Analyzing Intercontinental Circuitousness to Improve the Interconnection and Routing for ISPs

Pengcheng Du, Jessie Hui Wang, Jiahai Yang
Institute for Network Sciences and Cyberspace
Tsinghua University

Jianfeng Wang, Youjian Zhao
Computer Science Department
Tsinghua University

***Abstract*—Routing detours raise concerns both on network security of nations and efficiency of resource consumption. Geo-optimal routing can be viewed as a long term goal of Internet routing, and a careful designed Internet together with geo-optimal routing can produce the best performance for the Internet. In this paper, we focus on some representative intercontinental circuitous paths summarized from our measurement data set. With the help of PeerDB and CAIDA, we try to find possible reasons for the detour routing and propose some suggestions to improve interconnection and routing for ISPs.**

I. Introduction

The routing of the Internet is determined by all ISPs based on their technical, business and political reasons. Due to these technical, business and political considerations, there exist a lot of detour routes in the Internet, which raise concerns both on network security of nations and efficiency of resource consumption.

Roughly speaking, the intradomain routing based on OSPF tries to minimize congestion on all intradomain links, while the interdomain routing based on BGP tries to provide a way for ISPs to enforce their business agreements and political considerations [1]. Both routing protocols are dynamic to find the best route and deal with short-term congestions and link failures.

But what is the optimal Internet routing in a long term? In this paper, we propose that the optimal Internet routing should follow the geographically shortest path as much as possible, *i.e.*, *Geo-Optimal Routing*. In other words, we should try to avoid detour routes to be selected. We argue that routing protocols improve networking performance in a short term, while in a long term we should try to improve networking performance by careful network plannings. A careful network design together with geo-optimal routing would produce the best performance for the Internet. Thus geo-optimality, *i.e.* degree and popularity of circuitousness, can be used to evaluate a network planning and the routing resulted from this planning.

Based on this argument, in this paper we try to find intercontinental detour routes in the Internet, analyze the possible reasons, and propose suggestions to improve their network plannings. We are not pushing ISPs to follow geo-optimal routing now, but take geo-optimal routing as a final goal in a well-designed future Internet.

We conduct a measurement study from four looking glass servers in different continents to a set of destination addresses which covers different countries in the world. The result reveals that the interconnection in Europe and USA are much better than other continents. From $AS7018$, only 3.66% of routes are with circuitousness ratio larger than 1.5. The percentage is 3.71% in the routes from $AS3303$ (Europe) to the global Internet. In Africa the percentage is 28.46%; and in Oceania it is 20.32%. Both are much larger than Europe and USA.

Based on the measurement study and data analysis, we present three case studies, *i.e.*, from Europe ($AS3303$) to Asia, from USA ($AS7018$) to South America, and from Africa ($AS5713$) to the world. Based on PeerDB and CAIDA, we then try to find possible reasons for these detour routes, and propose some suggestions for ISPs to improve their network plannings.

The paper is organized as follows. In Section II, we introduce some previous works which study geographical properties of Internet routing. In Section III, we describe how to collect and analyze routing data. We present three case studies and analyze them in details in Section IV. Then we summarize possible reasons for routing detours and discuss challenges in this study and future work in Section V. Section VI concludes the paper.

II. Related Work

In 2002, Lakshminarayanan Subramanian *et. al.* published a paper on geographical properties of Internet routing [2]. They mainly focus on the circuitousness of routing in USA and Europe, and conduct a statistical study to understand the routing detours. They quantify the degree to which an ISPs routing policy resembles hot-potato or cold-potato routing and find that many tier-1 ISP networks may have poor tolerance to the failure of a single, critical geographic node.

Peter Matray *et. al.* also conduct a statistical study on spatial properties of internet routes characterize the length distribution of Internet links [3]. They mainly focus on three areas: USA, Europe and East Asia. Their measurement shows that the paths between Europe and East Asia is more likely to be circuitous. They also present an extreme circuitousness on a network route between Auckland, NZ and Tokyo, JP, which is similar as the case studies in our paper.

