

# How DNS Misnaming Distorts Internet Topology Mapping

Ming Zhang  
*Microsoft Research*

Yaoping Ruan  
*IBM Research*

Vivek Pai, Jennifer Rexford  
*Princeton University*

## Abstract

Network researchers commonly use reverse DNS lookups of router names to provide geographic or topological information that would otherwise be difficult to obtain. By systematically examining a large ISP, we find that some of these names are incorrect. We develop techniques to automatically identify these misnamings, and determine the actual locations, which we validate against the configuration of the ISP's routers. While the actual number of misnamings is small, these errors induce a large number of false links in the inferred connectivity graph. We also measure the effects on path inflation, and find that the misnamings make path inflation and routing problems appear much worse than they actually are.

## 1 Introduction

Network researchers commonly use reverse DNS lookups to infer router locations when extracting network topology and routing behavior, since large ISPs often embed topological or geographic information in the router's DNS name. Using these names, outsiders can infer information that would otherwise require explicit cooperation from the ISP. For example, decoding all of the router names seen during a traceroute can show what cities a network path visits. Previous research has used the information inferred from this approach to map ISP topology [12] and estimate network distance [4, 11, 13].

While this technique has been shown to be useful, errors can occur for a variety of reasons, affecting the conclusions drawn from this data. Router interfaces are often given DNS names manually by network operators, for troubleshooting convenience rather than as a primary addressing mechanism. As routers and line cards are moved, reconfigured, or cycled out of service for repairs or upgrades, and as IP addresses are reassigned across the ISP's network, the DNS information may not be updated, and inferences drawn from it become inaccurate. These naming errors may persist for long periods if they have no effect on normal network operations—the network operators may never need to perform troubleshooting on the incorrectly-named interfaces. However, external researchers attempting to analyze the ISP's network may be affected by these misnamed interfaces.

Without correcting for these DNS misnamings, researchers may get misleading or even conflicting re-

sults when applying inference techniques based on DNS names. We are unaware of any examination of the errors in this approach and their implications. In this work, we present the first systematic study on DNS misnamings, with validated results. Our contributions are as follow:

- We propose ways to detect misnamings, based on observing “abnormal” paths via traceroute. For example, we find stable paths that appear to visit the same point-of-presence (POP) multiple times.
- We develop heuristics for identifying and fixing misnamed addresses by correlating traceroutes from multiple vantage points. We analyze a large ISP and validate against the ISP's router configuration data.
- We examine the topological impact of DNS misnamings. Although DNS misnamings only occur in a small portion (0.5%) of IP addresses, their topological impact is disproportionately larger—we find that 20 out of 182 (11%) edges in a Rocketfuel-like network topology [12] are actually false edges.
- We find that DNS misnaming has an even greater impact on path inflation. Correcting the misnamed addresses reduces the number of unusually long paths by more than 50%.

In the rest of this paper, we describe the system we developed to map the ISP, how we find and resolve the naming problems, and how we determine the impact of these problems on the topology and routing measurements. We have performed these measurements on one large ISP, and have verified with them that the misnamings exist and that our solutions are correct.

## 2 Inferring POP-Level ISP Topologies

To understand how DNS misnaming affects researchers, we discuss how modern ISP networks are constructed, and what complicates the process of inferring their topology. At a high level, an ISP's network is a set of cities that have Points-of-Presence (POPs), and the links that connect these POPs. The POPs contains the routers that connect the ISP's links, and may also provide easy access to links of other peer ISPs and customers. These routers have multiple interfaces with separate IP address, and may also have DNS names configured for reverse lookups. A POP may also have multiple interconnected routers, rather than a single, larger router.

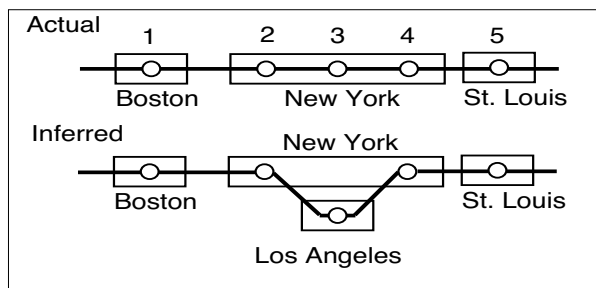


Figure 1: Mismatching causing a POP loop and extra edges. Circles are routers, and rectangles are POPs.

