-area overlay, and UUnet, a large ISP. We find that these techniques have complementary strengths and weaknesses and are best used in concert. The DNS and graph inference methods provide information where existing probe methods fail and are less dependent on router implementation choices. The existing probe methods can be made more effective in practice by using multiple vantage points and taking advantage of implementation synergies.

1 Introduction

Internet mapping is the process of discovering the topology of portions of the network, from corporate networks to ISPs, POPs, overlays, and even the overall Internet. The resulting maps provide knowledge of the underlying structure of the network and are valuable for managing its operation and understanding its properties, e.g., checking connectivity and identifying points of vulnerability to failure. The technology for Internet mapping has improved by strides over the past several years, and there are now several different approaches and mapping efforts [9, 11, 5, 3, 6, 4]. Of these, traceroute-based methods are the most broadly applicable as they can be used across multiple administrative domains where IP is the lowest common denominator and no special provision is made for mapping.

This paper focuses on a problem that is common to all Internet mapping efforts based on traceroute: IP alias resolution. Traceroute discovers the sequence of routers along an Internet path by sending packets with limited, consecutive time-to-live (TTL) values from a probe machine. When these packets expire in the network, routers return ICMP time-exceeded messages to the probe machine. The source address of these messages is typically that of the interface that received the packet. This implies that traceroute provides a list of interfaces but does not attempt to group those interfaces into routers. Alias resolution is the process of performing this grouping, removing IP aliases to reveal the true network topology.

Accurate alias resolution is an important though easily overlooked component of any traceroute-based mapping effort. Without it, the resulting map will not reflect the connectivity of the underlying network and can be misleading. For example, the path from A to B and the path from B to A may appear to be disjoint even when they follow the same sequence of routers in opposite directions and hence share properties such as propagation delay and capacity. More generally, alias resolution improves the utility of the recovered maps in two respects. First, while traceroutes reflect paths that were taken, alias resolution also exposes new paths through the network that exist and may be taken in the future, but were *not* taken during mapping, either due to current routing or because they were not directly measured for reasons of mapping efficiency. Second, grouping IP addresses into routers collapses virtual "interface-pair links" into real, router-to-router links. This is important because it reveals, for example, which paths will compete for available bandwidth. In the PlanetLab overlay that we have mapped there are 3,053 IP address pairs that represent "links" as seen by traceroute, but only 2,240 router-to-router links after alias resolution.

*{nspring,mirad,rodrig,djw}@cs.washington.edu. Department of Computer Science and Engineering, University of Washington, Seattle, WA 98195–2350.

Unfortunately, alias resolution is not straightforward because it is not supported as part of the IP protocol. Rather, the state-of-the-art is based on heuristics that take advantage of common router implementations. Several different techniques have appeared in recent years. Pansiot and Grad first introduced a method that relies on the practice of using the IP address of the outgoing interface as the source address of router-generated packets and the existence of a single dominant route from all router interfaces to a given remote destination [8]. This method was subsequently extended to use multiple vantage points as part of the Mercator project [5]. More recently, in our earlier work on Ally, we introduced a technique that relies on the common implementation of the IP identifier in certain packets generated by routers as a per-router counter. This technique, made efficient by clustering by DNS names and hop distance from a probe point (as approximated by return TTL values on packets), allows pairs of IP addresses to be tested to determine whether or not they are likely aliases [11].

In this paper, we present two new techniques for alias resolution and compare the performance of these and existing techniques on two real mapping experiments. The first technique is an extension of methods that recover geographic location from router names. It recognizes fragments of the DNS names assigned to router interface addresses to find those that identify a specific router. The second technique uses the graph of edges between IP addresses obtained from traceroutes and consists of two simple inference rules. One rule is that two adjacent IP addresses are likely to represent adjacent routers rather than the same router, given that routing does not contain loops. Conversely, the other rule is that IP addresses immediately preceding a merge point in the graph are likely to be aliases when point-to-point links are in use, for reasons that are elaborated in the paper. To evaluate these and existing techniques, we measure their performance while mapping the topology of PlanetLab, a wide-area overlay, and the topology of UUnet, a large network provider. These mapping tasks represent real yet diverse workloads (in terms of router equipment makeup, dense versus sparse topology, and scale) and allow us to gauge the strengths and weaknesses of the various methods. To the best of our knowledge, alias resolution techniques have not been systematically evaluated.

The results of these experiments allow us to provide tentative recommendations on how mapping efforts can best resolve IP aliases. We find that these methods have complementary strengths and weaknesses (none is redundant with the others) so that they are best used in concert when complete alias resolution is the goal. The new DNS and graph techniques are able to resolve aliases that are unresponsive to probes, thus finding up to a third more aliases than can be found with previous methods. However, they do not find a superset of the aliases found by existing methods. Having multiple methods is also useful to provide a check on the underlying assumptions made by individual methods, and hence improve overall accuracy. DNS resolution is generally accurate but requires knowledge of ISP naming conventions. Graph-based resolution relies on assumptions about ISP network design and so is less accurate by itself, but it has the advantage that it is largely not dependent on router implementation choices, unlike existing probe methods. We also find that existing probe techniques benefit from the use of multiple vantage points, which improves both their efficiency and effectiveness. Further, because probe packets can return multiple pieces of usable information there is a synergy in combining their implementation. Finally, we note that the effectiveness of the methods varies with the mapping task, suggesting that care is needed in applying them. In our case, DNS appears more appropriate for ISP mapping, while graph-based inferences are suitable for overlay mapping.

The rest of this paper is organized as follows. In Section 2, we present a classification of the various alias resolution techniques, including a description of DNS and graph-based alias resolution. In Section 3, we describe our methodology for comparing the various techniques by running head-to-head tests of different combinations as part of the task of mapping the Internet portions that underlie the PlanetLab overlay topology and UUnet, a large ISP. In Section 4, we present the results of these comparisons, highlighting the strengths and weaknesses of each method. Finally, in Section 5 we conclude with our recommendations for performing alias resolution in practice and propose future areas of study.

