

Inferring Subnets in Router-level Topology Collection Studies*

Mehmet H. Gunes
Department of Computer Science
University of Texas at Dallas
mgunes@utdallas.edu

Kamil Sarac
Department of Computer Science
University of Texas at Dallas
ksarac@utdallas.edu

ABSTRACT

Internet measurement studies require availability of representative topology maps. Depending on the map resolution (e.g., autonomous system level or router level), the procedure of collecting and processing an Internet topology map involves different tasks. In this paper, we present a new task, i.e., subnet inference, to advance the current state of the art in topology collection studies. Utilizing a technique to infer the subnet relations among the routers in the resulting topology map, we identify IP addresses that are connected over the same connection medium. We believe that the successful inclusion of subnet relations among the routers will yield topology maps that are closer, at the network layer, to the sampled segments of the Internet in router level topology measurement studies.

Categories and Subject Descriptors

C.2.1 [Communication Networks]: Network Architecture and Design — Network Topology

General Terms

Measurement

Keywords

Topology discovery, Router-level map, Subnet inference

1. INTRODUCTION

Internet measurement studies require availability of representative topology maps. Depending on the nature of measurement study, researchers may use different types of topology maps including autonomous system (AS) level, point-of-presence (POP) level, router level, or IP address level maps.

*Dataset used in this paper is available at <http://www.utdallas.edu/~mhg042000/I2>

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'07 October 24-26, in San Diego, CA, USA

Copyright 2007 ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

A POP-level topology map is often the most detailed information that Internet Service Providers (ISPs) make publicly available, if at all, about their network [5]. Due to various privacy and security reasons, ISPs keep their router level topology information confidential. On the other hand, router level Internet topology maps are considered useful in various contexts such as analyzing the topological characteristics of the Internet at the network level and designing topology generators that can produce Internet-like synthetic network topologies to be used in various simulation studies.

The confidentiality of router level topology maps introduces a practical challenge for the research community and requires them to use other means to collect this information. In order to facilitate router level topology measurement studies, several research groups and institutions developed tools and methodologies to collect the required topology information from the Internet [4, 7, 12, 14, 15, 18, 21, 22]. Most of these approaches utilize a well-known Internet debugging tool, called *traceroute* [11], to collect a large number of path traces from topologically diverse set of vantage points and make this data available to researchers.

Recent work on router level topology construction studies identified several tasks in building a topology map from a collected set of path traces. These tasks include (1) filtering erroneous traces [17], (2) resolving anonymous routers in path traces [23] and (3) identifying IP addresses, within the data set, belonging to the same router [8]. The accuracy and the completeness of these tasks affect the representativeness of resulting sample topology [10].

In this paper, we formulate a new task to build more accurate router level Internet maps from collected path traces. In this new task, we study the relation between the IP addresses in the data set to infer subnet relations among them. The successfully inferred subnet information helps in (1) improving the quality of the resulting map by annotating it with additional information, i.e., the resulting map includes subnet relations among the existing set of IP addresses, (2) increasing the scope of the map by adding new links into the resulting map, and (3) improving the IP alias resolution process. We believe that the successful inclusion of subnet relations among the routers will yield topology maps that are closer, at the network layer, to the sampled segments of the Internet.

Note that, the proposed subnet detection task presents similarities with that of the IP alias resolution task (the 3rd task of map construction). In IP alias resolution, the goal is to identify nodes that appear to be separate in the collected path traces and combine them into one single node (i.e.,




Figure 1: Sample Subnetwork

to detect IP addresses, in the data set, that belong to the same router). Similarly, the goal in subnet detection is to identify multiple links that appear to be separate and combine them to represent their corresponding single hop (e.g., point-to-point or multi-access) connection medium. As a result, the inclusion of this task will improve the accuracy and the completeness of the constructed topology map.

