

Macroscopic analyses of the infrastructure: measurement and visualization of Internet connectivity and performance

Bradley Huffaker, Marina Fomenkov, David Moore, and kc claffy

Abstract— The robustness and reliability of the Internet is highly dependent on efficient, stable connectivity and routing among networks comprising the global infrastructure. To provide macroscopic insights into Internet topology and performance, the Cooperative Association for Internet Data Analysis (CAIDA) has developed and deployed the *skitter* tool to dynamically discover and depict global Internet topology and measure performance across specific paths. We are developing a systematic approach to visualizing the multi-dimensional parameter space covered by *skitter* measurements aggregated on a daily basis. In this paper we discuss our techniques and apply them to selected daily *skitter* snapshots.

Keywords— *skitter* visualization IP AS geopolitical topology RTT

INTRODUCTION

The Internet has grown by several orders of magnitude since the original ARPAnet topology map fit on a single piece of paper. However, the need to visually understand Internet topology has not diminished. On the contrary, capturing macroscopic topology and routing information requires increasingly sophisticated techniques and analysis methodologies. CAIDA has developed specialized tools to capture, store, process, and visualize the performance and stability of a large sample of paths in the IPv4 address space.

This paper is organized as follows. Section I provides basic background on the data collected by CAIDA’s *skitter* tool, including selection of the hundreds of thousands of probed destinations. In section II we analyze properties of forward IP path hop count distributions, RTT distributions, and relationships between hop count and RTT. Section III examines relationships between geography and RTT performance and presents a visualization for comparing observed quality of performance to different countries. Section IV presents a visualization technique, *dispersion graphs*, for illustrating forward path topology relationships to hundreds of thousands of destinations based on administrative domain (ASN) and geography.

All authors are with CAIDA, San Diego Supercomputer Center, University of California, San Diego. E-mail: {bradley, marina, dmoore, kc}@caida.org.

Support for this work is provided by DARPA NGI Contract N66001-98-2-8922

I. BACKGROUND

A. The *skitter* project

*skitter*¹ is a tool that measures the forward path and round trip time (RTT) to a set of destination hosts by sending probe packets through the network. It does not require any configuration or cooperation from the remote sites on its target list. The main objectives of the *skitter* project are:

- Collect path and RTT data in a manner similar to traceroute. *skitter* increments the “time to live” field of each probe packet and records replies from each router (or hop) along the path to the destination host. Unlike the default of UDP used by traceroute, *skitter* uses ICMP echo requests as probes .
- Acquire infrastructure-wide global connectivity information, by measuring forward IP paths from a source to large sets of destinations. CAIDA currently uses *skitter* to probe hundreds of thousands of destination hosts around the world.
- Analyze the visibility and frequency of IP routing changes. Low frequency persistent routing changes are visible through analysis of variable links across specific paths.
- Visualize network-wide IP connectivity (viewed as directed graphs from a source). Revealing the nature of global IP topology is a primary goal of *skitter*. Probing paths from multiple sources to a set of destinations that carefully stratify the current IPv4 address space allows us to characterize a statistically significant fraction of macroscopic Internet connectivity.

An essential design goal of *skitter* is to execute its pervasive measurement while incurring minimal load on the infrastructure and upon final destination hosts. In line with this goal, *skitter* packets are 52 bytes in length, and we restrict the frequency of probing 1 packet every 2 minutes per destination and 300 packets per second to all destinations. To improve the accuracy of its round-trip time measurements, CAIDA added a kernel module to the FreeBSD operating platform used by its *skitter* boxes. Kernel timestamping does not solve the synchronization issue required for one-way measurements, but reduces variance caused by multitasking processing when making round-trip measurements. This feature helps to capture performance variations across the infrastructure more effectively. By comparing data from various sources, we can identify points of

¹<http://www.caida.org/tools/measurement/skitter/>

congestion and performance degradation or areas for potential improvements in the infrastructure.

The first four *skitter* monitors were deployed in July 1998. CAIDA gradually increased the number of monitors and currently there are 15 active *skitter* monitors probing various sets of destinations (Table I²). We describe the different destination lists below. Note that we do not explicitly seek permission to probe destinations, since the load presented is trivial - a few ICMP packets a day. We immediately delete from our target lists any sites that ask not to be probed.

B. Destination Lists

We created the **web servers** list in mid-1998 by collecting IP addresses of **web servers** from a variety of log file sources: NLANR's squid[11] caches[12], web servers, search engines. We also traversed parts of the IPv4 address space, using *in-addr.arpa* to get domain names and then adding *www* to the beginning of those domain names to test for existence of web servers. We used the resulting list on our first *skitter* monitors to probe thousands of web servers across the world. A year later, we augmented this list with additional destinations in the Asia-Pacific region. The current **web servers** list has remained unchanged since the end of 1999 and contains about 21,500 geographically diverse destinations. Unless otherwise specified, the data used in this paper corresponds to this **web servers** list and was collected from the end of 1999 to mid-2000. During that time, ten *skitter* monitors used this list. In the summer of 2000, we changed the destination list on several *skitter* boxes (see below), leaving only four still using the original **web servers** list.

Studying long term trends in the **web servers** list, we discovered a somewhat surprising result: destinations in our database become unreachable by *skitter* probe packets at the rate of about 2-3% per month. We are still investigating sources of this degradation. To find out whether non-leaf IP addresses (i.e., those that are not on the last hop of a path) exhibit a similar extinction rate, we developed the **routers** list. This list is composed of intermediate IP addresses seen in *skitter* traces from the Tokyo and Washington, DC (Table I) hosts running the **web servers** list on the 20th and 21st of June 2000. We plan to study the temporal decay characteristics of the **routers** list (containing nearly 31,500 destinations) by running it on two *skitter* monitors for an extended period of time.

The **small** list is a subset of less than 2000 destinations of the **web servers** list, which we run on one *skitter* host to support studies of path and performance dynamics on a finer scale. Each host on the **small** list is probed 10-11 times per hour.

