



AFRL-RI-RS-TR-2012-014

MAPPING THE WHOLE INTERNET WITH PASSIVE MEASUREMENTS

DUKE UNIVERSITY

JANUARY 2012

FINAL TECHNICAL REPORT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

STINFO COPY

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09. This report is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-RI-RS-TR-2012-014 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/s/

ROBERT L. KAMINSKI
Work Unit Manager

/s/

WARREN H. DEBANY JR., Technical Advisor
Information Exploitation and Operations Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) January 2012		2. REPORT TYPE Final Technical Report		3. DATES COVERED (From - To) July 2010 – July 2011	
4. TITLE AND SUBTITLE MAPPING THE WHOLE INTERNET WITH PASSIVE MEASUREMENTS				5a. CONTRACT NUMBER N/A	
				5b. GRANT NUMBER FA8750-10-2-0193	
				5c. PROGRAM ELEMENT NUMBER N/A	
6. AUTHOR(S) Bruce Maggs				5d. PROJECT NUMBER BYU1	
				5e. TASK NUMBER DU	
				5f. WORK UNIT NUMBER KE	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Duke University 2200 West Main Street, Suite 710 Durham, NC 27705-4677				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/Information Directorate Rome Research Site/RIG 525 Brooks Road Rome NY 13441				10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/RI/RRS	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-RI-RS-TR-2012-014	
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited. This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This final technical report describes an effort to develop a comprehensive and accurate map of the Internet using passive measurements, diverse data sets and statistical learning methods. The effort passively collected a comprehensive set of Internet traffic and topology information. It captured a wide range of sources, including the Akamai data, Honeypot traffic, Border Gateway Protocol (BGP) traffic, background radiation, and distributed packet capture at Web servers and other network choke points. The data was complemented with selective active (traceroute-like) measurements as needed to enhance and/or validate the inferred topology. The effort also fused these data sources into a coherent picture of Internet topology and geography in a real-time fashion.					
15. SUBJECT TERMS Internet mapping, Internet topology, Border Gateway Protocol, BGP					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 18	19a. NAME OF RESPONSIBLE PERSON ROBERT L. KAMINSKI
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code) N/A

TABLE OF CONTENTS

1.0 Introduction.....	1
2.0 Background: The Alidade Geolocation System.....	1
3.0 Alidade Architecture.....	2
4.0 Annotated Bibliography.....	6
5.0 Bibliography	8
6.0 Appendix: Summary of Data Sources.....	9
7.0 List of Acronyms	14

1.0 Introduction

This report summarizes progress made in improving the state of the art in Internet geolocation by the Principal Investigator (PI) under contract FA8750-10-2-0193. Geolocation is the process of predicting the geographical locations of hosts and routers connected to the Internet based on their Internet Protocol (IP) addresses. In collaboration with researchers at Cornell University and Akamai Technologies, the PI has laid the groundwork for this research by developing a geolocation framework called “Alidade,” which is described in more detail in the remainder of this report.

2.0 Background: The Alidade Geolocation System

Alidade is an Internet geolocation system. In common with other geolocation systems, Alidade’s goal is to accurately predict the geographical location of any given Internet IP address. For example, given a query for the IP address 128.2.205.42, Alidade should present an answer such as “Pittsburgh, Pennsylvania,” or perhaps an even more specific answer, such as Global Positioning System (GPS) coordinates.

Alidade differs from previous geolocation systems studied in academia in several important respects. First, Alidade must return an answer for any IP address without initiating measurements after the query is presented. In particular, Alidade is not permitted to direct any network traffic probes toward the target address in order to determine its location. Virtually all previous geolocation systems reported in the academic literature, including Octant (Wong, Stoyanov, & Sirer, April 2007), Spotter (S., Mátray, Hága, Sebók, Csabai, & Vattay, April 2011), etc., probe the target address to estimate its position. Geolocation systems that are more comparable to Alidade are commercial “off-line” systems such as Akamai’s Edgescap and MaxMind’s GeoLite City databases.