As an example, consider three routers A, B, and C, in Fig. 1-a, that are connected to each other via a multi-access link. Assume that a collected set of path traces include A-to-B link and B-to-C link and no path trace at hand includes the A-to-C link. In this case, a router level map that does not consider the subnet relation among these IP addresses will result in a topology map as shown in Fig. 1-b. On the other hand, a careful study of the IP addresses may detect the subnet relation between the routers and therefore improve the resulting map.

The rest of the paper is organized as follows. Next section presents the related work. Section 3 presents our approach on inferring subnets between IP addresses in a set of path traces. Section 4 presents our preliminary evaluation results on the proposed subnet inference approach. Finally, Section 5 concludes the paper.

2. RELATED WORK

In this section, we outline the challenges in router level topology mapping and present the current state-of-the-art tools and mechanisms to overcome these challenges.

Obtaining an accurate network map requires several important tasks including (1) reducing the number of probes, (2) verifying the accuracy of collected path traces, (3) resolving anonymous routers that are represented by ‘*’s in traceroute output, (4) resolving IP addresses belonging to the same router, an operation referred to as IP alias resolution, and (5) identifying physical subnets correctly.

The first task is to collect the topology information with minimum number of probes. *Mercator* system [7] performs traceroute by starting the probing far from the source for IP addresses that have an already traced IP with the same prefix. For such destination IP addresses, it starts the TTL value at the highest hop count of a responding router seen on previous paths. Likewise, *doubletree* algorithm [6] benefits from the tree-like nature of Internet paths (as seen from a vantage point) to prune redundant probes to nodes that are close to both the vantage point and the destination.

The second task involves making sure that the obtained trace corresponds to a real path. This need arises especially due to traffic engineering practices. That is, when an AS implements load balancing between some paths in their network, traceroute probes may traverse alternating paths and the returned IP addresses may not belong to neighboring routers. *Paris traceroute* [2] tries to minimize the

effect of load balancing. *Paris traceroute* eliminates per-flow load balancing by controlling packet headers, but may not always succeed in eliminating the effect of load balancing when per-packet load balancing is implemented. In addition, *sidecar* [19] detects changes in traversed paths by enabling record route option of probe packets.

The third task is to resolve anonymous routers. This is required because not all routers respond to traceroute probes all the time [23]. A router may ignore responding with ICMP error messages due to its policy or when it is loaded. Hence, we need to group observed ‘*’s that correspond to the same router to prevent an inflated network graph. In [3], Bilir et.al. presented a practical approach where anonymous routers are combined when their upstream and downstream neighbors are same.

The fourth task is on IP alias resolution. This task is an artifact of the traceroute-based topology collection procedure. Routers have multiple interfaces each one having its own IP address. A router may appear on multiple path traces with different IP addresses. Therefore, we need to identify and group IP addresses belonging to the same router. Several approaches have been proposed for alias resolution including DNS based approach of [16], source IP based *mercator* [7], IP identification based *ally* [20], record route based *sidecar* [19], and analytical approach APAR [9] to resolve IP aliases.

Finally, the last task is required for obtaining an accurate network map and is the main focus of the work presented in this paper.

3. INFERRING SUBNETS

Internet is a network of interconnected networks that belong to different entities or organizations such as academic, businesses, or governments institutions. Each entity has its own approach to network design, and therefore there are different network layouts [21]. Moreover, due to different requirements at the core and at the edge of a network, there are differences between the layouts of network core and network edge [13]. Our focus, in this work, is mainly on the network core rather than the edge.

Devices at the core of a network are connected to each other using a connection medium such as a point-to-point link or a multi-access link. Devices that are on the same connection medium are said to form a *subnet* and they can directly communicate with each other at the link layer. On the other hand, the communication between devices on different subnets requires routing support at the network layer. Moreover, based on the global addressing and efficient routing needs, devices on a subnet are assigned IP addresses from a specific address range that is explicitly reserved for the particular subnet [5]. In practice, the term *subnet* is used to refer to both the connection medium as well as the IP address range given to that medium. For the sake of clarity, in this paper, we use the term *subnet* to refer to the connection medium and the term *subnet address* to refer to the specific address range assigned to a connection medium or a subnet.

