

Geography-Based Structural Analysis of the Internet

Shiva Kasiviswanathan Stephan Eidenbenz Guanhua Yan
CCS-3, Los Alamos National Laboratory
New Mexico, USA
{eidenben,kasivisw,ghyan}@lanl.gov

Abstract—In this paper, we study some geographic aspects of the Internet. We base our analysis on a large set of geolocated IP hop-level session data (including about 300,000 backbone routers, 150 million end hosts, and 1 billion sessions) that we synthesized from a variety of different input sources such as US census data, computer usage statistics, Internet market share data, IP geolocation data sets, CAIDA’s Skitter data set for backbone connectivity, and BGP routing tables. We use this model to perform a nationwide and statewide geographic analysis of the Internet. Our main observations are:

(1) There is a dominant coast-to-coast pattern in the US Internet traffic. In fact, in many instances even if the end-devices are not near either coast, still the traffic between them takes a long detour through the coasts.

(2) More than half of the Internet paths are inflated by 100% or more compared to their corresponding geometric straight-line distance. This circuitousness makes the average ratio between the routing distance and geometric distance *big* (around 10).

(3) The weighted mean hop count is around 5, but the hop counts are very loosely correlated with the distances. The weighted mean AS count (number of ASes traversed) is around 3.

I. INTRODUCTION

Owing to its great importance, the Internet, has been a subject of a large number of studies. Much of the previous work has focused on studying topology of the Internet at the network level, without any regard to geography. In this paper, we perform a geography-based analysis of the Internet. Our main focus is on understanding the geographic properties of routing and the geographic structure of autonomous systems. Our conclusions provide new insights into the structure and functioning of the Internet.

Our results are obtained using a very high fidelity model of the US Internet infrastructure that we create by combining various datasets. Our background topology is derived primarily from the CAIDA’s Skitter dataset¹. We use the telegeography colocation database to obtain all the major point of presence locations in the US. We then simulate millions of end-devices and also billions of session-level traffic between these end-devices. The end-devices and the session traffic are generated in consultation with US census data, computer usage surveys, and market shares of various Internet service providers. For routing, we use an AS (autonomous system) path inference algorithm that uses realistic BGP tables to derive inter-domain paths. The level of authenticity captured by our model has rarely been achieved before.

¹The Skitter dataset graph is not connected, so we add few extra links based on other auxiliary datasets to make the graph connected.

It is a well known fact that the Internet routes could be highly circuitous [1], [2]. In this paper, we ask the question: How *geographic* is the Internet routing? We compute the *travel distance* between two end-points as the sum of the geometric (geographic) distance between the end-points of the various links on the path. For example, if the path from an end-device in Los Angeles to one in New York goes through San Francisco and Miami, the travel distance for this path is the sum of geometric distance from Los Angeles to San Francisco, from San Francisco to Miami, and from Miami to New York. Our experiments show that a large fraction of the traffic travels through the east and/or the west coasts of the US. Consider two end-devices A and B and the traffic flowing from A to B . Let s and t be the locations of A and B , respectively. What we observe is that for many such pairs A and B , the packets from A travels (possibly multiple times) to the east and/or the west coast before reaching B and this is true even if neither A nor B are near either coasts. We observe this phenomenon both at the national level (entire US traffic) and the state level (traffic originating from some particular state).

Looking at the ratio between travel and geometric distance, we observe more than 50% of the traffic has this ratio greater than 2 (i.e., the travel distance is at least twice the geometric distance) and about 20% of the traffic has this ratio greater than 4. One observes a similar behavior even if the traffic volume (number of bytes flowing across) is taken into account. For example, about 46% of the traffic volume our model generates are between end-devices that are less than 1000 miles apart, whereas, only 13% of the traffic volume have their travel distance less than 1000 miles.

Another related question that we investigate is the spread of the hop and AS counts and their relationship with distance. Majority of the paths have hop count less than 6, and we found that the average hop count is near 5. The AS count (the number of ASes passed on the way) is almost always less than 3 and for most of the traffic it is around 2. Also, a bit surprising is the fact that the hop count is very loosely correlated with the geometric distance. For example, it is almost equally likely two end-devices that are 500 miles or 2000 miles apart will have a hop count of 5. A similar lack of correlation also holds between the hop count and travel distance.

II. RELATED WORK

Over the past decade, there have been numerous efforts on analyzing the structural properties of the Internet topology. Much of the work has focused on studying topology at the

network level. We refer the reader to a recent survey of Willinger *et al.* [3] for more details on network topology generation schemes.

Much of the work on Internet routing has mainly focused on measuring properties like end-to-end performance, routing convergence, etc., or on modifying certain aspects of routing to get an improved performance. Our main focus is on understanding geographic properties of Internet routing. It is well known that the Internet route can be highly circuitous. This was first suggested by Tangmunarunkit *et al.* [1], who used a simplified routing model to show that the routing policies significantly increases the shortest hop distance. The paper by Tangmunarunkit *et al.* considered just the network path taken by the routes and ignored the geographic information. Subramanian *et al.* [2] were the first to study geographic properties of Internet routing. They used the GeoTrack [4] tool to determine the geographic path of the routes. They suggested that the circuitousness of Internet paths depends on the geographic and network locations of the end-host, and tends to be greater when paths traverse multiple ISP. Their dataset, however is quite small (it had only about 84,000 end-to-end paths). Spring *et al.* [5] documented some root causes of this circuitousness. Lakhina *et al.* [6] studied a wide range of geographic properties of the Internet, focusing on routers, links, and autonomous systems. Yook *et al.* [7] studied the fractal dimension of routers, ASes, and population density. They argued that the fractal dimension of all these parameters is around $3/2$.

We undertake the first large-scale study of the redundancy in Internet routing. A lot of models have also been proposed to characterize the routing and traffic in the Internet [8], [9], [10], [11]. Instead of relying on inter-domain routing models, we use an AS path inference algorithm to derive the actual inter-domain paths that are used in the Internet. By combining many real-life datasets we generate synthetic end-to-end sessions for the entire US population. The traffic we generate statically follows the traffic distribution observed in the US.