Another important distinction between Alidade and previous systems is that Alidade also does not initiate any active measurements *prior* to being presented with queries. Instead, Alidade relies on data collected from other sources for other purposes. For example, Alidade can process network measurement data such as traceroutes already collected by Akamai, the Cooperative Association for Internet Data Analysis (CAIDA), and other organizations, but does not launch its own traceroutes. As a consequence, it is not possible to detect whether Alidade is being used to predict the location of a network address by monitoring network traffic.

Yet another distinction is that Alidade is designed to operate on data at a very large scale. Alidade has been developed in collaboration between Akamai, Duke University, and Cornell University. The design is derived from Cornell's previous Octant geolocation system (Wong, Stoyanov, & Sirer, April 2007), but Alidade differs in that it is designed to operate on a Hadoop cluster of servers, allowing it to process much larger data sets. Furthermore, unlike Octant, Alidade has no notion of a single "target" address that it is to locate. Instead, based on the entirety of its input data, Alidade attempts to predict the location of every IP address for which it has some measurements. (In future work, to be carried out by the PI, Alidade will also infer the locations of addresses for which it has no measurement data.)

Alidade has reached a state of maturity where it can process a variety of different types of data on a large scale. The next phase of the development will be to evaluate the quality of the predictions that Alidade is making. Among the goals are to determine which data sets are most useful, and to fine tune the algorithms that are used for making predictions. Evaluating a geolocation system is non-trivial because of the general scarcity of "ground truth," i.e., IP addresses for which the true geographical location is known.

The remainder of the report provides a high-level description of the Alidade software architecture and describes how it operates. It concludes with an annotated bibliography, focusing on recent work that is seen as most significant or relevant to the research.

In an Appendix, the report reviews the various input data sources that Alidade processes. Much of this data is proprietary and belongs to Akamai Technologies. It covers the detailed features of each data source, and the efforts that have been undertaken to collect and in some cases improve the reliability and accuracy of these sources. It also describes efforts in building up long-term archives of data.

3.0 Alidade Architecture

A "run" of Alidade proceeds through three stages – data import, evaluation, and aggregation. The separation into stages is the consequence of a modular design that allows the stages to be run independently to a large extent. Alidade schedules these stages as *map-reduce* jobs to enable processing of large volumes of data in an efficient way.

Data Import

Alidade can ingest latency information from diverse datasets viz., *traceroutes*, *pings*, and *TCP stats*. The import is often followed with a transformation to an internal canonical format. Implicit in this transformation process is the task of inferring the "best latency" for any given landmark-target pair, where a landmark is the source of a measurement, and a target is the destination. Experience has shown that the simple approach of choosing the minimum latency between a pair of addresses can cripple the system if the data is

corrupted in any way. Alidade attempts to detect such inconsistencies and makes an educated guess at the best and “safe” measurement using the *median* and *variance* of the data distribution. The transformation helps retain just the key details that can be used to make informed decisions about the use of the data. For instance, by recording whether an observation is *direct* – explicitly measured observation with a *beacon* (an address with *a priori* information on its location) as its source or *indirect* - an observation not originating at a beacon and inferred from other direct data, Alidade can *weight* the measurements and control how they contribute to the answer.

Data Evaluation

The evaluation phase is an *iterative* estimation process. The predicted locations of the targets improve gradually as the iterations move them closer to the true location. To best understand the evaluation phase, we can visualize it as an orchestration of three different tasks – *selection of data points*, *generation of shapes*, and *finalization of result*, that encapsulate the actions performed during this stage.

Selection of Data Points

There are scenarios where processing all available observations between a pair of landmark and target nodes might simply not be practical. Further, a large quantity of observations doesn’t necessarily guarantee good results. Hence, selection of data points in such circumstances plays a vital role in avoiding wasting resources processing unnecessary data. In general, given a large volume of latency measurements Alidade first classifies them, as already mentioned, into two categories – direct and indirect observations. If there are a large number of short measurements, each say less than a millisecond, then it combines such good measurements into one to generate a tight area indicating the target’s location.

