

In Search of Path Diversity in ISP Networks

Renata Teixeira, Keith Marzullo, Stefan Savage, and Geoffrey M. Voelker
Computer Science and Engineering, University of California, San Diego

{teixeira,marzullo,savage,voelker}@cs.ucsd.edu

ABSTRACT

Internet Service Providers (ISPs) can exploit path diversity to balance load and improve robustness. Unfortunately, it is difficult to evaluate the potential impact of these approaches without routing and topological data, which are confidential. In this paper, we characterize path diversity in the real Sprint network. We then characterize path diversity in ISP topologies inferred using the Rocketfuel tool. Comparing the real Sprint topology to the one inferred by Rocketfuel, we find that the Rocketfuel topology has significantly higher apparent path diversity. We evaluate heuristics that improve the accuracy of the inferred Rocketfuel topologies. Finally, we discuss limitations of active measurements techniques to capture topological properties such as path diversity.

Categories and Subject Descriptors

C.2.1 [Computer Communications and Networks]: Network Architecture and Design—*network topology*

General Terms

Measurement

Keywords

Path diversity, Internet topology

1. INTRODUCTION

Path diversity is a metric that reflects the number of routes available to a packet to transit between two points in a network. In the Internet, the route is chosen hop-by-hop by a routing protocol among the available routes. Therefore, the path diversity of a network directly impacts its resilience to failure (how many alternatives are available if a link or router fails), as well as its ability to support traffic engineering. In this paper we consider the particular issue of IP-level path diversity in Internet Service Provider (ISP) networks. We focus on the IP level because it represents the

lowest protocol layer that is exposed to the users¹. Moreover, it is exactly this environment that is faced by designers of value-added high-availability network services – such as routing overlays and peer-to-peer networks [1, 2, 3].

By representing the network topology as a graph, in which routers are nodes and links are edges, one can use standard graph algorithms to calculate the number of node-disjoint and edge-disjoint paths between all pairs of nodes. In practice this is rarely possible. ISPs consider such detailed topology information both confidential and proprietary. As a result, there is considerable literature that simply treats the network as a “black box” and infers aspects of the topology through active measurements [1, 4, 5, 6]. However, Spring et al. have recently developed a tool, Rocketfuel, that derives *precise* topological maps of ISP networks using active measurements [7]. These maps are widely considered to be the most accurate IP-level topologies available to the broad research community.

In this paper, we use Rocketfuel-derived topologies to evaluate the IP-level path diversity between Points of Presence (PoPs) in a various ISPs. Overall, we find that the Rocketfuel topologies fall into two categories, those that have very limited path diversity among PoPs and those that have significant path diversity. For example, the Rocketfuel map for Tiscali, an European ISP, reveals only one PoP-disjoint path among 80% of its PoPs, whereas the Sprint topology inferred by Rocketfuel has at least two paths between all its PoP pairs.

However, our attempts to validate these results revealed at least one instance of substantial inaccuracies in Rocketfuel-derived data. We were able to obtain an exact IP-level topology of the tier-one Sprint ISP network and found that all pairs of PoPs have at least two PoP-disjoint paths between them, and 90% of pairs of PoPs have a minimum of four link-disjoint paths between them. Computing these same metrics over the Rocketfuel-derived topology for Sprint suggests significantly higher diversity. Rocketfuel overestimates PoP-disjoint paths by as much as a factor of 2.5, and link-disjoint paths by as much as a factor of 4.4. This discrepancy results from errors associated with topological inference methodology – the inclusion of links that do not exist and the omission of links that are actually present. Path diversity is particularly sensitive to such inaccuracies since each individual link impacts the results. By comparison, many traditional topological metrics, such as network diameter and average pairwise shortest path [8], are relatively stable in the absence of gross errors. We discuss the potential sources of inaccuracy in topological inference, and evaluate some initial heuristics for improving its accuracy. However,

¹Path diversity at the IP level is distinct from path diversity of the underlying switching or fiber infrastructure. In this work, we do not address physical level path diversity.