III. METHODOLOGY AND MODELING

In this section, we describe the various aspects of our modeling setup. As mentioned earlier, we use many different datasets such as the US census data, the US computer usage statistics, and the Internet market shares of various service providers to construct a large-scale realistic model of the US Internet infrastructure. The Internet model that we use in this paper was introduced by Yan *et al.* [12], and we refer the reader to that paper for a complete description of the model. In this paper, we summarize only the relevant features of this model. The motivation behind the Yan *et al.* paper was to perform a criticality analysis and assessment of the US Internet infrastructure. Yan *et al.* used a variety of different network analysis tools to identify critical Internet infrastructure facilities with the US. The various geometric and geographic analysis that we do in this paper are of different flavor from the analysis done by Yan *et al.* [12].

Backbone Topology. The Internet backbone consists of routers and links that are owned and operated by the major Internet service providers. We extract 18,000 backbone equipment locations housing approximately 291,000 unique IP (Internet Protocol) addresses in the US from the Skitter dataset collected by the CAIDA project (available at <http://www.caida.org/tools/measurement/skitter/>). Since the original Skitter dataset is not connected, we add two types of *virtual* links to make the resulting graph connected.

Each IP address corresponds to a network interface at a backbone router and multiple IP addresses can belong to the same physical backbone router. This is the well known *IP alias-resolution* problem [13]. To overcome this problem, we use the alias clustering data provided by the iPlane project (available at http://iplane.cs.washington.edu/data/alias_lists.txt). Using the iPlane data, for any two IP addresses in the Skitter dataset that belong to the same physical router, we create a virtual link between them. These are the first type of virtual links that we add to the Skitter graph.

If two IP addresses belong to the same autonomous system (AS) and are located at the same place, it is unlikely that traffic between them traverses through a different location. The geographical position of each backbone IP address, in the form of its longitude and latitude, is derived from the ip2location dataset (available at <http://www.ip2location.com>). Geolocation is generally considered to work fairly well to a city-level resolution, which is our main concern for analysis. Street-address level accuracy is much harder to achieve. We use a star structure to connect co-located IP addresses owned by the same AS (the star structure is used to prevent addition of too many links). We call the center of such a star topology a *hub IP*. These are the second type of virtual links that we add to the Skitter graph.

With these virtual links, we are able to produce a connected Internet backbone that covers more than 99.7% of the IP addresses in the Skitter dataset.

Internet Point of Presence. An Internet PoP (point of presence) is an access point to the Internet backbone, which is typically owned by an ISP (Internet Service Provider), or located in an Internet exchange point. We obtained a set of 543 PoPs from the telegeography colocation database (available at: <http://www.telegeography.com/>), which lists the operators present in each PoP. We then use the Skitter dataset to populate the PoPs with backbone IP addresses. For this we use the following heuristic: if the latitude and longitude of a backbone IP address agrees with that of a PoP, we assign it to that PoP. This simple assignment scheme, however, leads to inconsistency: for an AS \mathcal{A} present in PoP \mathcal{P} according to the telegeography colocation database, it may not have any of its backbone IPs assigned to that PoP. To circumvent this problem, we create a virtual backbone IP address that belongs to the AS \mathcal{A} in PoP \mathcal{P} ; moreover, if there exists a geolocated backbone IP address that belongs to AS \mathcal{A} within 15 miles from PoP \mathcal{P} , we connect the virtual IP address to the geolocated IP address. In total, this process created only 1247 virtual backbone IP

addresses. These virtual backbone IP addresses form only a small fraction (about 0.4%) of all backbone IP addresses.

Virtual Point of Presence. As the Skitter dataset is generated from traceroute output, inter-AS links derived from it are incomplete and biased [14]. To mitigate this problem, we assume that all the ASes present in the same PoP are internally connected. We create a virtual PoP IP inside each PoP and connect it to every hub IP in that PoP. Note that this process may introduce artificial inter-AS links that may not actually exist. For example, if AS \mathcal{A} and AS \mathcal{B} do not have any business relationships, there may not be physical AS links between \mathcal{A} and \mathcal{B} , but the above process might add links between \mathcal{A} and \mathcal{B} . The routing scheme in our model (explained below), however, uses inferred AS-level relationships obtained from realistic BGP data to compute AS-level paths and these artificial links will then not be used for routing.

End Devices. We generate synthetic end devices, including both residential and business computers, and then connect them to the Internet backbone topology. We distinguish residential and business computers in our model. In total, we generated 73,884,296 residential computers and 58,923,964 business computers in the US. These numbers are a rough approximation to the number of computers in the US. To generate residential computers, we use the latest US census bureau data (like income distribution, percentage of residential computer usage for each annual family income category, etc). This gives us a census-block level population in each 250×250 square meter grid in the US for the entire 24 hour duration [15]. To generate business computers, we use the Dun & Bradstreet (D&B) dataset, which provides information about all companies in the US, including their headquarter locations, numbers of employees, and SIC (Standard Industrial Classification) codes. A SIC code has four digits and indicates the business type of a company. The US census data presented in [16] gives us computer penetration ratios in different business categories.

Access routers. Internet access routers are used to connect end devices to the Internet backbone. We consider three types of Internet access services, dial-up, DSL, and Cable, as they are three mostly widely used Internet access methods in the US. We use the Home Broadband Adoption 2006 report by Pew Internet & American Life Project to randomly assign the Internet access type of each end device.

For the dial-up service, we collect a list of aggregators for each zip code from the Internet Service Provider directory (available at <http://www.findanisp.com>) and for each of these aggregators we create an Internet access router. For DSL and cable services, we use the subscriber numbers of the top companies. For both these services, the top nine companies collectively cover more than 50% of the market, see, e.g., <http://www.leichtmanresearch.com/press/081108release.html>. For each of these companies, we collect a list of zip codes where the company provides DSL or Cable services and use that information to add an Internet access router for each zip code within its service coverage.

If the chosen access service is dial-up, we randomly assign the end device to an aggregator for the zip code where the device is located. If the access type is DSL or Cable, we randomly choose an Internet broadband access router based on the market shares of the top broadband companies. After an Internet access router is chosen, we create a link between it and the end device.

