Geography-Based Analysis of the Internet Infrastructure

Shiva Prasad Kasiviswanathan† Stephan Eidenbenz‡ Guanhua Yan‡
†IBM T.J. Watson Research Center ‡Los Alamos National Laboratory
kasivisw@gmail.com {eidenben,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, 130 million end hosts, and one 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 dataset. 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.

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 Angles 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.* [10] 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

Model component	Data sources
Backbone topology	Skitter dataset: http://www.caida.org/tools/measurement/skitter/ Alias clustering data from the iPlane project: http://iplane.cs.washington.edu/data/alias_lists.txt IP geolocation dataset: http://www.ip2location.com/
Internet Point of Presence	Telegeography colocation database: http://www.telegeography.com/
Internet End Devices	US census data: census-block population in each $250{\times}250m^2$ grid in US for the 24-hour duration [3] US business information (headquarter locations, number of employees, SIC codes): Dun & Bradstreet (D&B) dataset Computer penetration ratio per business category: US census data [4]
Internet access routers	Dial-up service aggregators per each zip code: http://www.findanisp.com Broadband ISP market share: http://www.leichtmanresearch.com/press/081108release.html DSL central office locations: the LERG (Local Exchange Routing Guide) dataset from Telcordia Cable company service locations: Dun & Bradstreet (D&B) dataset
Internet sessions	Top 100 servers: http://www.alexa.com/ Internet traffic measurement results [5], [6], [7], [8]
Internet routing	BGP routing information from the University of Oregon Route Views Project: http://www.routeviews.org/ AS prefix sets: http://www.fixedorbit.com/ AS-level path inference: Qiu and Gao's algorithm [9]

TABLE I: Data sources used in our Internet infrastructure model

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 [11] 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.* [12] documented some root causes of this circuitousness. Lakhina *et al.* [13] studied a wide range of geographic properties of the Internet, focusing on routers, links, and autonomous systems. Yook *et al.* [14] 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 been proposed to characterize the routing and traffic in the Internet [15], [16], [17], [18]. Instead of relying on inter-domain routing models, we use an AS path inference algorithm to derive the actual inter-domain paths used in the Internet. By combining a number of 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.* [19], and we refer the reader to that paper for a complete description of the model. In Table I we summarize the list of data sources we used to generate our Internet model.

In total, we have generated 73,884,296 residential computers and 58,923,964 business computers in the entire US (except Hawaii and Alaska). We also model Internet access routers of three types, dial-up, DSL and Cable, based on the market share of top US broadband companies and dial-up service aggregators, and these access routers connect to the backbone topology at Internet PoP (Point or Presence) locations based on AS peering relationships. In addition, we have generated a total of 1.14 billion sessions, which include HTTP, email, P2P, and streaming traffic from every computer for a period of 24 hours. With this comprehensive, high-fidelity, Internet model, we shall investigate the geographic aspects of the US Internet infrastructure in the following section.

IV. Internet Routing Analysis

We analyze the paths generated by our experiments. For a session between a source at location s and a destination 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 $c = 2{\times}\text{arctan2}(\sqrt{a}, \sqrt{1-a})$ 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 1 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 1, 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 2. It is very easy to find such

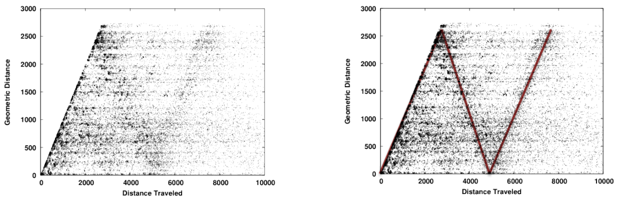

Fig. 1: Each point represents an Internet session. The X-axis represents the route distance between the source and the destination of the session. The Y-axis represents the geometric distance between the source and the destination of the session. Only a uniform $1/10000th$ fraction of all sessions we generated are shown 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.

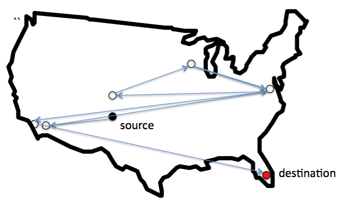

Fig. 2: A real traceroute path, generated from a computer in Santa Fe, New Mexico to Palm Beach, Florida. Notice the long route traverses both coasts.

long paths (that go from cost 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-B).