Unlike network layout design, the global naming of device interfaces (e.g., IP address assignment) adheres to IETF guidelines that are mostly respected in practice. In the following, we first present a summary of IP address assignment practices in the Internet and then show a methodology to use it to infer subnets in a set of collected path traces.

3.1 IP Address Assignment Practices

One basic requirement in the Internet is that each device (or each device interface) has a globally unique identifier. The unique identification of devices in the Internet is achieved by the Internet Protocol (IP) where unique IP addresses are assigned to device interfaces. Being a scarce commodity, IP addresses are assigned systematically adhering to the Internet Registry IP Allocation Guidelines (RFC-2050). Systematic assignment mechanism also helps in reducing the routing information at routing tables. Basically, each network domain (e.g., each Autonomous System or AS) gets a range of IP addresses for assignment. This address range is then divided into smaller address ranges, called subnet addresses, each of which is assigned to different subnets. The individual IP addresses in each subnet address range is then assigned to an interface connected to that subnet.

The smallest subnet in the Internet is built by using a point-to-point link that connects two device interfaces. A /30 or a /31 subnet address (the latter is introduced in RFC 3021) is defined and used to assign IP addresses to the interfaces in this type of networks. Larger subnet addresses (/29 or larger) are not used for point-to-point links as they cause waste of IP addresses. Multi-access links are used to connect several device interfaces to form a subnet as in Fig. 1-a. In general, these subnets include more than two devices connected to them. A number of technologies can be used to build multi-access links including Ethernet, FDDI, token ring, etc. When building a subnet, one chooses a subnet address range that has enough IP addresses for unique address assignment for each interface on the subnet. As an example, if a subnet is to include five device interfaces, one defines a /29 subnet address to assign unique IP addresses to each interface. In a /29 subnet address, we have $2^{(32-29)} - 2 = 6$ IPv4 addresses for assignment.

In general, a subnet with n devices needs a / x subnet address where $x = 32 - \lceil \log_2(n + 2) \rceil$. The first x bits of the assigned IP addresses denote the subnet address and the last $32 - x$ bits identify the device interfaces within the subnet. For example, if a subnet has a subnet address of 192.168.0.0/28, then the last 4 bits are used to identify the individual IP addresses of the device interfaces in this subnet. These four bits can identify at most 14 device interfaces. The remaining two IP addresses, namely 192.168.0.0 and 192.168.0.15, have special meanings and are not typically used for assignment.

3.2 Subnet Formation

In order to build an accurate router level topology map from collected path traces, we need to analyze the data set to infer subnet relations among the IP addresses. The above mentioned IP address assignment scheme introduces a relation between the IP addresses that are assigned to device interfaces on a subnet. This relation can be used to infer the existence of subnets among a number of IP addresses in the collected set of path traces. That is, we can form candidate subnets where some set of IP addresses can be grouped into a subnet address range under an address prefix of length / x . Note that, any two IP addresses can be grouped into an address range for a sufficiently large address range that can be represented using a sufficiently small / x prefix length. Therefore, we need to analyze path traces to collect evidence to eliminate candidate subnets that may not correspond to real subnets in the underlying network. Moreover, we may

perform additional probing of the network to increase our confidence in the accuracy of the inferred subnets.

Due to IP address assignment practices, each observed IP address belongs to some subnet where all interfaces on the subnet have IP addresses with the same maximal x bit prefix, i.e., subnet address, and interfaces on other subnets have different x bit prefixes, i.e., different subnet addresses. Based on this observation, we use an iterative approach to identify candidate subnets. We first form all candidate / x subnets from the data set by combining the IP addresses whose first x bits match. Next, we recursively form smaller subnets (e.g., / x , / $(x+1)$, ..., /31 subnets).