In the summer of 2000, we developed the **IPv4 space** list to try to cover one responding destination for each reachable /24 segment (256 addresses) of IPv4 address space. Stratifying the IPv4 space in this way should provide us with comprehensive topology coverage. We are still in the

²see also <http://www.caida.org/tools/measurement/skitter/lists> and <http://www.caida.org/tools/measurement/skitter/monitors.xml>

process of building this list, and have used a wide range of methods, e.g., *tcpdumps* from the UCSD-CERF link, collecting hostnames from web search engines, adding intermediate IP addresses seen in *skitter* traces. We have collected more than 313,000 destinations, each on a separate /24 segment. This is our largest destination list so far, which we use on two *skitter* monitors. A *skitter* monitor takes approximately 2.5 days to traverse this entire list once. Note that there are over 16 million potential /24 segments in the IPv4 address space, and about 4 million of them are currently in the routable address space. Thus, our coverage is still far from complete.

We created the **DNS clients** list to study the connectivity and performance of root DNS servers. We collected IP addresses seen in passive data obtained from a number of root servers, and selected one IP address per routable prefix. If many IP addresses in the same prefix were found, we used the address that made the largest number of requests to a DNS root server. The resulting list covered 46,844 prefixes (out of approximately 87,408 currently globally routable ones as of 8 August 2000 [10]). To increase the coverage of prefixes, we added addresses from the **IPv4 space** list. The six root server *skitter* boxes run the current **DNS clients** list of more than 58,000 destinations.

Figures 1-4 illustrate the geographical make-up of our destination lists. Each pie chart shows individual countries that contribute more than 3.3% of the destinations to a list, the rest being clumped by continents. The total number of destinations in each list is also shown.

C. Data collected

A 24-hour *skitter* data set typically contains from 300,000 to 500,000 traces, with the number of traces per destination dependent on the size of the destination list. Each trace consists of an RTT to the ultimate destination, and the addresses of intermediate routers that responded. Thus, *skitter* measurements yield a large volume of data, which we categorize along the following dimensions:

- number of hosts running *skitter* and collecting data
- number of destinations probed by each monitor
- timestamp of each probe
- number of distinct forward IP paths observed to each destination
- number of times each particular path was observed per time interval
- length of a path (measured in IP hop count)
- round-trip-time for each trace

We consider a trace to be *complete* when it contains addresses of all intermediate routers. A trace is *incomplete* if a few intermediate addresses are missing, but the probe still reached the destination and got a corresponding RTT. Both complete and incomplete traces are called *responsive*. A trace is *nonresponsive* if the probe failed to reach the destination. The analysis described below pertains to responsive traces only.

TABLE I
LIST OF ACTIVE *skitter* MONITORS AS OF OCTOBER 1, 2000.

Hostname	Controlling organization	Location	Destination List
apan-jp.skitter.caida.org	APAN	Tokyo, JP	Web servers
iad.skitter.caida.org	ABOVE.NET	Washington, DC, US	Web servers
nrt.skitter.caida.org	ABOVE.NET	Tokyo, JP	Web servers
riesling-ether.caida.org	SDSC	San Diego, CA, US	Web servers
lhr.skitter.caida.org	ABOVE.NET	London, UK	IPv4 space
waikato.skitter.caida.org	University of Waikato	Hamilton, NZ	IPv4 space
sjc.skitter.caida.org	ABOVE.NET	San Jose, CA, US	Routers
yto.skitter.caida.org	CANET	Ottawa, CA	Routers
champagne.caida.org	VBNS	Urbana/Champaign, IL, US	Small
a-root.skitter.caida.org	Verisign	Herndon, VA, US	DNS clients
e-root.skitter.caida.org	NASA	Moffet Field, CA, US	DNS clients
f-root.skitter.caida.org	VIX	Palo Alto, CA, US	DNS clients
k-peer.skitter.caida.org	RIPE	Amsterdam, NL	DNS clients
k-root.skitter.caida.org	RIPE	London, UK	DNS clients
l-root.skitter.caida.org	ISI	Marina del Ray, CA, US	DNS clients

II. IP PATH LENGTH AND ROUND TRIP TIME

A. The hop count distribution

IP hop count is a natural connectivity metric that characterizes the proximity (in IP-space topology) of a *skitter* source to the set of destinations it probes. Figure 5 shows the hop count distribution for four *skitter* sources located in California, Japan, and Canada for August 4, 2000. Note that the *skitter* source in Marina del Ray was probing a different set of destinations than the other three monitors shown here. The overall shape of the hop count distribution is similar for all four sources regardless of their destination list. The position of its center on the x-axis, however, primarily depends on the connectivity of the source. For example, the California monitors are both near major exchange points, and have lower IP path lengths to their destinations. The distributions for the Canada and Japan monitors are shifted to the right, implying that they tend to be further away from most of the network.

The data we collected over the last two years suggests that the hop count distribution does not change dramatically for a given monitor over time. Figures 6 - 8 compare the hop count distributions observed by different *skitter* monitors for a weekday in February and August 2000. In all three cases, the shape does not change much over this six-month period. For the San Diego *skitter* monitor (Figure 6), the distribution shifted slightly to the right, indicating some increase in the average path length, likely due to a topology change close to the monitor itself. The Tokyo monitor distribution (Figure 8) shifted slightly to the left, suggesting the opposite type of change.

B. The round trip time distribution

Round trip time (RTT) is a simple Internet performance metric whose value depends on the geographic and topological position of the *skitter* source host with respect to the destinations it probes, as well as on the conditions of

the Internet along paths to those destinations. Figures 9 and 10 show the distribution of median RTTs for a number of different *skitter* sources for June 27, 2000.

If the database of destinations is internationally diverse, as in our data, then a few prominent peaks are usually present in the RTT distribution, corresponding to major geographical clusters of destinations: east and west coasts of the United States, and Europe. RTTs to destinations in Asia, Australia, South America, and Africa generally fall in the tail of the distribution.