Arpit Gupta *et. al.* take a first look at ISP interconnectivity between various regions in Africa [4]. They reported that many circuitous Internet paths that should remain local often detour through Europe. They also investigate the causes of circuitous

Internet paths and propose to increase peering and cache proxy placement for reducing latency to popular Internet sites.

The authors of [5] study the circuitousness in USA. They report a dominant coast-to-coast pattern in the US Internet traffic. Their study is based on a set of geolocated IP hop-level session data they synthesized from a variety of different input sources using the model proposed in [6]. In [7], the authors refer this phenomenon as "Internet Boomerang Routing". Motivated by the security concerns raised by boomerang routing, they take Canadian boomerang routing as a case study, and try to show its extent and recurring patterns. In [8], the authors study path inflation by comparing against the paths that are shortest in terms of different metrics including geographic distances. Based on experience from simulations, they argue it is necessary to include an explicit notion of geographic distance in the routing information to implement geo-optimal routing.

Our work is different from above works in several ways. We mainly focus on intercontinental detours, which may have more significant influence on Internet routing. We use real routes collected by traceroutes. Instead of conducting a statistical analysis, we focus on case studies, and try to propose some suggestions for ISPs. Besides understanding the circuitousness phenomenon of Internet routing, the goal of our study also includes proposing a framework to evaluate and improve routing and planning of one network.

III. Data Collection and Data Sources

A. Data Collection and Preparation

Our data collection is based on traceroute from various looking glass servers to all destination addresses all over the world [11] [12]. To study the geographical features of global routing, we select 4 looking glass servers [13] distributed in different continents. One are located in Europe, *i.e.*, 164.128.251.1 of Swisscom in Switzerland. One is located in North America, *i.e.*, 12.0.1.202 of AT&T in USA. The looking glass server in Africa is 196.25.9.45 of Telkom SA in the Republic of South Africa. In Oceania, we select the looking glass server 203.202.125.3 of Optus in Australia.

We then generate a set of destination addresses. In order to cover the global Internet as much as possible, we uniformly select addresses from the whole IPv4 address space. In total, we select about 20,000 destination addresses. Let us denote the set as D.

From the four looking glass servers, we traceroute all destinations in D. In order to control the load we incur on those looking glass servers and make sure that the traceroute to one destination is not interrupted by the traceroute to the next address, we submit one traceroute command every 50 seconds. It takes us about 11 days to finish the data collection from one looking glass server.

We save the result of traceroute for further analysis. Each entry in our results is a path to one destination address, including the IP address, domain name (if exist), and AS number of each router on the path. We also save the delay to each router on the path.

Some routers do not reply traceroute commands, which makes some paths incomplete. If the available last hop of one path is not in the same AS as the destination, we filter the path out since the information of this entry is not enough for our study.

B. Other Data Sources

First, we need to determine the geographical location of each router on each path. In our work, this is done based on *GeoLiteCity*[14], and we use the screenshot of June 2014, the month in which we conduct the data collection of routing paths. When we find a suspicious detour route, we then check the locations of each hop from multiple databases [15] [16] [17]. After that, we further look into the delay time of each hop and check if the suspicious intercontinental link brings sharp rise in delay.

During our analysis, we need to know the geographical coverage of ISPs, *i.e.* PoPs, and also the business relationship between two ISPs, *i.e.* peer-to-peer and transit. We depend on PeerDB [18] for the former, and depend on CAIDA [19] for the latter.

IV. Case Studies of Global Internet Routing

In order to understand the severity of routing detours, we define a metric, *i.e.* circuitousness ratio. Let us denote the path from the source s to d as $P_s^d = (p_0, p_1 ... p_i ... p_n, p_{n+1})$, where $p_0 = s$ and $p_{n+1} = d$. The circuitousness ratio is defined as follows:

$$C_s^d = \frac{\sum_{i=0}^{n} G(p_i, p_{i+1})}{G(s, d)}, \quad (1)$$

wherein $G(p_i, p_j)$ gives the geographical distance between p_i and p_{i+1}. We calculate it using the algorithm in [20].

