Measured Impact of Crooked Traceroute

Matthew Luckie
Computer Science
University of Waikato
Hamilton, New Zealand
mjl@wand.net.nz

Amogh Dhamdhere, kc claffy
CAIDA
University of California, San Diego
La Jolla, CA, USA
{amogh,kc}@caida.org

David Murrell
Computer Science
University of Waikato
Hamilton, New Zealand
dmurrell@waikato.ac.nz

ABSTRACT

Data collected using traceroute-based algorithms underpins research into the Internet's router-level topology, though it is possible to infer false links from this data. One source of false inference is the combination of per-flow load-balancing, in which more than one path is active from a given source to destination, and classic traceroute, which varies the UDP destination port number or ICMP checksum of successive probe packets, which can cause per-flow load-balancers to treat successive packets as distinct flows and forward them along different paths. Consequently, successive probe packets can solicit responses from unconnected routers, leading to the inference of false links. This paper examines the inaccuracies induced from such false inferences, both on macroscopic and ISP topology mapping. We collected macroscopic topology data to 365k destinations, with techniques that both do and do not try to capture load balancing phenomena. We then use alias resolution techniques to infer if a measurement artifact of classic traceroute induces a false router-level link. This technique detected that 2.71% and 0.76% of the links in our UDP and ICMP graphs were falsely inferred due to the presence of load-balancing. We conclude that most per-flow load-balancing does not induce false links when macroscopic topology is inferred using classic traceroute. The effect of false links on ISP topology mapping is possibly much worse, because the degrees of a tier-1 ISP's routers derived from classic traceroute were inflated by a median factor of 2.9 as compared to those inferred with Paris traceroute.

Categories and Subject Descriptors

C.4 [**Performance of Systems**]: Measurement techniques

General Terms

Measurement

Keywords

Traceroute, Internet topology

1. INTRODUCTION

Data collected using traceroute-based algorithms underpins research into router-level aspects of the Internet's topology and is vital to developing fine-grained models of Internet reachability, performance, structure, and growth. Macroscopic Internet topology mapping projects use traceroute-based algorithms [1] to infer forward IP paths toward mil-
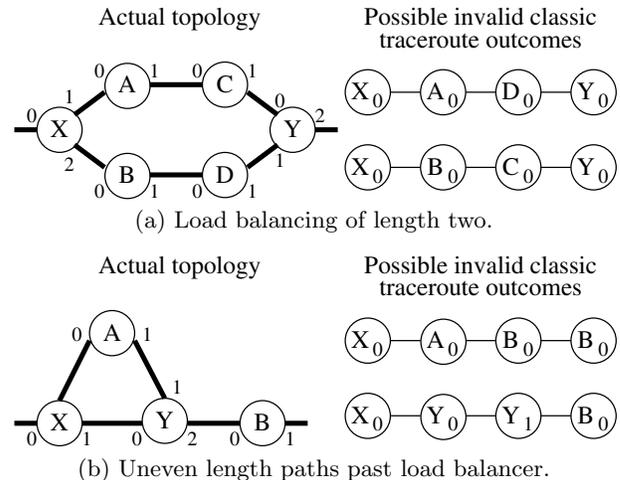

Figure 1: Load balancing can cause false inference of router-level links. In 1(a) links $A_0 \to D_0$ and $B_0 \to C_0$ are invalid as the routers are not adjacent. In 1(b) $A_0 \to B_0$ is invalid as the routers are not adjacent, and $Y_0 \to Y_1$ is invalid as the addresses are aliases.

lions of destinations [2, 3, 4], and ISP-mapping tools use traceroute to develop maps of individual ISPs [5]. To infer the forward IP path, traceroute algorithms use a series of probe packets, incrementing the TTL value with each probe to elicit ICMP time exceeded responses from consecutive routers in the path. An *IP-interface graph* derives from inferring IP links between interface addresses in packets that time out at adjacent hops. An implicit assumption is that the addresses represent distinct neighbouring routers.

Classic traceroute can infer false links if a device that load-balances traffic is in the path [6]. The most common methods of load-balancing traffic are: (1) per-destination, where the path depends only on the destination IP address in the packet; and (2) per-flow, where the path depends on the source and destination IP addresses, IP protocol type, and the first four bytes of the transport header [7] – the source and destination ports in TCP and UDP, and the type, code, and checksum in ICMP. Because classic traceroute sends UDP probes with different destination port values, or ICMP probes with different checksum values, per-flow load balancers can forward consecutive packets on different paths leading to inference of false links. Figure 1 illustrates the two cases of load-balancing that will induce such false

link inference: (a) a path after a load-balancer is at least two hops in length before the paths converge; or (b) the length of paths after a load-balancer is uneven. In each case there are four possible topology inferences, two containing false links, so classic traceroute has a 50% chance of inferring a false link in the topologies illustrated.