The data for three North American monitors are presented in Figure 9. The distributions are not normalized. Note that the Mountain View monitor had a different destination list at the time. In the Mountain View and San Jose monitor curves, the first peak corresponds to west coast destinations, the second one to east coast destinations, and the third peak and tail to European, Asian, and other destinations. For the Canada monitor, the first peak (RTTs to east coast destinations) is shifted to the right and merges together with the west coast peak. The third peak (RTTs to European destinations) is shifted to the left relative to that of west coast sources.

Depending on the location of the *skitter* source, the order of peaks may change. For the London monitor (Figure 10) the leftmost peak corresponds to European destinations, followed by the US destinations, and then by those European destinations that are reached via the US.

For the Tokyo monitor (same Figure), a small number of Asian destinations are reached in less than 100 ms. However, paths between Asian sources and Asian destinations often go via the US. Therefore, the rise of the Japan monitor's prominent leftmost peak does not start until 150 ms along the x-axis, approximately the time it takes to get to and return from a destinations in the US. The other two peaks are similarly shifted to the right.

We suppose that in all cases the differences between peak positions are determined primarily by geographical




Fig. 1. **Web servers** list: 21,577 destinations.




Fig. 2. **Routers** list: 34,223 destinations.




Fig. 3. **IPv4 space** list: 298,821 destinations.

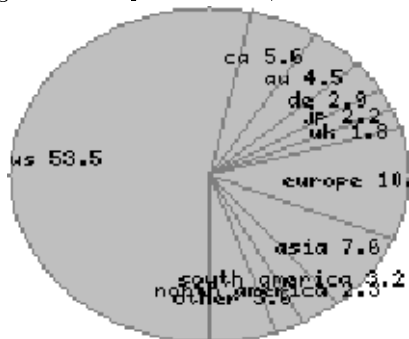


Fig. 4. **DNS clients** list: 58,317 destinations.

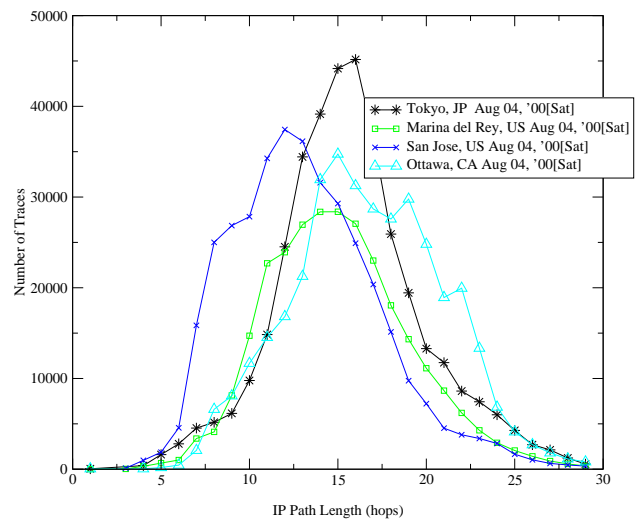


Fig. 5. Hop count distributions for four *skitter* sources, August 4, 2000.

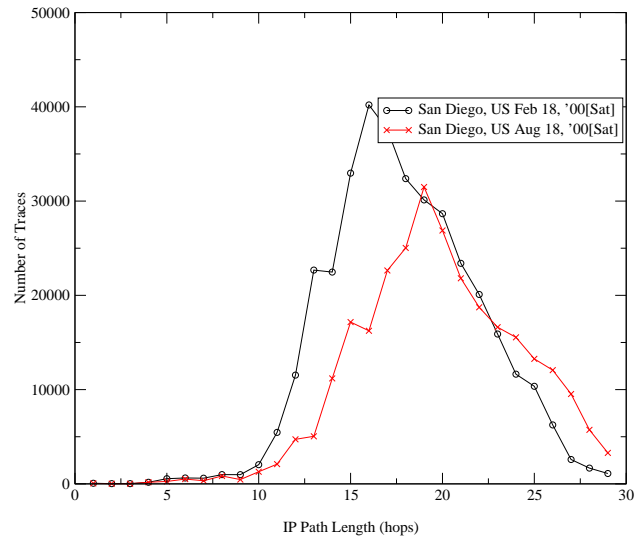


Fig. 6. Hop count distributions for the San Diego, US *skitter* monitor, February 18 and August 18, 2000.

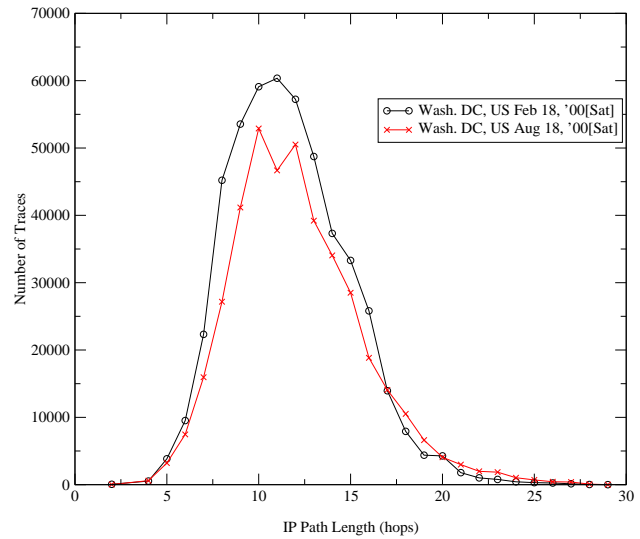


Fig. 7. Hop count distributions for the Washington DC, US *skitter* monitor, February 18 and August 18, 2000.