A. Routing between Europe and Asia

We first try to conduct a measurement study on the routing of Europe to the global Internet. $AS3303$ (Swisscom) provides a looking glass server and we take it as a representative of European ISPs. From this server, we find out its paths to different countries using traceroute command, and see if there can be any improvement.

In the data set of all routing paths sourced from $AS3303$, we find out the paths to countries in East Asia are problematic, *i.e.*, with routing detours. In order to further investigate the routing, we category the routes from $AS3303$ to all destinations in East Asia, and show them in Figure 1. Please note one path in the Figure is not a single route to one destination IP address, but represents a group of routes to a same network in a same country.

From CAIDA, we can see that $AS3303$ has direct interconnections with all destination ISPs in this figure, except $AS4538$ (CERNET). However, it does not mean that $AS3303$ has geo-optimal routes to these destinations. Both the routing to $AS9318$ (Korea Hanaro Telecom) and $AS2497$ (Japan Internet Initiative) have fairly big circuitousness ratios, *i.e.*, $C_{3303}^{9318} = 2.15$ and $C_{3303}^{2497} = 1.89$.

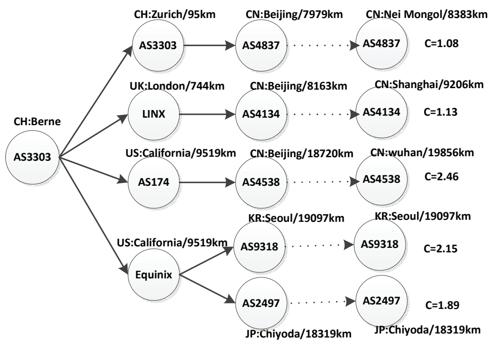

Fig. 1. Routing between One European AS and Countries in East Asia

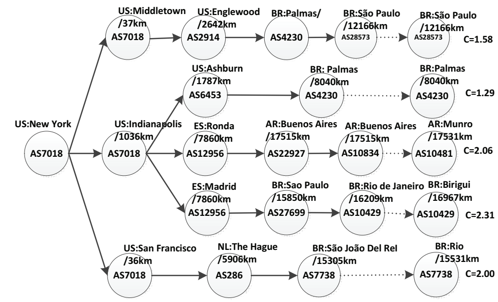

Fig. 2. Routing between One USA ISP and Countries in South America

To compare with these detour routes, we also list two categories of direct routes, *i.e.* the upper two routes in Figure 1. The traffic flows from $AS3303$ are exchanged to their destination ASes in two European cities, Zurich and London. The result circuitousness ratios are $C_{3303}^{4837} = 1.08$ and $C_{3303}^{4134} = 1.09$ respectively. Obviously, they can have geo-optimal routing because the source ISP and the destination ISP set up interconnection links in European cities (confirmed by peerDB and CAIDA), and thus traffic flows can follow geo-optimal paths.

Contrast to the upper two routes, the lower two routes are traversing USA, making the paths detour routes. According to PeerDB, $AS9318$, $AS2497$ and $AS3303$ share at least one PoP in Europe, *i.e.*, LINX London, at least one PoP in Asia, *i.e.* Hong Kong, and at least one PoP in USA, *i.e.*, Equinix Palo Alto. Our measurement results show that in fact two ISPs decide to exchange traffic flows at Equinix Palo Alto in USA, and the path is with a delay of about 300ms, which is worse than the route to $AS4837$ (China Unicom). Intuitively, both Hong Kong and London would be better than current choice. It is reasonable to suspect that they do not peer with each other in Europe and Asia, although they are colocated at these PoPs, or they have a special routing consideration which results in the current routing decision.