2 Alias Resolution Techniques

The alias resolution problem we study is illustrated in Figure 1. To map a portion of the network, many traceroutes are run from diverse vantage points to collect a set of overlapping paths through the network. Each path consists of a series of IP addresses representing the order of visited router interfaces. These paths are the input for the alias resolution process that merges interfaces that belong to the same router. The output is the topology of a portion of the network that identifies individual routers and the links between them, i.e., what is traditionally meant by a network map.

In the rest of this section, we describe existing techniques (Mercator and Ally) and present two new tech-

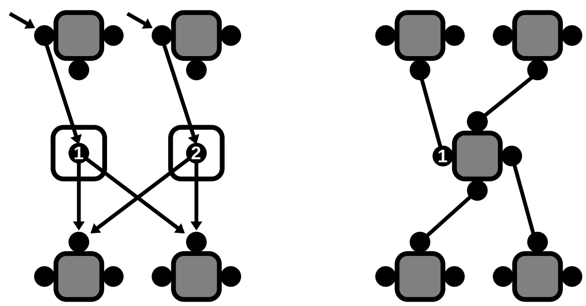

Figure 1: Boxes represent routers and circles represent interfaces. Traceroute lists input interface addresses from paths (left). Alias resolution clusters interfaces into routers to reveal the true topology. Interfaces ❶ and ❷ are aliases (right).

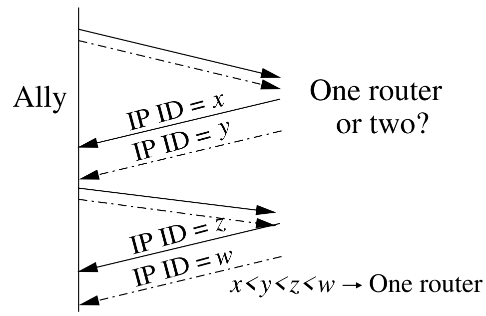

Figure 2: In-order IP identifiers from two different IP addresses suggest aliases.

niques for alias resolution (DNS and Graph). These are the four techniques that we evaluate experimentally in the later sections of the paper. We describe these techniques in two broad classes: fingerprint-based methods, which actively probe routers and compare responses, and inference-based methods, which interpret patterns drawn from the traceroute data. This classification is useful because it highlights the difference between the existing fingerprint methods and the new inference methods, and because it provides a framework for classifying new methods that may emerge over time.

2.1 Fingerprint Methods

Existing techniques for alias resolution send "probe" packets into the network and study the responses to find evidence of shared, underlying identity. We term these *fingerprint* techniques because they implicitly compare the signatures or fingerprints of routers to find matches. They are applicable only when routers are responsive to probe packets, which we found excludes their use on 10 to 50% of the IP addresses in our experiments.

2.1.1 Common Source-Address (Mercator)

Pansiot and Grad [8] introduced an alias resolution technique based on comparing the source address of messages sent by the router's host processor. While the source address of ICMP time-exceeded messages (in the middle of the traceroute) is set to that of the input interface, ICMP port-unreachable messages (at the end of the traceroute) typically use the output interface address.[1] Since IP addresses are unique[2] and the output interface is constant when routing is stable, the source address of ICMP port-unreachable messages can be used to group aliases together. Finding aliases by source address requires only one probe packet per address, and so is quite efficient.

Govindan and Tangmunarunkit [5] refine Pansiot's technique in two ways. To account for unstable routing, they repeatedly probe addresses to expose likely aliases. To account for unreachable routers, they use source routing to inject probe packets into other parts of the network that may reach the probed address.

The "Mercator" tool we will evaluate is based on Govindan's approach without the use of source routing. Instead, it sends probes to each IP address from different vantage points in the network. Each pair of interface addresses that share a source address as seen by any vantage point is considered an alias.

2.1.2 Common IP-ID Counter (Ally)

Rocketfuel's alias resolution component, Ally, builds on Mercator and introduces a technique that inspects the IP identifier (IP-ID) field of responses [11]. The original purpose of the IP identifier is to uniquely identify packets for reassembly after fragmentation. It is commonly implemented as a counter that is incremented after sending each packet. Thus, packets that are sent one after the other should have consecutive IP identifiers. Ally uses this observation to identify aliases. Other recent work also uses this observation to count hosts behind a NAT [2] and measure reordering [1].

Ally sends a series of probe packets to two candidate IP addresses as shown in Figure 2 to solicit IP identifiers in response packets. When responses have in-order

[1] The choice of source address is not standardized and appears to depend on router implementation.

[2] Use of private IP address space on public portions of the Internet is not significant in our experience.

IP identifiers, it suggests they were generated from a single counter. This technique can identify more aliases than Mercator because a single IP ID counter is more common than a single source address. However, a naive implementation of the scheme would be inefficient because it requires $O(n^2)$ pairwise tests. To guide the search, Rocketfuel clustered IP address pairs by return TTL and sorted by piecewise-reversed DNS name. Other clustering metrics are possible, for example, latency from different vantage points. We report on the effectiveness of these orderings on efficiency in Section 4.3. Further tuning is required to make the scheme accurate. The test will yield false-positives when the counters of different routers happen to appear synchronized, and so a verification phase is needed to confirm aliases at a later time. We report on the effectiveness of the verification phase in Section 4.1.1. Some leeway must also be made for the impact of cross-traffic arriving at the router and reordering along network paths.

The "Ally" tool we evaluate later is based solely on the IP identifier technique, with the TTL-based clustering required to make it practical, and does not include the Mercator method. We separate these methods to provide a clean comparison for evaluation.

2.2 Inference Methods

The two new techniques we present in this paper are based on drawing inferences by looking for patterns in the database of traceroute paths and supplementary data, instead of probing routers. First, we describe resolution based on reverse-DNS mappings. Second, we describe resolution based on graph inference rules. These techniques are applicable even when routers do not respond to direct probe packets.

2.2.1 DNS-based Alias Resolution

If an ISP uses systematic naming conventions for its routers, then information can be gathered by decoding names using this convention. This approach is the basis of earlier work on recovering geographic location [7, 11, 10]. In our work, we have logically extended it to determine whether two IP addresses are aliases. For example, sl-bb21-lon-14-0.sprintlink.net and sl-bb21-lon-8-0.sprintlink.net are aliases for the same backbone router in London (the 14-0 and 8-0 appear to refer to slot or port numbers.) We extended Rocketfuel's DNS name decoder, which previously extracted geographic location and role information, to extract the fragments of the DNS name that uniquely identify a router. Mapping a new ISP requires reverse-engineering a new expression to extract these unique fragments, a process that is simplified by the examples given by the probing methods above, but limits the applicability of using DNS.