Tools like `traceroute` only report the IP addresses of the interfaces on the forwarding path, but not the POPs traversed. To derive POP-level topology, the interface IP addresses must be mapped to their corresponding POPs. While the network operators have this information readily available to them, external researchers do not, and must use some other means to infer it.

The commonly-used mapping method is to perform reverse DNS lookups on the returned IP addresses, and then extract the city name or POP code that many large ISPs embed in the DNS name. For example, 12.122.12.109 reverse-resolves to `tbr2-p012601.phlpa.ip.att.net`, indicating it is an AT&T router in Philadelphia (`phlpa`), and 144.232.7.42 reverse-resolves to `sl-bb22-nyc-6-0.sprintlink.net`, indicating it is a Sprint router in New York City (`nyc`). By mapping from IP addresses to POPs, researchers can then extract other information, such as what cities are visited along a path and how many routers are traversed in each POP.

DNS misnaming can cause severe errors in inferred topologies, as shown in Figures 1 and 2. In Figure 1, the actual path has one router in Boston and three routers in a POP in New York. The inferred topology has a POP loop because the DNS name of `IP3` is misconfigured with a name that suggests the interface is located in Los Angeles. Figure 2 shows a simple topology consisting of many routers in a large POP in San Francisco, with connections to Seattle and Salt Lake City. Reverse DNS lookup of `IP3` suggests the router is within Seattle while it is actually in San Francisco. DNS misnaming causes four major effects on topology inference:

- **Path inflation:** In Figure 1, the misnamed router induces the POP-level “loop,” making the path appear needlessly inflated, since the Los Angeles round-trip is unnecessary. The effect on inferred path inflation can be severe, particularly for short paths.
- **False edges:** If the NYC POP in Figure 1 does not have any real links to LA, the misnamed router suggests that these POPs are directly connected, adding a false edge to the inferred topology.

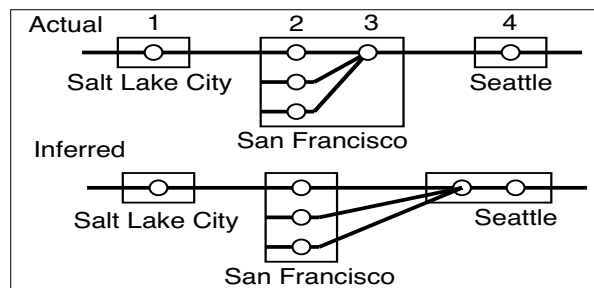


Figure 2: Mismatching can shift routers across POPs, yielding multiple edges between a pair of POPs

- **Extra inter-POP links:** In Figure 2, both ends of the SF–Seattle link are labeled as being in Seattle, causing the dense intra-POP links in SF to appear as multiple links to Seattle. Though technically possible, such redundant links are unlikely, since a smaller number of higher-capacity links would require less hardware and less expense.
- **Missing edges:** If router 3 in Figure 2 were misnamed as another city, such as Los Angeles, then the traceroute path would not contain a direct SF–Seattle connection, causing the inferred topology to miss a real link between the two POPs.

### 3 Data Collection

To map the ISP topology, we perform distributed traceroutes that traverse many paths of the network under study. The reason for distributed traceroutes is not only to improve coverage of the ISP’s links, but also to view mislabeled IP addresses from multiple vantage points.