Similar as the lower two paths, the traffic flows to $AS4538$ are exchanged in USA, and $C_{3303}^{4538} = 2.46$. In this case, $AS4538$ does not have any direct interconnection agreement with $AS3303$, and the traffic flows go through $AS174$ (Cogent). We then look into PeerDB, and find that CERNET only has PoPs in Asia and USA, and no PoP in Europe. Therefore, $AS3303$ has to sent the traffic flows to Asia or USA, and it selects USA in our case study. In fact, $AS3303$ also has PoP in Hong Kong, which might be a better choice for the traffic flows between them. Therefore, two ISPs can consider to interconnect with each other in Hongkong, which might be good for both sides.

As a summary, we have the following suggestions:

1) *We suggest that $AS9318$ and $AS2497$ examine the routing between them and $AS3303$, and see if the current configuration and policy are what they expect. They should fully exploit their PoPs in European, try to peer with their peers at multiple geographical locations and use multiple interdomain links efficiently.*
2) *We suggest that $AS4538$ try to set up interconnections with European ISPs in Hong Kong, or extend its network to Europe to set up interconnection links with European ISPs directly in Europe if the cost is bearable.*

B. Routing between USA and South America

The looking glass server in North America used in this study is from $AS7018$ (AT&T). Its routes to Europe, Asia and Africa are near geo-optimal, *i.e.*, with small circuitousness ratios. Figure 2 illustrates its routes to different locations in South America. To facilitate reading, we have categorized all paths into several categories and only describe the representatives of each category.

Different from Figure 1, $AS7018$ does not have direct interconnections with any destination AS in Figure 2. Therefore, all traffic flows between them should be transited by one or more intermediate ASes.

The upper two routes are from USA to South America directly. The second route is with a fairly small circuitousness ratio, *i.e.*, $C_{7018}^{4230} = 1.29$. For the first route, $C_{7018}^{28573} = 1.58$, which is slightly larger than the second route. In fact, the first route also traverses $AS4230$ (Embratel), but the intermediate AS is $AS2914$ (NTT) instead of $AS6453$ (TATA), which makes the first route longer than the second route.

From the upper two routes, we can see that it is possible for traffic flows from $AS7018$ to be delivered to South America without big circuitousness. However, all of the remaining three routes take a detour through Europe, and thus result in larger circuitousness ratios.

Both of the third route to $AS10429$ (Telefonica S.A.) and the forth route to $AS10481$ (Prima S.A.) go through $AS12956$ (Telefonica Backbone) in Spain, and we have $C_{7018}^{10429} = 2.318$ and $C_{7018}^{10481} = 2.06$. It is a little weird since $AS12956$, as a global tier-1 ISP, has a lot of PoPs in all three areas, *i.e.*, USA, Europe and South America. $AS12956$ is capable to transmit traffic flows from USA to South America directly. We look into the delay to each hop in the path, and find two sharp rises, which should be caused by two intercontinental links. The detour route might be a result of traffic engineering conducted by $AS12956$. If it is true, it is reasonable for us to conjecture that $AS12956$ should increase its backbone capacity between its PoPs in USA and PoPs in South America in a near future.

The last route goes through Netherlands, and $C_{7018}^{7738} = 2.001$. Based on information from CAIDA, we try to check if it is possible for $AS7018$ to reach $AS7738$ (Telemar) through the intermediate ASes in the upper two routes, *i.e.*, $AS6453$ and $AS2914$. We find that $AS6453$ is in fact one provider of $AS7738$, and $AS7018$ has a peer-to-peer relationship with $AS6453$. Therefore, it is really a feasible path between $AS7018$ and $AS7738$ to go through $AS6453$, and this path would have a smaller circuitousness ratio. A further investigation shows that the reason that $AS7018$ does not select this path might be that $AS7018$ prefers a route to its customer $AS286$ (KPN).