The DNS technique can only be as accurate as the ISP's database, which must be updated as addresses are reassigned and ISPs are merged, and may include the occasional typo. It is also incomplete, as some ISPs only name their backbone (core) routers and addresses used at exchange points and peering links may not have such structured names, for example, att-gw.sea.cw.net. Nonetheless, we have found it to be a valuable source of information.

The DNS tool we evaluate in this paper includes rules for Abilene, AT&T, CalREN, Cogent, Exodus, Geant, Genuity, PSI, Qwest, Telstra, Sprint, UUnet and Verio. These rules were generated by hand by observing the pattern of aliases measured by the fingerprinting approaches above. The DNS tool assumes that two ISPs do not "fight" over the same router — interface names are given so that each router has names from only a single domain.

2.2.2 Graph-based Alias Resolution

The traceroute data collected as part of mapping can also guide the search for aliases. We construct a directed graph using the IP addresses as nodes and the pairs of IP addresses seen by traceroute as edges. We then look for patterns in this graph to suggest likely and unlikely aliases. If inferences based on the graph can be made sufficiently accurate, they can provide a "best-guess" for unresponsive addresses. Our graph-based techniques are based on the following observations:

1. Two addresses that directly precede a *common successor* are aliases, assuming point-to-point links are used.

2. Addresses found in the *same traceroute* are not aliases, assuming there are no routing loops.

Common successor When links between routers are point-to-point, and the input interface is used as the source address for time-exceeded messages, this interface address implicitly identifies the router at the opposite end of the point-to-point link. So, edges from different nodes that are incident on the same node in the traceroute graph suggest an alias as shown in Figure 3.

When multiple-access or switched networks are used, this heuristic may fail — the successor address does not identify a single router at the other end of the link.[3] Similarly, this technique requires traceroute paths to overlap to

[3]However, this suggests a method for finding switched or multiple-access networks, which we defer to future work.

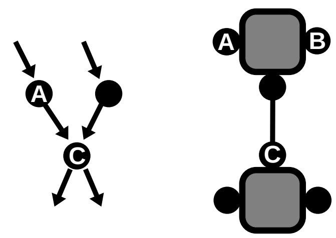

Figure 3: The common successor technique identifies aliases using the assumption of point-to-point links. At left is the IP address graph: nodes A, B, and C represent interface addresses. At right, A and B are shown to represent interfaces on the same router, connected by point-to-point link to C.

find aliases, which is not always the case, e.g., upstream and downstream traceroutes may have no IP addresses in common.

Same traceroute In addition to finding likely aliases, traceroute data can disprove aliases and obviate the need for probing. If routing loops are not present during mapping (or traces containing routing loops are recognized and discarded) different addresses that appear in the same trace cannot be aliases. If successful, this observation can be used both to re-prioritize or reduce the workload of IP-ID probing and to correct the common successor technique in the presence of switched networks. Each trace of length n disproves $\binom{n}{2}$ aliases, and so provides a rich source of data.

2.3 Summary

The existing, fingerprint-based techniques for alias resolution (Mercator and Ally) probe routers and inspect the responses to find evidence of shared, underlying identity. These approaches fail when routers are unreachable or are implemented differently than expected. The new, inference-based techniques that we have presented (DNS and Graph) work by drawing inferences from the database of traceroutes and supplementary information. These approaches are applicable even when routers are unresponsive, but have other limitations in their completeness and accuracy. In the following sections we evaluate the four methods, both individually and in combination.

3 Experimental Methodology

In this section, we introduce three metrics for assessing the performance of alias resolution: accuracy, completeness, and efficiency. To compare Mercator, Ally, DNS, Graph and reasonable combinations of techniques along these axes, we collect two datasets — the first is a map of the PlanetLab overlay, the second is a (much larger) ISP map of UUnet.

3.1 Performance Metrics

Guided by our earlier experience with alias resolution, we develop three performance metrics that can be used to assess any alias resolution technique: *accuracy* measures how often discovered (or disproven) aliases are correct; *completeness* measures how many aliases are discovered; and *efficiency* measures the amount of probe traffic used to discover the aliases. We describe the methodology for estimating each metric in turn.

3.1.1 Accuracy

Since we are mapping networks that we do not control, we cannot compare the results of our alias resolution techniques to the true map. Instead, we estimate accuracy by measuring the agreement between methods. Since the four methods are based on different sources of information, we believe that agreement between them indicates accurate alias resolution. We compute two agreement measures, a "false" positive rate and a "false" negative rate.

Each of the four techniques can be seen as making a series of statements about pairs of IP addresses A, B, and C of the form: $A \equiv B$ or $B \not\equiv C$, where $\equiv$ represents "is an alias for." These pair-wise statements are explicit for Ally, which uses pair-wise tests, and Graph, which finds pairs of addresses with common successors, but implicit in the groups of addresses found by Mercator and DNS. Intuitively, we wish to compare each pair of techniques for consistency over these statements. However, since each approach provides an incomplete picture of the aliases found, we can compare only over the aliases for which the pairs of methods have made a statement.