Recall that there are 543 PoPs in the backbone topology and each of them has a list of backbone IPs. Also, each PoP IP is associated with an AS number. Given an Internet access router, we use the following heuristic to decide which PoP IP it connects to. First, we sort all PoPs according to their distances from the Internet access router. Then, starting from the closest PoP, we check whether it has a PoP IP that peers with the ISP company owning that Internet access router. This can be done by checking whether the AS number of the PoP IP connects to any one of the AS numbers owned by the ISP company in the AS-level graph. If we cannot find one, we try the second closest PoP. This process repeats until one such PoP IP is found. Thereafter, we create a link between it and the Internet access router.

Sessions. We generate synthetic sessions, including HTTP, email, P2P, and streaming traffic, for every computer for a period of 24 hours. We generated a total of 1.14 billion sessions. For each Internet session, we assign it a session type and choose its origin and destination. For an email or P2P session, we assume that the end-points of the session are end-devices residing either at a home or a business location. For HTTP or streaming traffic, we assume that the source of the session is an end-device whereas the destination of the session is a server. To pick an end device as either the source or destination, we pick a device from a (US) state based on the percentage of devices in that state. When the end-device is a server (for HTTP and streaming sessions), we pick a server from one of the top 100 servers that are most visited, based on the proportion of web access hits they receive (information available at <http://www.alexa.com/>). Many web servers are located in the technological centers of Silicon Valley and Washington D.C., as well as a few smaller centers mostly in metropolitan areas. We ignore the effects of content distribution networks (as they are difficult to model). Figure 1 gives a more visual representation of the sessions that our simulation generates. For example, Figure 1(b) shows that more than 20% of traffic originating in the mid-west states of North Dakota, South Dakota, Kansas, Oklahoma, and Texas travel goes to destinations less than 1000 miles away, which roughly makes it to either coast. Similarly, from Figure 1(d) we see that about 40-50% of the traffic originating from California goes a distance greater than 2000 miles (so to east coast).

In our simulations, 40% of the sessions are HTTP (with a split of 25.14% and 4.86% among home and business computers), 30% of the sessions are email (with a split of 11.87% and 18.13% among home and business computers), 20% of the sessions are P2P (with a split of 18.71% and 1.29% among home and business computers), and remaining 10% of

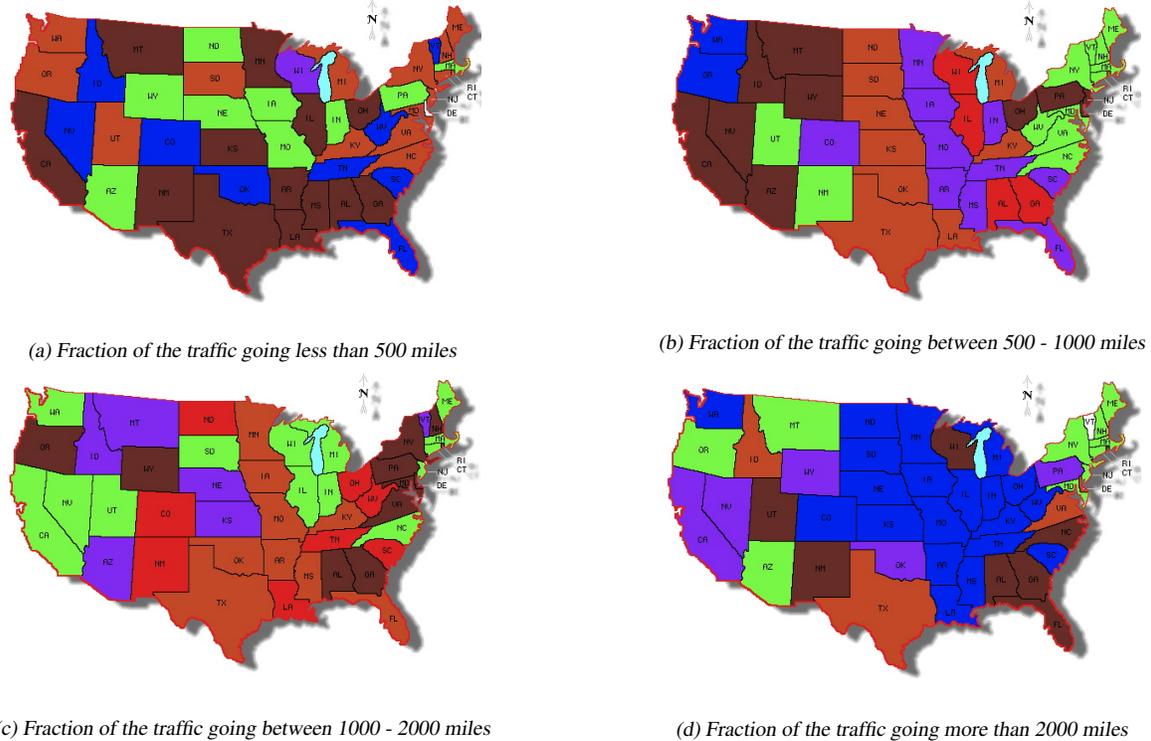


Fig. 1: Traffic distribution generated by our model. Figure (a): Fraction of the traffic going to a geometric distance less than 500 miles. Figure (b): Fraction of the traffic going to a geometric distance between 500 and 1000 miles. Figure (c): Fraction of the traffic going to a geometric distance between 1000 and 2000 miles. Figure (d): Fraction of the traffic going to a geometric distance greater than 2000 miles. Each state is colored by the fraction of the traffic of that particular type originating from that state. Blue means between 0-10%, Brown means between 10-20%, Green means between 20-30%, Orange means between 30-40%, Purple means between 40-50%, and Red means between 50-60%. For example, if in Plot (a) a state is colored blue then between 0-10% of total traffic originating from that state goes less than 500 miles.

the sessions are streaming (with a split of 6.29% and 3.71% among home and business computers). The above traffic mix was generated in consultation with the information available from [17], [18], [19], [20].