When we look into the PeerDB, we notice that $AS7738$ has some PoPs in USA. $AS7018$ is a Tier-1 ISP and it also has a lot of PoPs in USA. It is also a reasonable choice for two ISPs to set up interconnection links in a suitable location, and then they can exchange traffic flows directly.

In summary, we have the following suggestions:

1) *We suggest that* $AS7018$ *review its interdomain routing to different ISPs in South America and see if its current routing is what it expects, especially the routes to* $AS10481$ *and* $AS10429$*.*
2) *We suggest that* $AS12956$ *examine its intradomain routing among PoPs in different areas, and see if the current detour routing is what it expects. We also suggest to investigate a long-term solution of network planning which increases the backbone capacity of related links.*
3) *We suggest that* $AS7738$ *and* $AS7018$ *examine the possibility of interconnecting with each other in some co-located PoPs.*

C. Routing between Africa and the World

The looking glass server in Africa used in this study is located in $AS5713$ (Telkom SA). Africa is a developing area in the Internet. In [4], the authors have observed the phenomenon that the traffic flows between two Africa ISPs are often exchanged in Europe. Our measurement results also demonstrate this phenomenon.

In Figure 3, the upper two routes are from $AS5713$ to countries in Africa. One is transited at LINX in Europe, and the other is transited by $AS174$ in USA. Both routes are with large circuitousness ratios.

We also find the routes to countries in South America are similar as Africa. Although South America and Africa are geographically close, there is only two cables between two continents [21]. Therefore, Africa ISPs also rely on IXPs in Europe and USA to exchange the traffic flows to South America, which results in routes with large circuitousness ratios.

Figure 4 plots several routes from $AS5713$ to Asia countries. $AS6453$ (TaTa) is global tier-1 ISP headquartered in India, and it provides transit service in Asia-Pacific areas. Intuitively, $AS5713$ can rely on this ISP to reach Asia countries, *e.g.*, the upper two routes. However, there are still a lot of

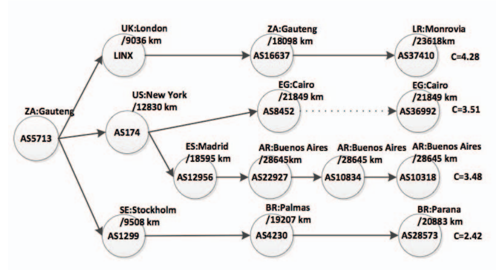

Fig. 3. Routing between One Africa ISP and Countries in Africa and South America

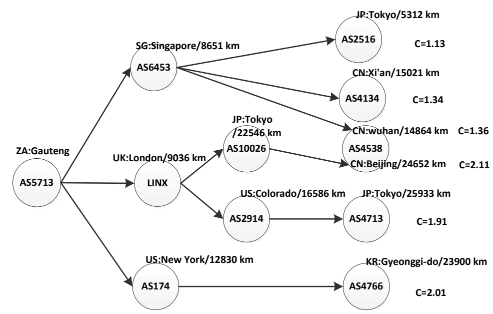

Fig. 4. Routing between One Africa ISP and Countries in Asia

traffic flows going through $AS174$ (USA) or LINX in Europe. Although their circuitousness ratios vary a lot, there is no much difference between the delay of these routes, which suggests the performance of links in TaTa Asia should be improved.

The routing to Australia is similar as the routing to Asia, and we plot several representative routes in Figure 5. Roughly speaking, there are three paths from Africa to Australia, via Asia, Europe or USA. The route via Asia is with a smallest circuitousness ratio. However, in our measurement, its delay is longest among all three routes. The route via Europe, *e.g.* the second route, is best in terms of delay time.

We summarize our suggestions from the above observations as follows.