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

IMC'03, October 27–29, 2003, Miami Beach, Florida, USA.
Copyright 2003 ACM 1-58113-773-7/03/0010 ...\$5.00.

even with these improvements we find that substantial differences remain and therefore conclude that additional work is needed before Rocketfuel-derived topologies will be appropriate for evaluating sensitive topological metrics.

The remainder of this paper is organized as follows. Section 2 presents related work. Section 3 characterizes path diversity in the Sprint topology and the Rocketfuel ISP topologies. Section 4 discusses sources of inaccuracy for path diversity in the inferred topologies, and evaluates heuristics for improving their accuracy. Section 5 summarizes our experiences and results of the paper.

2. RELATED WORK

There has been extensive research in the area of mechanisms that could take advantage of path diversity. However, little work has been done to characterize path diversity in Internet topologies. There have been a number of efforts to map Internet topologies using active measurements [9, 10, 7]. As part of these studies, these efforts highlight some of the challenges to measuring network topologies. They do not, however, evaluate the accuracy of the resulting topologies as compared to a known correct network topology or the effects of active measurement techniques on the resulting path diversity of the inferred topologies.

Tangmunarunkit et al. [8] identify a set of metrics to evaluate network topologies. One of these metrics is topology resilience, which is defined as the size of the cut set of a balanced bi-partition. Resilience is related to path diversity, but it is too coarse of a metric to characterize path diversity. In their work, they use these metrics to discuss the accuracy of topology models when compared to network topologies generated by active measurements. But, they do not contrast measured topologies to actual network topologies.

Giroire et al. analyze the robustness of IP backbones [11]. They study the resilience of a backbone network to fiber cuts and propose solutions for finding an optimal mapping of a given IP topology to a fiber infrastructure.

Akella et al. [12] study the performance and reliability benefits of using multihoming for high-volume servers and receivers in the Internet. Using traceroute measurements among servers in the Akamai content distribution network, they explore the performance benefits of multihoming where traffic uses first-hop diversity of routing to different ISPs on the first hop to the same destination. Using traceroute measurements between Akamai servers and Keynote Systems nodes in geographically diverse network locations, they also construct trees of paths from multiple Akamai servers to individual Keynote nodes to explicitly represent the path diversity between them. This work focuses on the path diversity enabled just using multihoming and standard IP routing across ISPs, whereas our goal is to characterize the path diversity of complete ISP topologies.

In [13], we state the problem of studying path diversity in Internet topologies and illustrate it with two case studies: inside an ISP and across multiple ISP networks. In this paper, we focus on path diversity inside an ISP network. We expand our definition and analysis of path diversity, and discuss the limitations of computing path diversity using topologies generated by active measurements.

3. PATH DIVERSITY IN ISP TOPOLOGIES

Routing inside an ISP network is performed by Interior Gateway Protocols (IGP) such as OSPF [14] and IS-IS [15]. In general, IGPs forward packets using the shortest path to the destination. Despite this prevalent policy of using a single path, there are some opportunities for leveraging path diversity inside a network. Routing protocols use multiple paths between two points in the network either

when there is a failure or when there are multiple paths with equal cost. In the former case, traffic from a failed path is switched to a backup path after the failure. In the latter, multiple paths are used simultaneously to spread out traffic to a destination to balance load across paths and potentially provide higher bandwidth to the destination.

Equal-cost multipath routing is available both in OSPF [14] and IS-IS [15]. When there are multiple paths to a destination, routers typically choose which interface to forward a packet based on a hash function of some of the fields that identify a TCP flow. The implementation of this hash function is vendor-specific. ISPs are responsible for assigning weights to links to determine which paths have equal cost, but they have no control over which specific path a packet is forwarded on.

