

On Measuring the Geographic Diversity of Internet Routes

Attila Csoma, András Gulyás, László Toka

Abstract

Route diversity in networks is elemental for establishing reliable, high-capacity connections with appropriate security between endpoints. As for the Internet, route diversity has already been studied at both Autonomous System- and router-level topologies by means of graph theoretical disjoint paths. In this paper we complement these approaches by proposing a method for measuring the diversity of Internet paths in a geographical sense. By leveraging the recent developments in IP geolocation we show how to map the paths discovered by traceroute into geographically equivalent classes. This allows us to identify the geographical footprints of the major transmission paths between end-hosts, and building on our observations, we propose a quantitative measure for geographical diversity of Internet routes between any two hosts.

Index Terms

geodiversity; traceroute; geolocation; disjoint routes

I. INTRODUCTION

The value of knowledge of the Internet topology is arguably immense. In the last decades we have witnessed a myriad of stories in which topology-related information about the Internet was directly compiled into more efficient architectures, services and more appropriate business decisions. Content Delivery Networks (CDNs) [1], in which global topological peculiarities are highly exploited for e.g. surrogate server and cache placement strategies or request routing mechanism, are just a narrow segment of the whole spectrum. Peer-to-peer networks [2, 3], data center placement [4], traffic engineering [5], business-based AS peering strategies [6], just to mention a few, are all receivers of Internet topology related knowledge. With this non-comprehensive list of receivers in mind, it should come at no surprise that many researchers have contributed to our current understanding of the topology of the Internet.

An elemental metric of Internet topology is the diversity of routes between end-hosts, as multiple uncorrelated routes can provide better throughput, resiliency and security. In [7], [8] and [9] authors describe how to increase the resilience of future networks and the role of multipath communication in it. A detailed description about network security can be found in [10] and authors in [11] propose a method which improves network security; however it requires multiple path between end-hosts.

Existing studies of IP-level route diversity usually focus on extracting routes between end-hosts, e.g. by using `traceroute`, and on computing their diversity by means of edge or node disjointedness in a graph theoretical sense [12, 13, 14]. Such analysis provides an interpretation of route diversity in a microscopic level where each node in the route is a router interface having a particular IP address. In [15] the authors propose to interpret route diversity at a higher level, namely at the level of PoP's (Point of Presence): the interfaces residing in the same building or campus are grouped together, forming a PoP, and finally route diversity is computed at the level of these PoP's. In this paper we interpret route diversity at an even higher, geographical level. We propose grouping routers in a given geographic vicinity and compute route diversity at the level of geographical regions (e.g. the level of cities) independently from ASes. Our contribution is threefold: *first we describe a method for identifying the geographically equivalent routes in traceroute outputs; second, we show the efficiency of the method in terms of successfully merged traceroute routes and present their possible applications; finally we define a metric which can capture the geo-diversity of Internet routes between endpoints and compute this metric for our measurement dataset.* Such characterization of routes' "geo-diversity" is clearly beneficial if one is curious about connectivity between end-hosts in case of e.g. power outages affecting larger geographical areas.

The rest of the paper is organized as follows. In Sec. II we overview the corresponding related work. Sec. III describes our `traceroute` measurements and our algorithms for extracting geographically equivalent routes from those. In Sec. IV we define and evaluate a metric called Geographic Diversity Index (GDI) that captures how Internet routes differ from each other. In Sec. V we validate our framework and present its performance. Finally we draw the conclusions and list the possible applications and future work in Sec. VI.

A. Csoma is with Budapest University of Technology and Economics, Hungary and with HSNLab, Dept. of Telecommunications and Media Informatics. e-mail: csoma@tmit.bme.hu

A. Gulyás is with Budapest University of Technology and Economics, Hungary, with HSNLab, Dept. of Telecommunications and Media Informatics and with MTA-BME Information Systems Research Group. e-mail: gulyas@tmit.bme.hu

L. Toka is with Budapest University of Technology and Economics, Hungary, with HSNLab, Dept. of Telecommunications and Media Informatics and with MTA-BME Information Systems Research Group. e-mail: toka@tmit.bme.hu

The work of László Toka was supported by the János Bolyai Research Scholarship of the Hungarian Academy of Sciences (MTA).