#### 3.1 Traceroute Measurements

We perform traceroutes from 132 nodes on Planet-Lab [7], across sites in the US, Canada, South America, Europe, Middle East, and Asia. From each node, we perform traceroutes to all 265,448 prefixes in the BGP tables of RouteViews [5], RIPE-NCC [8], and Route-Server [9]. Some of these prefixes are either partially or completely superseded by more specific subnets. To discard these prefixes, we use the algorithm from Mao *et al.* [3] to extract 259,343 routable address blocks. We randomly pick one destination IP address in each block to traceroute and we remove unstable paths caused by routing changes. We modify traceroute to probe only a single destination port to reduce the chance of being accused of port scanning. We also use a blacklist to avoid known prefixes that easily trigger alarms. Data collection spanned 20 hours on March 30, 2005.

To study the misnaming of a specific ISP, we first pick the traceroutes that traverse the target ISP. We use the

BGP tables to map IP addresses to their autonomous systems (ASes). The mapping is constructed by inspecting the last AS, termed the *origin AS*, in the AS path for each prefix [1]. Some IP addresses may map to multiple origin ASes (MOAS) [15], in which case we consider it part of the target ISP if one of the origin ASes is that ISP. With the IP-to-AS mapping, we can then identify all the traceroutes that intersect with the target ISP.

### 3.2 IP-to-POP Mapping

To obtain POP-level information, we perform the reverse DNS lookups of the IP addresses encountered by traceroute, and then use the parsing rules of the *undns* tool [10] to extract POP-level information. Version 0.1.27 of *undns* has parsing rules for 247 ASes. For our target AS, we added four new city names for POP names that were not present in *undns*.

With the POP-level information of an IP, we use the longitude and latitude of the city as an estimate of the geographic location of that POP. We acquire the geographic location through Yahoo maps, by requesting a map of the city/state pair; the latitude and longitude of the city are embedded in the HTTP response. This enables us to calculate the geographic distance between two POPs. We will discuss this in more detail in Section 5.3, where we quantify the impact of misnaming on path inflation.

## 4 Identifying and Correcting Mismappings

In this section, we present our algorithms for identifying misnamed router interfaces and associating them with the correct POPs. We propose two heuristics for detecting and correcting misnamed interfaces.

### 4.1 POP-Level Loop

Normally, a path inside an ISP should not contain a POP-level loop, because ASes typically employ intradomain routing protocols that compute shortest paths using link weights. Inter-POP link weights are usually much larger than those of intra-POP links, to reduce propagation delay and avoid overloading expensive long-haul links. Therefore, for stable paths, the traffic that passes through a POP should not return to the same POP again.

To determine which IP address in a POP-level loop has been mislabeled, we leverage our distributed traceroutes. Misnamed IPs are likely to appear repeatedly in the abnormal paths when we combine the traceroutes from multiple locations. Assuming we have a collection of stable paths with POP-level loops, a simple strategy is to count how many times each IP appears and pick the ones that appear most frequently. However, this strategy may not work well, because it treats all the IPs equally. For example, a correctly-named IP address may appear frequently, simply because it is close to a misnamed IP.

To handle this problem, we assume that most DNS entries are correct and mismappings are infrequent, which we see is true for the ISP we study in Section 5. Therefore, we could resolve all the POP-level loops by fixing only a small number of misnamed IP addresses. We devise a greedy algorithm to solve this problem.

For each abnormal path with a POP-level loop, we have several possible candidates for mismappings. For each interface in the path, we check if we can resolve the loop by mapping this address to a different POP. If so, we consider this IP possibly misnamed. For example, in the inferred path in Figure 1, the second and the fourth IPs are candidates, since we can break the loop by mapping either of them to the Los Angeles POP. The third IP is also a candidate, because we can resolve the loop by mapping it to New York. In this way, we can obtain a set of candidate mismappings for each abnormal path. To select the most likely candidate, we consider all abnormal paths together. Our goal is to identify the minimum set of IPs that needs to be relabeled to resolve all the loops.