This limitation hinders not only efforts to create accurate IP-level graphs, but also impedes the accuracy of several topology inference techniques based on traceroute data, from statistical metrics such as router degree distribution [8] and path diversity [9], to methodologies for extracting router and AS-level graphs [2, 10, 11, 12, 13]. Although we cannot identify – or even estimate the proportion of – false links in data previously collected with classic traceroute without corresponding ground truth data, we can design an experiment to estimate their impact. We simultaneously collected topology data using classic and other traceroute techniques [6, 7] cross-validating different sources of data to identify false inferences specific to classic traceroute.

2. IDENTIFYING AND ANALYSING CLASSIC TRACEROUTE ARTIFACTS

Our method has two steps: identifying links that are inferred but not actually travelled by packets in flows, which we call *link artifacts*, and assessing whether such links actually exist. More precisely, we define a link to be a measurement artifact of classic traceroute if it is not also inferred with a traceroute algorithm that enumerates per-flow load-balanced paths toward the destination. Augustin, *et al.*'s multipath detection algorithm (MDA) attempts to detect load-balanced paths between a source and destination by varying the first four bytes of the transport header in consecutive traceroute packets [7]. Per-flow load-balancers using these fields in forwarding decisions will select different next hops for consecutive packets. The MDA traceroute algorithm assumes packets are distributed evenly amongst next hops, and subsequently sends enough probes to obtain a level of confidence that it has observed all next hops from a given router interface in the path. We use the MDA traceroute algorithm with a 99% confidence level to enumerate per-flow load balanced paths toward a destination. To avoid our MDA traceroute packets appearing like a port scan, we vary the source port and hold the destination port constant; this also has the benefit of accommodating firewalls [14] that permit traceroute probes to a small range of ports above the base port (33435) that classic traceroute uses [1]. To minimise triggering intrusion detection systems, we halt probing after three consecutive unresponsive hops, send at most two probes per hop, and wait at least five seconds between completing one traceroute and beginning the next to the same destination.

Identifying a link artifact using this method does not prove the link does not exist – it just proves that router interfaces were visited in a sequence that is not visited by packets in flows towards the destination. The disparity arises because most routers use the address of the incoming interface as the source address in an ICMP time exceeded response; the same sequence of routers might have been visited by MDA traceroute but via different interfaces. Therefore, a measurement artifact may represent a valid router-level link.

We can use IP address alias resolution [15, 16] techniques to resolve such conflicts. Specifically, if we can establish that the same sequence of routers in a link artifact was visited by MDA traceroute, the link artifact actually exists. Figure 2 illustrates this heuristic; the first graph is an interface graph inferred with MDA traceroute, and the second graph is an interface graph inferred with classic traceroute. In the second graph, $A \to E$ is a measurement artifact because it is not also in the MDA graph, which attempts to reflect all possible next hops after A to a 99% confidence level. If A and B are aliases as in the third graph, or if C and E are aliases as in the fourth graph, or if C and F are aliases, then $A \to E$ is a valid router-level link. If we do not find a corresponding router-level link in the MDA graph, a link artifact is declared a false router-level link, provided A and E are responsive to alias resolution probes; that is, we received responses from them and their IP-ID appears to use a counter [5]. We can also declare a false link if the IP addresses of the artifact are aliases, caused by a combination of classic traceroute sending packets belonging to different flows and uneven length paths past a per-flow load balancer as in figure 1(b).

Because we might also incorrectly infer a link as a result of a routing change, we collect our classic and MDA traceroute data for each destination between measurements of topology using Paris traceroute [6]. We infer that routing was stable for our classic and MDA traceroutes if neither Paris traceroute reports a link not found by MDA traceroute. The sequence of measurements to each destination is thus: (1) initial Paris traceroute, (2) classic traceroute to identify measurement artifacts, (3) MDA traceroute to enumerate all links to 99% confidence, (4) final Paris traceroute to infer if a path has changed. Previous work [17] found 9% of paths to be unstable on a time scale of tens of minutes; we use Paris traceroute to infer unstable paths, so we can focus on artifacts due to load balancing.

We use Ally [5] and Mercator [18] to resolve aliases.[1] Ally infers aliases if a sequence of probes sent to alternating IP addresses yields responses with incrementing, interleaved IP-ID values. Mercator infers aliases when a router responds with a different source address than that probed. To classify a pair of addresses as aliases with Ally, we require the IP-ID values in five consecutive responses to be in sequence. Ally's default is to send three probes up to 400ms apart, to require that the IP-ID values be within a small range, and to allow for some reordering by allowing IP-ID values to be returned out of sequence but close together. We send our probes one second apart to reduce the chance of rate-limiting, and of false positives due to IP-ID values being out of sequence but close enough to pass the allowance made for reordering. We try Ally up to four times if we do not obtain five responses. Ally is susceptible to false positives if the counters of two routers happen to be in sequence at the time of probing [19], so we repeat the Ally measurement every 20 minutes for two hours for address pairs classified as aliases.