Another surprising feature from Figure 1 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 west coast of the US). We now try to explain why there is this concentration near these lines. We look at each of the 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 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. *Distance Ratio Analysis*

To better understand how far apart these distances could be, we look at various “ratio plots”. In Figure 3(a), we study the ratio of route distance to geometric distance. Let us define *stretch*, or *distance ratio* as referred to in [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 detour from the geometric path. The average stretch (over all sessions) is 10.25.

In Figure 3 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

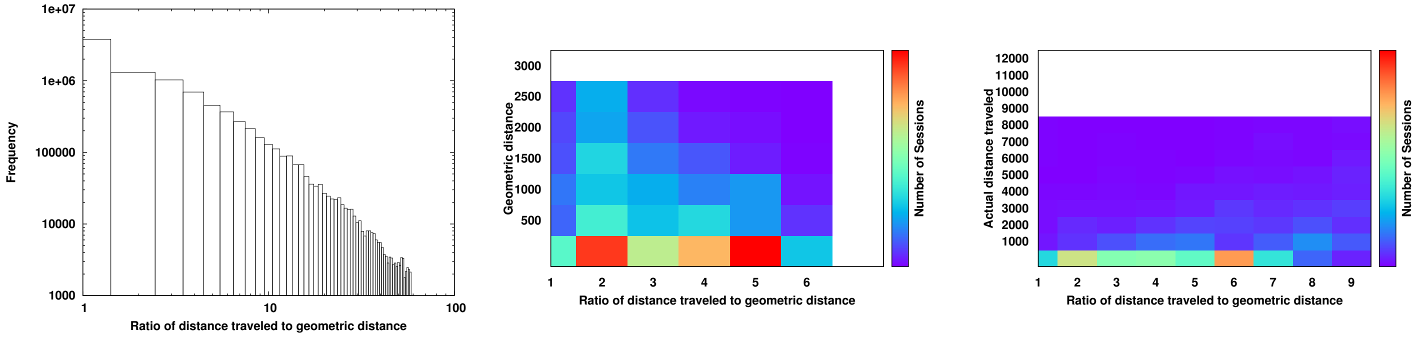

Fig. 3: (a): Histogram of the ratio of distances (stretch) generated using a uniform random sample of about 8 million sessions (in log-log scale). (b): Heat map representing stretch vs. geometric distance. The color scheme used in the heat map is shown adjacent to the plot. A reddish region has more number of points (sessions) than a bluish region. (c): Heat map representing stretch vs. travel distance.

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.

C. 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 4(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 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 4(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 4(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.

D. 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 4(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 5(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 Pearson correlation coefficient 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.* [20] by analyzing the CAIDA dataset for the Asia-Pacific region.

In Figure 5(b), we plot the variation of the hop count against the travel distance. As in Figure 5(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 number[1] 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 5(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.

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

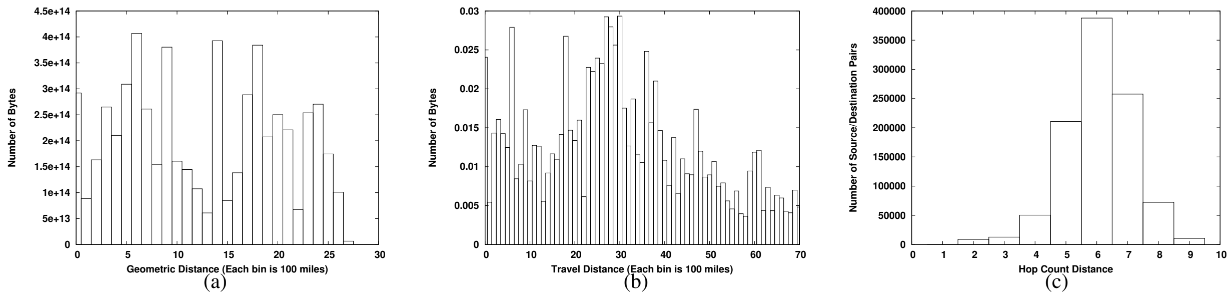

Fig. 4: (a): The distribution of the traffic according to geometric distance using a uniform sample of around 8 million sessions. The X-axis represents the geometric distance broken into bins of 100 miles. The Y-axis represents the volume of traffic (measured in number of bytes). (b): The distribution 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.