The pseudocode of our greedy algorithm for identifying misnamed IPs is shown below. We first compute the candidate set for each abnormal path. Then we greedily pick a candidate IP address that helps to resolve loops for many paths, while at the same time it seldom appears in a path where renaming does not resolve its loop. Finally, we remove the paths whose loops can be resolved by the selected IP and output this IP. This process continues until there are no abnormal paths.

```
For each abnormal path
  Compute the candidate set of misnamed IPs;
While the set of abnormal paths is not empty
  Compute the union of all candidate sets;
  For each candidate IP in the union set
    Count the # of paths where it is in
      their candidate set, CountCandidate;
    Count the # of paths where it appears
      but not in their candidate set,
      CountNotCandidate;
  Pick CandidateIP with the max value of
    CountCandidate - CountNotCandidate;
  Remove all the abnormal paths whose loop
    can be resolved by fixing CandidateIP;
  Output CandidateIP;
```

After identifying the misnamed IP addresses, the next question we want to answer is: can we find the correct POPs of those misnamed IPs by only examining the traceroute data? If so, we can then resolve the mismappings without the ISP's internal data, and supplement the existing topology mapping systems with this DNS name auto-correcting mechanism to achieve higher accuracy.

As we just described, we test if we can resolve a loop by mapping an IP to a different POP. We often have multiple choices—for example in Figure 1, we can map  $IP_4$  to Los Angeles, St. Louis, or any other POP that does not appear in the path to resolve the loop. However,  $IP_4$

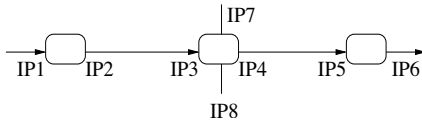


Figure 3: Misnaming leads to router-level discrepancy.

is more likely to be in Los Angeles or St. Louis than in some other random POP because it is connected to both POPs. Therefore, we assign a misnamed IP to a POP by voting based on its neighbors [2]. If the majority of them map to the same POP, we consider it the correct POP for that IP. We assume that routers have more intra-POP links than inter-POP links. Given that inter-POP links span much longer distances and are more expensive, we believe this assumption is true for most major ISPs.

## 4.2 Router-Level Discrepancy

Traceroute usually reports the IP address of the incoming interface of each router on the forwarding path. For example in Figure 3, the traceroute only reports  $IP_1$ ,  $IP_3$ , and  $IP_5$  along the path. Sometimes, we can infer the IP of the outgoing interfaces from that of the incoming interfaces. We take advantage of the fact that the inter-POP links of many major ISPs are high-speed point-to-point links (e.g., Packet-Over-SONET links). This means the IP addresses at the opposite ends of a link are in the same /30 subnet. Among the four IPs in a /30 subnet, the two ending with 01 and 10 are used as router interface addresses while the two ending with 00 and 11 are used as network and broadcast addresses respectively. So, if we know that both  $IP_3$  and  $IP_5$  (144.232.9.149) are backbone routers, we can infer that  $IP_4$  is 144.232.9.150 and obtain its DNS name. Since  $IP_3$  and  $IP_4$  are on the same router, their names should map to the same POP; if not, we call this a *router-level discrepancy*.

We collect all such abnormal IP pairs and assign each IP address to a router. For example, in Figure 3, suppose there are three such pairs,  $(IP_3, IP_4)$ ,  $(IP_3, IP_7)$ , and  $(IP_3, IP_8)$ . We will assign  $IP_3$ ,  $IP_4$ ,  $IP_7$ , and  $IP_8$  to the same router. Then for each router, we decide its correct POP by voting. If the majority of its interfaces map to the same POP, we consider it the correct POP of that router, and the IP that maps to a different POP a misnamed IP. For example, suppose  $IP_4$ ,  $IP_7$ , and  $IP_8$  map to Chicago while  $IP_3$  maps to Detroit, we infer that  $IP_3$  is misnamed and it should map to Chicago.