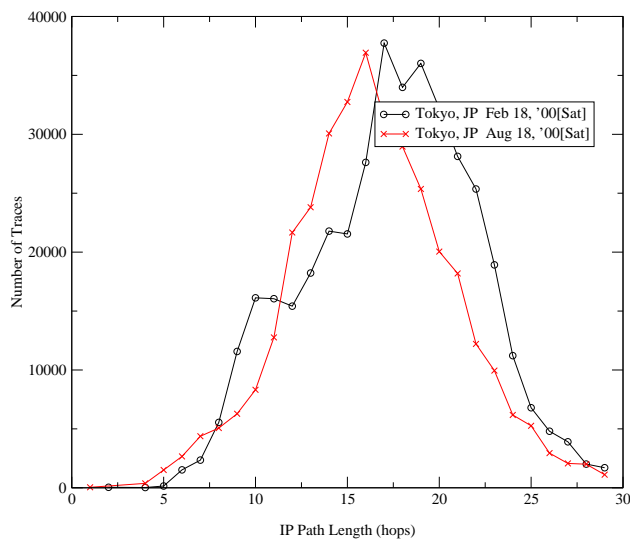


Fig. 8. Hop count distributions for the Tokyo, JP *skitter* monitor, February 18 and August 18, 2000.

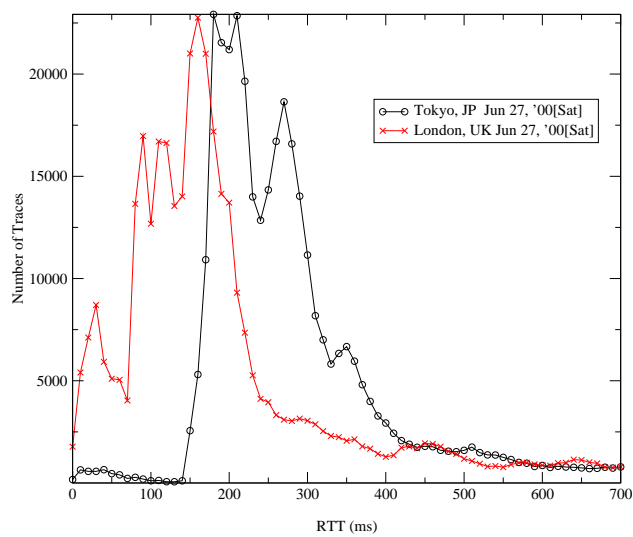


Fig. 10. RTT distributions for the London, UK and Tokyo, JP *skitter* monitors, June 27, 2000.

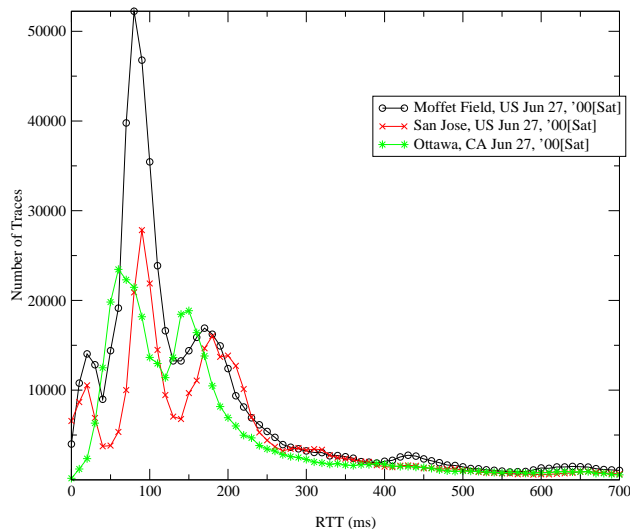


Fig. 9. RTT distributions for North-American *skitter* monitors, June 27, 2000.




Fig. 11. RTT distributions for the San Diego, US *skitter* monitor, February 18 and August 18, 2000.

distance between the source and large sets of destinations.

We showed earlier that the hop count distribution did not significantly change over time for our monitors. What about the RTT distribution? Figures 11 - 13 compare the RTT distributions for three *skitter* monitors observed on a weekday in February and August 2000. Our data for the San Diego monitor (Figure 11) shows that the RTT distribution also did not change significantly between February and August of 2000. The RTTs have dropped somewhat (a macroscopic indication of performance improvement from this monitor), while the number of hops (seen earlier in Figure 6) has slightly increased!

The RTT distributions for the *skitter* monitor in Washington, DC (Figure 12) show a similar small shift to the left, implying a performance improvement, but in this case the distribution of hops (Figure 7) did not change. Finally, the RTT distribution for the Japan *skitter* monitor (Fig-

ure 13) remained stable over the half-year period despite a slight decrease in the number of hops observed by this monitor (Figure 8).

Does RTT show any correlation with hop count? Figures 14 - 15 show percentiles of RTT as a function of IP path length. Generally, we find a rather weak correlation between RTT and IP path length. This result is unsurprising given the complex nature of many layer 3 architectures, where a packet may traverse many IP interfaces in a single machine room, as well as layer 2 connectivity, which may 'hide' hops (from the layer 3 measurement methodology) that still incur non-negligible delay. The drop near the end of the London graph is caused by insufficient statistics due to too few paths.

The minimum RTT between two nodes depends on physical characteristics of a connection such as the geographic distance to the destination, bandwidth, and processing




Fig. 12. RTT distributions for the Washington DC, US *skitter* monitor, February 18 and August 18, 2000.




Fig. 13. RTT distributions for for the Tokyo, JP *skitter* monitor, February 18 and August 18, 2000.

characteristics of intermediate switching equipment. While the speed of light in cable media of the links along the path is a fixed component of RTT, delays in routers due to forwarding lookups, queuing, and other processing vary. For short distances, propagation delay should be negligible relative to time spent in routing hardware. For longer distances, a packet may go through different types of media and many competing factors contribute to delay at each hop. Thus, the RTT to even the same destination may have significant variation during a day. Median and higher RTTs tend to reflect the extent of congestion along the path; maximum RTTs often are due to unpredictable network anomalies.




Fig. 14. Dependence of RTT on the length of IP path for the San Diego, US *skitter* monitor, August 18, 2000.