A not-so-good situation is one where there are a lot of long latencies between the landmarks and the targets. In such cases, Alidade looks for at least one direct measurement less than 75 milliseconds (ms) and picks up the top 50 indirect measurements to be used for geolocation. The rest, if any, are discarded. The direct measurements in this case help to generate the first estimate of the target’s location and the indirect ones help to refine this estimate. In general, direct measurements are weighted higher than the indirect measurements; the justification is that the direct measurements are more reliable compared to indirect. Hence, the only case where Alidade is likely to fall short in geolocating a target is when there are no quality direct measurements, say less than 75ms from any landmark to a given target.

Generation of Shapes

Using the list of observations selected at the previous step, Alidade proceeds to generate shapes corresponding to the measurements and evaluate them. The shape is Alidade’s estimate of the location of a target. To enable a fair comparison with existing geolocation systems, Alidade picks a point within the area as the target’s “exact” location. Even for beacons, where the “true location” or point is known, Alidade computes a point

representing the “calculated location” of the same. In short, Alidade plays a fair game – every target’s location is computed and *nothing is hardcoded*.

A direct measurement implies that the origin of the measurement should be a point with an associated *latitude* and *longitude* coordinate. With the point as the center, a circle of radius proportional to the measurement is drawn. Shape generation for indirect measurements follow a slightly different approach, since they need not have a point associated with the landmark. Inferring a landmark, or in other words geolocating the landmark gives us the center of the circle.

Although it looks like a “chicken and egg problem”, with a target’s geolocation requiring that the landmark be geolocated first, which may require geolocation of some other landmark and so on, progress is still possible since the entire evaluation phase proceeds in a iterative fashion. In each iteration, more addresses are be geolocated and for ones already geolocated the estimates improved. Hence, its possible to infer a point for the landmark associated with the indirect measurement using results from a previous iterative evaluation step. If that fails, Alidade can leverage other datasets viz., *HostParser answers* (described in the Appendix), to guess a point where a circle of appropriate radius can be pinned down. An indirect measurement is discarded if all attempts to generate a point for the landmark fail.

The shapes can further be refined using inputs from various datasets viz., *water files*, and *registry information*. With shapes generated for all measurements from multiple landmarks to a single target, calculating the common overlap or intersection of the shapes provides a good estimate of the target’s location. If the intersection of the shapes is a single region, then there is nothing more to be done. The more common case of intersection giving rise to multiple fragments or regions, a composite of the fragments is used as the target’s estimated location. Further, by weighting the shapes and using an additive weighting scheme for the intersections and taking only the top x percent of the weighted fragments, Alidade can track just the fragments that provide most information regarding the target’s location.

Finalizing the result

This phase is an *optional* step, usually carried out for benchmarking purposes. It involves a series of tests to evaluate the quality of results computed by Alidade. The first test is to understand the contribution from the HostParser answers. It proceeds by loading all HostParser answer points and checking for containment of these points in the corresponding target’s calculated location (area). For targets that are beacons, tasks include calculating the *error distance* – distance between true and calculated locations and *area size* – radius of a circle of equivalent area. By plotting the area size against the error distance factors contributing to bad results, which can be either corrupt datasets or true geolocation errors, potential problems with either the input data or Alidade’s algorithms can be highlighted.

Alidade also provides the option of looking up the city “closest” to the geolocated target. It finds all cities contained within the target’s shape, and the city closest to the answer point is used as the city location. A variation in the approach allows for cities to be weighted, based on the estimated population, with the weights used to choose the closest city. The higher the population of the city, the greater the weight and the more likely it is a candidate for the city closest to the target’s location.