This heuristic may not work if a router is moved to another POP with none of the DNS names of its interfaces being updated. In practice we have never seen such a case. However, even if this case does occur, it will be most likely to be detected by POP-level loops since there will be many misnamed interfaces. We can resolve it by voting based on its neighbors as described in Section 4.1.

| IP | Wrong POP  | Correct POP | Method |
|----|------------|-------------|--------|
| 1  | WA         | CA          | Loop   |
| 2  | MA         | CO          | Loop   |
| 3  | FL         | CO          | Loop   |
| 4  | CA         | CO          | Loop   |
| 5  | VA         | DC          | 01/10  |
| 6  | VA         | DC          | 01/10  |
| 7  | City A, CA | City B, CA  | Missed |
| 8  | City A, CA | City B, CA  | Missed |
| 9  | City C, PA | City D, PA  | Missed |

Table 1: Summary of all misnamed IPs. Loop: POP-Level Loop, 01/10: Router-Level Discrepancy

We could have also used the IP alias check [12] to detect misnamed IPs, but we may not know the right IP pairs to compare in advance. For example, in Figure 3,  $IP_7$ ,  $IP_8$ , and  $IP_4$  may not appear in the traceroute measurements without using the 01/10 rule. Even if they do appear, we may not know to check IP aliases between  $IP_3$  and  $IP_7/IP_4/IP_8$  because their DNS names look unrelated. The 01/10 rule helps us to quickly identify a small number of abnormal IP pairs and focus on them.

## 5 Case Study on a Large ISP

In this section, we validate our algorithms for identifying and fixing misnamed IPs against the router configuration data for a large ISP. We then study the impact of misnamed interfaces on the inferred topology and path inflation. Although ISPs' naming conventions may be different, the techniques we describe are applicable to other ISPs as well. We plan to study other ISPs in the future.

### 5.1 Validation With Configuration Data

The ISP under study (kept anonymous by agreement) has hundreds of routers and dozens of POPs at different cities around the United States. We first select the traceroutes that traverse the ISP. As described in Section 3.1, we traced to 265,448 prefixes from 132 nodes on PlanetLab. After applying the IP-to-POP mapping, we discovered 113 POPs, which cover most of the ISP's POPs.

Among the traceroutes that traverse the ISP, we find 1,957 paths with non-transient POP-level loops. Using the algorithm described in Section 4.1, we are able to identify four misnamed IPs, which are listed as  $IP_1$ ,  $IP_2$ ,  $IP_3$ , and  $IP_4$  in Table 1. By comparing with the router configuration data, we confirm that these four IP addresses are indeed misnamed. In addition, the voting algorithm in Section 4.1 is able to map those misnamed interfaces to their correct POPs.

Since the ISP is a large backbone provider, most internal links are point-to-point links. We use the router-level discrepancy heuristic described in Section 4.2 to look for

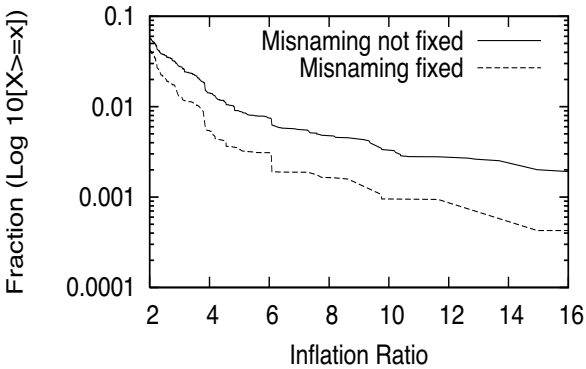


Figure 4: CCDF of path inflation ratio before and after fixing misnamings.

misnamed interfaces in all the non-transient traceroute results. This heuristic allows us to identify two more misnamed interfaces— $IP_5$  and  $IP_6$  in Table 1. We again confirm that these interfaces are misnamed and that our voting algorithm maps them to the correct POPs.

Finally, we check the completeness of our algorithms. Although we are able to identify six misnamed IPs, we fail to detect three misnamings, which are  $IP_7$ ,  $IP_8$ , and  $IP_9$  in Table 1. A closer look at the traceroute data reveals that each of the three IPs has only one neighboring POP and is misnamed to its neighboring POP.  $IP_9$  actually resides in City  $D$ , which is a nearby suburb of the larger City  $C$  in its name; similarly,  $IP_7$  and  $IP_8$  are located in a small City  $B$  near a large POP in City  $A$  in California. There is no way that we can identify these misnamed interfaces based on traceroute measurements. Arguably, this type of misnaming has very limited impact on topology mapping and path inflation, since these are small POPs with a degree of 1 and are misnamed as a big POP that is very nearby.