The false positive rate for a technique being tested T_t relative to a "reference" technique T_{ref} is the conditional probability that T_{ref} disagrees with T_t when T_t asserts that a pair of IP addresses are aliases. It is calculated as the number of address pairs (A,B) for which T_t asserts $A \equiv_t B$ when T_{ref} asserts $A \not\equiv_{ref} B$ divided by the number of address pairs in which T_t asserts $A \equiv_t B$ and T_{ref} asserts either $A \equiv_{ref} B$ or $A \not\equiv_{ref} B$. That is, the false positive rate of T_t is the relative number of address pairs "in dispute" given that T_t has declared those pairs to be aliases and T_{ref} has declared either way. The false negative rate is simply the complement: the relative number of address

False Positives	$\frac{\#\ \text{pairs}(A,B)\ \text{s.t.}\ (A \equiv_t B \wedge A \not\equiv_{ref} B)}{\#\ \text{pairs}(A,B)\ \text{s.t.}\ (A \equiv_t B \wedge (A \equiv_{ref} B \vee A \not\equiv_{ref} B))}$
False Negatives	$\frac{\#\ \text{pairs}(A,B)\ \text{s.t.}\ (A \not\equiv_t B \wedge A \equiv_{ref} B)}{\#\ \text{pairs}(A,B)\ \text{s.t.}\ (A \not\equiv_t B \wedge (A \equiv_{ref} B \vee A \not\equiv_{ref} B))}$

Table 1: The false positives are calculated as the number of pairs in dispute when a technique T_t asserts alias and T_{ref} asserts not alias. The false negative rate is the complement.

pairs "in dispute" given that T_t has declared those address pairs to be non-aliases. These equations are summarized in Table 1.

We separate false positives from false negatives because each technique can play a different role. A low false negative rate but high false positive suggests a technique may be useful for disproving aliases to guide a search. Conversely, a high false negative rate but low false positive rate suggests a technique that may be an efficient component but incomplete on its own.

3.1.2 Completeness

Our second measure of an alias resolution technique is how completely it can collapse IP address aliases into routers. Again, since we are mapping a live network outside of our control, we do not know how many aliases (or routers) are actually present. Instead, we compare each technique to the union of all techniques and use the relative number of aliases found by each technique to represent its completeness. Because the individual techniques are incomplete on their own, we will explore several combinations of techniques to more precisely understand their strengths and weakness and how they would compose in an integrated alias resolution approach.

When measuring completeness, the pair-wise comparisons used for measuring accuracy are less important. That is, incomplete discovery of pairs of IP addresses can still result in a complete grouping of IP addresses to routers. For example, if $A \equiv B$ and $B \equiv C$, the pair-wise test between IP addresses A and C is unnecessary and would not change the result. Instead, we define "aliases" to be the number of additional IP addresses that belong to a router. (By analogy, Samuel Clemens and Superman each have only one alias (Mark Twain and Clark Kent)).

The completeness of a technique is measured as follows. Each statement of aliases $A \equiv B$ from every technique is considered in grouping IP addresses into routers. The "total" number of aliases in the topology is the sum of the number of "aliases" of each router. Again, this "total" may be incomplete or even an overestimate if undetected false positives are present. The completeness of a technique is thus the ratio of the number of aliases it finds to the "total" number of aliases.

In a second set of measurements, we consider the "total" number of aliases in the topology to be the sum of the number of "aliases" of each *responsive* router. The aliases for responsive routers are the only ones that probing techniques can discover, thus the completeness of probing techniques increases when only responsive routers are considered. Inference techniques can discover aliases for responsive as well as unresponsive routers, so their relative completeness should remain unchanged.

3.1.3 Efficiency

Our final concern is the efficiency of alias resolution, which relates to how much work each technique requires and hence how quickly it completes. We measure the efficiency of each technique by counting the packets sent in the alias resolution process after mapping is complete. A packet count ignores the complexity of local computation (such as the graph search) and result storage (such as lists of discovered and disproven aliases) which we do not consider to be bottlenecks based on our earlier mapping experiences.

Tradeoffs between approaches may gain efficiency, and we characterize only a few points in the design space. For example, using more vantage points for TTL clustering (described in Section 4.3) initially costs packets from each vantage point, but may result in overall savings because fewer pairs will be tested by Ally.

3.2 Mapping Tasks

Now that we have presented four techniques and three metrics, we discuss the two datasets we will use as Internet mapping workloads. Our approach is to map two Internet structures that represent different extremes: PlanetLab, a wide-area overlay, and UUnet, a large ISP.

3.2.1 PlanetLab (Overlay)

PlanetLab is a wide-area overlay that consists of roughly 127 nodes at 53 geographically distinct sites, most educational. We collected the equivalent of 3,012 traceroutes between these sites to map its topology on May 6, 2003, using the reverse path tool running on Scriptroute [12] at each site.[4] The resulting map consists of 1,815 IP ad-

[4] The reverse path tree tool does not take complete traceroutes for efficiency reasons: each trace is stopped when a branch merges with the rest of the tree. For more detail, see [12].

Map	Tests performed	Confirmed aliases	Disproven aliases	Test false positive rate
PlanetLab	265 million	557	120 (22%)	5×10^{-7}
UUnet	0.2 million	2,782	64 (2%)	3×10^{-4}

Table 2: False alias statistics from using Ally on PlanetLab and UUnet. While there are many aliases that were initially believed but disproven, the error rate of the test itself is very low, indicating that a single verification pass is sufficient.

dresses and 3,053 interface-to-interface links that resolve to 983 routers and 1,347 router-to-router links.

The PlanetLab map is useful for our evaluation because it is relatively small, diverse, and sparse. It is a real mapping task whose importance will likely increase as overlay networks gain popularity. However, it is small enough that we can run alias tests exhaustively to create a dataset suitable for evaluating whether heuristics that guide the search miss valid aliases. It is also made up of several ISPs, which reduces the influence of any particular hardware vendor's implementation choices or ISP topology design and router configuration choices. Finally, overlay maps are much sparser than ISP maps, so it represents one extreme for testing mapping techniques.

3.2.2 UUnet (ISP)

The second topology that we mapped is that of UUnet, a very large network service provider. We used the BGP-directed probing methodology from Rocketfuel [11] with PlanetLab servers as measurement vantage points. To gather a map of this scale, we used the flexibility of Scriptroute to construct a modified traceroute that stopped as soon as it left address space originated by AS701 (UUnet in North America), thereby reducing the volume of measurement traffic. We collected two million traceroutes from 49 PlanetLab sites to map the topology on May 9, 2003. The resulting map of UUnet and its periphery (the adjacent routers of customers and peers) consists of 10,812 IP addresses and 25,015 interface-to-interface links that resolve to 7,391 routers and 8,074 router-to-router links.

We chose ISP mapping as a tractable subset of whole Internet mapping, and UUnet as the canonical example of a large and well-known ISP. Compared to the overlay workload, this map is larger, denser, and less diverse in geography and network design. It allows us to investigate the scalability of the different alias resolution techniques as well as to observe whether the differences between overlay and ISP mapping affect the success of the techniques.