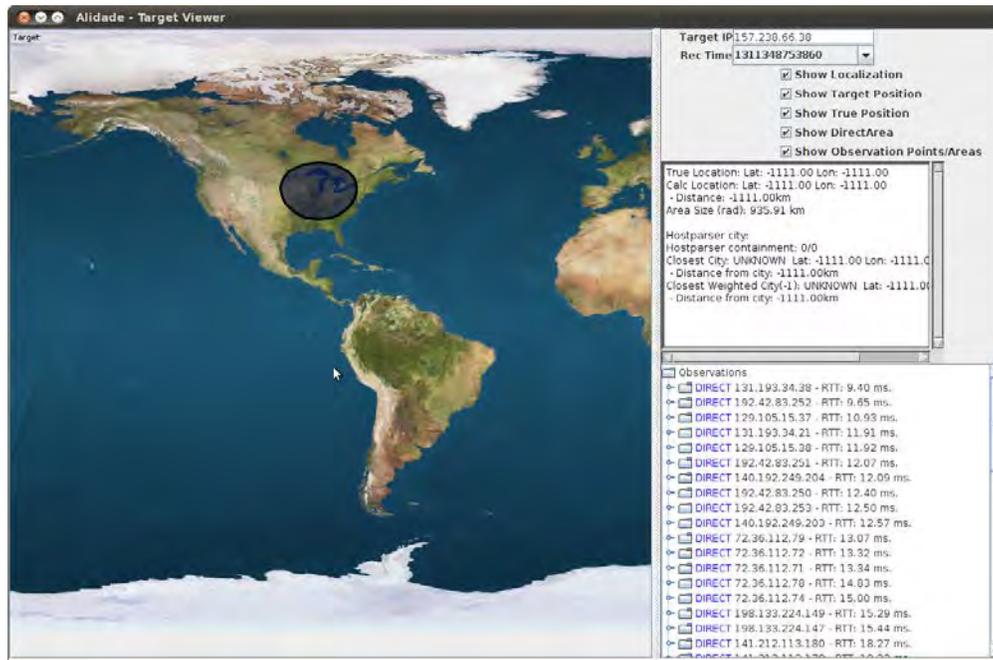


Figure 1: ViewResults tool in action

Finally, to help in debugging or evaluating results of a specific target, Alidade provides an option of writing out detailed results to a database called *HBase*, which can be queried by the *ViewResults* tool, shown in Figure 1. The *ViewResults* tool helps to visualize the iterative location estimation process and the contributions of various measurements to the final computed location.

Data Aggregation

To accommodate cases where a target that needs to be geolocated has absolutely no associated measurements, Alidade allows the solutions computed to be aggregated. Aggregation merges the answers computed for IPs (with some associated measurements) within a prefix. If the aggregation is done carefully and represented in a meaningful fashion, it can help geolocate other IPs in the same Classless Inter-Domain Routing (CIDR) block for which there are no measurements. The current implementation is still just a prototype towards achieving this goal.

4.0 Annotated Bibliography

This section reviews recent technical literature pertaining to IP geolocation including proposed algorithms, the measurements which are typically involved in their support, and related matters such as assessments of dynamic network topography. The organization is by research groups and their major papers. The order is arbitrary.

Model Based Geolocation with Emphasis upon Locating Routers

Laki, *et al*, build a system, Spotter (S., Mátray, Hága, Sebök, Csabai, & Vattay, April 2011), which combines probability density functions constructed from spatial densities giving lower and upper bounds on ranges from observing locations. This achieves statistical power because they pursue a hunch that range versus latency follows densities which are independent of location. Their model for distance versus latency is universal in the sense that *any* IP geolocation system ought to be able to use it.

Geolocation Using Probability Models for Latency Measurements

Where Spotter fits a regression line to latency measurements in order to calibrate a conversion from latency to distance, GeoWeight (Arif, Karunasekera, & Kulkarni, 2010) uses range probability estimates from a PlanetLab and iPlane latency-distance calibration to directly calculate probability mass estimates of position. In particular, whereas Spotter assumes that the density of distance given a particular latency value is unimodal and symmetric (as well as Gaussian), GeoWeight constructs an empirical density model for latency bins which have equal numbers of calibrating observations. This is essentially an equalized range histogram for each latency bin. The prediction of target placement, then, for a particular landmark is a series of concentric annuli each having an appropriately calculated weight.

Extreme Priors

Sources of information regarding IP address density and geographic position are increasing in number and size. These are derived from search engine queries, lists of WiFi locations, monitoring traces of cell phones and correlating these with TCP connections, war-driving, as well as cartographically registered demographic and economic data (Guo, Liu, Shen, Wang, Yu, & Zhang, 2009). In the limit, if a target is known to be using and near a pre-surveyed Internet access point, the geographic position of that point might serve as a good proxy for the target's position¹. This is essentially the approach of Wang, Burgener, Flores, Kuzmanovic, and Huang (Wang, Burgener, Flores, Kuzmanovic, & Huang, 2011). They exploit leakage of geographic information from

¹ In the simplest case, if the user is in an Internet café and the café's position is known, the position of the user is known.