Fig. 15. Dependence of RTT on the length of IP path for the London, UK *skitter* monitor, August 18, 2000.

III. GEOGRAPHY AND ROUND TRIP TIME

A. Geographic patterns in the Internet

Clustering of RTTs in Figures 9 - 13 is determined by the density of the human population in different regions of the globe and by its level of technological advancement. Thus, we should not study RTTs without considering geographical layout of the network. The distance between points of communication is a major component in the performance of a network as measured by round-trip times. Several measures of distance are applicable in trying to understand the relationship between RTT and geographical location of destinations:

- great circle (the shortest distance between two points on a sphere)
- difference in longitude coordinates \times
- difference in longitude coordinates + difference in lati-




Fig. 16. Median RTT vs great circle distance for the San Diego, US *skitter* monitor, August 18, 2000.

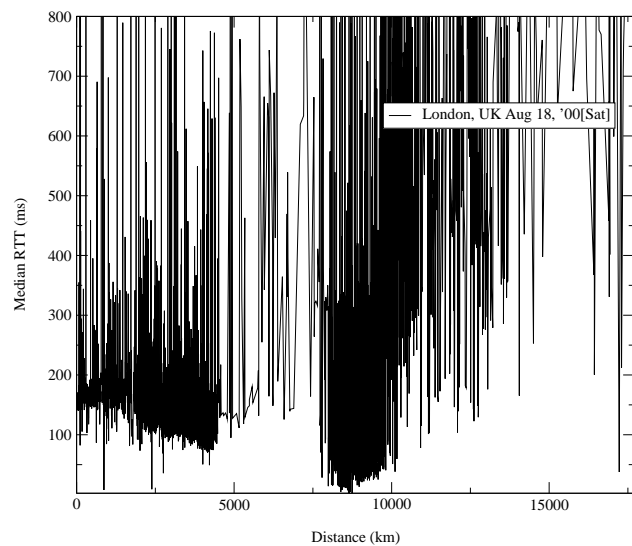


Fig. 17. Median RTT vs great circle distance for the London, GB *skitter* monitor, August 18, 2000.

tude coordinates X+Y

- distance via the US

We use CAIMIS's **This is now ipmapper get bibi from david when he has time** *IPMapper*[?] tool to determine the geographical coordinates of our IP addresses. Finding the geographic location of a host from its IP address is currently non-trivial and imprecise. Some host names legitimately indicate geographic location, but it is not a reliable or universal method. (Note that the DNS system supports the LOC record as a mechanism for a site to register geographic location information for their IP addresses but unfortunately few organizations use this functionality.) Inaccuracy in our geographical mapping of IP addresses is obvious in Figures 16 - 20. Some RTT values appear in geographic locations that would require packet transmission faster than the speed of light. We presume these points do not represent violations of causality, but rather result from inaccurate determination of the geographical coordinates of the corresponding IP addresses. Those hosts are in reality simply closer to the *skitter* host than our geographical database placed them.

Figures 16 - 18 show the median RTT as a function of geographic distance from monitors in San Diego, London, and Tokyo, using the great circle measure of distance. Large scattering of points in those figures suggests that the geographical distance between hosts is only one out of many factors influencing the latency of a connection. Other factors depending on the state of the network along the path may significantly increase the RTT between two hosts separated by a given distance. Figure 18 is particularly interesting as it shows that the actual distance between the monitor and many of the destinations it probes does not correspond to the sum of per-hop distances traveled by packets. Many paths to European or other destinations go from Japan to the United States first. Thus this graph suggests paths that are significantly longer than the great circle distance between the source and destination. Such




Fig. 18. Median RTT vs great circle distance for the Tokyo, JP *skitter* monitor, August 18, 2000.

data reflects a market (financial) reality of international transit, and limits the integrity of straight great-circle distance as a useful predictive metric of RTT. A metric such as 'from the source to the US, and then from the US to the destination' could provide better predictive power.

It is known that the bulk of global telecommunications infrastructure (e.g. transoceanic and transcontinental links) is deployed east-west, not north-south. Therefore, the difference in longitude between source and destination appears to be a useful distance metric. A visualization that plots median RTTs by longitude (Figure 19 - 20) conveniently clumps data by continents. The longitude of the *skitter* host is indicated with a vertical line. Different colors (grey scales) show major clusters of destinations in North America (from 50 to 150 degrees West), Europe (20 degrees West to 30 degrees East), Asia (50 to 140 degrees East) and Oceania (Australia at 100 to 170 degrees East,

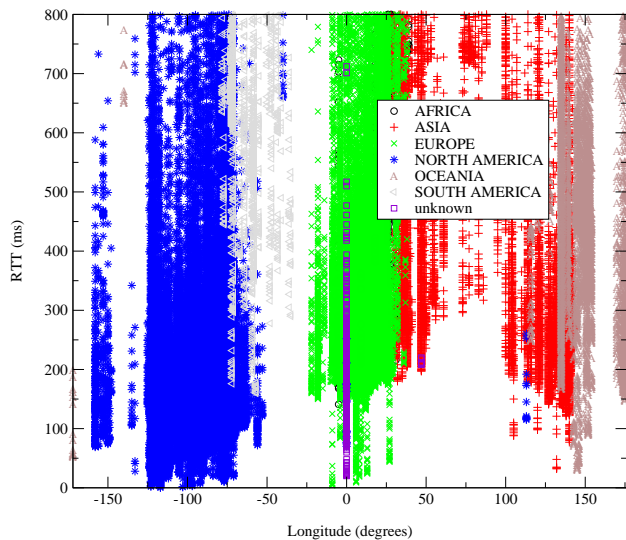


Fig. 19. Dependence of RTT on the longitude of destinations for the San Diego, US *skitter* monitor, August 18, 2000.