Previous works have suggested that approximately 80% of the web document transfers are less than 100 kilobytes in size [21], but this distribution has a heavy tail [21], [22]. We choose 25 kilobytes as the average size of the HTTP sessions. This number is computed by downloading a number of webpages and finding the average size of these downloaded web pages. We choose the average size of the email sessions as 100 kilobytes. This number is based on the average size of all emails in the inbox of various employees at a large institution. The average streaming rate of streaming sessions is 200 kilobytes per sec [23] and the average duration of streaming session is approximately 125 sec [24]. This gives an average size of approximately $250 \times 125 \approx 30$ megabytes for streaming sessions. The average size of a P2P session is computed by observing the history of already completed transfers in a P2P client [25]. Now for each session its size (based on whether it is HTTP, email, P2P, or streaming) is drawn from an exponential distribution with mean (average size) as given above.

Routing. Internet routing is strictly hierarchical: inter-domain

routing protocols (e.g., border gateway protocol) regulate Internet traffic among different ASes, and intra-domain routing protocols (e.g., Open Shortest Path First and Routing Information Protocol) specify how traffic is routed within the same AS. Due to complexity of BGP and the fact that commercial relationships between ASes are generally unavailable to the public, we use AS-level paths inferred from existing BGP routing tables for inter-domain routing. We use the AS path inference algorithm from [26], which is able to infer AS-level paths with 95% accuracy. For the intra-domain routing, we simply use the shortest path algorithm. For some destination PoP IPs, the algorithm fails to infer AS-level paths to them. In such circumstances, we derive their AS numbers and use www.xedorbit.com to obtain a list of prefixes for each of them. We then use the AS path inference algorithm from [26] to infer AS-level paths to these prefixes. These derived AS-level paths are further used to compute the IP-level paths to these destination PoP IPs. To compute the route between any two PoP IPs we use an algorithm similar to [27] which has been shown to achieve more than 78% accuracy. A more thorough description of the routing scheme is given in [12].

IV. INTERNET ROUTING ANALYSIS

We analyze the paths generated by our experiments. For a session between a source at location s and a destination

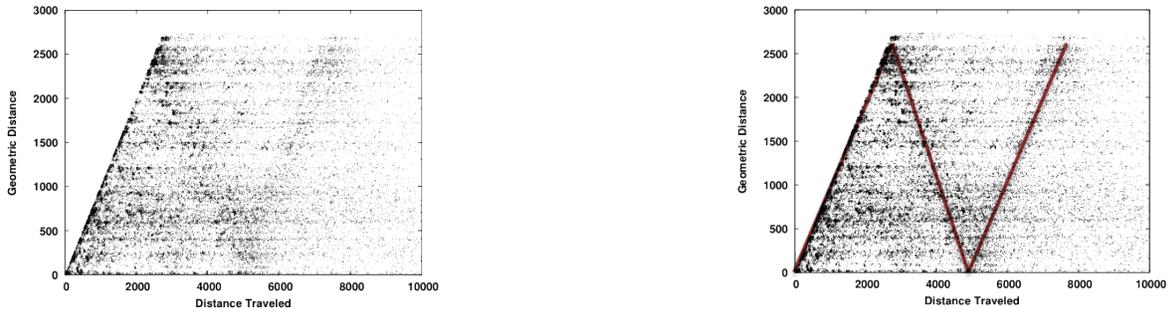


Fig. 2: Each point represents an Internet session. The X-axis represents the route distance between the source and the destination of the session. The Y-axis represents the geometric distance between the source and the destination of the session. Only a uniform 1/10000th fraction of all sessions we generated are represented in this plot. The plot on the right is same as that on the left except that lines $y = x$, $y = -x + 5100$, and $y = x - 5100$ are drawn for visual aid.

at location t , we use the *Haversine* formula to compute the *geometric distance* between s and t . The Haversine formula takes as input the latitude and longitude of the end-points. Let (lat_s, lon_s) and (lat_t, lon_t) be the latitude and longitude of the locations s and t . The Haversine distance d between s and t equals, $d = R \times c$. Here, $R = 3961$ miles is the radius of the earth and $c = 2 \times \arctan2(\sqrt{a}, \sqrt{1-a})$ where $a = \sin^2((lat_t - lat_s)/2) + \cos(lat_s) \times \cos(lat_t) \times \sin^2((lon_t - lon_s)/2)$.

A. Nationwide Analysis

The *route (travel) distance* between s and t is computed by summing up the lengths of the link on which the packets travel from s to t in our simulation. Figure 2 plots the geometric distance against the route distance. Each point here is a session in our simulation. Notice that since the route distance is always greater than the geometric distance, there are no points above the $y = x$ line. In the following, we refer to geometric and route distance of a session to mean the geometric and route distance between the end points of the session.

In Figure 2, we notice that for many sessions the route distance is far greater than their corresponding geometric distance. In order to validate the feasibility of our synthesized data, we geolocated a few trace route exercises. One such example is shown in Figure 3. It is very easy to find such long paths (that go from coast to coast multiple times), in fact it is more the rule than the exception and we encourage the reader to try this little experiment at home. The fact that there exists sessions whose route distance is significantly more than their corresponding geometric distance may not be all that surprising given that there are many economic and engineering aspects that drive the Internet routing and this observation has been made before (see, e.g., [1], [2]). But what is a bit surprising is the number of such sessions (we elaborate this point in Section IV-C).

Another surprising feature from Figure 2 is the heavy concentration of points near the lines $y = x$, $y = -x + 5100$, and $y = x - 5100$ (in form of a triangular strip). Note that, $5100 = 2 \times 2550$ is approximately the round-trip distance between the east coast and the west coast (2550 is approximately the average distance between the east and the

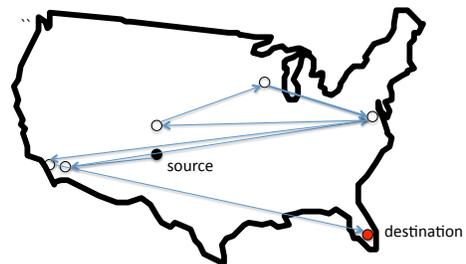


Fig. 3: A real traceroute path. The traceroute was sent from a computer in Santa Fe, New Mexico to Palm Beach, Florida. Notice the long route going through both coasts.

west coast of the US). We now try to explain why there is this concentration near these lines. We look at each of these three lines separately.