1) *We suggest Africa ISPs should cooperate with each other and try to exchange local traffic flows locally.*
2) *We suggest Africa and South America examine the utilization rate of the intercontinental cables. If the utilization rate is low, they may want to improve its usage and decrease their reliance on ISPs or IXPs in Europe and USA. Otherwise, they may need to have*

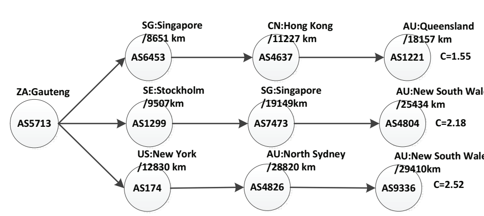

Fig. 5. Routing between One Africa ISP and Countries in Australia

more intercontinental cables.

3) *We suggest $AS6453$ examine the performance of its backbone from Africa to Asia-Pacific. $AS6453$ has had a PoP in Africa, but a lot of traffic flows to Asia-Pacific still go through Europe and USA possibly due to performance considerations.*

V. Remarks and Future Work

Our measurement results show that Europe and USA, especially Europe, are still the center of the Internet. The traffic flows in and between these two continents would not go through other places, while a lot of traffic flows are exchanged there although their sources and destinations are in other continents. This observation is consistent with our intuition.

We calculate the circuitousness ratio C for all routes we collect from four looking glass servers. From $AS7018$, only 3.66% of routes are with circuitousness ratio larger than 1.5. The percentage is 3.71% in the routes from $AS3303$ (Europe) to the global Internet. In Africa the percentage is 28.46%; and in Oceania it is 20.32%. Both are much larger than Europe and USA.

A. Remarks

It is well known that routing decisions are affected by a lot of factors such as technical, business and politic considerations. Therefore, it is infeasible to assert the reasons for each detour route. However, detour routes may bring a lot of concerns on security and resource consumption. In a short term, detour routes may improve performance and avoid some bottleneck links. However, in a long run, network congestions should be solved by planning networks carefully instead of depending on detour routes. This is why we propose the concept of *Geo-Optimal Routing*, and try to use it to evaluate the global routing ecosystem and point out possible ways to improve the global interconnection and routing performance for ISPs.

Based on observations from our measurement study, we summarize the possible reasons and potential solutions as follows.

1) *detour due to business considerations*
   As shown in Figure 6.a, the upstream ISP has two paths to the destination, wherein the path via a peer is shorter than a path via a customer ISP. The upstream ISP prefers the longer path since the customer ISP would pay for the traffic flows between them. It results in a detour route. This kind of detour routing is determined by the upstream ISP on purpose. The route from $AS7018$ to $AS7738$ illustrates this kind of detour routing. According to PeerDB and CAIDA, there might be a short path $AS7018, AS6453, AS7738$ which goes from USA to South America directly. But $AS7018$ prefers the route via its customer $AS286$. Generally said, in this case the destination ISP may be more eager to shift the traffic flows to the short path. However, it is not easy for a downstream ISP to affect the routing decisions of upstream ISPs.
2) *detour due to improper inter-domain routing*
   Figure 6.b illustrates this kind of detours. Two ISPs interconnect with each other at multiple locations, and they use improper location to exchange the traffic flows. In this case, two ISPs should examine the configuration of their routing protocols.
3) *detour due to the lack of an inter-domain link*
   Figure 6.c illustrates this kind of detours. Two ISPs are co-located at a city, but they do not set up any interconnection link in this city. As a result, some traffic flows have to traverse a long path to be exchanged. The solution to this kind of detours is to set inter-domain interconnections at co-located cities.
   It is not easy to tell whether a detour route is caused by Figure 6.b or 6.c. For example, in Figure 1, the routes from $AS3303$ to $AS9318$ and $AS2497$ are circuitous through USA. From PeerDB and CAIDA, we can see that all three ISPs share at least one PoP in Europe, *i.e.*, LINX London, at least one PoP in Asia, *i.e.* Hong Kong. Obviously, a route via London or Hong Kong is shorter than current route. We are not aware of whether the current route is selected because they do not have interconnections in these two cities or because their routing configurations are not optimized.
4) *detour due to the lack of interconnection agreement*
   Figure 6.d illustrates this kind of detours. Take the detour route from $AS3303$ to $AS4538$ shown in Figure 1 as an example. Two ISPs are co-located in Hong Kong, but they do not have any business relationship. Their traffic flows have to be exchanged via other ISPs, which may result in detours.
   We understand that setting up relationships with more ISPs would increase the complexity of routing and management. Our investigation and suggestion only serve as a first step to improve the interconnection of one ISP.
5) *detour due to the lack of cables*
   Sometimes detour routing is caused by the limitation of hardware cables. For example, there is only two cables between Africa and South America, and we suspect the intercontinental cables are often congested. Therefore, ISPs select detour routes to avoid congestions. Obviously, the solution to this kind of detours is to invest more money in cables.