In this section, we characterize the path diversity in the real Sprint topology and the inferred Rocketfuel ISP topologies. We identify two kinds of path diversity. One can consider one ISP full topology, i.e., all IP paths available in the absence of failure or routing constraints (including backup links). We call this type of path diversity *complete*. Complete path diversity, however, is not always available to traffic traversing the ISP network because of link failures and network engineering practices (some paths may be used only for backup purposes). We call the path diversity available at a given moment *active*.

3.1 Path Diversity in the Sprint Network

We were fortunate to have access to the exact continental Sprint US PoP-level topology under a non-disclosure agreement². This network is composed of 17 PoPs scattered throughout the US. PoPs are interconnected by multiple high-bandwidth optical links (OC-48 and OC-192), not all of which are active at a given time. Links represent IP links between PoPs. Sprint provided annotations in the topology that determine which links are used only as backup. In this topology, we compute active path diversity by ignoring backup links. Even though this is a static topology and consequently, its path diversity represents the amount of path diversity available in the Sprint network in the absence of failures.




Figure 1: Path diversity in the Sprint Network.

Figure 1 shows the cumulative distribution function of the number of link- and PoP-disjoint paths for all pairs of PoPs in the Sprint network. The set of PoP-disjoint paths for a given pair of PoPs is a subset of the link-disjoint paths for the same pair. Consequently, the number of PoP-disjoint paths is no greater than the number of

²An approximate topology can be found at <http://www.sprint.net>

link-disjoint paths. Since PoPs can fail (e.g., from power failures), PoP-disjoint paths represent the highest amount of availability.

We present both active and complete path diversity. The solid line represents the maximum possible number of paths for each pair of PoPs as determined by the minimum between the out-degree of the source PoP and the in-degree of the destination PoP. Measured path diversity close to this limit indicates that path diversity is limited only by the degree of the PoPs.

We can see in Figure 1 that the Sprint network has a large path diversity. When considering both link- and PoP-disjoint paths, all pairs of PoPs have more than one path. The significant difference between the number of link- and PoP-disjoint paths is expected because there are multiple links connecting a pair of PoPs. The number of link-disjoint paths is relatively large. Almost 90% of pairs of PoPs have at least four link-disjoint paths between them, and 40% of pairs of PoPs are connected through eight or more link-disjoint paths (these are generally the most important paths as defined in [11]). Active path diversity represents the potential for spreading traffic in the Sprint network, whereas complete path diversity represents the limits of potential availability (i.e., how many links need to fail before the network is partitioned).

Since ISP topologies are confidential, and we have only been able to obtain access the Sprint topology, we need to use topologies generated by end-to-end measurements for other ISPs. We now consider those generated by the Rocketfuel tool.

3.2 Path Diversity in Rocketfuel Topologies

Rocketfuel topologies are constructed from traceroute-like probes, which identify the incoming interfaces of the routers traversed by probes [7]. After collecting all links found by probes, Rocketfuel applies techniques to map different interfaces in the same router to a single node in the topology.

Because the Sprint topology used above was at the PoP level, we focus here on inter-PoP path diversity in Rocketfuel topologies. Since PoPs represent physical locations, we group all routers located in the same city in the Rocketfuel topologies into a PoP and consider only links between routers in different PoPs.

We have measured the path diversity for all the topologies described in [7]. Based on our results, we divide this set of topologies into two groups: *low* contains topologies with a significant percentage of pairs of PoPs with only one path between them, and *high* contains topologies with a higher level of path diversity. Here we only present results for one representative topologies of each group. We select Tiscali, an European ISP, to represent the first group and Sprint to represent the other. To distinguish the Sprint topology inferred by Rocketfuel from the real Sprint topology, we call the former *Sprint_{inf}* and the latter *Sprint_{real}*.

For the sake of comparison, we eliminate PoPs in *Sprint_{inf}* that are not present in *Sprint_{real}*. We also remove non-European PoPs from the Tiscali topology so that both topologies contain only the main PoPs of their networks. The resulting *Sprint_{inf}* topology has 17 PoPs and the Tiscali topology has 49 PoPs. We compute path diversity between every pair of PoPs in both topologies.