- (1) **Line $y = x$:** The points close to the line $y = x$ represent the sessions where the source and destination are at distance y apart, and the routes taken by the packets have lengths almost y (i.e., sessions where geometric distance is very close to the route distance). These points represent the best-case scenario as the routing is almost perfect.
- (2) **Line $y = -x + 5100$:** The points close to the line $y = -x + 5100 \equiv x = 5100 - y$ represent the sessions where the source and destination are at a geometric distance of y , whereas the route distance is $5100 - y = 2 \times 2550 - y$. Most of the points that lie close to this line have the property that if source is at location s and destination at location t then *roughly* either of the following happens: (a) if s is to the west of t , then the route taken by packet in going from s to t involves going s to the west coast, from the west coast to the east coast, and from the east coast to t , or (b) if s is to the east of t , then the route taken by packet in going from s to t involves going s to the east coast, from the east coast to the west coast, and from the west coast to t . To quantify the above statement, we picked all the points that lie close to this line (between the lines $y = -x + 4900$

and $y = -x + 5300$) and analyzed the paths that produce these points. We noticed that more than 90% of these points satisfied either the property (a) or (b).

(3) Line $y = x - 5100$: The points close to the line $y = x - 5100 \equiv x = 5100 + y$ represent the sessions where the source and destination are at a geometric distance of y , whereas the route distance is $5100 + y = 2 \times 2550 + y$. Most of the points that lie close to this line have the property that if source is at location s and destination at location t then roughly either of the following happens: (a) if s is to the west of t , then the route taken by packet in going from s to t involves going s to the east coast, from the east coast to the west coast, and from the west coast to t , or (b) if s is to the east of t , then the route taken by packet in going from s to t involves going s to the west coast, from the west coast to the east coast, and from the east coast to t . Again to quantify this statement, we picked all the points that lie close to this line (between the lines $y = x - 4900$ and $y = x - 5300$) and analyzed the paths that produce these points. We noticed that 98.5% of these points satisfied either the property (a) or (b).

From Figure 2 and the above discussion, we conclude that there is a very interesting coast to coast shuttling of traffic even when the source and destination are close to each other. In particular, peering agreements among ASes are typically structured such that geometric distance is not the main cost driver.

B. Per State Analysis

Figure 2 shows that at a national level there are some interesting phenomena happening. So the natural question would be to see if some similar happens when we restrict ourselves to traffic originating from some particular state in the US. To answer this question, we repeated the simulation on a state level. For each state in the US, we looked at the sessions originating from that particular state and used the route distance that we get out of this to redraw the previous plot between the geometric distance and the route distance. In the full version, we present the analysis of all the 48 states (excluding, Hawaii and Alaska which were not part of our model). For now, we concentrate on 6 states, that form a good representation of our results. We now analyze some interesting trends appearing in these six state plots.

(1) California: The plot for California (Figure 4) follows the national plot (Figure 2). We observe that there are lots of points near the lines, $y = x$, $y = -x + 5100$, and $y = x - 5100$. The reason for this behavior is same as in the national case. Also, compared to other states there are far more points in the California plot because our session generation model takes into account population, income, economy of a region, etc, which are much higher in California than other states.

(2) Illinois: The plot for Illinois (Figure 5) has two thick horizontal stripes. The first one is approximately between $y = 550$ to $y = 800$. A majority of the points that lie in

this stripe represent sessions originating from Illinois and going to the east coast. Observe that even for the same geometric distance there is a wide discrepancy in the route distance. The second stripe is approximately between $y = 1700$ and $y = 1800$. A majority of points that lie in this strip represent sessions originating from Illinois and going to the west coast.

(3) New York: The plot for New York (Figure 6) has lots of points between lines $y = x$ and $y = x + 400$. About, 80% of the points that lie between these two lines represent sessions originating from New York and taking a small detour within the east coast (like going to Washington DC) before heading to the destination. For larger values of y , this destination is the west coast. In the New York plot, one also notices a concentration of points around the lines $y = x + 2500$ and $y = x + 5000$.

(4) Florida: The plot for Florida (Figure 7) is similar to Illinois in that it has two dominant horizontal stripes. The first stripe is approximately between the lines $y = 700$ and $y = 1000$, and it is produced by the traffic originating from Florida and going to the Northeast of the US. The second stripe is approximately between $y = 2400$ and $y = 2600$, and this stripe is produced by the traffic originating from Florida and going to the west coast.

(5) Texas: The plot for Texas (Figure 8) has lots of points in and around the circles depicted in the plot. This circle is on the left is caused by the traffic originating from Texas and going either to the east or the west coast (which are roughly at the same distance). The circle on the right is again caused by the traffic originating from Texas and going either to the east or the west coast, but this time there is a detour through the other coast, i.e., the traffic to the west coast takes a detour through east coast and vice-versa.

(6) New Mexico: The plot for New Mexico (Figure 9) is quite sparse because population, income, etc, are quite low in New Mexico. There is some concentration of the points in and around the circle depicted in the plot. This is mainly due to the traffic originating from New Mexico and going to the east coast.

C. Distance Ratio Analysis

Figure 2 and the various state plots Figures 4 to 9 showed that in many cases the routing distance is far greater than the geometric distance. To better understand how far apart these distances could be, we look at various “ratio plots”. In Figure 10(a), we study the ratio of route distance to geometric distance. Let us define *stretch*² of a path as the ratio between length of the route and the geometric distance between the end-points of the path. For about 45% of the paths the stretch is between 1 and 2. An ideal stretch of exactly 1 was never achieved, but this is to be expected, since some small amount of detour compared to geometric distance will always exist. About 81% of the paths have stretch less than 4. So still a significant fraction (about 19%) of the paths suffer a large

²In [2], stretch is referred to as the *distance ratio*.

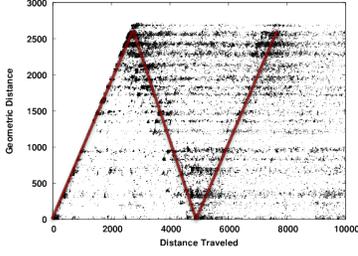


Fig. 4: Plot for the state of California.