B. Challenges and Future Work

The study on the geo-optimality and possible improvement of global routing faces a lot of challenges. We list some of the challenges as follows.

- *Accuracy and Completeness of Data Collection*
  We rely on responses of routers to the traceroute command to collect data about global routing. Unfortunately, not all routers reply to the traceroute command. Therefore, we frequently see incomplete paths in which the information of many hops is missing. These incomplete paths cannot be used for further analysis.

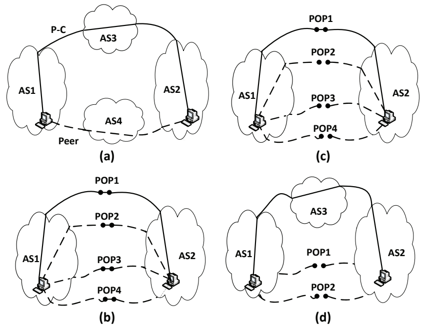

Fig. 6. Possible Reasons of Detour Routes

Another challenge during data collection is that the Internet routing is dynamic. Some routes are temporally selected due to congestions or other issues. Our goal is to analyze the global routing ecosystem and propose some long-term suggestions. Those temporal routes are not suitable for our goal and should be removed from the data set.

In the future, we will try to collect more snapshots, compare the routes in different snapshots and then remove temporal routes.

- *Mapping IP addresses to Geographical Locations*
  Mapping one IP address to its geographical location has drawn a lot of attentions of researchers [22] [23], but the problem is not well solved yet. Industry and researchers have published several databases, but no one can guarantee the accuracy of all entries of one database. In this study, when we find a suspicious detour route, we then check the locations of each hop from multiple databases. After that, we further look into the delay time of each hop and check if the suspicious intercontinental link brings sharp rise in delay. It takes us a long time to verify a detour route, and this way is not applicable in a systematic large scale study.
- *Accuracy of PeerDB and CAIDA*
  We depend on PeerDB and CAIDA to analyze reasons of detour routes. Unfortunately, sometimes the related ISP does not publish its information in PeerDB, and we can not guarantee the accuracy of these two databases.

VI. Conclusion

In this paper, we study intercontinental circuitous paths based on the data set we collect using traceroute from several looking glass servers to the whole IPv4 address space. We summarize the circuitous paths and extract several representative routes for each source-destination pair. With the help of the information about business relationship between ISPs from CAIDA and the information about PoP coverage of each ISP from PeerDB, we analyze these representative routes in details. Based on the analysis, we summarize possible reasons for Internet detours and propose suggestions for some ISPs. Our work is a first step towards a framework to improve network planning for both the Internet and each individual ISP.

VII. Acknowledgment

This work is supported by the National Natural Science Foundation of China under Grant No. 61202356.

References

[1] A. Di Menna, T. Refice, L. Cittadini, and G. Di Battista, "Measuring route diversity in the internet from remote vantage points," in *Proc. International Conference on Networks (ICN 2009)*, 2009, pp. 24–29.