This approach assumes that Ally renders few false negatives; that is, a router uses a single shared counter if it uses a counter at all. Sherwood, *et al.* [21] tested this assumption on 374,337 aliases inferred by DisCarte, a system that infers router-level topology through alignment of traceroute and record-route data; they report that routers *usually* record

---

[1] More recent tools for alias resolution [19, 20] are more appropriate for constructing a complete router-level graph; Ally lets us carefully probe small sections of the topology to try to identify false links.
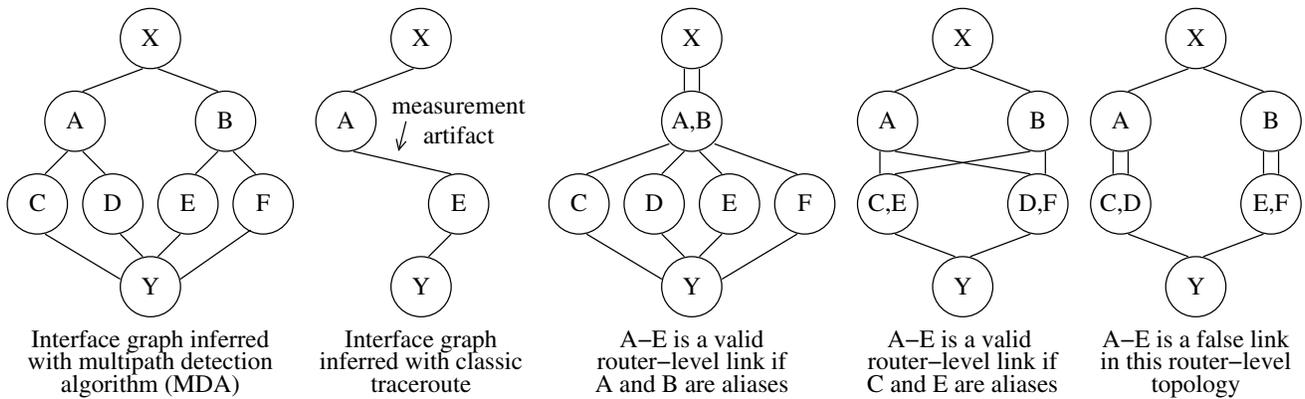
Figure 2: Classification of measurement artifacts in an interface graph. A measurement artifact introduces error into the router-level topology if it connects two routers that are not neighbours.

traceroute	record–route
192.107.171.142	192.107.171.51
192.107.171.49	130.217.2.5
130.217.2.6	203.167.234.86
203.167.234.85	218.101.61.194
218.101.61.193	10.65.0.65
203.98.50.1	203.98.50.1
203.98.50.251	203.167.233.14
203.167.233.10	202.84.219.122
202.84.219.121	202.84.219.121

Figure 3: Inference of /30 subnets and candidate aliases to check with Ally. Dashed lines associate interfaces inferred to be in common /30 subnets; solid lines associate interfaces we believe are aliases.

the egress interface in an RR option, and traceroute responses (ICMP time exceeded messages) *usually* record the ingress interface. These two sources of topology information allow aliases to be inferred using logic rules that correspond to network engineering practices. For a set of 374,337 aliases they tested, Ally measurements rejected 3.8% as not aliases, a low disagreement rate that is a combination of Ally limitations and imperfect topology data provided to DisCarte.

We repeat this cross-validation using a smaller set of likely aliases and analyse Ally's behaviour with ICMP-echo, TCP-ack, and UDP probes. Figure 3 illustrates how we derive our set of likely aliases. Given a sequence (n_1^m, n_2^m), (n_1^{m+1}, n_2^{m+1}) where n^m is an inferred /30 subnet, n_1^m is an address returned in a record-route probe, n_2^m is the address returned to a traceroute probe, and n^m, n^{m+1} are adjacent in the topology data, we infer n_2^m, n_1^{m+1} to be aliases of the same router because there are few other likely network engineering explanations for the use of these addresses. Using record-route and traceroute probes from 22 Archipelago [22] vantage points on August 13th 2010, we assembled a set of 26,510 likely IP alias address pairs using this approach. For 20,762 (78.3%) pairs, both interfaces have an incrementing IP-ID for at least one probe method; 17,799 pairs for UDP, 7,918 pairs for TCP, and 5,598 pairs for ICMP. Nearly all responses to UDP probes have an incrementing IP-ID. ICMP probes are more likely to elicit a response, but most re-

sponses echo the probe's IP-ID, and 37% of interfaces probed with TCP packets reply with randomly generated IP-ID values; both response types are not useful for alias resolution.