4 Results

In this section, we evaluate the alias resolution techniques along the axes of accuracy, completeness, and efficiency. For each metric, we discovered that relatively small engineering fixes to Mercator and Ally can provide a large benefit in practice, which we report on before providing comparisons across all of the techniques.

4.1 Accuracy

An alias resolution technique is accurate when its statements about whether IP addresses represent the same router are correct. In this section, we first describe how to remove false positives when using Ally. We then compare the results of each technique to those of the rest, identifying relative false positives and negatives.

4.1.1 IP Identifier False-Positives

The IP identifier technique infers the existence of a single counter shared between two aliases. By random chance, some counters may be temporarily synchronized and appear as a single counter. A verification phase that tests these addresses at a later time establishes whether they actually represent aliases. The pairs that are reclassified as aliases are considered "confirmed" while the rest are "disproven." These disproven aliases represent inaccuracy that would appear in the resulting map if the verification phase had not been run.

In Table 2 we show the false positive rate relative both to the "confirmed" aliases and to the total number of tests performed. We perform many more tests than necessary for PlanetLab, as described in Section 3, so despite its small size, 265 million alias-pairs are tested. The 120 pairs that were falsely believed to be aliases yield an error rate of 1 in 2 million tests. Fewer pairs were tested over the topology of UUnet, and the test showed a false positive rate of 1 in 3 thousand. This shows that the false positive rate inherent in the approach is very low, but not that there are no systematic errors in the method; we defer the latter to the next subsection

However, while the number of false positives appears insignificant relative to the number of tests performed, it

Tested	Reference Technique		
Technique	Ally	DNS	Graph
Mercator false +	0/ 382 (0%)	0/ 185 (0%)	0/ 105 (0%)
Ally false +	-	0/ 334 (0%)	0/ 281 (0%)
Ally false -	-	0/12,881 (0%)	22/2,686 (0.8%)
DNS false +	0/ 334 (0%)	-	0/ 154 (0%)
DNS false -	0/12,881 (0%)	-	1/2,772 (0.04%)
Graph false +	22/ 303 (7.3%)	1/ 155 (0.6%)	-
Graph false -	0/ 2,664 (0%)	0/ 2,771 (0%)	-

Table 3: PlanetLab: Error rate of alias resolution techniques. We compare aliases discovered using each method to those inferred by Ally, DNS, or Graph.

Tested	Reference Technique		
Technique	Ally	DNS	Graph
Mercator false +	3/ 1,293 (0.2%)	23/ 410 (5.6%)	9/1,345 (0.7%)
Ally false +	-	11/ 965 (1.1%)	6/2,933 (0.2%)
Ally false -	-	6/17,633 (0.03%)	116/ 190 (61.1%)
DNS false +	6/ 960 (0.6%)	-	0/2,330 (0%)
DNS false -	11/17,638 (0.06%)	-	319/4,603 (6.9%)
Graph false +	116/ 3,043 (3.8%)	319/ 2,649 (12.0%)	-
Graph false -	6/ 80 (7.5%)	0/ 4,284 (0%)	-

Table 4: UUnet: Error rate of alias resolution techniques. We compare aliases discovered using each method to those inferred by Ally, Graph, or DNS.

is quite substantial when compared to the number of confirmed aliases and can lead to inaccuracy in the resulting map. A verification phase to verify the initial set of aliases and discard false positives from the set is thus essential. Fortunately, the low false-positive rate ensures that a single verification phase is sufficient to detect and discard false positives.

We also investigated the cause of false positives and found that the individual IP-ID test is less effective when routers rate-limit responses so that nearby identifiers are difficult to observe. That is, seeing two or three IP identifiers that are nearby or in order is not as convincing as four, but when ICMP rate-limiting is used, retrying to obtain four samples is futile. Instead, generous thresholds that favor false positives (over false negatives) that are then caught by the verification phase work well to resolve rate-limiting routers accurately.

4.1.2 Comparative Evaluation

To measure the relative accuracy of each technique we compare it with each of the others. We compute the error rates as the percentage of cases in which a pair of techniques disagree on a classification of an alias or a non-alias pair, as defined in Section 3.1. The "false positive" rate of a technique is the likelihood that its assertion that a pair of addresses are aliases will be disputed by a reference technique. The "false negative" rate is the complement, or the likelihood that a pair classified as not-aliases will be disputed by a reference technique.

Table 3 summarizes the error rates of each technique for the mapping of PlanetLab. Mercator, Ally, and DNS show zero false positives when compared to the other three techniques. We do not quantify "false negatives" with Mercator, as the technique does not disprove aliases: there is always some chance that another vantage point would show an unproven pair of addresses to be aliases. The graph-based technique that disproves aliases (the "same traceroute" inference rule) has no false negatives. However, the common successor inference rule that finds likely aliases has a false positive rate of 7.3% relative to Ally, and 0.6% relative to DNS. This inconsistency between Graph's aliases and Ally's and DNS's not-aliases also appears in the false negative rate of the two approaches relative to Graph. The graph-based technique incorrectly classified 23 IP addresses pairs as aliases (one compared to

DNS and 22 compared to Ally), making it the least accurate approach.

In Table 4, we show the error rate of the four techniques in the mapping of UUnet. Those alias pairs that represent false positives for Mercator and Ally relative to DNS are consistent — both tools assert that some addresses are aliases while DNS disagrees, suggesting incorrect or simply out of date DNS names. The false positive rates of Mercator and Ally when compared with Graph are likely the result of undetected loops in the traceroute data — when collecting two million traceroutes, eventually some of these will experience strange routing. Using the "same traceroute" inference rule, Graph falsely asserts 15 IP address pairs are not-aliases. The 4-12% false positive rate of the graph technique is a consequence of the topology design of UUnet: the assumption of point-to-point links is not accurate due to the use of switched networks in the UUnet topology. Ally's high false negative rate of 61% relative to Graph is a consequence of the small overlap between techniques (the 190 pairs Ally claims are not aliases that are also classified by Graph). In a separate experiment, we confirmed many of Graph's aliases using Ally, so we expect this rate would decrease as more tests are performed. We found that nearly all of UUnet's routers are responsive to fingerprinting methods, so graph's inaccuracy is only a small concern for mapping this ISP. However, this demonstrates that the "common successor" technique depends on its assumptions and should be used with care.