Web sites regarding organizational entities associated with them. In addition, Wang *et al* use a rich set of carefully screened *passive* landmarks in combination with a pre-surveyed indirect inference of their *network distance* from a target² to determine the closest passive landmark and then give that landmark's location as their estimate of target position.

Similarly, Li, Chen, Guo, Liu, Zhang, Zhang, and Zhang use a rich set of landmarks and outbound HTTP/Get probing of them from a target to develop a geolocation system called *GeoGet* (Li, et al., 2009). They simply map the address with unknown location to the landmark having the shortest delay, arguing this approach may be good for contexts where landmarks and users are not "richly connected".

Structon is a massive effort to mine 500 million Web pages in China to derive their geographic positions (Guo, Liu, Shen, Wang, Yu, & Zhang, 2009). Guo, Liu, Shen, Wang, Yu, and Zhang report that through simple data mining, two thirds "of Web server IP addresses can be correctly mapped to correct cities, and 73 percent mapped to correct provinces". But the inference algorithms of Guo *et al* "significantly [improve] the accuracy to about 87.4 percent at the city level, and 93.5 percent at the province level".

² This uses what can potentially be a massive number of traceroutes, and so is an active technique.

5.0 Bibliography

Arif, M. J., Karunasekera, S., & Kulkarni, S. (2010). GeoWeight: Internet host geolocation based on a probability model for latency measurements. *Proceedings of the 33rd Australasian Computer Science Conference*. Brisbane, Australia.

Guo, C., Liu, Y., Shen, W., Wang, H., Yu, Q., & Zhang, Y. (2009). Mining the Web and the Internet for accurate IP address geolocations. *Proceedings of INFOCOMM*.

Li, D., Chen, J., Guo, C., Liu, Y., Zhang, J., Zhang, Z., et al. (2009). *IP-geolocation mapping for involving moderately-connected Internet regions*. Microsoft.

S., L., Mátray, P., Haga, P., Sebok, T., Csabai, I., & Vattay, G. (April 2011). Spotter: A Model Based Active Geolocation Service. *Proceedings of IEEE INFOCOM 2011*. Shanghai, China.

Wang, Y., Burgener, D., Flores, M., Kuzmanovic, A., & Huang, C. (2011). Towards street-level client-independent IP geolocation. *Proceedings of the Eighth USENIX Conference on Networked Systems Design and Implementation*.

Wong, B., Stoyanov, I., & Sirer, E. G. (April 2007). Octant: A Comprehensive Framework For The Geolocalization Of Internet Hosts. *Proceedings of the Symposium on Networked System Design and Implementation*. Cambridge, Massachusetts.

6.0 Appendix: Summary of Data Sources

This section contains a list of data sources that Alidade can process, as well as additional sources under consideration. In some of the data sets, we have “ground truth” information for IP addresses. In this document, the definition of ground truth for an IP is that we have high confidence that we know IP’s geographical location, which is typically represented by latitude and longitude.

Traceroute data

All traceroutes in the different datasets listed below are collected with standard traceroute tools using ICMP or UDP packets, with each of the hop in a traceroute containing the IP address (but not the hostname) and the corresponding round-trip latency for that hop.

PlanetLab to PlanetLab and Akamai traceroute data

- These traces were initiated by the Alidade project solely for the purpose of evaluating accuracy. In actual operation, Alidade does not initiate any measurements.
- Traceroutes were collected from PlanetLab³ nodes to both PlanetLab and Akamai IPs.
- 16 different sets of about 15 million traceroutes in each set were collected by Cornell around early 2010 over a period of about two weeks.
- The traceroutes were collected with a standard traceroute tool using UDP packets from about 500 PlanetLab nodes to about 1,000 PlanetLab IP addresses and 25,000 Akamai IP addresses.
- About 70% of the traceroutes reached their destinations.
- Ground truth⁴ is known for each traceroute source.
- Ground truth is known for each destination.
- A new collection is underway with updated PlanetLab and Akamai IP addresses using ICMP packets and the Paris traceroute tool⁵.