Overall, 93.8% of pairs where both interfaces respond with an incrementing IP-ID are inferred by Ally to be aliases; that is, they share a counter. Of the 1,266 pairs of interfaces inferred by Ally to not be aliases, in all but 60 pairs the probe method used is UDP. Six IP addresses in two origin ASes account for 1,108 of the 1,144 pairs, suggesting six outliers rather than a systematic problem with UDP probes. In 64 pairs, the TCP method inferred aliases where the UDP method inferred not aliases, suggesting the TCP method sometimes reveals a shared counter where the UDP method does not. We also noticed 27 pairs that, with Ally tests every 20 minutes, sometimes returned out of sequence IP-ID values, suggesting that responses are sometimes generated from different counters even if the addresses are aliases; 78% of these pairs were also with UDP probes. Therefore, when classifying artifacts with Ally we used ICMP where available, then TCP, and finally UDP. To summarise, we note that alias resolution can introduce classification errors, but this data and the data from [21] show the error rate is low.

3. MACROSCOPIC-LEVEL IMPACT

To assess the negative impact of false links from classic traceroute on statistics derived from Internet-scale topology data sets, we collected macroscopic topology data with classic, Paris, and MDA traceroute techniques[2] from 22 CAIDA Archipelago vantage points[3]. We supplied each vantage point with a list of about 16,189 randomly generated IP addresses based on prefixes found in a Route Views [24] BGP table from 4th of September 2010. The complete set of 356,154 addresses contains an address in each /16 for prefixes not enclosed in any other prefix of length 16 or shorter, and an address in each other prefix of length 24 or shorter, but only one destination in any /24. We excluded addresses in the Team Cymru bogon prefix list [25] and the CAIDA do-not-probe list.

---

[2]The tools used and data collected are available at the scamper website [23].
[3]syd-au, yow-ca, yto-ca, scl-cl, lej-de, her-gr, nap-it, mty-mx, ams2-nl, hlz-nz, bma-se, tpe-tw, amw-us, eug-us, hnl-us, jfk-us, laf-us, lax-us, san-us, sea-us, sql-us, wbu-us.

	ICMP	UDP
Total IP-level links		
Classic	427,211	413,040
Paris	422,130	398,179
MDA	455,965	462,762
Artifact links in classic traceroute	13,313	32,233
Fraction of classic links	3.12%	7.80%
Inferred valid	45.9%	38.6%
Inferred false	24.3%	34.7%
Unclassified	29.9%	26.7%
Artifact links seen in other MDA traces	4,840	6,756
Fraction inferred valid	52.8%	34.9%
Fraction inferred false	13.5%	4.5%
Fraction unclassified	29.7%	22.2%
False links in graph	3,230	11,179
Fraction of classic links	0.76%	2.71%
Fraction violating no-loop	15.2%	9.67%
Spurious AS-level links	23	58

Table 1: Summary of our data measuring the macroscopic impact of classic traceroute on our global data set. 0.76% of links inferred with classic ICMP and 2.71% links inferred with classic UDP are inferred to be false. Most false links are between distinct routers but some are a loop caused by visiting aliases of the same router.

We collected this data over three days, beginning on the 4th of September 2010. We inferred 427k links from ICMP classic traceroute data, and 413k links from UDP classic traceroute data; ICMP traceroutes are able to probe further along a path than UDP traceroutes, and therefore infer more topology, because fewer edge networks filter ICMP than UDP [26]. We inferred 455k links from ICMP MDA data, and 463k links from UDP MDA data; we infer more links with UDP even though we are able to probe fewer edge networks because more routers per-flow load-balance UDP packets than do with ICMP packets [7]. In total, we inferred 13,313 artifact links with ICMP probes and 32,233 artifact links with UDP probes that were not in the corresponding MDA path, representing 3.1% and 7.8% of links in the classic traceroute data. Table 1 summarises our data and the macroscopic impact of using classic traceroute, which we discuss in more detail in the following sections.

3.1 Fraction of traceroutes with artifacts

Because more routers per-flow load-balance UDP packets than do ICMP packets, our dataset has more UDP traceroutes with link artifacts (20.0%) than ICMP traceroutes with link artifacts (9.5%), but most of those are repeated artifacts on topology nearby the vantage points [27]. Only 6.6% of UDP traceroutes and 2.8% of ICMP traceroutes yielded new link artifacts. Most classic traceroutes that contain an artifact have a single artifact: 60% of paths inferred with ICMP and 62% of paths with UDP. However the tail of the distribution is long, especially for paths inferred with classic UDP traceroute, where 1.4% have between 6 and 12 artifacts. In some cases the artifacts extend beyond the portion of the topology where load balancing is featured; these cases occur when the lengths of the paths after a load bal-

ancer are different, so the artifacts extend beyond where load balanced paths converge.

3.2 Classification of measurement artifacts

We were able to classify as either valid or false 70% of link artifacts derived from our ICMP data and 73% of link artifacts derived from UDP data; of these, the method defined in section 2 revealed that 24.3% of ICMP and 34.7% of UDP link artifacts were false, i.e., did not exist. But the fraction of link artifacts in the complete set of IP-level links inferred by classic traceroute is small, 3.12% and 7.80% for ICMP and UDP respectively. Therefore, the fraction of false links due to the classic traceroute algorithm is also small: 0.76% and 2.71% for ICMP and UDP respectively.