Figure 2 shows the cumulative distribution function of the number of link- and PoP-disjoint paths for every pair of PoPs. There is a large difference between the amount of path diversity in the two topologies. The *Sprint_{inf}* topology has at least two PoP-disjoint paths between all of its PoPs, whereas Tiscali has only one path between 80% of its PoPs. The difference is more pronounced for link-disjoint paths. *Sprint_{inf}* has at least five link-disjoint paths between all of its PoPs and Tiscali has less than five link-disjoint paths between almost 85% of its PoPs.

One can only compute active path diversity from Rocketfuel-




Figure 2: PoP-level path diversity in the Sprint and Tiscali topologies inferred by Rocketfuel.

generated topologies. Hence, the lack of path diversity in the Tiscali topology does not mean that a single link or router failure will necessarily disconnect the network. Most likely, Tiscali has backup links that can be used in the case of failure or use optical-level protection mechanisms. But the results do indicate that Sprint has a more robust network because it has a higher path diversity. Sprint has also more potential for balancing traffic between PoPs as indicated by the number of link-disjoint paths.

More surprising is the difference between the diversities of the *Sprint_{inf}* and *Sprint_{real}* topologies³. Because active measurements do not capture all the links in the network, we expected that the path diversity in the *Sprint_{inf}* topology would be lower than that of the *Sprint_{real}* topology. More specifically, we expected the set of disjoint paths between a pair of PoPs in *Sprint_{inf}* to be a subset of the set of disjoint paths (including backup links) found for the same pair of PoPs in *Sprint_{real}*. Comparing Figures 1 and 2, we see that, instead, *Sprint_{inf}* has higher path diversity. We discuss this issue in the next section.

4. ACCURACY OF PATH DIVERSITY IN INFERRERD TOPOLOGIES

Topologies generated from active measurements are an approximation of the real ISP network. There may be actual links that are not in the generated topology, and there may be links in the generated topology that do not actually exist. As a metric path diversity is particularly sensitive to such errors, since each error changes the path diversity by one for at least one pair of PoPs. This contrasts with other metrics [8], such as node diameter and average pairwise shortest path, which are relatively insensitive to a small number of erroneous links.

4.1 Comparison of Inferred and Real Topologies

Let $L_I(i, j)$ and $P_I(i, j)$ be the number of link- and PoP-disjoint paths, respectively, in *Sprint_{inf}* for a pair of PoPs (i, j) . Let $L_R(i, j)$ and $P_R(i, j)$ be the corresponding values for *Sprint_{real}*. We further investigate the discrepancies in path diversity between

³The *Sprint_{inf}* and *Sprint_{real}* topologies were collected approximately six months apart, but the Sprint topology did not change much during that period.




Figure 3: Comparison of the number of link-disjoint paths in $Sprint_{inf}$ and $Sprint_{real}$.




Figure 4: Comparison of the number of PoP-disjoint paths in $Sprint_{inf}$ and $Sprint_{real}$.

the two topologies by computing the ratios:

$$RL_I(i, j) = L_I(i, j)/L_R(i, j), \text{ and}$$

$$RP_I(i, j) = P_I(i, j)/P_R(i, j) \text{ for every pair of PoPs } (i, j).$$

Figure 3 presents the cumulative distribution of $RL_I(i, j)$ and Figure 4 presents the cumulative distribution of $RP_I(i, j)$ for all pair of PoPs (i, j) . When $RL_I(i, j)$ or $RP_I(i, j)$ is 1, the path diversity measured for (i, j) in $Sprint_{inf}$ matches that of $Sprint_{real}$. When $RL_I(i, j)$ or $RP_I(i, j)$ is greater than 1, $Sprint_{inf}$ overestimates path diversity, and when $RL_I(i, j)$ or $RP_I(i, j)$ is less than 1, $Sprint_{inf}$ underestimates path diversity.