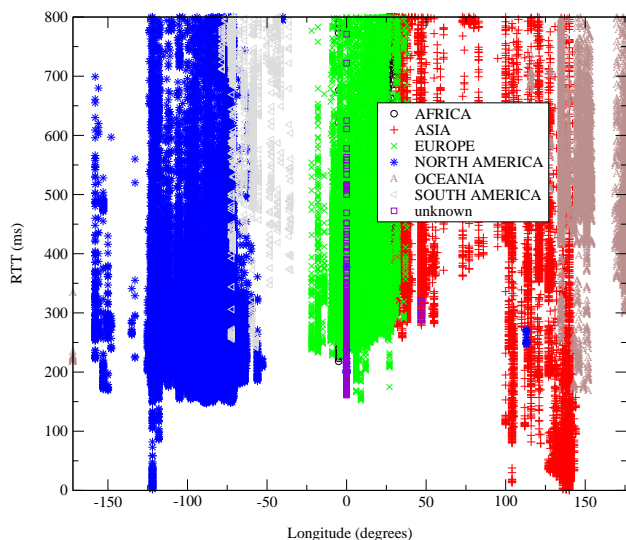


Fig. 20. Dependence of RTT on the longitude of destinations for the Tokyo, JP *skitter* monitor, August 18, 2000.

and New Zealand at 180 degrees East). Note that data for South America and Africa are shifted slightly incongruently with the rest of the world, which derives from the fact that connectivity to those regions travels north-south rather than east-west. The points in Africa are hardly visible due to the much greater density of European destinations at the same longitude.

Figure 20 shows that paths from the Tokyo monitor to North American, European and even some Asian destinations go through a node on the east coast of the US first. RTTs have a local minimum in the vicinity of this node located at about 90 degrees West. This node then becomes a secondary source of packets, with RTTs increasing both eastward and westward from it.

Visualizing RTT versus longitude can also reveal specific topological changes in connectivity between the *skitter* host and a certain subset of destinations. For example, between

30 and 31 August 1999, connectivity from the *skitter* source host in Korea to Australian destinations had drastically improved: the minimum RTT decreased from 600 ms to 300 ms. At the same time, the distributions of RTTs to other geographical groups remained practically the same. (Graph not shown.) RTT-vs-longitude graphs are an ideal tool for discovering such macroscopic trends or events.

B. Quality of service by geographic regions

Two parameters are important for characterizing the quality of Internet performance to a certain group of destinations: the speed of connection and of path stability. We can assess both parameters from *skitter*-measured RTT distributions by considering cumulative RTT distributions in various geographical destination groups (Figure 21 - 22). Cumulative RTT distributions show the percentage of destinations to which the median RTT is less than a given value. The curves are colored by continents. Areas with steep slopes reflect large clusters of destinations that are all reachable within a narrow interval of time. (They correspond to prominent peaks in the RTT distributions presented in Figures 9 - 13.) The steeper the curve, the more stable is the performance to this geographical group of destinations. The further left is that steep slope segment, the lower RTTs are, meaning faster performance.

Figure 21 shows a typical pattern of RTT distributions observed by the San Diego *skitter* monitor for various continents on 14 May 2000. Around 90% of all North American destinations are reachable in 50 to 170 ms, while most European destinations are reachable in the 150-300 ms range. Only two thirds of Asian destinations are reachable within a well-defined interval between 120 and 300 ms; the remaining one-third are highly variable. Continents with flatter curves (Oceania, South America) are characterized by poorer and less consistent performance in general. Note that in Figure 22 all curves are much less steep than in Figure 21. On that day, May 18, 2000, a problem in the connectivity of the San Diego monitor occurred, and RTTs to all destinations became considerably higher and more variant. Only 80% of all destinations were reached in less than 700 ms, while on a normal day this percentage is 95%.

Cumulative RTT distributions can also be used to detect groups of destinations that warrant further performance evaluation. For example, one can identify groups of destinations that consistently have exceptionally high RTTs. Such data may suggest areas in need of an infrastructure upgrade or possibly subject to misconfigured routing.

IV. ADMINISTRATIVE AND GEOPOLITICAL PATH TOPOLOGY

A. Dispersion by autonomous system

A common question of interest is how many and which ISPs/Autonomous Systems carry most Internet traffic or control the greatest amount of IP connectivity? We explore this question by visualizing the *AS dispersion* of paths observed by a *skitter* source. Each path contains IP addresses of the intermediate nodes between the source and




Fig. 21. Cumulative RTT distributions for the San Diego, US *skitter* monitor, May 14, 2000.

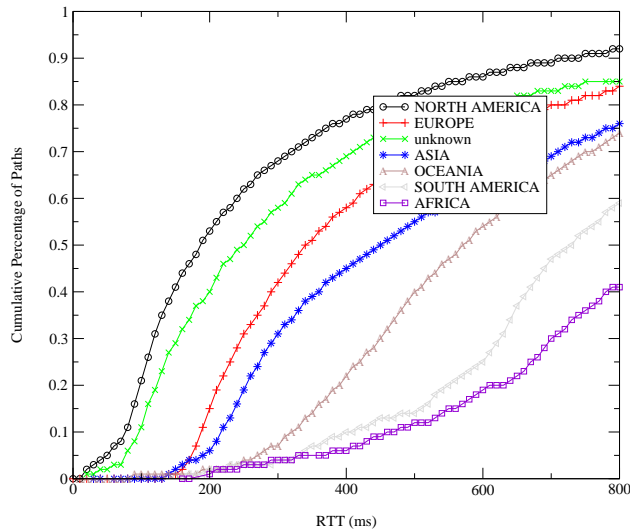


Fig. 22. Cumulative RTT distributions for the San Diego, US *skitter* monitor, May 18, 2000.

the destination. We use a routing table database to abstract these IP addresses into AS (Autonomous System) numbers, which approximately map to ISPs. In order to convert IP paths into these ‘forward’ AS paths, we use BGP (Border Gateway Protocol) routing tables collected by University of Oregon’s Route Views [10]

Figure 23 is an example of an *AS dispersion* graph for the San Diego *skitter* monitor. Complete traces to 21,574 different destinations observed for a 24-hour period on April 27th, 2000 are included. The x-axis represents the IP hop number along the path. The color (grey scale) and numeric label in the vertical bars at each hop identify the AS responsible for the IP address at this hop. The height of the bar represents the proportion of paths that passed through a particular AS at a given hop. Areas are gray when the set of paths disperse into too many distinct ASes to delineate clearly in the plot. The data is sorted from the bottom by

proportion of paths traveling through each AS. Black bars indicates paths that have ended in under 24 IP hops.