At this point, we need to identify the candidate subnets that correspond to real subnets in the Internet. That is, even though a given set of IP addresses can map to a, say, candidate /29 subnet, there may not be a corresponding real /29 subnet in the underlying network among these IP addresses. Instead, the addresses may belong to two separate /30 subnets in the Internet. Similarly, the candidate /29 subnet may be part of a larger real subnet. Therefore, we need to detect and prune candidate subnets that do not correspond to real subnets. We develop a set of complementary conditions that will help us do the necessary verification check during the pruning process as follows:

Accuracy: Given a loop-free path trace, two or more IP addresses from the same subnet cannot appear in any path trace without having a successor/predecessor relationship with each other. That is, IP addresses in a subnet should appear next to each other whenever they appear in the same trace. For instance, consider the sample topology in Fig. 2 where $H1$ and $H2$ are end-hosts and $R1$ and $R2$ are routers on a FDDI ring. Without the knowledge of network topology, a path trace from $H1$ to $H2$, i.e., (b, e, h) , will indicate that b and h can not be in the same subnet as they are two hops away of each other. In addition, aliases of IP addresses in a subnet should appear as successor/predecessor if they appear in the same trace as well. This condition arises from the fact that nodes within the same subnet are directly connected and should appear one hop away from each other in a loop-free path trace if they appear together. RFC 1812 states that ICMP error messages should be sent with the IP address of the outgoing interface toward traceroute source. In this case, all IP addresses on a trace will be from distinct subnets. In practice, however, IP addresses of other interfaces, e.g., incoming interface, might be returned and these differing practices may yield path traces with two IP addresses from the same subnet. In this case, such IP addresses will be at most one hop away of each other.

Distance: Given a candidate / x subnet, the IP addresses from within this subnet should be at similar distances to a vantage point. Hence, we determine the hop distance of IP addresses with respect to each vantage point from path traces and prune candidate subnets whose IP addresses ap-




Figure 2: A sample network between two end-hosts

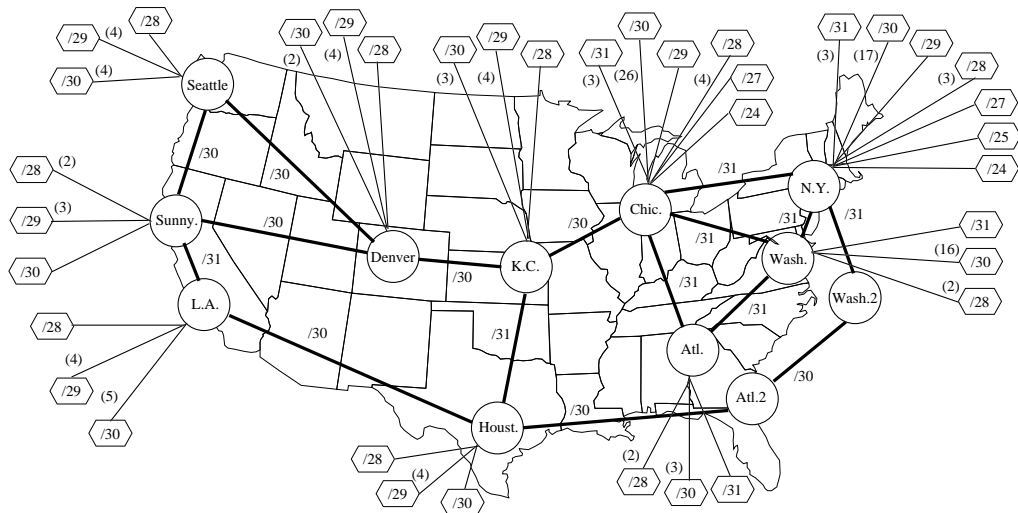


Figure 3: Internet2 backbone topology on 29 Apr 2007

pear more than one hop away of each other. In addition, since IP addresses of a subnet might not be observed from the same vantage point, we perform additional probing of the IP addresses of each candidate subnet. We send probe packets to each IP address with a non-existing port number. Then, from the ICMP response, we obtain the TTL value of the original packet when it reached the destination. We use dissimilarities in the TTL values to identify candidate subnets that do not correspond to real subnets. For instance, let's assume, for a sample topology of the network in Fig. 2, we send probes from vantage point at $H1$. Then, from probe responses, we will realize that h can not be in the same subnet with b or c since h is three hops away while b and c are one hop away of $H1$.