Without ground truth data our validation capability is limited. Instead, we searched MDA traceroute data from other vantage points and destinations for links we inferred as artifacts of classic traceroute. Because routing decisions are primarily based on the destination IP address, it is possible that a link artifact exists in the underlying topology towards destination D', even though we are 99% confident that the link artifact does not exist in the topology we inferred towards destination D. For classic traceroute with ICMP probes, 52.8% of the link artifacts which we classified as valid links were also observed in another trace. However, only 13.5% of link artifacts that we classified as false links were seen in another traceroute; i.e. we classified it incorrectly. Despite the limitations of our method, our classification of links was consistent with the best available data nearly four times as often as it was inconsistent.

We investigated methods to reduce the fraction of false router-level links in ICMP-traceroute-based graphs [2]. If we remove traceroutes where the same interface address appears twice in the path, our ICMP link count drops to 391k (91.6%) but the percentage of false router-level links inferred reduces to 0.52% from 0.76%. It is possible to reduce this fraction further to 0.35% by including only traceroutes that reach their destination, but at the cost of reducing the overall link count to 129k (30.2%).

3.3 Artifacts violating no-loop condition

Analytical alias resolution techniques assume any two interfaces observed in a traceroute path are not aliases in the absence of a forwarding loop [15, 10, 16], an assumption that figure 1(b) shows may be false. We tested link artifacts with Ally to infer whether the interfaces on both ends of the link artifact represented the same router. We found 1081 links in UDP-classic traceroute paths and 490 links in ICMP-classic traceroute paths where the interfaces were aliases; these undetected loops represented 9.67% and 15.2% of link artifacts we inferred to be false. Load-balancing packets across paths of different lengths is also a cause of false loops in classic traceroute [6], but filtering out classic traceroute paths with a loop only removed 431 (40%) of these artifacts for UDP-classic and 151 (31%) for ICMP-classic.

3.4 Impact on AS-link inference

Publicly available BGP data does not provide a complete AS-level graph for most ASes in the Internet [28]. A second method to infer AS-level links is to match adjacent IP addresses found in traceroute paths with the origin ASes of their longest matching prefix in public BGP data; an AS-link is inferred if the addresses correspond to different origin

ASes. However, this approach is prone to false inference of AS-level links [29, 30] primarily because of IP address sharing between peering BGP routers in neighbouring ASes [31]. Our focus is on AS-level links derived from classic traceroute that we infer to be false at the IP-level.

To measure the impact of false links derived from classic traceroute on AS-link inference, we derived an AS-level graph from the IP-level graphs assembled using the topology data collected with MDA traceroute, and then checked if links we inferred as being false at the router-level introduced spurious AS-level links, i.e, those not in the MDA graph. The UDP-based graph yielded 26733 AS-level links and the ICMP-based graph yielded 28591 AS-level links, reflecting less filtering of ICMP probes at the edges of the Internet [26]. The false links from classic traceroute with UDP probes yielded 58 spurious AS-level links, and the false links from classic traceroute with ICMP probes yielded 23 spurious AS-level links, a fraction of 0.22% and 0.8% respectively. This result is consistent with the hypothesis that load-balancing across AS-level paths is rare [7], so false links from classic traceroute should rarely cross AS boundaries.

Without ground truth data, we are limited in our ability to determine how many of these spurious AS-link inferences actually exist. Instead we searched the UCLA Internet topology collection [32], which is the largest collection of BGP data available. We found two of 23 spurious ICMP-derived links and three of 58 spurious UDP-derived links in their data; that is, we were mostly accurate at inferring these spurious links as false.

4. ISP-LEVEL IMPACT

In this section, we study the impact of false links on metrics for a single ISP (AS): router-level degree distribution and PoP-level path diversity. We used ground truth information from a tier-1 ISP network: the complete set of interfaces for their 1986 routers. The ground truth data does not include the set of router-level links in the ISP, but it does allow us to separate the impact of alias resolution errors from the impact of false links.

To obtain a set of router-level links, we used CAIDA's BGP-derived AS relationship data [33], and probed a random address in each prefix announced by customers of the ISP. On June 8th 2010, we probed the ISP network from 22 Ark vantage points using classic, Paris, and MDA traceroute methods with UDP packets. We observed interfaces from 819 distinct routers (41%) with 1242 router-level links with Paris traceroute, 1340 router-level links with MDA traceroute, and 2398 links with classic traceroute, i.e. 44% of classic traceroute links in our data are suspicious.