Figure 3 shows that the number of link-disjoint paths in $Sprint_{inf}$ varies from 1.25 to 4.4 times the number of link-disjoint paths in $Sprint_{real}$. On average, there are 158% more link-disjoint paths in $Sprint_{inf}$ than in $Sprint_{real}$. $Sprint_{inf}$ underestimates path diversity for 9% of the pairs of PoPs, but, averaged over all pairs of PoPs, it overestimates by 37%.

An inspection of both topologies shows that there is a considerable number of false links in the Rocketfuel topology. There are two kinds of false links in $Sprint_{inf}$: links between PoPs that are not connected in $Sprint_{real}$ and extra links between PoPs that are connected. Both kinds of false links explain the higher path diversity in the inferred topology. False links between PoPs that are connected in $Sprint_{real}$ only increase the number of link-disjoint

paths, whereas false links connecting PoPs that are not connected in $Sprint_{real}$ increase both the number of link- and PoP-disjoint paths.

We discuss the origins of these false links and other limitations of active measurements techniques for capturing path diversity in the next section.

Some of the false links are easy to identify as being false. For example, there are inter-PoP links in $Sprint_{inf}$ connecting an access router in one city to a backbone router in another city. Such links cannot exist in the Sprint network. In fact, in our PoP-level topologies, all inter-PoP links should be between backbone nodes. Therefore, we remove all links connected to access nodes from $Sprint_{inf}$. We call this new version of the topology $Sprint_{inf-bb}$. As another example, in [16], the authors of Rocketfuel present techniques for inferring IGP weights in ISP networks using the Rocketfuel topologies. For this follow-up work, they used a modified version of the Rocketfuel topologies. In this new version that we call $Sprint_{inf-rev}$, they eliminated all links for which there was not a corresponding reverse link.

To quantify the impact of these heuristics on improving the accuracy of path diversity in the Rocketfuel topologies, we compute four new ratios for the link- and PoP-disjoint paths in the corresponding $Sprint_{inf-bb}$ and $Sprint_{inf-rev}$ topologies:

$$RL_{I-bb} = L_{I-bb}(i, j)/L_R(i, j), \text{ and}$$

$$RP_{I-bb} = P_{I-bb}(i, j)/P_A(i, j) \text{ for } Sprint_{inf-bb} \text{ and}$$

$$RL_{I-rev} = L_{I-rev}(i, j)/L_R(i, j) \text{ and}$$

$$RP_{I-rev} = P_{I-rev}(i, j)/P_R(i, j) \text{ for } Sprint_{inf-rev}.$$

Figures 3 and 4 present the cumulative distributions of these ratios. From these figures, we see that both heuristics reduce the difference in the amount of path diversity, but even after applying them they overestimate the path diversity for the vast majority of pair of PoPs. It is interesting to note that, after correcting some of the false links, path diversity is underestimated for some pairs of PoPs. Path diversity lower than that of $Sprint_{real}$ is a reflection of links that Rocketfuel was unable to capture with measurements. $Sprint_{inf-rev}$'s path diversity is the closest to the path diversity found in $Sprint_{real}$, yet it only matches the number of link-disjoint paths of $Sprint_{real}$ for approximately 20% of the pairs of PoPs and 28% of the pairs when considering PoP-disjoint paths.

4.2 Active Measurements and Path Diversity

Clearly, the difference in path diversity presented in the previous section is specific to the Sprint topology inferred by Rocketfuel. Even though Rocketfuel is an improvement in measuring ISP topologies, active measurement techniques still have limitations that prevent them from accurately capturing path diversity. These limitations are an obstacle for using such topologies as basis for studying traffic engineering or robustness of ISP networks.

