Computer Networks 81 (2015) 178–200

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Graph Based Induction of unresponsive routers in Internet topologies

Hakan Kardes [a], Mehmet Hadi Gunes [b,*], Kamil Sarac [c]

[a] Inome Inc., United States
[b] University of Nevada Reno, United States
[c] University of Texas at Dallas, United States

ARTICLE INFO

Article history:
Received 8 July 2013
Received in revised form 27 December 2014
Accepted 16 February 2015
Available online 23 February 2015

Keywords:
Internet topology measurement
Internet mapping
Unresponsive router resolution

ABSTRACT

Internet topology measurement studies utilize traceroute to collect path traces from the Internet. A router that does not respond to a traceroute probe is referred as an unresponsive router and is represented by a '*' in the traceroute output. Unresponsive router resolution refers to the task of identifying the occurrences of '*'s that belong to the same router in the underlying network. This task is an important step in building representative traceroute-based topology maps and obtaining an optimum solution, i.e., minimal graph under trace and distance preservation conditions, is shown to be NP-complete. In this paper, we first analyze the nature of unresponsive routers and identify different types of unresponsiveness. We also conduct an experimental study to understand how the responsiveness of routers has changed over the last decade. We then utilize a novel graph data indexing approach to build an efficient solution to the unresponsive router resolution problem. The results of our experiments on both synthetic and sampled topologies show a significant improvement in accuracy and effectiveness over the existing approaches. Our experiments also demonstrate that applying IP alias resolution along with unresponsive router resolution results in a final topology closer to the original topology.

1. Introduction

Internet has become an important part of our daily interactions and keeps expanding around the world. As the largest man-made complex network, it has been growing with no central authority. Each network is built by a different AS for possibly different purposes, e.g., small local campuses to large transcontinental backbone providers. As each AS grows its own network based on local economic and technical objectives, a single or few ASes are not representative of the Internet.

Understanding Internet topology is valuable for commercial, social, and technical purposes. Knowledge of the network graph helps in understanding the large scale characteristics and dynamics of the Internet [1]. For instance, such graphs are needed in analyzing the topological characteristics of the Internet and designing topology generators that can produce realistic synthetic topologies [2]. Researchers have also indicated that the deployment of networks by content providers has a flattening effect on the hierarchical AS structure [3]. Knowledge of the underlying networks help in path exploration using new approaches [4] and optimize data delivery [5]. Additionally, analysis of the Internet topology is useful to develop failure detection measures, network planning, and optimal routing algorithms [6]. Researchers test new protocols and systems using simulations or emulations

* Corresponding author.
E-mail addresses: hkardes@cse.unr.edu (H. Kardes), mgunes@unr.edu (M.H. Gunes), ksarac@utdallas.edu (K. Sarac).

http://dx.doi.org/10.1016/j.comnet.2015.02.012
1389-1286/

[7,8], but more realistic results can be obtained when real topologies are utilized in the analysis [9,10].

Internet service providers (ISPs) keep their topology information confidential due to commercial and security reasons. This policy introduces a practical challenge for the research community and requires them to generate measurement probes to sample the Internet topology at various levels, including IP, link, router, point of presence (PoP), and autonomous system (AS) levels. Most measurement studies utilize *traceroute* [11] or its variations to collect a large number of path traces from topologically diverse set of vantage points. Traceroute relies on the TTL expiration mechanism of routers to elicit responses from routers between source and destination (see RFC 1393: Traceroute Using an IP Option).