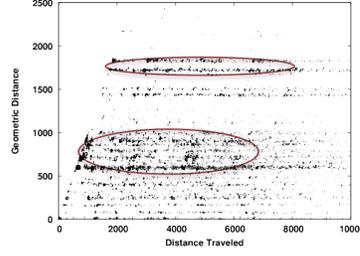


Fig. 5: Plot for the state of Illinois.

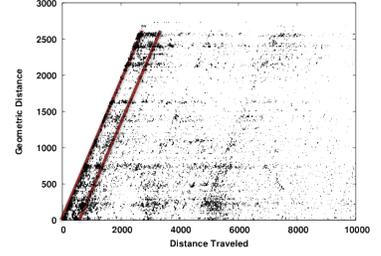


Fig. 6: Plot for the state of New York.

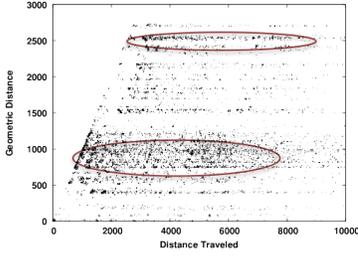


Fig. 7: Plot for the state of Florida.

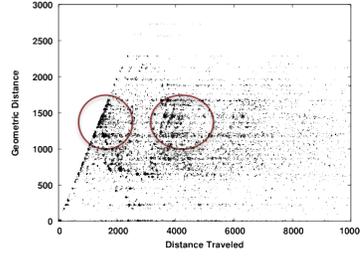


Fig. 8: Plot for the state of Texas.

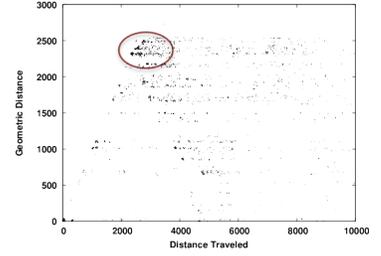


Fig. 9: Plot for the state of New Mexico.

For each state, only a uniform 1/1500th fraction of all sessions that we generated for that particular state are represented in this plot. The red markers are drawn only for visual aid.

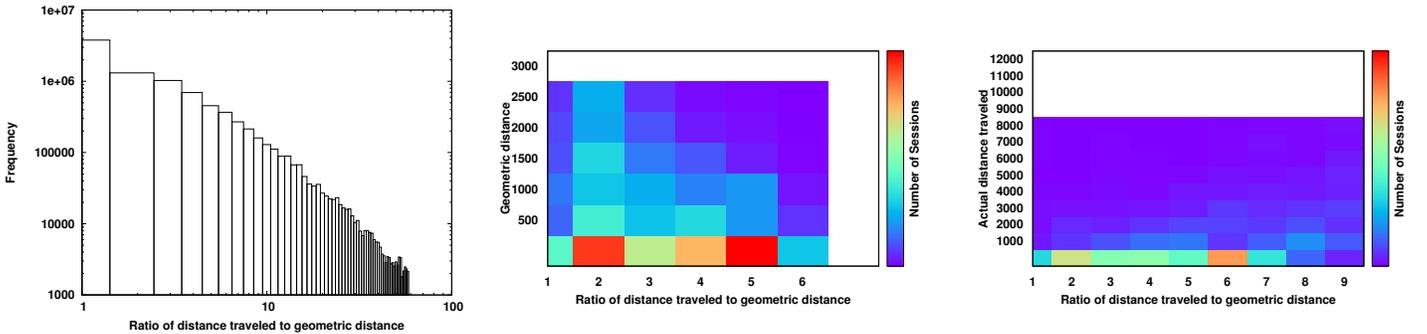


Fig. 10: (a): Histogram of the ratio of distances (stretch) generated using a uniform random sample of about 8 million sessions (in log-log scale). (b): Heat map representing stretch vs. geometric distance. The color scheme used in the heat map is shown adjacent to the plot. A reddish region has more number of points (sessions) than a bluish region. (c): Heat map representing stretch vs. travel distance.

detour from the geometric path. The average stretch (over all sessions) is 10.25.

In Figure 10 through heat maps we also show: (i) stretch vs. geometric distance, and (ii) stretch vs. travel distance. In these heat maps the number of sessions decrease gradually as we go from a red to blue region. An observation is that generally the *short* distance traffic (i.e., traffic going between source and destination which are geometrically close) have large stretch. For example, if restricted to traffic that goes less than 500 miles (in geometric distance) then the average stretch is as big as 35. The simple reason for this is that if the geometric distance is small, then even a small detour (relative to the geometric distance) will result in large stretch. Most of the *long* distance traffic (i.e., traffic going between source and destination which are geometrically far) have small stretch.

For example, if considers traffic that goes more than 2000 miles (in geometric distance) then the average stretch is only around 1.8. The reason for this being that since the geometric distance is big, even a reasonably big detour (relative to the geometric distance) will not lead to a big stretch.

D. Traffic Distribution Analysis

The previous plots were only considering distances, and were completely ignoring the volume of traffic (measured in number of bytes) that go across various source-destination pairs. As one would expect there is a lot of asymmetry in the volume of traffic among different source-destination pairs. In Figure 11(a), we compare the volume of traffic against the geometric distance. This plot just depends on our model of traffic (session) generation and is independent of the routing