The implication of these tables is that different alias resolution techniques can serve different purposes, and these tables provide an order in which to compose the statements made by each technique: in order of increasing error. A tool integrating these approaches would likely consider Mercator's aliases authoritative, then add statements from graph's not-aliases (same traceroute), DNS aliases, Ally aliases, Ally not-aliases, DNS not-aliases, and finally graph's aliases (common successor).

4.2 Completeness

Completeness measures the fraction of aliases discovered by a technique out of the total number of aliases in the network. In the absence of a true map that shows us how many aliases are in the network, we consider the union of aliases identified by all techniques as the total, and compare the number of aliases discovered by a single technique or combination of techniques to this total. In this section, we first study the improvement offered by multiple vantage points to source-address based alias resolution, and then compare the completeness of the different techniques.

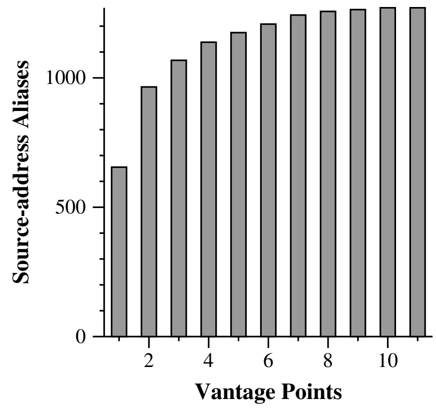

Figure 4: Each additional vantage point contributes to a source-address based alias-resolution technique in the mapping of UUnet

4.2.1 Improving Source-Address-based Completeness

In Figure 4, we show the additional aliases found when using a Mercator-inspired approach and looking for aliases from up to eleven vantage points. While a single vantage point matched 655 aliases, using all eleven vantage points found 1,271, almost doubling the number of discovered aliases. This shows that there is an appreciable gain in using additional vantage points. These extra vantage points serve the same purpose as source-routed probes in Mercator: some router addresses can only be reached by certain vantage points. Beyond eleven vantage points, however, we reach the point of diminishing returns where adding an additional vantage point does not contribute sufficiently to the method's completeness to make it worthwhile.

4.2.2 Comparative Evaluation

The completeness of each alias resolution approach is shown in Tables 5 and 6 for PlanetLab and UUnet respectively. While individual approaches find at most 80% of the aliases in the network, using them in combination completes the picture. We also find that there are nearly a third more aliases in the network, according to DNS and graph approaches, than were found previously by Ally and Mercator. These extra aliases primarily represent unresponsive routers. We investigated the 20% of aliases from responsive routers in PlanetLab that were missed by Ally and found that most of these were the result of addresses

Technique group	Of 832 Overall	Of 694 Responsives
Mercator	345 (42%)	345 (50%)
Ally	557 (67%)	557 (80%)
DNS	331 (40%)	258 (37%)
Graph	332 (40%)	238 (34%)
Mercator ∪ Ally	608 (73%)	608 (88%)
Graph ∪ DNS	547 (66%)	409 (59%)
Mercator ∪ Ally ∪ Graph	756 (91%)	662 (95%)
Mercator ∪ Ally ∪ DNS	727 (87%)	654 (94%)

Table 5: PlanetLab: Completeness of techniques. We define the union of aliases found by all techniques to be 100%.

that were only temporarily responsive: these did not respond when probed by Ally. The vast majority of routers in UUnet were responsive, so there is very little difference between the completeness of techniques when considering only responsive addresses, and hence we only present one column of completeness results for UUnet.

We show a few combinations of techniques. Mercator with Ally shows the completeness of existing fingerprinting methods. Graph with DNS shows the completeness available by inference methods alone. Mercator, Ally, and Graph represent a group that would be effective for Internet mapping in which there are too many ISPs for DNS to be practical. Mercator, Ally, and DNS represent a group that would be effective for ISP mapping. Unstated in the table is the union of all techniques, which we use to define 100% completeness. We removed Graph from the UUnet completeness analysis due to the relatively high false positive rate it exhibited, and consider the union of aliases discovered by Mercator, Ally, and DNS as 100% completeness for this ISP. Nevertheless, the conclusion to draw from Tables 5 and 6 is that these techniques can be used in combination to find more aliases in the network than any technique alone.

4.3 Efficiency

In this section, we evaluate the efficiency of alias resolution approaches in the context of network mapping. Our metric of (in-)efficiency is the number of packets sent in the process. We count packets instead of time to completion as there is the potential to exploit some parallelism to improve speed.

Technique group	Of 3,421 Overall
Mercator	1,271 (37.2%)
Ally	2,782 (81.3%)
DNS	1,290 (37.7%)
Mercator ∪ Ally	3,086 (90.2%)

Table 6: UUnet: Completeness of techniques. We define the union of aliases found by the three techniques to be 100%. Nearly all routers were responsive (3,378 of 3,421), so a second column is not shown. We removed Graph from this analysis for concern about its false positive rate.

4.3.1 Using TTLs and DNS for IP-ID Efficiency

The IP identifier field exposes an alias when responses to probes are returned in-sequence from two different interface IP addresses. To solicit these responses, pairs of interface addresses must be tested individually, and this process can require many packets. To better guide the search for aliases, we apply two heuristics. First, we test only those addresses whose responses include similar TTLs – addresses that have paths of the same length (in hops) back to the measurement source. In this section, we show how this can be used to prune the search space from the all-pairs of the naive approach to a manageable subset. Our second heuristic is to test addresses having similar names first, relying on the implicit structure of DNS names to expose most aliases quickly.

In Figure 5, we show the distribution of return TTLs as seen by our measurement host. The return TTL is the value in the outer IP header of the ICMP error message sent by the router, as opposed to the outgoing TTL in the packet header encapsulated by the ICMP error message. We now concern ourselves with the pairs of addresses that share a return TTL.