The process of discovering the topology of an ISP network can be divided into *probing* the ISP network to discover links and *processing* those links to infer the ISP topology. Links that are kept only as backups are not going to be traversed by probes. Topologies generated from active measurements cannot contain such links unless there is a failure during the measurement period that causes probes to be forwarded through backup links. Consequently, inferred topologies capture mostly active path diversity, and in some cases part of the complete path diversity. In the case of Sprint network, complete path diversity is only 5% more than active on average. Therefore, in this section, we chose the topology of the Sprint network without backup links as a basis of comparison.

Below we discuss some sources of inaccuracy introduced in each step of generating ISP topologies from active measurements and how they might be addressed. As an example, we analyze the percentage of links that were missed and added both during probing the Sprint network and processing the probed links. Table 1 summarizes the percentage of inter-PoP links missed and added at each step, computed over the total number of links in $Sprint_{real}$. The last column presents the overall percentage of missed and added links by combining both steps. We computed this difference by comparing inter-PoP links from $Sprint_{real}$ to $Sprint_{inf}$.

Links	Probing	Processing	Overall
Missed	58.57%	-	8.57%
Added	37.14%	717.14%	608.57%
- Reverse links	-	80%	-
- PoPs not connected	5.71%	285.71%	117.14%
- Connected PoPs	31.43%	351.43%	491.43%

Table 1: Percent of missed and added links in each step of generating the Rocketfuel PoP topology relative to the number of links in the $Sprint_{real}$ topology.

4.2.1 Inaccuracies introduced during probing

We compare $Sprint_{real}$ without backup links with the topology induced by the raw Rocketfuel traces⁴ to quantify inaccuracies introduced during probing. We select links from the traces that connect different PoPs and that are also present in $Sprint_{inf}$. We call the topology induced by the set of links discovered by probes as $Sprint_{traces}$. Table 1 shows that the Rocketfuel probes missed almost 60% of the inter-PoP links in the Sprint network. Active measurements miss links for two main reasons:

- **Lack of vantage points:** End-to-end measurements only capture links traversed by probes. A *vantage point* is a host that is the source of probes. Increasing the number of vantage points and varying their location increases the probability of finding links. If there are two paths with the same cost between two points in the ISP network and probes are always hashed to the same path, then all links in the alternate path are missed. Probes from another vantage point may find alternative equal-cost paths. Missing links may cause the resulting topology to understate the active path diversity.

To capture all inter-PoP links, it is important that probes enter the ISP network in every PoP and exit to all the other PoPs. Capturing all the path diversity inside a PoP is harder. Probes should enter and exit the ISP network through each access router in a PoP.

- **Incomplete traceroutes:** Probes may be incomplete, i.e., there may be some unknown hops (“*”) in the traceroutes. In this case, Rocketfuel conservatively ignores incomplete parts of the traces. Examining the traces for the Sprint network, we observed that there were not many instances of non-responses to probes in the segment of the trace that traverse the core of the Sprint network. Consequently, we do not believe that is this a main source of missing links in $Sprint_{traces}$.

More surprising, $Sprint_{traces}$ contain a number equivalent to 37.14% of the total number of links in $Sprint_{real}$ that do not exist in the real network even when considering backup links. Almost all

⁴Traceroutes used to generate the topologies are available at: <http://www.cs.washington.edu/research/networking/rocketfuel/>.

of those extra links are between a pair of PoPs that are connected in the real network. This may arise because of network topology changes during the probes. A smaller fraction of links, equivalent to 5.71% of all links in $Sprint_{real}$, connects PoPs that have no connection in the real network. Those extra links could have been caused by incorrect DNS resolution.

The two main causes of finding extra links during probing are:

- **Changes in the path of a probe:** The process of probing takes time during which its path may change. The path a probe takes may change for two reasons: a change in the network topology itself, or a routing policy change that impacts the best forwarding path. ISPs are constantly maintaining and upgrading their networks by adding and removing links from the topology. Failures cause some links to be removed from the topologies and may cause some backup links to become active.