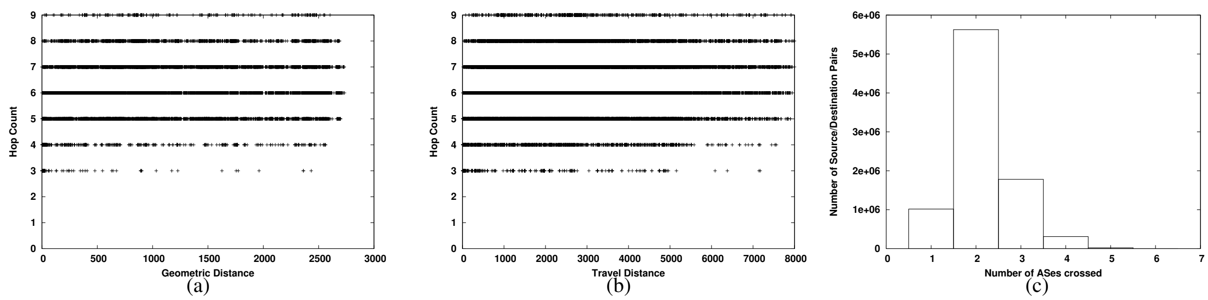

Fig. 5: (a): Variation of hop count with geometric distance. Sample size: around a million sessions. (b): Variation of hop count with travel distance. Sample size: around a million sessions. (c): Frequency distribution of the AS count. Sample size: around 8 million sessions.

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 many cases the traffic bounces multiple times between the two coasts before reaching the destination. 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 and Jian Qiu, Lixin Gao for their AS path inference algorithm. Most of the work of the first author was done while he was a postdoc at LANL

References

[1] H. Tangmunarunkit, R. R. Govinda, S. Shenker, and D. Estrin, "The impact of routing policy on internet paths," in *IEEE INFOCOM'01*.
[2] L. Subramanian, V. N. Padmanabhan, and R. H. Katz, "Geographic properties of internet routing," in *Proceedings of the General Track of the annual conference on USENIX Annual Technical Conference*, 2002.
[3] T. N. McPherson and M. J. Brown, "Estimating Daytime and Nighttime Population Distributions in U.S. Cities for Emergency Response Activities," in *Bulletin of the American Meteorological Sociaty*.
[4] "Computer and Internet Use in the United States: 2003," Available at: http://www.census.gov/prod/2005pubs/p23-208.pdf.
[5] http://www.internettrafficreport.com/.
[6] http://www.readwriteweb.com/archives/p2p_growth_trend_watch.php.
[7] http://www.dslreports.com/shownews/85022, 2007.
[8] J. Charzinski, "Internet Client Traffic Measurement and Characterization Results," in *Proceedings of the 13th International Symposium on Services and Local Access (ISSLS 2000)*.
[9] J. Qiu and L. Gao, "AS Path Inference by Exploiting Known AS Paths," in *Proceedings of Globecom'06*, 2006.
[10] 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.
[11] 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.
[12] N. Spring, R. Mahajan, and T. Anderson, "The causes of path inflation," in *Proceedings of ACM SIGCOMM Conference*. ACM, 2003, p. 124.
[13] 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.
[14] 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.
[15] J. Leguay, M. Latapy, T. Friedman, and K. Salamatian, "Describing and simulating internet routes," in *Proceedings of the 4th IFIP-TC6 Networking Conference*, pp. 2–6.
[16] 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.
[17] M. Zukerman, T. D. Neame, and R. G. Addie, "Internet traffic modeling and future technology implications," in *Proceedings of IEEE Infocom'03*.
[18] V. Paxson and S. Floyd, "Wide-area traffic: The failure of poisson modeling," *IEEE/ACM Transactions on Networking*, vol. 3, pp. 226–244, 1995.
[19] G. Yan, S. Eidenbenz, S. Thulsidasan, P. Datta, and V. Ramaswamy, "Criticality analysis and assesment of national internet infastructure," *Computer Networks*, vol. 54, no. 7, pp. 1169–1182, 2010.
[20] B. Huffaker, M. Fomenkov, D. Moore, and E. Nemeth, "Measurements of the Internet Topology in the Asia-Pacific Region," in *Proceedings of INET00*.