

TraceNET: An Internet Topology Data Collector

M. Engin Tozal
Department of Computer Science
The University of Texas at Dallas
Richardson, TX 75080 U.S.A.
engintozal@utdallas.edu

Kamil Sarac
Department of Computer Science
The University of Texas at Dallas
Richardson, TX 75080 U.S.A.
ksarac@utdallas.edu

ABSTRACT

This paper presents a network layer Internet topology collection tool called **tracenet**. Compared to **traceroute**, **tracenet** can collect a more complete topology information on an end-to-end path. That is, while **traceroute** returns a list of IP addresses each representing a router on a path, **tracenet** attempts to return all the IP addresses assigned to the interfaces on each visited subnetwork on the path. Consequently, the collected information (1) includes more IP addresses belonging to the traced path; (2) represents “being on the same LAN” relationship among the collected IP addresses; and (3) annotates the discovered subnets with their observed subnet masks. Our experiments on Internet2, GEANT, and four major ISP networks demonstrate promising results on the utility of **tracenet** for future topology measurement studies.

Categories and Subject Descriptors

C.2 [COMPUTER-COMMUNICATION NETWORKS]:
Network Architecture and Design

General Terms

Measurement

Keywords

Internet, Network, Subnet, Topology, Traceroute

1. INTRODUCTION

Many successful research projects and efforts have been introduced attempting to derive an accurate and large scale topology map of the Internet [17, 18, 22, 16]. These efforts focus on different but correlated topology maps: *IP level* maps show IP addresses that are in use on the Internet; *router level* maps group the interfaces hosted by the same router into a single unit (via alias resolution); *subnet level*

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'10, November 1–3, 2010, Melbourne, Australia.

Copyright 2010 ACM 978-1-4503-0057-5/10/11 ...\$10.00.




Figure 1: Traceroute vs tracenet on a path trace. Traceroute collects a single IP address vs tracenet collects a subnet at each hop.

maps enrich the router level maps with subnet level connectivity info; and *AS level* maps demonstrate the adjacency relationship between ASes.

Adverse to the benefits of having a network topology map, the main tools used to collect router or IP address level topology data are a few and operationally limited. **Traceroute** [12] and **ping** are the main data collection tools. **Traceroute** collects a list of IP addresses one for each router on the path between two hosts and **ping** is mainly used to check whether an IP address is in use or not. Almost all topology mapping projects use data collected by **traceroute** from multiple vantage points [19].

In this study, we propose a new end-to-end topology collection tool called **tracenet**. An accurate and complete Internet topology map at the router level requires identifying all routers and subnets among them. **Traceroute** attempts to collect an IP address at each router on a path between two hosts whereas **tracenet** attempts to collect a subnet at each router on the same path. In the worst case, **tracenet** returns the exact path that would be returned by **traceroute**, and, in the best case, it collects the complete topology of each subnet visited on the path. Consequently, a single session of **tracenet** (1) discovers new IP addresses that are missed by **traceroute**, (2) marks multi-access and point-to-point links, (3) reveals subnet relationship among IP addresses on the path, and (4) annotates the subnets with their observed subnet masks. **Traceroute**'s ability to collect a similar data is often limited in practice due to the difficulties of obtaining a reverse path trace and due to the dynamics of the underlying routing behavior between the two systems. As an example, consider the use of **traceroute** and **tracenet** to collect router level topology info between a vantage point