Changes in the network topology may lead to false connectivity. If the forwarding path to the destination of the probe changes during the traceroute, then it may appear that there is a link between two routers that are not connected in the real network. This is one possible reason for false links in the Rocketfuel version of the Sprint topology. Eliminating links that do not have corresponding reverse links, as was done for the $Sprint_{inf}$ -rev topology, should greatly reduce such false connectivity.

- **Incorrect DNS:** DNS names are used both to determine which links belong to a particular ISP network and to map IP addresses to PoPs. Some addresses may be incorrectly mapped due to misconfiguration or change in the ISP topology that have not been reported to DNS, thereby leading to the inclusion of a false link in the traces.

4.2.2 Inaccuracies from processing probed links

After capturing as many links as possible using active measurements, tools like Rocketfuel process the resulting traces to produce a more accurate map of the ISP network. In order to assess the effects of processing the traces on the final topology, we compare the percentage of inter-PoP links connecting every pair of PoPs in $Sprint_{traces}$ with the percentage in $Sprint_{inf}$. Table 1 shows that processing the traces adds more than seven times as many links as the real network for two main reasons:

- **Alias resolution:** After collecting links, measurement tools perform a series of tests to resolve aliases. If two interfaces that belong to the same router are not accurately resolved to the same node, then the resulting topology will have a larger number of router-disjoint paths than in the real network. On the other hand, if two interfaces in different routers are mistakenly mapped to the same router, then the resulting topology will have more link-disjoint paths.

Rocketfuel introduced a number of new techniques to resolve aliases. One of these techniques is to use information in DNS names. An interface’s DNS name sometimes encodes the router that has the interface, and at other times encodes the router at the other end of the link. This difference in convention can lead to incorrect grouping of interfaces. Grouping two interfaces in different routers will result in a router with more connectivity, thereby overstating the number of link-disjoint paths. In Table 1, extra links between PoPs that are connected in $Sprint_{real}$ represent interfaces that were not correctly mapped to the same router in a PoP, whereas links

in $Sprint_{inf}$ between PoPs that have no connection in the real network represent either interfaces in different routers that were mistakenly collapsed or a router that has an incorrect DNS name.

- **Adding reverse links:** Once a probe found a link between two routers, it is tempting to add a reverse edge between them. Indiscriminately adding reverse links, however, may inflate the number of links in the inferred topology as compared to the number of links in the real network. Traceroute reports the incoming interfaces of routers. Hence, if two different vantage points probe the same link from different directions, then it will appear as two different links. If the alias resolution mechanism was perfect, then it would map all the interfaces to the same router and considering all links to be bi-directional would be a safe assumption.

After the alias resolution step, Rocketfuel completes the topology by adding reverse links. This step introduces a number of false links to $Sprint_{inf}$ that is equivalent to 80% of the links in the real network. This result suggests that, until we have more accurate techniques for alias resolution, links should be considered directed in the traces.

Given how sensitive path diversity is to false or missing links, accurately represent path diversity in ISP networks using active measurements remains an open problem. We are working on improving the accuracy of path diversity in the Rocketfuel topologies, and determining the limits to which active measurement techniques can capture such topological properties.

5. CONCLUSION

In this paper we describe our efforts to characterize path diversity in ISP network topologies. First, using the real Sprint IP topology, we find that Sprint has significant path diversity. We then characterize path diversity in a larger set of ISP networks using the recent Rocketfuel ISP topologies inferred from active measurements. However, comparing the real Sprint topology to the one inferred by Rocketfuel, we find that the Rocketfuel topology has significantly higher apparent path diversity. We then evaluate two heuristics that improve the accuracy of capturing path diversity in the topologies, but substantial differences remain.