[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*, ser. ATEC '02. Berkeley, CA, USA: USENIX Association, 2002, pp. 243–259. [Online]. Available: http://dl.acm.org/citation.cfm?id=647057.713857

[3] P. Mtray, P. Hga, S. Laki, G. Vattay, and I. Csabai, "On the spatial properties of internet routes," *Computer Networks*, vol. 56, no. 9, pp. 2237 – 2248, 2012. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1389128612000989

[4] A. Gupta, M. Calder, N. Feamster, M. Chetty, E. Calandro, and E. Katz-Bassett, "Peering at the internets frontier: A first look at isp interconnectivity in africa," in *Passive and Active Measurement*, ser. Lecture Notes in Computer Science, M. Faloutsos and A. Kuzmanovic, Eds. Springer International Publishing, 2014, vol. 8362, pp. 204–213.

[5] S. Kasiviswanathan, S. Eidenbenz, and G. Yan, "Geography-based analysis of the internet infrastructure," in *INFOCOM, 2011 Proceedings IEEE*, April 2011, pp. 131–135.

[6] G. Yan, S. Eidenbenz, S. Thulasidasan, P. Datta, and V. Ramaswamy, "Criticality analysis of internet infrastructure," *Computer Networks*, vol. 54, no. 7, pp. 1169 – 1182, 2010. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1389128609003545

[7] A. Clement and J. A. Obar, "Internet boomerang routing: Surveillance, privacy and network sovereignty in a north american context," Tech. Rep., 2013, available at SSRN: http://ssrn.com/abstract=2242593.

[8] W. Mhlbauer, S. Uhlig, A. Feldmann, O. Maennel, B. Quoitin, and B. Fu, "Impact of routing parameters on route diversity and path inflation," *Computer Networks*, vol. 54, no. 14, pp. 2506 – 2518, 2010. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1389128610001088

[9] G. S. N. Chatzis and A. Feldmann, "On the importance of internet exchange points for todays internet ecosystem," 2013.

[10] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger, "Anatomy of a large european ixp," in *Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, ser. SIGCOMM '12. New York, NY, USA: ACM, 2012, pp. 163–174.

[11] S. Garcia-Jimenez, E. Magana, D. Morato, and M. Izal, "Pamplona-traceroute: Topology discovery and alias resolution to build router level internet maps," in *Global Information Infrastructure Symposium, 2013*, Oct 2013, pp. 1–8.

[12] A. Abdelkefi, Y. Eftekhari, and Y. Jiang, "Locating disruptions on an internet path through end-to-end measurements," in *Computers and Communications (ISCC), 2013 IEEE Symposium on*, July 2013, pp. 000 648–000 653.

[13] "Looking glass," http://lg.eastlink.ca/.

[14] "Geolitecity," http://dev.maxmind.com/geoip/legacy/geolite/.

[15] "ip-geo-location," http://www.tcpiputils.com/ip-geo-location.

[16] "db-ip," https://db-ip.com/db/.

[17] "whatismyipaddress," http://whatismyipaddress.com/.

[18] "Peerdb," https://www.peeringdb.com/private/.

[19] "Caida," http://www.caida.org/home/.

[20] "Calculate distance, bearing and more between latitude/longitude points," http://www.movable-type.co.uk/scripts/latlong.html.

[21] "cable-map-2014," http://www.telegeography.com/.

[22] A. Lakhina, J. W. Byers, M. Crovella, and I. Matta, "On the geographic location of internet resources," in *Proceedings of the 2Nd ACM SIGCOMM Workshop on Internet Measurment*, ser. IMW '02. New York, NY, USA: ACM, 2002, pp. 249–250. [Online]. Available: http://doi.acm.org/10.1145/637201.637240

[23] G. Wang, C. Zhang, X. Qiu, and Z. Zeng, "Modelling a tractable and annotated isp's router-level topology based on statistical data and geolocation mapping," in *Communication Software and Networks (ICCSN), 2011 IEEE 3rd International Conference on*, May 2011, pp. 31–35.