Figure 24 shows the same graph four months later. Notice the increase in the number of paths transferred to AS 7018 at the 10th hop. Also notice the more diverse peering at the 13th hop in the later plot. In both plots, SDSC’s routing policy did not change - most packets use the CERFnet (AS 1740) link, with the rest going via the vBNS (AS 145) or CalREN (AS 11422).

The AS dispersion graphs for the *skitter* monitors at root server locations can be used to answer several important questions about placement of the root servers:

- is the server near the edge of its network?
- does the location have rich peering?
- does the location have diverse upstream transit?

Figures 25 - 28 show AS dispersion graphs for four DNS root *skitter* monitors observed on October 9, 2000. Each server was using the DNS Clients list of 58,312 destinations. Two of our root server monitors, F-Root and L-Root (Figures 26 and 28, respectively), are three hops away from a major exchange point (MAE-West). In both cases splintering into a variety of distinct ASes at the 4th hop is a sign of extensive peering at this hop. The A-root monitor (Figure 25) appears to be even closer to an exchange point. Note how the dispersion patterns for the four root server monitors differ from that for the San Diego monitor. The latter exhibits a limited upstream transit: most paths travel through the same ASes and significant fanout does not occur until hop 11.

B. Country dispersion graph

We visualize an even higher level of path abstraction by mapping IP addresses to countries (instead of autonomous systems) and considering *country dispersion* of paths. Here again we use CAIDA’s *NetGeo*[?] to obtain geographical information.

Figures 29 and 30 show country dispersion graphs for two *skitter* monitors. It is not surprising that most paths in the Tokyo graph cross the US at some point, since North American networks still play a major role in providing Internet connectivity to the rest of the world.

CONCLUSION

We have presented several preliminary examples of visualization techniques that we are investigating in the analysis of *skitter* data. They have provided insight into correlations among metrics and diversity of infrastructure, and pointed us in several fascinating directions. At the same time, they have generated more questions than they have answered. We are also looking into more advanced visualization techniques, in particular, exploring three-dimensional visualization of topology and RTT performance across segments of infrastructure.

Managing dynamically changing data that is geographically and logically diverse is a challenge. Mapping many hundreds of thousands of IP addresses (nodes) to even approximate geographic location information, much less precise latitude/longitude coordinates, is a non-trivial task,

often requiring knowledge of company-specific heuristics or common data formats. The accuracy of this process is not as high as we need.

While the benefits of the described project are easy to understand, the methodology for accomplishing this analysis is hardly straightforward. First, it is critical to ensure that the measurements themselves do not impact the operation of the networks being measured. Second, we need to determine how much data needs to be gathered, and improve methods for collecting, reducing, aggregating, and mining gigabyte and terabyte datasets. Finally, we need to refine our techniques that analyze, interpret, and correlate these and other data sets in order to better visualize events, anomalies, and trends.

REFERENCES

- [1] Govindan R., Tangmunarunkit H.: Heuristics for Internet Map Discovery, <ftp://ftp.usc.edu/pub/csinfo/tech-reports/papers/99-717.ps.Z>
- [2] Govindan R., Reddy A.: An analysis of Internet Inter-Domain Topology and Route, Stability. In *Proc. IEEE INFOCOM '97*, Kobe, Japan, Apr 1997.
- [3] Braun H.-W., Claffy K.: Global ISP Interconnectivity by AS number, <http://moat.nlanr.net/AS>
- [4] Cheswick, W., Burch, H.: Internet Mapping Project, <http://cm.bell-labs.com/who/ches/map/index.html>
- [5] Paxson V.: End-to-End Routing Behavior in the Internet, *IEEE/ACM Transactions on Networking*, 5(5), 601-615, 1997.
- [6] Matthews W., Cottrell L.: The PingER Project: Active Internet performance Monitoring for the HENP Community, submitted to *IEEE Communications Magazine on Network Traffic Measurements and Experiments*.
- [7] Kalidindi S., Zekauskas M.: Surveyor: an infrastructure for Internet performance measurements, <http://telesto.advanced.org/kalidindi/papers/INET/inet99.html>
- [8] McGregor A.J., Braun H.-W., Brown J.A.: The NLANR Network Analysis Infrastructure, <http://moat.nlanr.net/Papers/TonyM-IEEE-comms.ps>
- [9] Moore D., Periakaruppan R., Donohoe J.: Where in the World is netgeo.caida.org?, http://www.caida.org/outreach/papers/inet_netgeo
- [10] Meyer D.: University of Oregon route Views project, <http://www.antc.uoregon.edu/route-views/>
- [11] D. Wessels: <http://squid.nlanr.net/>
- [12] <http://ircache.nlanr.net/>
- [13] D. Moore, R. Periakaruppan, J. Donohoe, and K. Claffy: Where in the World is netgeo.caida.org?, http://www.caida.org/outreach/papers/inet_netgeo/
- [14] IPMapper. Geographic Location Service. Caimis Geo, Inc. <http://www.ipmapper.com>




Fig. 23. AS dispersion for the San Diego, US *skitter* monitor, April 27, 2000.




Fig. 24. AS dispersion for the San Diego, US *skitter* monitor, August 27, 2000.




Fig. 25. AS dispersion for the Herndon, US (A-root) *skitter* monitor, October 9, 2000.