Completeness: Ignore candidate subnets that have less than one quarter of their IP addresses present in the collected data set. A $/x$ subnet can include up to $2^{32-x} - 2$ IP addresses and we require that at least one quarter of these addresses to appear in our data set. This requirement helps us increase our confidence in the accuracy of the inferred subnets. Without this requirement, it would be easy to form a candidate subnet (likely a large one) using a few IP addresses falling into the same subnet address range. However, the existence of a small number of IP addresses within the candidate subnet makes it difficult to verify the existence of the corresponding real subnet. Hence, if possible, we try to identify routers with IP addresses from the subnet address range of candidate subnets that do not satisfy this condition and perform additional traces to the new IP addresses.

MaxFit: Ignore smaller candidate subnets that are subset of a bigger candidate subnet after assessing the previous conditions. Since candidate subnets are formed recursively, from bigger ones to smaller ones, IP addresses of a $/x$ subnet will, most likely, appear in smaller (i.e., $/(x+1)$, $/(x+2)$, etc.) subnets. In such cases, we accept IP addresses to be in the biggest candidate subnet. Only exception is when an IP address appears to be in both a $/30$ and a $/31$ candidate subnets. In this case, $/31$ candidate subnet is chosen as the valid subnet.

During our subnet inference process, the above rules col-

lectively work to identify subnets in the data set and use this information to build the final map by explicitly showing point-to-point links as well as multi-access links.

4. EVALUATIONS

In this section, we present a preliminary evaluation of our subnet inference approach. Note that, evaluating the accuracy of our presented approach requires the knowledge of the router level topology of the underlying network. Currently, we have access to Internet2 backbone topology information from their web site. Therefore, for our evaluation, we run path traces to collect Internet2 backbone topology and use our subnet inference approach during the construction of the topology map. After inferring the topology map, we compare it with the genuine Internet2 topology and quantify the accuracy of our subnet inference approach. In addition, we present the utility of subnet inference in improving the map and in helping alias resolution task.

4.1 Obtaining Internet2 Backbone Maps

In this section, we describe how we obtain the genuine Internet2 backbone topology and how we collect path traces to map Internet2 backbone.

First of all, we obtain the topology map of Internet2 backbone from Abilene - Visible Backbone [1] on 29 April 2007. Fig. 3 presents the core routers of Internet2 and subnets connected to each router. The subnet size of point-to-point links connecting core routers is indicated on each link. In addition, hexagons in the figure indicate other $/x$ subnets connected to each router. The numbers shown on top/left of the lines indicate the number of such $/x$ subnets if more than one $/x$ subnet is connected to the router. For instance, Los Angeles router is connected to Houston and Sunnyvale routers over $/30$ and $/31$ subnets, respectively. Los Angeles router is, also, connected to five $/30$, four $/29$, and a $/28$ subnets. Overall, in the original Internet2 map, there are 547 routers with 793 IP addresses, and 150 subnets ranging in size from $/24$ to $/31$. Note that, we might not have all of the routers that are connected to the Internet2 core.

Next, we try to construct the same topology using collected path traces. In constructing the map, we benefit from

8 vantage points at various universities that are connected to the Internet2. First, we choose 360 IP addresses from the list of 793 IP addresses that we obtained from Internet2 web site and consider them as our probe destinations. In order to map the verifiable subnets, we try to select IP addresses from each subnet that is connected to Internet2 backbone routers. Next, we run traceroute queries from our vantage points to each of these IP addresses to collect our topology data. In addition to 360 IP addresses, during the subnet inference process, we choose 58 new IP addresses from some of the subnet address ranges and probe them to resolve ambiguities during the subnet inference phase.