Because the fraction of routers observed using the Ark vantage points was 41% of the ground truth set, we repeated the probing using PlanetLab vantage points to try and observe more of the ISP's topology. On July 17th 2010, we obtained data from 265 PlanetLab vantage points – one per site for sites that were operational when we ran our experiment. For the data collected in our Ark experiment, the router degree distribution for the Paris and MDA data were practically indistinguishable, so to reduce measurement load we restricted probing to classic and Paris traceroute using UDP and ICMP probes. We staggered the data collection across 24 hours, so the ISP's customers would receive, on average, one traceroute every 80 seconds. Despite using an order of magnitude more vantage points, we observed 889 distinct
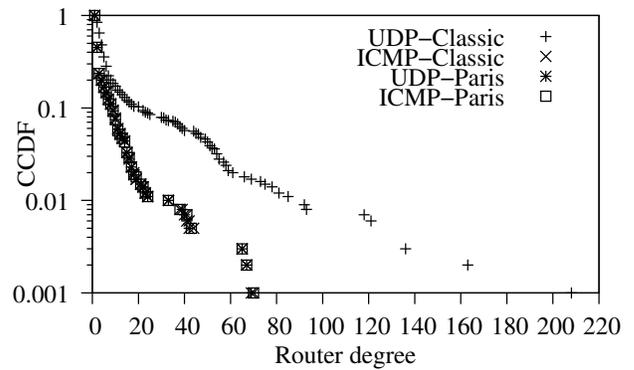

Figure 4: CCDF of the router-level degree distribution inferred by classic and Paris traceroute.

routers (45%) from PlanetLab. UDP-Paris, ICMP-classic, and ICMP-Paris observed 1437–1444 router-level links, and UDP-classic observed 4346, i.e. 67% are suspicious. The two experiments demonstrate that even if we choose probing destinations in prefixes originated by the ISP's customers, our probes are not guaranteed to discover the complete topology of the ISP. We emphasise that the results we report next are for the subset of the tier-1 ISP's topology that we were able to discover with active probes.

4.1 Impact on router degree distribution

Figure 4 shows the CCDF of the degree distribution of the router-level topology as inferred by Classic and Paris traceroute with ICMP and UDP probes. Classic traceroute with UDP probes significantly skewed the degree distribution toward higher degrees. False router-level links, inaccurately inferred from classic traceroute data, make the ISP appear to have much richer router-level connectivity than actually exists. Specifically, only 1% of routers had a degree larger than 20 with UDP-Paris, while 10% of routers had a degree larger than 20 with UDP-Classic. We also found that router degrees were inflated by a median factor of 2.9 in the graph inferred by UDP-Classic, as compared to degrees inferred using UDP-Paris. The ISP does not appear to load balance ICMP packets on a per-flow basis, as the degree distribution obtained using ICMP-Classic closely follows that obtained with ICMP-Paris.

4.2 Impact on path diversity

Path diversity is typically measured in the number of edge and node-disjoint paths between pairs of PoPs [9]. To identify PoPs, we collapsed the router-level graph into a PoP-level graph using information about the city in which each router is located[4]. To quantify the impact of false links on the inferred path diversity of the ISP network, we compute the maximum number of edge-disjoint and node-disjoint paths between each pair of PoPs, as inferred by each probing method. Figure 5 shows the CDF of the number of edge-disjoint paths between each pair of PoPs in the PoP-level graph inferred by classic and Paris traceroute. Figure 6 shows the same result for the number of node-disjoint paths

---

[4]The ground truth provided by the tier-1 ISP does not include PoP-level information, so we assume that all routers in the same city are at the same PoP.
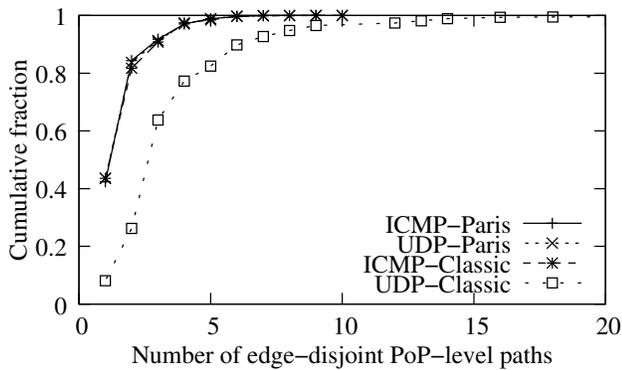
Figure 5: Cumulative frequency of the maximum number of edge-disjoint paths between each pair of PoPs in the PoP-level topologies inferred by classic and Paris traceroute.
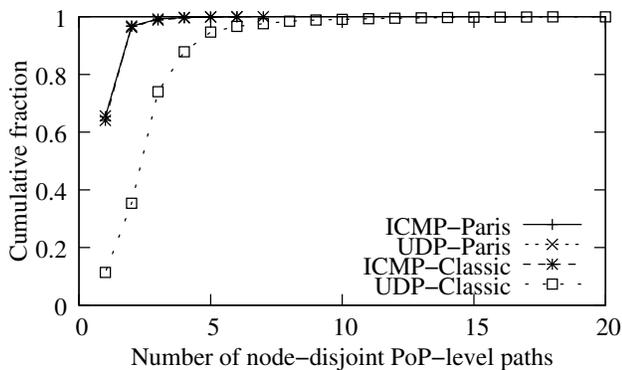