In Figure 6, we show the distribution of differences in return TTL. That is, those that share a value have distance 0, and if one interface's response has TTL value 250 and anothers has value 251, they have a distance of 1. From the CDF, we observe that fewer than 10% of the all-to-all alias probes are required if matching TTLs. However, we found one alias pair with a distance of 1; to catch this alias would require 25% of the all-to-all probes.

By adding more measurement points from which to capture the return TTL, the pairwise testing approach becomes feasible without sacrificing completeness. In Figure 7, we show the cumulative fraction of address pairs with increasing Euclidean distance for one to five vantage points. These additional vantage points permit a small search to tolerate noise. In our tests over the PlanetLab

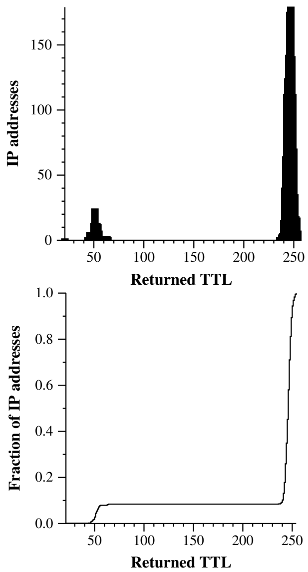

Figure 5: Distribution of return TTLs as seen from our site. The two modes represent routers that use an initial TTL of 255 and those that use an initial TTL of 64.

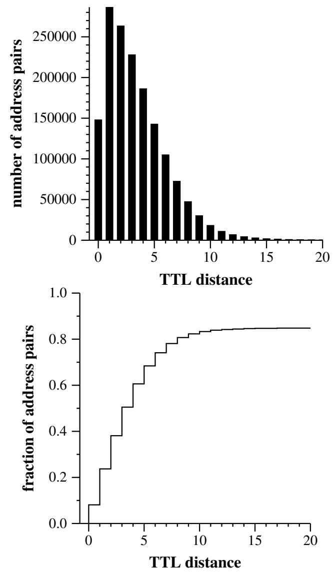

Figure 6: Distribution of the distance between return TTLs as seen from our site. Fewer than 10% of address-pairs share the same return TTL, but the median distance is only 3 hops.

topology, there was a single alias pair at distance three; to catch that alias would only require 2% of all-pairs alias probes.

In Figure 8, we present a different view of efficiency. While choosing pairs to test, the DNS can be used to help find aliases quickly. We sort DNS names "piece-wise reversed" to preserve the hierarchical structure of the names, without actually decoding this structure as we do in the DNS technique. Figure 8 shows the aliases discovered when considering only those addresses that are near each other in this sorted list. Most aliases can be found by considering only those addresses with adjacent names. Of course, selecting pairs to test using DNS does not shorten the alias resolution process, as each pair must be tested, it simply helps make initial data available more quickly.

4.3.2 Comparative Evaluation

In Table 7, we show the number of packets required for each alias resolution approach. We sent 25,822 packets in the course of mapping the PlanetLab overlay before alias resolution; this is an efficient mapping based on the reverse path tree tool from Scriptroute [12]. For the two fingerprint approaches, we show results when using both one and five sources. Mercator uses additional sources to find more aliases, but faces diminishing returns. Ally uses additional sources to narrow the search and gain efficiency. We choose TTL distance in this table to encompass all discovered aliases. To find this set for the experiment, we compared many more pairs of aliases than was necessary; the table reflects a count of only the packets that would be needed in practice. While Ally uses a four packet technique to detect aliases, most candidate alias pairs can be disproven with only two, so on average 2.2 packets are sent for each tested pair. The table assumes that only a single packet is required to look up a hostname; a few more will be needed initially to populate the local cache with referral records.

This table shows that IP identifier-based alias resolution can be much more resource intensive than the rest of map construction, requiring ten times as many packets as mapping itself, and twenty times as many packets as source-address- or DNS-based techniques. With several vantage points, however, the cost can be kept manageable without sacrificing completeness.

4.4 Summary

Table 8 summarizes the relative strengths and weaknesses of the techniques presented in this paper. We add two columns to the list of metrics. The "Unresponsives" column represents the potential to resolve aliases when IP addresses are unresponsive to probe traffic: the inference-based methods succeed where the fingerprinting methods

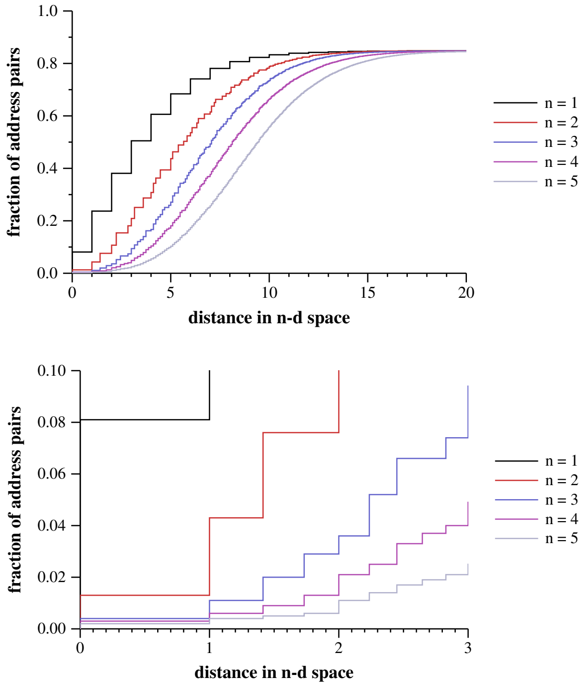

Figure 7: The distribution of the Euclidean distance between return TTLs when measured from one to five vantage points. The lower graph shows the lower left corner of the upper; additional vantage points help tolerate some error in the TTL measurement without many pairwise tests.