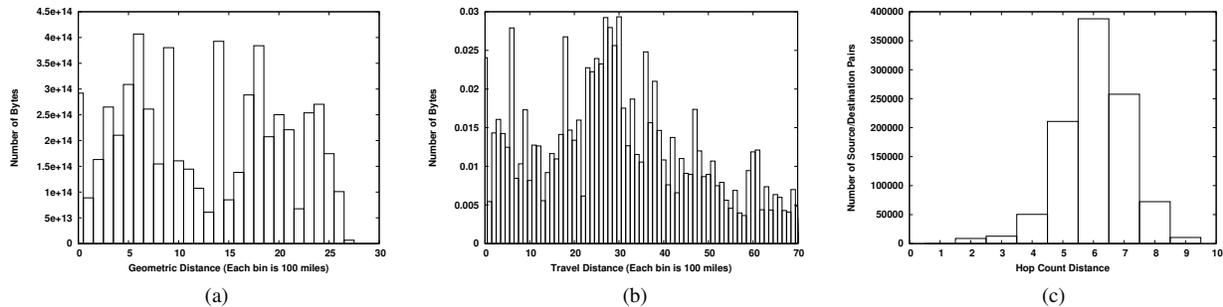


Fig. 11: (a): The distribution of the traffic according to geometric distance using a uniform sample of around 8 million sessions. The X-axis represents the geometric distance broken into bins of 100 miles. The Y-axis represents the volume of traffic (measured in number of bytes). This plot is closely related to Figure 1. (b): The distribution of the traffic according to travel distance. The X-axis represents the route distance broken into bins of 100 miles. The Y-axis represents the volume of traffic (measured in number of bytes). (c): Frequency distribution of the hop count. Sample size: around a million sessions.

strategy used. About 22% of the traffic volume our model generates goes less than 500 miles, about 46% of the traffic volume goes less than 1000 miles, and about 76% of the traffic volume goes less than 2000 miles. The farthest source-destination pair in our model was around 2650 miles apart. What is also interesting to observe is the multi-modality of this plot that arises due to distances between various metropolitan areas in the US. Due to large population density in the metropolitan areas a large fraction of sessions we generate are between these metropolitan areas.

In Figure 11(b), we plot the volume of traffic against the route distance. Because of the routes being far from geometric, only about 13% of the traffic volume has route distance less than 1000 miles, about 26% of the traffic volume has route distance less than 2000 miles, and about 76% of the traffic volume has route distance less than 5000 miles. Comparing this to the Figure 11(a), one notices that about 76% of the traffic volume has geometric distance less than 2000 miles, whereas, to get the same percentage in the route distance one has to go 5000 miles. So again, one notices the discrepancy between the properties of travel and geometric distances.

E. Hop Count and AS Count Analyses

We now analyze the hop and AS distribution of the routing paths. The hop count between a source and destination is defined as the number of hops that a packet takes in going from the source to the destination. In Figure 11(c), we plot the distribution of the hop count. We observe that a large fraction of paths (about 38.3%) have a hop count of 6. Also, about 20.8% of the paths have a hop count of 5 and 25.4% of the paths have a hop count of 6. The weighted mean hop count is 5. The plot also suggests that the hop count distribution is tightly concentrated around its mean.

In Figure 12(a), we plot the variation of the hop count against the geometric distance. The plot suggests that the geometric distance has very little effect on the hop count. For example, if we look at the paths with hop count 6, we see that there is a uniform spread of these paths independent of the distance between source and destination. That is, there

are almost equal numbers of *close* and *far* source-destination pairs with a hop count of 6. The same observation holds for other hop counts too. So we conclude that the number of hops is dependent more on the commercial relationships between ASes, and less on the geometric distance. To make this conclusion more formal, we measure the correlation coefficient³ between hop count and geometric distance. The correlation coefficient turned out to be quite small (about 0.15) suggesting that the hop count and geometric distance are almost independent. A similar conclusion was obtained by Huffaker *et al.* [28] by analyzing the CAIDA dataset for the Asia-Pacific region.

In Figure 12(b), we plot the variation of the hop count against the travel distance. As in Figure 12(a), we observe little correlation between the hop count and the travel distance (the correlation coefficient is only 0.136).

The AS count between a source (s) and a destination (t) is defined as the number⁴ of ASes crossed by a packet traveling from s to t (including, the source and destination ASes). If the entire path remains within a single AS, then the AS count is 1. In Figure 12(c), we plot the distribution of the AS count. We observe that a large fraction of paths (about 75.8%) have an AS count less than 2. About 96.2% of paths have an AS count less than 3, which can be used to conclude that almost all routing paths cross at most 3 ASes.

V. CONCLUSIONS

In this paper, we studied a variety of geographic properties of the US Internet infrastructure. To perform our study, we combined many different data sources to create a realistic model of the US Internet infrastructure. We show that a large fraction of the traffic gets routed through the coasts and in

³Given n measurements of variables F and G , $(f_1, g_1), \dots, (f_n, g_n)$, the (Pearson) sample correlation coefficient between F and G is defined as $\sum_{i=1}^n (f_i - \bar{f})(g_i - \bar{g}) / ((n-1)\sigma_f\sigma_g)$, where \bar{f} and \bar{g} are the sample mean of F and G and σ_f and σ_g are the sample standard deviation of F and G , respectively.

⁴We count the number of crossings between ASes, so if a path enters the same AS twice, we count it as two crossings.

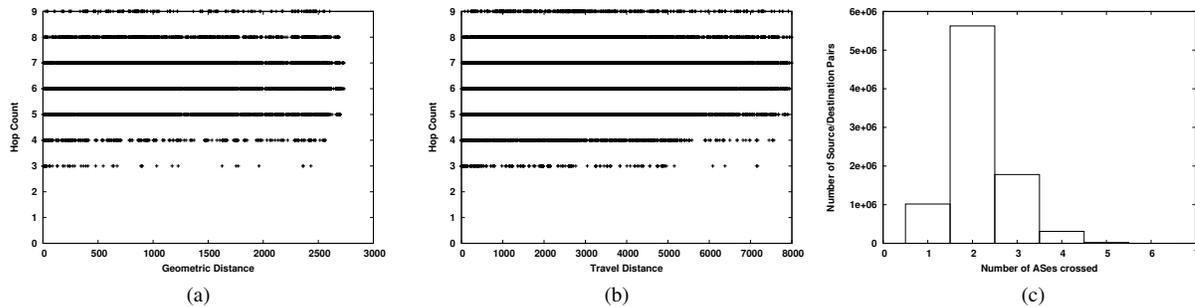


Fig. 12: (a): Variation of hop count with geometric distance. Sample size: around a million sessions. (b): Variation of hop count with travel distance. Sample size: around a million sessions. (c): Frequency distribution of the AS count. Sample size: around 8 million sessions.

many cases the traffic bounces multiple times between the two coasts before reaching the destination. We also investigated the geographic structure of the ASes, and confirmed the observation made by Lakhina *et al.* [6]. The contributions made in this paper extend our knowledge of the geographic aspects of the Internet. The results can be used for various policy decisions and for designing better generative models for the Internet.