## 5.2 Impact on Topology Mapping

As discussed earlier in Section 2, misnamed interfaces may lead to false edges in topology mapping. Using the mapping techniques in [12], we find that the six misnamed interfaces ( $IP_1$  to  $IP_6$  in Table 1) lead to *twenty* false edges which do not exist in the real topology. This corresponds to 11% of the total number of inferred edges. We can see that although misnamed IPs are rare, they have a significant influence on topology inference.

Past work relies on the speed-of-light rule to identify false edges [12]. To determine whether a link is false, we first infer the geographic location of the two endpoints of the link from their DNS names. Based on this information, we calculate the shortest time it takes for light to traverse the distance between the two endpoints. Then we estimate the one-way latency of the link using the actual RTT measurements in traceroute. If the latency

estimated from the traceroute is smaller, we know the inferred location of at least one of the endpoints is wrong and the link is false. However, the speed-of-light rule has some limitations. First, it can only identify false edges whose actual distance is shorter than the distance inferred from DNS names. Second, the one-way link latency estimated from traceroute measurements may be inaccurate because of Internet routing asymmetry, queuing delay, or delay in router response. Third, it can only detect misnamed IPs but cannot assign the IPs to the correct POPs. In our dataset, the speed-of-light rule only identifies 1 misnamed IP. In comparison, we discover and fix 6 out of the 9 misnamed IP addresses.

## 5.3 Impact on Path Inflation Studies

Misnamed IPs may inflate the linearized geographic distance of a path, as we explained briefly in Section 2. We now study to what extent misnamings may affect path inflation. As in [13], we compute the *inflation ratio* of a path as the ratio of the linearized distance of a path to the geographic distance between the source and the destination. This ratio reflects how much a path is inflated because of network topology constraints [11].

We calculate the inflation ratio for every possible IP-level path inside the ISP. Figure 4 compares the complementary cumulative distribution function (CCDF) of the inflation ratios, before and after correcting the misnamed IPs. The curves are plotted with a logarithmic scale on the y-axis to emphasize the tail of the distribution. The inflation ratio on the x-axis starts from 2 because we want to focus on the paths that are severely inflated. We can clearly see that a small number of misnamings introduce many unusually long paths. For the paths with inflation ratio over four, more than 50% of them are miscalculated due to misnamed interfaces. We also examined the length of these severely inflated paths. Among the paths with inflation ratio over 2, roughly 60% of them have a direct distance longer than 500 miles. This means their inflated distance is longer than 1000 miles.

## 6 Related Work

The pioneering work of Rocketfuel provides techniques for inferring detailed ISP topology using traceroutes [12]. In their work, they tried to filter out false edges by removing the links whose distance to latency ratio exceeds the speed of light. Although this heuristic helps to remove certain false edges, it may still miss those less obvious ones due to the reasons we discussed in Section 5.2. In a later work, Teixeira *et al.* found that the Rocketfuel topology of Sprint has significantly higher path diversity than the real topology because of extra false edges [14]. Since path diversity directly impacts the resilience of a network to failures, such overestimated path diversity may severely mislead network

designers and operators. They suspected this is due to imperfect alias resolution. However, this still cannot explain the POP-level false edges. Our work complements these existing works by identifying that DNS misnamings could be a major source of POP-level false edges. We also propose ways to fix the misnamings.

## 7 Conclusion & Discussion

We have shown that DNS misnaming, a relatively harmless problem from the network operator's standpoint, can be a much more serious problem for network researchers. A small fraction of misnamed router interfaces gets magnified, leading to a larger fraction of false links in the inferred connectivity graph. These links then cause errors in the path inflation metrics, leading to a mistaken belief that the routing decisions are worse than they actually are. The approaches we have developed to identify and correct the misnaming are able to resolve the problems that have significant impact on topology mapping and path inflation and we have verified them in consultation with a major ISP. Our future plans include conducting similar study on other major ISPs, and to expand the scope of the problems examined.