Although the Rocketfuel topologies have benefitted from some evaluation and validation, this is the first time that they have been systematically compared to a real topology and evaluated from the perspective of path diversity. Due to constraints faced by active measurements, these topologies both contain false links and miss actual links. We show that the post-processing of the traces (in particular, the alias resolution step) introduces a huge number of false links in the topology. As a metric, path diversity is particularly sensitive to such errors since each error changes the path diversity by one for at least one pair of PoPs. We are working on improving the accuracy of path diversity in the Rocketfuel topologies, and determining the limits to which active measurement techniques can capture such topological properties.

Acknowledgments

We thank Christophe Diot, Antonio Nucci, Konstantina Papagianaki, and Nina Taft for generously providing the Sprint topology. We also thank Neil Spring and Ratul Mahajan for the insightful discussions on the Rocketfuel topologies. Support for this work was provided by DARPA FTN Contract N66001-01-1-8933. Renata Teixeira was also supported by a fellowship from Capes/Brazil.

6. REFERENCES

- [1] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris, "Resilient Overlay Networks," in *Symposium on Operating Systems Principles*, pp. 131 – 145, Oct. 2001.
- [2] A. Keromytis, V. Misra, and D. Rubenstein, "'sos: Secure overlay services'," in *Proceedings of ACM SIGCOMM*, (Pittsburgh, PA), Aug. 2002.
- [3] I. Clarke, O. Sandberg, B. Wiley, and T. W. Hong, "Freenet: A Distributed Anonymous Information Storage and Retrieval System," in *Proceedings of the ICSI Workshop on Design Issues in Anonymity*, (Berkeley, CA), July 2000.
- [4] P. Francis *et al.*, "IDMaps: A Global Internet Host Distance Estimation Service," *IEEE/ACM Trans. on Networking*, Oct. 2001.
- [5] T. Ng and H. Zhang, "'towards global network positioning'" in *Internet Measurement Workshop*, Nov. 2001.
- [6] A. Rowstron and P. Druschel, "Storage management and caching in PAST, a large-scale, persistent peer-to-peer storage utility," in *Proceedings of ACM SOSP*, (Lake Louise, Alberta), Oct. 2001.
- [7] N. Spring, R. Mahajan, and D. Wetherall, "Measuring ISP Topologies with Rocketfuel," in *Proceedings of ACM SIGCOMM*, (Pittsburgh, PA), Aug. 2002.
- [8] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "Network Topology Generators: Degree-Based vs. Structural," in *Proceedings of ACM SIGCOMM*, (Pittsburgh, PA), pp. 147 – 159, Aug. 2002.
- [9] R. Govindan and H. Tangmunarunkit, "Heuristics for Internet Map Discovery," in *Proceedings of IEEE INFOCOM*, (Tel Aviv, Israel), Mar. 2000.
- [10] Cooperative Association for Internet Data Analysis, "Skitter." <http://www.caida.org/tools/measurement/skitter/>, Jan. 2000.
- [11] F. Giroire, A. Nucci, N. Taft, and C. Diot, "Increasing the Robustness of IP Backbones in the Absence of Optical Level Protection," in *Proceedings of IEEE INFOCOM*, (San Francisco, CA), Mar. 2003.
- [12] A. Akella *et al.*, "A Measurement-Based Analysis of Multihoming," in *Proceedings of ACM SIGCOMM*, (Karlsruhe, Germany), Aug. 2003.
- [13] R. Teixeira, K. Marzullo, S. Savage, and G. M. Voelker, "Characterizing and Measuring Path Diversity of Internet topologies," in *Proceedings of ACM SIGMETRICS (extended abstract)*, (San Diego, CA), June 2003.
- [14] J. Moy, "OSPF Version 2." RFC 1247, July 1991.
- [15] D. Oran, "OSI IS-IS Intra-domain Routing Protocol." RFC 1142, Feb. 1990.
- [16] R. Mahajan, N. Spring, D. Wetherall, and T. Anderson, "Inferring Link Weights using End-to-End Measurements," in *Proceedings of Internet Measurement Workshop*, (Marseille, France), Nov. 2002.