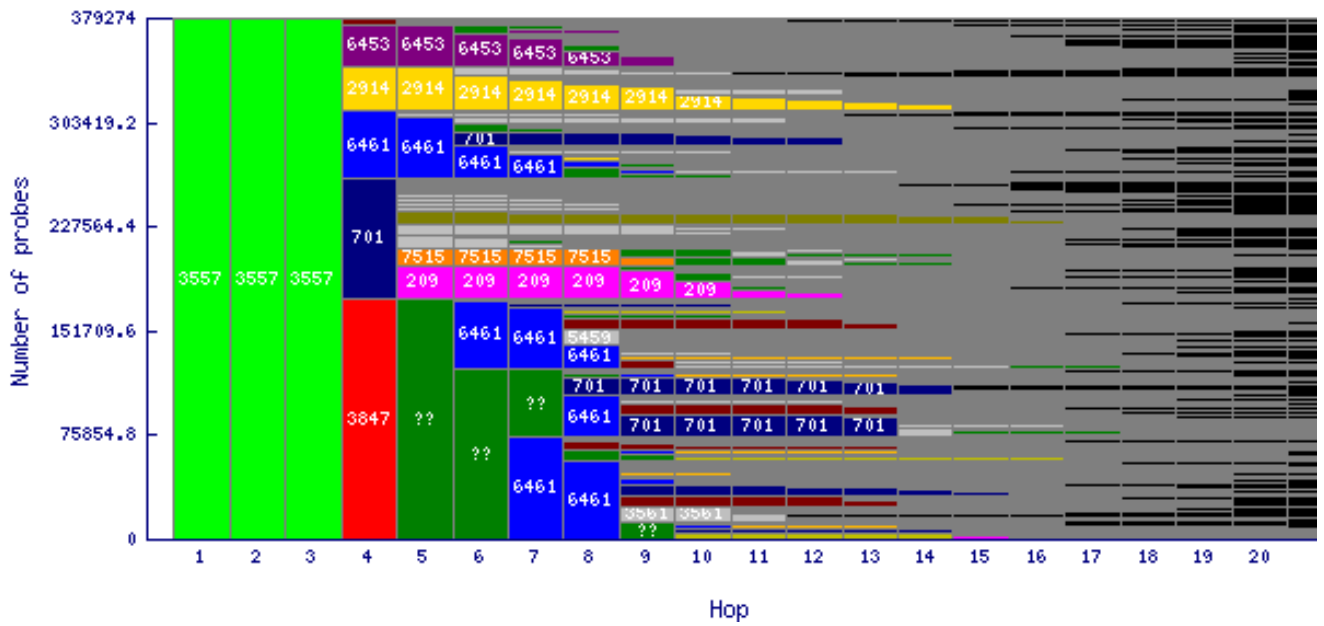


Fig. 26. AS dispersion for the Palo Alto, US (F-root) *skitter* monitor, October 9, 2000.




Fig. 27. AS dispersion for the London, UK (K-root) *skitter* monitor, October 9, 2000.




Fig. 28. AS dispersion for the Marina del Rey, US (L-root) *skitter* monitor, October 9, 2000.




Fig. 29. Country dispersion for the Palo Alto, US *skitter* monitor, October 9, 2000.

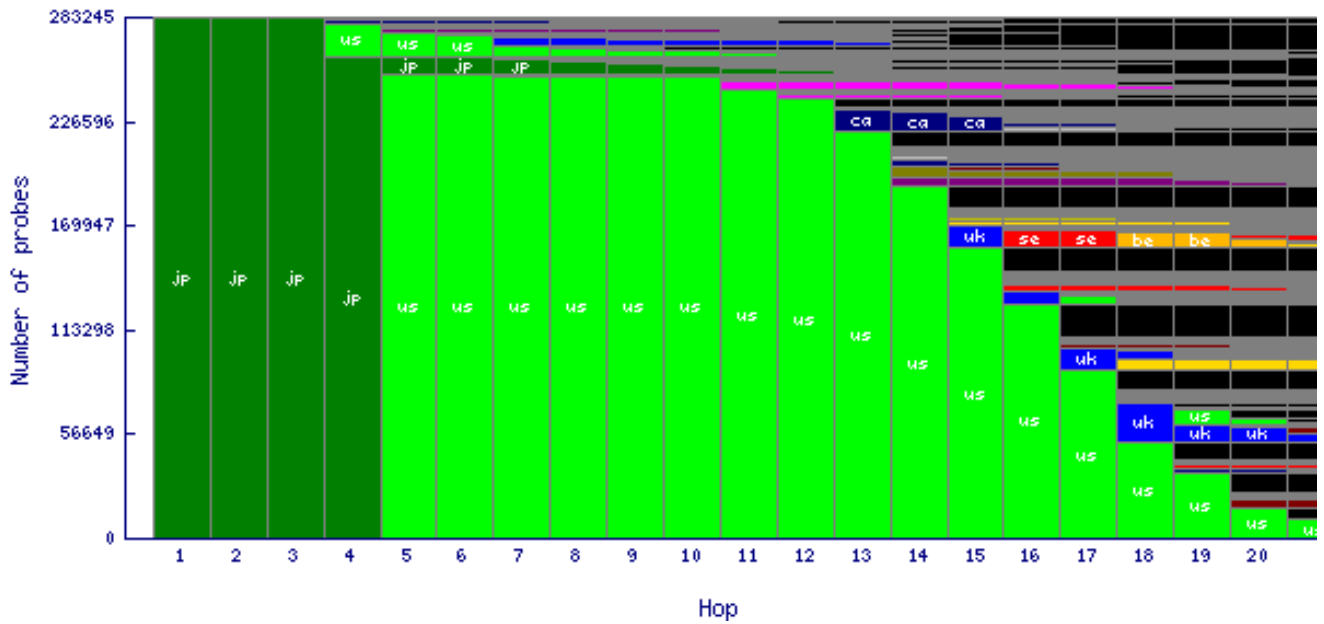


Fig. 30. Country dispersion for the Tokyo, JP *skitter* monitor, October 9, 2000.