Akamai to Akamai and PlanetLab traceroute data

- The Akamai to PlanetLab traces were collected for the purpose of evaluating accuracy, and would not be initiated as part of Alidade’s normal operation.
- Traceroutes were collected from all Akamai locations to both PlanetLab nodes and all Akamai locations.

³ <http://www.planet-lab.org/>

⁴ PlanetLab has ground truth data on their nodes, however, the data is not necessarily all accurate and needs to be verified (see <http://www.onelab.eu/index.php/media-centre/news/446.html>)

⁵ <http://www.paris-traceroute.net/>

- A data set of about 6 million traceroutes was collected around the end of October, 2010.
- The traceroutes were collected from about 1,500 Akamai locations to two IP addresses from each of 1,500 Akamai locations and to about 1,000 PlanetLab IP addresses.
- The traceroutes were collected using an internally developed traceroute program similar to the standard tool using ICMP packets.
- Over 90% of the traceroutes reached their destinations.
- The ground truth locations of all sources and destinations are known.
- A new collection is planned using updated Akamai and PlanetLab addresses and locations.

Mapper traceroute data

- Akamai's "mapping" system, which determines which clients should be directed to which Akamai servers, performs a large number of traceroutes in order to model the Internet topology. Alidade does not initiate these traceroutes, but uses them in its normal course of operation.
- Traceroutes are collected from 200+ Akamai locations to about 280K IP addresses. Most of these IP addresses belong to resolving name servers.
- About 5-6 million traceroutes are collected on a daily basis and this data has been archived since February 2011.
- Traceroutes are performed to the 280K name servers that sent Akamai the most DNS requests on the previous day.
- About 30 traceroutes are collected to each NS from Akamai locations that are in the same continent. Determination of name server continents is performed using Akamai's commercial geolocation tool called EdgeScape⁶.
- Traceroutes are collected using an internal developed traceroute program similar to the standard tool, using ICMP packets.
- About 50% of the traceroutes reached their destinations.
- The ground truth locations of all sources, but not destinations, are known.

EdgeScape traceroute data

- Akamai's commercial geolocation product, EdgeScape, periodically performs traceroutes to randomly select client IP addresses.
- Traceroutes are collected from about 10 Akamai locations to about 20 million end-user IPs.
- About 20 million traceroutes are collected roughly every 2 weeks and one set of data per month has been archived since January 2010.
- One traceroute was collected from an Akamai location to each IP addresses.
- Traceroutes are collected using an internally developed traceroute program similar to the standard tool using ICMP packets.
- About 30% of the traceroutes reached the destination.

⁶ <http://www.akamai.com/html/technology/products/edgescape.html>

- Ground truth is known for sources but not destinations.

iPlane⁷ traceroute data

- The iPlane project, which is independent of our efforts, collects a large amount of traceroute data and makes this data available to other researchers.
- Traceroutes are collected by iPlane from various academic institutions (including a large number of PlanetLab nodes) to about 140K IP addresses.
- About 20 million traceroutes are collected daily and this data has been archived by us since February, 2011.
- The traceroutes are collected from about 200 locations to 140K IP addresses.
- The traceroutes are collected using a standard traceroute tool, but it is unclear if the collection used ICMP or UDP packets.
- About 45% of the traceroutes reach their destinations.
- iPlane does not provide ground truth locations for sources, although they may not be difficult to determine.
- Ground truth is not available for destinations.

Ping data

Ping data is similar to traceroute data except that instead of recording a route from a source to a destination, the ping tool only collects the round-trip latency.

Mapper ping data

- In addition to collecting traceroute data, Akamai's mapping system also collects a large amount of ping data.
- ICMP ping data is collected from over 1000 Akamai locations to about 30K IP addresses (mostly routers and servers).
- Data is collected around the clock with a ping to each IP roughly every few minutes.
- Data has been archived roughly every 15 minutes since August 2010.
- Ground truth is known for all source locations.
- Ground truth is known for some destinations (mostly Akamai machines), but not the vast majority of destinations.