Figure 6: Cumulative frequency of the maximum number of node-disjoint paths between each pair of PoPs, in the PoP-level topologies inferred by classic and Paris traceroute.

between pairs of PoPs. There were significantly more edge-disjoint and node-disjoint paths between pairs of PoPs in the topology inferred by classic traceroute, due to the false router-level links inferred. Specifically, 18% of PoP pairs have more than 5 edge-disjoint paths in the topology inferred by classic traceroute, while only 1% of PoP pairs have more than 5 edge disjoint paths in the topology inferred by Paris traceroute. We emphasise again that our probes have not inferred the complete topology of this ISP. In particular, we found that 66% of PoPs have no node-disjoint paths between them when using data from Paris traceroute. This situation is unlikely in a tier-1 ISP network, which would be designed to provide redundancy against PoP-level failures.

5. RELATED WORK

The popularity of traceroute-like tools for measuring Internet topology at the router and AS-level led to interest in studying the accuracy of the traceroute-derived topology maps. Lakhina, *et al.* [27] found that sampling bias in traceroute-like measurement can cause a non-heavy tailed degree distribution to appear heavy-tailed. Mao, *et al.* [34], Hyun, *et al.* [35] and Zhang, *et al.* [31] studied incongruities between traceroute and BGP AS paths. Mao, *et al.* [34]

provided an algorithm to correct IP-to-AS mappings; Hyun, *et al.* [35] and Zhang, *et al.* [31] quantify the incongruities of third-party addresses and address-sharing between BGP-speakers respectively. Teixeira, *et al.* [9] found that a topology of the Sprint network derived from Rocketfuel significantly over-estimated PoP-level path diversity. They conjectured this was due to false links in the Rocketfuel topology. Augustin, *et al.* [6] argued that classic traceroute can lead to anomalies such as false links, loops, and cycles, and proposed Paris traceroute which crafts packets so a sequence of packets will follow a single path through a load-balanced topology. To the best of our knowledge, we are the first to quantify the false links inferred from classic traceroute, and to measure the effect of these false links on properties such as degree distribution and path diversity.

Sherwood, *et al.* presented DisCarte [21], an approach to infer and cross-validate inferred topology using traceroute and the record-route (RR) IP option. Most routers use the address of the interface at which a packet expires when sending a time exceeded message, but use the address of the transmitting interface when inserting an address into the RR option. DisCarte uses both data sources to reveal a more complete topology graph, including hidden routers that do not send time exceeded messages in response to traceroute probes; DisCarte therefore reduces the number of false router-level links caused by hidden routers.

Load-balancing has not always been common. In a 1995 study on routing dynamics, Paxson [17] found two out of 33 vantage points whose hosting campus network exhibited per-packet load-balancing, a pathology he called *"fluttering"* since it can reduce TCP performance. He also found evidence of five routers he called *"tightly-coupled"* – that is, routers that likely load-balanced packets on a per-flow basis. Augustin, *et al.* [7] measured load balancing on 771,795 Internet paths in 2007 and found that 39% had at least one router that balanced traffic load on a per-flow basis, while 70% had at least one router that balanced on a per-destination basis.

6. CONCLUSION

The state of the art in Internet topology measurement is essentially and necessarily a set of hacks, which introduce many sources of possible errors. This paper investigated the impact of one type of inferential error – assuming a single path between a source and destination. Our results suggest that the impact of this error on macroscopic topology data collected with the ICMP method is minor – in our data, link artifacts we infer to be false router-level links are 0.76% of the total set. The impact of using classic traceroute for ISP topology mapping is more significant; two thirds of the links collected for a targeted ISP were suspicious, and false links in classic traceroute led to substantially inflated router degrees and path diversity in the ISP network. Other errors arising from traceroute measurement are still invisible, though acknowledged [21]. Hidden nodes, stale interface IP addresses, and routing dynamics still have an unquantified presence in the Internet topology graphs inferred with traceroute.

Acknowledgements

7. REFERENCES

[1] Van Jacobson. traceroute.
    `ftp://ftp.ee.lbl.gov/traceroute.tar.gz`.

[2] Bradley Huffaker, Daniel Plummer, David Moore, and k claffy. Topology discovery by active probing. In *Symposium on applications and the Internet (SAINT)*, January 2002.

[3] Yuval Shavitt and Eran Shir. DIMES: Let the Internet measure itself. *ACM SIGCOMM Computer Communication Review*, 35(5), October 2005.

[4] Harsha V. Madhyastha, Tomas Isdal, Michael Piatek, Colin Dixon, Thomas Anderson, Arvind Krishnamurthy, and Arun Venkataramani. iPlane: An information plane for distributed services. In *OSDI*, November 2006.

[5] Neil Spring, Ratul Mahajan, and David Wetherall. Measuring ISP topologies with Rocketfuel. In *ACM SIGCOMM*, August 2002.