After collecting path traces, we have 3,092 path traces which contain 808 unique IP addresses and 61 ‘*’s. Note that, collected path traces may contain segments that do not belong to Internet2 backbone. In order to construct the router level map from the collected traces, we perform the following tasks:

Filter inaccurate path traces: We first filter erroneous traces and combine the same path traces. The filtering step identifies 6 traces with routing loops and reduces 3,092 path traces to 2,465 unique paths.

Resolve anonymous routers: We perform anonymous router resolution using our approach presented in [3]. That is, we combine ‘*’s in different traces that have the same next hop upstream and downstream routers. This step reduces the number of ‘*’s from 61 to 34.

Identify alias IP addresses: We utilize *ally* tool [20] and our *APAR* tool [9] to resolve aliases. Note that, *ally* requires $O(n^2)$ probes for n IP addresses. We reduce the number of probes by limiting probes to pair of IP addresses that are one hop away of an identified subnet. This approach reduces the number of probes significantly as compared to graph reduction approach of [20] (see Section 4.4). *APAR* and *ally* tools identify 198 alias IP pairs by probing observed IP addresses from vantage point at UT-Dallas.

Identify subnets: We infer subnets using our approach presented in Section 3.2. At this step, the mapping system determines additional destinations that are needed to verify the existence of some of the candidate subnets and performs additional traces/probing. Overall, we infer 176 subnets using our subnet inference approach.

4.2 Verifying Inferred Internet2 Subnets

Among the 176 subnets identified in the collected topology, 116 subnets are part of the Internet2 backbone topology that we can verify. Table 1 compares the genuine subnets and subnets inferred from the collected data. First column indicates the number of each $/x$ subnet in the genuine topology. Following columns indicate the number of observed $/x$

	Σ	$/24$	$/27$	$/28$	$/29$	$/30$	$/31$	\emptyset	n/o
$/24$	2	1	1						
$/25$	1								1
$/27$	2			2					
$/28$	18		4	5	4	4		1	
$/29$	25		8	9			1	4	3
$/30$	86		4	1		67		7	7
$/31$	16				2		13	1	

Table 1: Comparison of genuine and inferred Internet2 subnets

subnet for each $/y$ subnet, i.e., (x, y) cell corresponds to the number of subnets that are $/y$ in the genuine topology but observed as $/x$. For instance, out of two $/24$ subnets, one is observed as $/24$ and the other as $/27$. In addition, \emptyset column indicates the number of genuine subnets whose IP addresses are not in any of the inferred subnets. Finally, n/o indicates the number of genuine subnets whose IP addresses are not observed in the collected path traces.

There are two types of mismatches in the Table 1. First, some subnets appear to be smaller than they are. This happens when the data set does not have IP addresses from the subnet address range that would necessitate the bigger one. For instance, for one of $/29$ subnets in the genuine topology only two IP addresses are observed in the collected data set and those IP addresses form a subnet of $/31$. In addition, except a single case, only one IP address of the genuine subnets in \emptyset set were observed in collected path traces. Similarly, some subnets, i.e., n/o set, have none of their IP addresses present in the collected data set. Inferring a subnet as a smaller one, in general, does not incorporate error into the map but annotates links with smaller size subnets.

On the other hand, there are subnets which are inferred to be bigger than they are. For instance, IP addresses of one $/29$ and two $/30$ subnets form a candidate subnet of $/27$ and there are no traces that invalidate the $/27$ candidate subnet. This happens when two subnets share a common router and have a common prefix. For instance, sample network in Fig. 4 shows two such subnets. That is, $R2$, $R3$, and $R4$ are connected over a $/29$ subnet and $R4$, $R5$, and $R6$ are connected over another $/29$ subnet. Note that, IP addresses of both subnets also belong to a $/28$ subnet address range. Then, if the only vantage point is $R1$, we will, at most, observe two IP addresses of subnets in a trace where they will be one hop away of each other. Hence, neither accuracy condition nor distance condition that we presented in Section 3.2 will be able to detect the error and we will assume both subnets as a single $/28$ subnet.