TCP data

Akamai web servers (caches) collect statistics about the TCP connections made with client machines. Only a small fraction of the TCP connections are sampled, but for each one, several useful pieces of data are logged. This data is referred to as "TCP stats."

⁷ <http://iplane.cs.washington.edu/>

TCP stats from Akamai machines

- TCP stats are collected for TCP connections between over 1000 Akamai locations and (potentially) any client/user IP addresses requesting content from an Akamai server.
- Only a very small percentage of all client requests are logged.
- The round-trip time between the sending of a SYN/ACK by an Akamai server and the receipt of an ACK is included in the log.
- The TTL (time to live) of the SYN packet received from a client is included in the log.
- Data has been archived since 2009.
- Ground truth is known for source locations but not destinations.

HostParser data

Akamai has developed a tool called “HostParser” that translates the reverse DNS names that ISPs register for their routers or other resources into city names. HostParser is similar to the well known undns tool but has much more coverage and is kept up to date by the EdgeScape product maintainers.

HostParser data from EdgeScape

- HostParser extracts locations from reverse DNS host names.
- A new set of Hostparser data is produced roughly every 2 weeks.
- We have been archiving this data for several months.
- Ground truth is not known for this data.

Shapefile data

Shapefile data on coastlines, islands, ponds, and lakes of the world

- Obtained from public source (<http://mapaspects.org/gis/data/index.php?dir=vector%2Fworld%2F>)
- Alidade can use shapefiles to eliminate water body and other areas that are considered unlikely locations.

Shapefile data on land mass of the world

- Obtained from public source (http://mapaspects.org/gis/data/index.php?dir=vector%2Fworld%2FGlobal_Admin_Areas%2F)
- The shapefile data includes country level data (level 0), state/region level data (level 1), and county level data (level 2).
- Alidade can use this data together with HostParser output (a city location) to pull in an appropriate shapefile for area processing.

Registry data

Registry data from EdgeScape

- Registry data that EdgeScape gathers from various registries (ARIN, RIPE, APNIC, LACNIC, AFNIC, JPNIC, etc.)
- Unlike other IP-based data that Alidade uses that is per-IP, this

- data is represented as cidr blocks
- Alidade has interfaces to pull out registry information for any IP and can use such information in similar fashion as Hostparser data

Other data under consideration

Below are a few additional data sources being investigated to see how they can be used in Alidade to improve its geo-location capability, augment the ground truth data for evaluation, and development to support queries on IPs with no direct observations/data. The investigation mostly focus on if there are values in such dataset and aspects that can help Alidade, plus what Alidade development work may need to use such data sources.

World cities, landmarks and populations

- Data sources include the NGA, US Census, FIPS dataset, and GeoNames data
- GeoNames data seems to be most complete
- Potential usage in Alidade is to help position cities, as well as finding the area that an answer covers

Additional shapefile at city level resolution

- Other shapefile sources, like Natural Earth data and shapefiles at city or sub-city level resolutions, preferably with global coverage
- Potential usage in Alidade is to provide finer resolution shapefiles to help position cities and refinement to the answer area

Icecast GPS locations

- Internal Akamai data that may contain GPS locations from mobile phone users to Icecast service
- Data needs to be validated
- Data can be used to augment ground truth data in Alidade to help expand the volume and diversity of ground truth data

Insight Zip code information

- Internal Akamai data that may contain zipcode for end users IPs from online transactions
- Data needs to be validated
- Data can be used to augment ground truth data in Alidade to help expand the volume and diversity of ground truth data

BGP data

- Data may be useful to answer queries on IPs with no direct observations/data by focusing on relevant information (like direct observations of IPs from the same prefixes/cidrs)

7.0 List of Acronyms

AS	Autonomous System
CAIDA	Cooperative Association for Internet Data Analysis
DNS	Domain Name Service
GPS	Global Positioning System
ICMP	Internet Control Message Protocol
IP	Internet Protocol
ISP	Internet Service Provider
UDP	User Datagram Protocol