[6] Brice Augustin, Xavier Cuvellier, Benjamin Orgogozo, Fabien Viger, Timur Friedman, Matthieu Latapy, Clémence Magnien, and Renata Teixeira. Avoiding traceroute anomalies with Paris traceroute. In *IMC*, October 2006.

[7] Brice Augustin, Timur Friedman, and Renata Teixeira. Measuring load-balanced paths in the Internet. In *IMC*, October 2007.

[8] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the Internet topology. In *ACM SIGCOMM*, August 1999.

[9] Renata Teixeira, Keith Marzullo, Stefan Savage, and Geoffrey M. Voelker. In search of path diversity in ISP networks. In *IMC*, October 2003.

[10] M.H. Gunes and K. Sarac. Analytical IP alias resolution. In *IEEE International Conference on Communications*, June 2006.

[11] Kai Chen, David R. Choffnes, Rahul Potharaju, Yan Chen, Fabian E. Bustamante, Dan Pei, and Yao Zhao. Where the sidewalk ends: extending the Internet AS graph using traceroutes from P2P users. In *ACM CoNEXT*, 2009.

[12] Yihua He, Georgos Siganos, Michalis Faloutsos, and Srikanth Krishnamurthy. Lord of the links: a framework for discovering missing links in the Internet topology. *IEEE/ACM Transactions on Networking*, 17(2), 2009.

[13] Hyunseok Chang, Sugih Jamin, and Walter Willinger. Inferring AS-level Internet topology from router-level path traces. In *Proceedings of SPIE ITCom*, 2001.

[14] R-fx Networks. Advanced policy firewall (APF).

[15] Neil Spring, Mira Doncheva, Maya Rodrig, and David Wetherall. How to resolve IP aliases. University of Washington UW-CSE-04-05-04, 2004.

[16] Ken Keys. Internet-scale IP address alias resolution

techniques. *ACM SIGCOMM Computer Communication Review*, 40(1), January 2010.

[17] Vern Paxson. End-to-end routing behaviour in the Internet. *IEEE/ACM Transactions on Networking*, 5(5), October 1997.

[18] Ramesh Govindan and Hongsuda Tangmunarunkit. Heuristics for Internet map discovery. In *IEEE INFOCOM*, March 2000.

[19] Adam Bender, Rob Sherwood, and Neil Spring. Fixing Ally's growing pains with velocity modeling. In *IMC*, October 2008.

[20] Ken Keys, Young Hyun, Matthew Luckie, and kc claffy. Internet-Scale Alias Resolution with MIDAR. `http://www.caida.org/workshops/isma/1002/`.

[21] Rob Sherwood, Adam Bender, and Neil Spring. DisCarte: a disjunctive Internet cartographer. In *ACM SIGCOMM*, August 2008.

[22] CAIDA. Archipelago measurement infrastructure. `http://www.caida.org/projects/ark/`.

[23] Matthew Luckie. Scamper. `http://www.wand.net.nz/scamper/`.

[24] University of Oregon route views project. `http://www.routeviews.org/`.

[25] Team Cymru. Bogon bit notation list v5.0. `http://www.cymru.com/Documents/bogon-bn.html`.

[26] Matthew Luckie, Young Hyun, and Brad Huffaker. Traceroute probe method and forward IP path inference. In *IMC*, October 2008.

[27] Anukool Lakhina, John W. Byers, Mark Crovella, and Peng Xie. Sampling biases in IP topology measurements. In *IEEE INFOCOM*, April 2003.

[28] Ricardo V. Oliveira, Dan Pei, Walter Willinger, Beichuan Zhang, and Lixia Zhang. In search of the elusive ground truth: The Internet's AS-level connectivity structure. In *ACM SIGMETRICS*, June 2008.

[29] Z. Morley Mao, David Johnson, Jennifer Rexford, Jia Wang, and Randy Katz. Scalable and accurate identification of AS-level forwarding paths. In *IEEE INFOCOM*, March 2004.

[30] Ricardo Oliveira, Beichuan Zhang, and Lixia Zhang. Observing the evolution of Internet AS topology. In *ACM SIGCOMM*, August 2007.

[31] Yu Zhang, Ricardo Oliveira, Zhang Hongli, and Lixia Zhang. Quantifying the pitfalls of traceroute in AS connectivity inference. In *Passive and Active Measurement Conference (PAM)*, April 2010.

[32] UCLA. Internet topology collection. `http://irl.cs.ucla.edu/topology/`.

[33] The CAIDA AS relationships dataset. `http://www.caida.org/data/active/as-relationships/`.

[34] Zhuoqing Morley Mao, Jennifer Rexford, Jia Wang, and Randy H. Katz. Towards an accurate AS-level traceroute tool. In *ACM SIGCOMM*, August 2003.

[35] Young Hyun, Andre Broido, and kc claffy. On third-party addresses in traceroute paths. In *Passive and Active Measurement Conference (PAM)*, April 2003.