ACKNOWLEDGEMENTS

We thank the CAIDA project for their Skitter dataset, Jian Qiu and Lixin Gao for their AS path inference algorithm. We would also like to thank Pallab Datta, Sunil Thulsidasan, and Venkatesh Ramaswmy for their help in developing the model.

REFERENCES

- [1] H. Tangmunarunkit, R. R. Govinda, S. Shenker, and D. Estrin, "The impact of routing policy on internet paths," in *IEEE INFOCOM*, vol. 2, 2001, pp. 736–742.
- [2] L. Subramanian, V. N. Padmanabhan, and R. H. Katz, "Geographic properties of internet routing," in *ATEC '02: Proceedings of the General Track of the annual conference on USENIX Annual Technical Conference*. USENIX Association, 2002, pp. 243–259.
- [3] W. Willinger, D. Alderson, and J. Doyle, "Mathematics and the internet: A source of enormous confusion and great potential," *Notices of the American Mathematical Society*, vol. 56, no. 5, pp. 586–599, 2009.
- [4] V. Padmanabhan and L. Subramanian, "An investigation of geographic mapping techniques for Internet hosts," *ACM SIGCOMM Computer Communication Review*, vol. 31, no. 4, p. 185, 2001.
- [5] N. Spring, R. Mahajan, and T. Anderson, "The causes of path inflation," in *Proceedings of ACM SIGCOMM Conference*. ACM, 2003, p. 124.
- [6] A. Lakhina, J. Byers, M. Crovella, and I. Matta, "On the Geographic Location of Internet Resources," *IEEE Journal on Selected Areas in Communications*, vol. 21, no. 6, pp. 934–948, 2003.
- [7] S. Yook, H. Jeong, and A. Barabási, "Modeling the Internet's large-scale topology," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 99, no. 21, 2002.
- [8] J. Leguay, M. Latapy, T. Friedman, and K. Salamatian, "Describing and simulating internet routes," in *In 4th International IFIP-TC6 Networking Conference*, pp. 2–6.
- [9] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the self-similar nature of ethernet traffic (extended version)," *IEEE/ACM Transactions on Networking*, vol. 2, pp. 1–15, 1994.
- [10] M. Zukerman, T. D. Neame, and R. G. Addie, "Internet traffic modeling and future technology implications," in *Proceedings of IEEE Infocom'03*.
- [11] V. Paxson and S. Floyd, "Wide-area traffic: The failure of poisson modeling," *IEEE/ACM Transactions on Networking*, vol. 3, pp. 226–244, 1995.
- [12] G. Yan, P. Datta, S. Eidenbenz, S. Thulsidasan, and V. Ramaswamy, "Criticality analysis and assessment of national internet infrastructure," *Computer Networks*, vol. 54, no. 7, pp. 1169–1182, 2010.
- [13] N. Spring, M. Dontcheva, M. Rodrig, and D. Wetherall, "How to Resolve IP Aliases," Department of Computer Science and Engineering, University of Washington, Tech. Rep. UW-CSE-TR 04-05-04, 2004.
- [14] Y. He, G. Siganos, M. Faloutsos, and S. V. Krishnamurthy, "A Systematic Framework for Unearthing the Missing Links: Measurements and Impact," in *Proceedings of 4th USENIX Symposium on Networked Systems Design & Implementation (NSDI'07)*, 2007.
- [15] T. N. McPherson and M. J. Brown, "Estimating Daytime and Night-time Population Distributions in U.S. Cities for Emergency Response Activities," in *Bulletin of the American Meteorological Society*.
- [16] "Computer and Internet Use in the United States: 2003," Available at: <http://www.census.gov/prod/2005pubs/p23-208.pdf>.
- [17] <http://www.internettrafficreport.com/>.
- [18] http://www.readwriteweb.com/archives/p2p_growth_trend_watch.php.
- [19] <http://www.dsreports.com/shownews/85022>, 2007.
- [20] J. Charzinski, "Internet Client Traffic Measurement and Characterization Results," in *Proceedings of the 13th International Symposium on Services and Local Access (ISSLS 2000)*.
- [21] M. Arlitt and C. Williamson, "Internet Web Servers: Workload Characterization and Performance Implications," *IEEE/ACM Transactions on Networking*, vol. 5, no. 5, October 1997.
- [22] M. Crovella and A. Bestavros, "Self-similarity in World Wide Web Traffic: Evidence and Possible Causes," *IEEE/ACM Transactions on Networking (TON)*, vol. 5, no. 6, pp. 835–846, 1997.
- [23] L. Guo, E. Tan, S. Chen, Z. Xiao, O. Spatscheck, and X. Zhang, "Delving into Internet Streaming Media Delivery: A Quality and Resource Utilization Perspective," in *Proceedings of the 6th ACM SIGCOMM on Internet measurement*, 2006.
- [24] W. Tan, W. Cui, and J. Apostolopoulos, "Playback-Buffer Equalization for Streaming Media Using Stateless Transport Prioritization," *Packet Video*, 2003.
- [25] S. Saroiu, P. Gummadi, and S. Gribble, "A Measurement Study of Peer-to-Peer File Sharing System," in *Proceedings of Multimedia Computing and Networking*, 2002.
- [26] J. Qiu and L. Gao, "AS Path Inference by Exploiting Known AS Paths," in *Proceedings of Globecom'06*, 2006.
- [27] H. V. Madhyastha, T. Anderson, A. Krishnamurthy, N. Spring, and A. Venkataramani, "A Structural Approach to Latency Prediction," in *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*, Rio de Janeiro, Brazil, 2006.
- [28] B. Huffaker, M. Fomenkov, D. Moore, and E. Nemeth, "Measurements of the Internet Topology in the Asia-Pacific Region," in *Proceedings of INET00*.