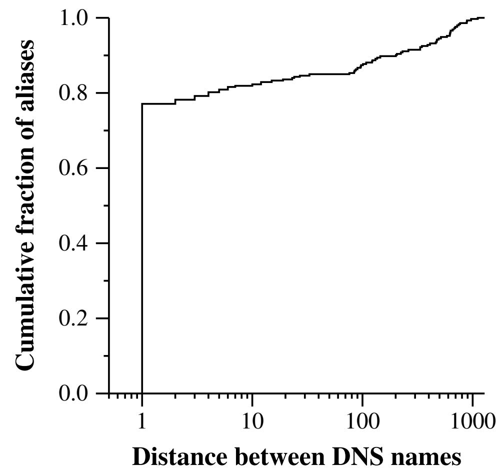

Figure 8: The distribution of aliases found as address pairs with increasingly different DNS names are considered. Most aliases are found when testing address pairs with the most similar names.

cannot. The "Simplicity" column represents the absence of implementation pitfalls. While the rest of the techniques have been straightforward, developing a practical alias test based on IP identifiers (Ally) has been a challenge.

Table 8 shows that each technique has a role. Mercator and DNS can provide efficient, accurate resolution of many aliases. Ally adds completeness for responsive routers at the cost of efficiency and simplicity. Graph adds completeness, especially for unresponsive routers, at some cost to accuracy. Graph's relative inaccuracy may be acceptable when it helps produce a more accurate map by resolving those aliases that cannot be resolved by any other method.

5 Conclusion

Alias resolution is an important component of all traceroute-based Internet mapping efforts — without it, the recovered map does not represent the IP level topology and can be misleading. In this paper, we have presented two new techniques for alias resolution. One is based on recovering information from DNS router names. The other is based on graph inference rules that predict when IP addresses are likely aliases and when they are likely not. Both techniques are of note because, in contrast with existing methods, they do not send probe traffic to routers and so are less dependent on router implementation choices.

We compared the performance of these and existing alias resolution techniques by running a mapping experiment over the topologies of PlanetLab, a wide-area overlay, and Uunet, a large and well-known ISP. The two represent real yet diverse mapping workloads in terms of router equipment makeup, dense versus sparse topology, and scale. We defined metrics for accuracy, completeness and efficiency to compare across methods. To the best of our knowledge, alias resolution techniques have not previously been systematically evaluated.

We hope that our results will help to guide future mapping efforts. Our overall finding is that all of the methods are best used in concert when complete alias resolution is the goal because they have complementary strengths and weaknesses and none is redundant with the others. The new DNS and graph techniques are able to resolve aliases that are unresponsive to probes, thus finding up to a third more aliases than can be found with previous methods. Having multiple methods is also useful to provide a check on the underlying assumptions made by individual methods, and hence improve overall accuracy. DNS

Technique	Intuition	Packets	Per Alias
Mercator	#Addrs × #Srcs		
One-source	#Addrs × 1	1,815	7.5
Five-source	#Addrs × 5	9,075	26.3
Ally	2.2 packets per test, plus #Addrs×#Srcs		
One-source	Test pairs of TTL distance ≤ 1	273,073	467.6
Five-source	Test pairs of TTL distance ≤ 3	52,813	90.4
DNS	#Addresses	1,815	5.5
Graph	No extra packets	0	-

Table 7: We show the efficiency of each technique for the mapping of PlanetLab. For comparison, 25,822 packets were sent in the process of collecting the reverse path trees for PlanetLab.

	Accuracy	Completeness	Efficiency	Unresponsives	Simplicity
Mercator	+	-	+	-	+
Ally	+	+	-	-	-
DNS	+	-	+	+	+
Graph	-	+	+	+	+

Table 8: Summary of strengths (+) and weaknesses (-) of each technique.

resolution is generally accurate but requires knowledge of ISP naming conventions. Graph-based resolution relies on assumptions about ISP network design and so is less accurate by itself, but it has the advantage that it is largely not dependent on router implementation choices, unlike existing probe methods. We also find that existing probe techniques benefit from the use of multiple vantage points, which improves both their efficiency and effectiveness. Further, because probe packets can return multiple pieces of usable information there is a synergy in combining their implementation. Finally, we note that the effectiveness of the methods varies with the mapping task, suggesting that care is needed in applying them. In our case, DNS appears more appropriate for ISP mapping, while graph-based inferences are suitable for overlay mapping.

In the future, we hope to improve the techniques in several respects. Further inference rules or checks may improve the accuracy of the graph technique to the point where it could be used by itself. This would be particularly welcome because the existing probe based methods are reliant on router implementation choices that can easily be altered. Checks between the different methods may result in new techniques for identifying point-to-point versus switched networks, e.g., MPLS and Ethernet switches. Ideally, we would like to reduce the configuration needed to use the DNS technique, perhaps by automatically checking for common naming conventions. Finally, we would like to package an implementation of these techniques as a reusable tool for other network mapping efforts.

Acknowledgements

We thank Ratul Mahajan, Mic Bowman and the PlanetLab team, and Kate Deibel.

References

[1] J. Bellardo and S. Savage. Measuring packet reordering. In *ACM SIGCOMM Internet Measurement Workshop*, 2002.

[2] S. M. Bellovin. A technique for counting NATted hosts. In *ACM SIGCOMM Internet Measurement Workshop*, 2002.

[3] Y. Breitbart, *et al.* Topology discovery in heterogeneous IP networks. In *IEEE INFOCOM*, 2000.

[4] k. claffy, T. E. Monk, and D. McRobb. Internet tomography. In *Nature*, 1999.

[5] R. Govindan and H. Tangmunarunkit. Heuristics for Internet map discovery. In *IEEE INFOCOM*, 2000.

[6] B. Lowekamp, D. R. O'Hallaron, and T. Gross. Topology discovery for large ethernet networks. In *ACM SIGCOMM*, 2001.

[7] V. N. Padmanabhan and L. Subramanian. An investigation of geographic mapping techniques for Internet hosts. In *ACM SIGCOMM*, 2001.

[8] J. Pansiot and D. Grad. On routes and multicast trees in the Internet. In *ACM Computer Communication Review*, pp. 41–50, 1997.

[9] C. Partridge, *et al.* Using signal processing to analyze wireless data traffic. In *ACM Workshop on Wireless Security (WiSe)*, 2002.

[10] N. Spring, R. Mahajan, and T. Anderson. Quantifying the causes of path inflation. In *ACM SIGCOMM*, 2003. (to appear).

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

[12] N. Spring, D. Wetherall, and T. Anderson. Scriptroute: A public Internet measurement facility. In *USENIX Symposium on Internet Technologies and Systems (USITS)*, 2003.