This case, however, will alter the network map by adding non-existent links. Accuracy and distance conditions can prevent this case when there are vantage points at routers other than $R1$ and $R4$. For instance, if our vantage point is $R7$, accuracy condition will prevent the error when it observes a path trace through $R2$, i.e., $(R6, R4, R2, \dots)$, or through $R3$, i.e., $(R6, R4, R3, \dots)$, where $.11$ will not be a successor/predecessor of $.2$ or $.3$, respectively. Likewise, distance condition will prevent the error when it observes that $.11$ is one hop away while $.2$ and $.3$ are three hops away of the vantage point.

Overall, among 116 verifiable subnets, we inferred 95 subnets correctly with the exact size. In addition, we inferred




Figure 4: Sample network

12 subnets smaller than they were due to the lack of other IP addresses. That is, all observed IP addresses of those subnets were accurately clustered but they formed smaller $/x$ subnets, i.e., longer same x bit prefixes, than the actual. On the other hand, there were 9 subnets inferred bigger than they were. In 2 cases, 3 subnets were combined into a bigger one, and, in 6 cases, 2 subnets were combined. All of the inference errors were due to subnets as in Fig. 4. Some of the errors could have been detected if we were to have vantage points that we could send probes as required by distance condition. We were only able to obtain distances of IP addresses with respect to the vantage point at UT-Dallas.

4.3 Improving Topology Map

Identifying subnets among IP addresses help in adding new links, that are not observed in path traces, to the map. For instance, identifying the subnet involving A , B , C in Fig. 1-b, we will be able to add the missing link between A and C . In the collected data, we had 1,727 links between routers. However, through subnet inference, we added 3,359 links between IP addresses in the inferred subnets. The number of new links is high especially due to a $/24$ subnet, which had 73 of its IP addresses present in the data set, that we identified. Identifying this subnet, we added 2,627 new links between its IP addresses because only one link existed in collected path traces.

4.4 Improving IP Alias Resolution

The current state-of-the-art probing based tool for IP alias resolution is *ally*. A straightforward application of *ally* requires a prohibitively large number of probes in the network. One approach proposed to reduce the number of probe pairs, is a graph reduction approach [20]. This approach identifies IP addresses that can not be aliases by observing that they appear on the same path trace. This approach reduced the number of pairs, for 808 IP addresses in our data set, from 362,028 to 280,229. We reduce the number of probes to 35,537 pairs using the inferred subnets where we probe IP address pairs that appear one hop away of a subnet. This approach may cause missing some alias pairs but reduces the number of required probes significantly.

5. CONCLUSIONS AND FUTURE WORK

In this paper, we present a new task into the process of building an accurate topology map from a set of collected path traces. We utilize a technique to infer the subnets in a set of collected path traces and improve the constructed topology map by identifying IP addresses that are connected over the same connection medium. The new task, subnet inference, is similar to the alias resolution, which identifies IP addresses belonging to the same router. The successful inclusion of subnet relations among the routers will yield topology maps that are closest, at the network layer, to the sampled segments of the Internet. In our experiments, we present the utility of identifying subnets in a topology map. Besides, we show how identified subnets can be utilized to reduce the number of probes while resolving IP aliases.

To our knowledge, this is the first study to promote an approach for inferring subnets in a topology mapping study. In this work, we showed the effectiveness of our subnet inference approach on Internet2 backbone. We are currently building a mapping system to obtain topology maps of Internet core using PlanetLab servers [4]. We are also working on improving the subnet inference so that errors are minimized.