One of the other inferred metrics that is likely to be affected by these misnamings is path asymmetry [6]. Even if packets traverse the exact same set of links in both directions, the addresses reported by traceroute will differ in the two directions, so a misnaming of a single interface will give the appearance of asymmetric paths. While we are interested in determining how much false asymmetry arises from misnamed interfaces, it requires cooperation at both endpoints to generate and compare traceroute traffic in both directions. Our current infrastructure does not provide this capability, since we do not control the destination endpoint. It may be possible to model a large ISP and use intra-AS routing information to separate the causes of perceived asymmetry, but this effort requires more explicit data from the ISP than we currently have. Our current approach only uses explicit information from the ISP for verification, not for problem identification.

Additionally, misnaming may provide a false sense of security when inferring shared fate of links—mislabeled may give someone the mistaken impression that two paths with the same source and destination traverse different cities, and would therefore not use the same physical POPs. Especially in the cases where real links exist between the cities, even a moderately careful inspection would provide a false impression that the paths do not share fate. In this scenario, misnaming could affect an organization's disaster recovery planning, rather than affecting the analyses of external researchers.

Our larger goal is to raise awareness of this kind of problem so that network researchers performing

inference-based analysis become aware of the possibility that a large number of anomalous results may stem from a small number of input errors, instead of automatically assuming that the network itself is anomalous. Beyond just prodding other researchers to re-examine their approach in using DNS names for topological or geographic data, our longer-term goals are to stimulate new research into automatically detecting and resolving these problems, as well as to identify other research areas where this kind of mislabeling may exist. Given how easily unchecked DNS errors can cause serious misinterpretations of traceroute data, we believe that other network measurement may be similarly affected.

## 8 Acknowledgments

We would like to thank Nick Feamster for his comments on an earlier draft of this paper, and our shepherd Geoff Voelker and the anonymous reviewers for their useful feedback. This work was supported in part by NSF Grants ANI-0335214, CNS-0439842, and CNS-0520053.

## References

- [1] P. Barford and W. Byrd. Interdomain routing dynamics. *Unpublished report*, June 2001.
- [2] N. Feamster, D. G. Andersen, H. Balakrishnan, and M. F. Kaashoek. Measuring the effects of Internet path faults on reactive routing. In *Proc. ACM SIGMETRICS*, June 2003.
- [3] Z. Mao, J. Rexford, J. Wang, and R. H. Katz. Towards an accurate AS-level traceroute tool. In *Proc. ACM SIGCOMM*, 2003.
- [4] V. N. Padmanabhan and L. Subramanian. An investigation of geographic mapping techniques for Internet hosts. In *Proc. ACM SIGCOMM*, Aug. 2001.
- [5] U. of Oregon RouteViews Project. <http://www.routeviews.org>.
- [6] V. Paxson. End-to-end routing behavior in the Internet. In *Proc. ACM SIGCOMM*, Aug. 1996.
- [7] PlanetLab. <http://www.planet-lab.org>.
- [8] RIPE. <http://www.ripe.net>.
- [9] RouteServer. <http://www.bgp4.net/>.
- [10] ScriptRoute. <http://www.scriptroute.org/>.
- [11] N. Spring, R. Mahajan, and T. Anderson. Quantifying the Causes of Path Inflation. In *ACM SIGCOMM*, Aug. 2003.
- [12] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP topologies with Rocketfuel. In *Proc. ACM SIGCOMM*, Aug. 2002.
- [13] L. Subramanian, V. N. Padmanabhan, and R. H. Katz. Geographic Properties of Internet Routing. In *Proc. USENIX Annual Technical Conference*, June 2002.
- [14] R. Teixeira, K. Marzullo, S. Savage, and G. Voelker. In Search of Path Diversity in ISP Networks. In *Proc. Internet Measurement Workshop*, Oct. 2003.
- [15] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. Wu, and L. Zhang. An analysis of BGP multiple origin AS (MOAS) conflicts. In *Proc. Internet Measurement Workshop*, Nov. 2001.