

Visual Traceroute and Archiving

Geolocating end hosts has improved dramatically in the last few years with tools such as [geoiptools.com](https://www.geoiptools.com) (<https://www.geoiptools.com>) and the Maxmind database (https://www.maxmind.com/en/geoip2-services-and-databases?gclid=EAlaIqobChMlpqra2bry4glVIMVCh1Z7QopEAAAYAAEglUyFD_BwE&rlid=google). However, when it comes to router nodes, the results typically show them as being located at the owner's head office. For example, ESnet routers which are scattered over the US and also in Europe, are identified as located at the ESnet headquarters in Berkeley, California. This makes it hard to understand and visualize the router a packet takes between a source and destination. The paper at <https://conferences.sigcomm.org/imc/2017/papers/imc17-final96.pdf> shows that the existing databases are not reliable for geolocating routers and that there is room to improve their country- and city-level accuracy. They present a set of recommendations to researchers concerning the use of existing geolocation databases to geolocate routers.

Locating nodes

We propose to research and develop a suite of tools to use trilateration (https://www.e-education.psu.edu/natureofgeoinfo/c5_p12.html) techniques to locate hosts. In particular, we will focus on the needs for locating Internet routers. We plan to use Round Trip Time (RTT) measurements made by a ping-based tool from Measurement Agents (MAs) located around the world to the targets, in particular, the routers at traceroute hops. The RTTs are then converted to distance between the MA and target to input to the trilateration.

Ping Based Tool

This will typically be the `traceroute.pl` web-based traceroute/ping tool developed by SLAC. Others may be added should we find a set of accessible ping MAs that we can access to request pings to a target.

Locations of Measurement Agents

We will start with perfSONAR MAs and PingER MAs. A map of the initial MAs can be seen below. The red dots are MAs in the PingER project, the green dots are perfSONAR MAs.



We will maintain a database of these MAs including their name, IP address, web site, physical address, and location, etc.. We will also test these MAs at regular intervals to ensure they are working and update the database of reachable active MAs accordingly. We will also contact associates, particularly in Pakistan, to add more MAs.

Populating the router location database

We will maintain a database of router locations. To populate this we will identify representative targets in regions of the world and carriers, and make regular traceroutes from the MAs to the targets and use the geolocation to locate the unknown positions of the routers at each hop. Having a database of router locations will mean that interactively locating the routes in a traceroute will be much faster (no need to wait while we ping the router at each hop from multiple MAs). If the router location is not in the database when a user requests a visual traceroute then the result will be added to the database.

To assist in populating this database, we will also develop a tool to provide a traceroute archive. The traceroutes will be made from an MA at Stanford in the US and a second in Pakistan in Islamabad to over multiple end hosts in over 160 countries (see <http://www.slac.stanford.edu/comp/net/wan-mon/viper/pinger-coverage-gmap.html> for a map of the target hosts). The measurements will be made regularly (initially daily) and archived. The measurements will be accessible via the web, together with an analysis including: a comparison of the traceroutes for 3 selected days (by default the most recent three days), identifying routing routes and comparing traceroute RTTs to hops with those for ping RTTs. This data may later be analyzable to identify significant changes in the traceroutes.

Traceroute visualization tool

This will be based on, evolve from and extend the TULIP traceroute visualization tool (see <http://www.slac.stanford.edu/comp/net/tulip/> and <https://confluence.slac.stanford.edu/pages/viewpage.action?pageId=160007381> which has now no longer supported and fallen into disuse).

Outcomes

A web based tool to visualize traceroutes.

A study of the accuracy of the predictions and the impact of different numbers of MAs, their RTTs and distance to the target.

A database of the locations of Internet routers.

A publicly accessible archive of traceroutes.