After collecting the path traces, the information needs to be processed to build the corresponding topology map. In particular at link, router and PoP levels, we need to *(i)* verify correctness of the collected paths [12], *(ii)* identify underlying tunnels [13], *(iii)* resolve IP addresses belonging to the same router [14,15], *(iv)* infer IP addresses that are connected over a common subnet link [16,17], and *(v)* resolve unresponsive routers that are represented by '*'s in traceroute outputs [18]. The first task is due to the fact that certain traffic engineering practices may cause traceroute to return IP addresses that do not correspond to a real end-to-end path in the network. Augustin et al. [19] propose *Paris traceroute* tool to address this problem. The second task helps in identifying and addressing inaccuracies due to tunneling techniques such as MPLS. Researchers have developed approaches to identifying underlying tunnels [20,13]. The third task is due to the fact that routers have multiple IP addresses that may appear in different path traces. Without any resolution, each IP address is treated as a potentially different router in the initial graph and IP alias resolution is applied to identify the IP addresses which belong to the same router. Several tools exist to resolve alias IP addresses [14,15,21]. The third task, subnet resolution, identifies subnet relations among IP addresses and reveals connectivity that is not observed in collected path traces [16,22,17]. The last task emerges due to the fact that not all routers respond to traceroute probes during topology collection and is the main focus of this paper.

All of these tasks are important for obtaining accurate samples of the underlying network. The accuracy and completeness of these tasks significantly affects the accuracy of the resulting topology maps [23–25]. Hence, topology measurement studies should handle these tasks to obtain representative topology maps. Moreover, when handling these tasks, one needs to make decisions based on observations to infer the underlying topology. As the earlier decisions affect the later ones, obtaining the most likely topology under various conditions has been shown to be NP-hard [26]. Several approaches have been proposed to reduce the set of hypotheses in the decision making of the resolution tasks [27,18,14,28,20].

A router that is unresponsive to measurement probes is represented by a '*' in traceroute output and the same router may be probed multiple times in different traces. Even though routers are expected to generate an ICMP message indicating TTL expiration, it is not mandated (see RFC 1812: Requirements for IP Version 4 Routers). Moreover, different protocols have different success rates in eliciting a response [29]. Without any resolution, each occurrence of '*' is treated as a potentially different router in the graph. Unresponsive router resolution task focuses on identifying the '*'s belong to the same unresponsive router. In this paper, we use the term *unknown node* to refer to a '*' in the traceroute output and *unresponsive router* to refer to the actual router that is represented by this unknown node (i.e., by this '*'). Similarly, we use the term *known node* to refer to an observed IP address in the traceroute output and *responsive router* to refer to the actual router that is represented by this known node (i.e., by the IP address(es)).

Depending on the number of unresponsive routers and the topology collection scenario, the collected set of path traces may include a large number of '*'s. For instance, let us consider the Internet2 backbone presented in Fig. 1a where squares represent the vantage points attached to each router shown as a circle. Assume that the routers at *Washington, DC* and *SaltLake* are configured to be unresponsive and path traces are collected between all vantage points. Under these assumptions, the topology that is constructed from the 36 path traces between 9 vantage points would be as in Fig. 1b (with no resolution). The common approach of pruning (see Section 4 for details) will produce a graph as seen in Fig. 1c and leave artificial nodes that needs to be further processed.

In this paper, we first analyze the nature of unresponsive routers and identify different types of unresponsiveness. We observe that temporary unresponsiveness is an important case that is omitted in previous studies. We also investigate the responsiveness of routers to active network measurements. Expanding our earlier work [29] on skitter [30] data sets, we present the prevalence of unresponsive routers in historical Ark [31] and iPlane [32] measurement data sets.

We then introduce a *Graph Based Induction* (GBI) approach, based on our earlier work [18], to resolve unresponsive routers in traceroute based topology mapping studies. In our work, we define a new induction approach from the practical context of the unresponsive router resolution problem and develop an efficient implementation. To this end, we first examine topology maps that are constructed from traceroute data with unresponsive routers and identify a number of graph structures that are formed among unknown nodes and their known neighbors. We then use *Structural Graph Indexing* (SGI) to detect these structures in the graph and reduce the unknown nodes (i.e., the occurrences of '*'s) into their corresponding unresponsive routers. We enhance our SGI [33] approach to efficiently query constructed topology graphs.

In our evaluations on synthetic topologies, we observe that GBI has better resolution than previously proposed Initial Pruning [34] and Neighbor Matching [35] approaches. Due to their high algorithmic complexity (see Section 2), we did not compare GBI with graph minimization [36], dimensionality reduction [37], and spectral