6. REFERENCES

- [1] Abilene - Visible Backbone. <http://pea.grnoc.iu.edu/Abilene/>.
- [2] B. Augustin, X. Cuvelier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira. Avoiding traceroute anomalies with Paris traceroute. In *Proceedings of IMC*, Rio de Janeiro, Brazil, October 2006.
- [3] S. Bilir, K. Sarac, and T. Korkmaz. Intersection characteristics of end-to-end Internet paths and trees. In *IEEE ICNP*, Boston, MA, November 2005.
- [4] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman. PlanetLab: An overlay testbed for broad-coverage services. *ACM/SIGCOMM Computer Communication Review*, 33(3):3-12, 2003.
- [5] M. Crovella and B. Krishnamurthy. *Internet Measurement: Infrastructure, Traffic and Applications*. Wiley, 2006.
- [6] B. Donnet, P. Raoult, T. Friedman, and M. Crovella. Efficient algorithms for large-scale topology discovery. In *Proceedings of ACM/SIGMETRICS*, pages 327-338, New York, NY, 2005. ACM Press.
- [7] R. Govindan and H. Tangmunarunkit. Heuristics for Internet map discovery. In *IEEE INFOCOM*, March 2000.
- [8] M. Gunes and K. Sarac. Analytical IP alias resolution. In *IEEE ICC*, Istanbul, Turkey, June 2006.
- [9] M. Gunes and K. Sarac. Resolving IP aliases in building traceroute-based Internet maps. Technical report, University of Texas at Dallas, December 2006.
- [10] M. Gunes and K. Sarac. Importance of IP alias resolution in sampling Internet topologies. In *IEEE Global Internet*, Anchorage, AK, May 2007.
- [11] V. Jacobson. *Traceroute*. Available from <ftp://ee.lbl.gov/traceroute.tar.Z>.
- [12] S. Kim and K. Harfoush. Efficient estimation of more detailed Internet IP maps. In *Proceedings of IEEE ICC*, Glasgow, Scotland, June 2007.
- [13] L. Li, D. Alderson, W. Willinger, and J. Doyle. A first-principles approach to understanding the Internet's router-level topology. In *Proceedings of ACM/SIGCOMM*, pages 3-14, New York, NY, 2004. ACM Press.
- [14] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: An information plane for distributed services. In *OSDI*, November 2006.
- [15] D. McRobb, K. Claffy, and T. Monk. *Skitter: CAIDA's macroscopic Internet topology discovery and tracking tool*, 1999. Available from <http://www.caida.org/tools/skitter/>.
- [16] J. Pansiot and D. Grad. On routes and multicast trees in the Internet. In *ACM/SIGCOMM Computer Communication Review*, 28(1):41-50, 1998.
- [17] V. Paxson. End-to-end routing behavior in the Internet. In *Proceedings of ACM/SIGCOMM*, pages 25-38, New York, NY, 1996. ACM Press.
- [18] Y. Shavitt and E. Shir. DIMES: Let the Internet measure itself. *ACM/SIGCOMM Computer Communication Review*, 35(5):71-74, 2005.
- [19] R. Sherwood and N. Spring. Touring the Internet in a TCP sidecar. In *Proceedings of the ACM/SIGCOMM IMC*, pages 339-344, New York, NY, 2006. ACM Press.
- [20] N. Spring, M. Dontcheva, M. Rodrig, and D. Wetherall. How to resolve IP aliases. Technical report, University of Washington, May 2004.
- [21] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson. Measuring ISP topologies using rocketfuel. *IEEE/ACM Transactions on Networking*, 12(1):2-16, February 2004.
- [22] N. Spring, D. Wetherall, and T. Anderson. Scriptroute: A public Internet measurement facility. In *Proceedings of USITS*, March 2003.
- [23] B. Yao, R. Viswanathan, F. Chang, and D. Waddington. Topology inference in the presence of anonymous routers. In *IEEE INFOCOM*, San Francisco, CA, March 2003.