II. RELATED WORK

Numerous existing studies apply geographical information to uncover non-trivial aspects of the Internet topology. In [16] the authors use the geographical positions of routers to estimate route circuitousness, route length distribution and geographic fault tolerance from an end-to-end and from an ISP perspective. However, the DNS name-based geolocation method used in their work has its own limitations and may create a false spatial distribution as described in [17]. The distance and angle between consecutive IP hops are investigated in [18]. In [19] the authors use geographic information to construct the hyperbolic map of the Internet and prove the navigability of the AS level topology using greedy algorithms. Points of presence are detected using delay constraints on an IP interface graph and the distribution of PoPs around the globe is visualized in [15]. Authors of [20] used PoP detection to evaluate the accuracy of some IP geolocation database. Inter-AS route diversity is examined through the network of Sprint in [14]. In [21] the authors study the route diversity of multi-homed and overlay networks as seen from multiple vantage points using graph theoretical methods exclusively. However none of these works capture route diversity in the pure geographical sense on the router-level.

III. MEASUREMENT FRAMEWORK

We built a framework that determines the extent of geographical heterogeneity of end-to-end routes that are being used to carry traffic between any two points in the Internet. The method that we implemented is the following: first, we run `traceroute` measurements to collect the IP-level routes between the selected endpoints; second, we use *MaxMind* [22], a localization tool that determines the geographical position of the recorded IP addresses; third, we group those IP-level routes that we consider equivalent from a geographical perspective; finally, we calculate a route diversity index for the selected endpoints. In this section we present these steps in details.

A. Route measurements with `traceroute`

Usually `traceroute` is used to discover end-to-end routes between two endpoints in the Internet. Network operators use it for detecting network errors, researchers use it to build Internet topology models. Although most in-network routers and endpoints support its operation, `traceroute` has a number of well-known shortcomings. On one hand it can be easily deceived by load-balancers, on the other hand it is an active measurement tool and due to the extra data traffic it generates, certain network equipment are configured to disable reactions to `traceroute` (and `ping`).

Several projects exist that collect `traceroute` measurements and make them publicly available. These data sources differ significantly based on their vantage point types, their vantage point location, the `traceroute` implementation they run and their endpoint selection methods. Two such projects are IPlane [23] and CAIDA's ITDK (Internet Topology Data Kit) [24]. IPlane offers a route performance prediction service and periodically runs `traceroute` measurements from PlanetLab [25] nodes to a set of endpoints changed bi-weekly. ITDK data sets are produced from measurements collected by CAIDA's Archipelago project: `paris-traceroute` [26] is run from 89 vantage points spread over 37 countries to randomly selected endpoints. In order to measure route diversity between two endpoints, we need to detect as many routes between those two endpoints as possible. Although the `paris-traceroute` output of ITDK is more reliable than that of IPlane's `traceroute`, the random selection of endpoints implemented by CAIDA hinders the collection of routes between the same vantage- and endpoints. Therefore we used the data of IPlane's `traceroute` measurements.

B. IP localization and filtering of routes

Once we have the IP-level routes, we determine the geographic position of each node appearing in them. Naturally the accuracy of the positioning is of paramount importance. As it is also noted in [27, 17], the use of DNS names and contents of various registries leads to unacceptable inaccuracy. Instead, we use the freely available geolocation tool MaxMind GeoLite [22] in order to establish the position for the IP addresses recorded in the measurements. As pointed out in [28], it is one of the most reliable, freely available geolocation database.

We filter the `traceroute` dataset from IPlane in order to remove measurements of vantage- and endpoint pairs between which only 1 IP-level route was observed. After localizing the IP hops, we further removed those vantage- and endpoint pairs between which only 1 geographic path was observed. With this second filter we eliminated traces differing only due to IP-aliasing and load balancing. The remaining set of traces contained ~ 0.5 million discovered routes; interestingly $\sim 80\%$ of vantage- and endpoint pairs had only 1 geographic path in the measurements.

C. Clustering of routes

After establishing *geo-paths* for the remaining routes by localizing their IP hops, we set out to decide which routes can be considered to be the same and which ones are different in a geographical sense. We make this decision by clustering the geo-paths on a hop-by-hop basis and by defining geodiversity, a mutual metric, between geo-paths. Iterating through all the geo-paths from a given vantage point to a given endpoint, we choose an appropriate cluster for each geo-path: if the geo-path

satisfies the geographical equality with all cluster members then it is assigned to that cluster if not, then a new cluster is created for it. At the end, all geo-paths in each cluster are considered to be geographically the same.

Two geo-paths are geographically equal if they are not farther from each other than 50 km: the distance between any of their nodes and the closest one of the straight lines determined by consecutive nodes of the other geo-path is not larger than 50 km. We used the arbitrary threshold of 50 km to reflect a large city's diameter [29]. Note, that this threshold also allows for fiber duct curves in the physical network, alleviating the mismatch between the location of IP-level nodes of traceroute measurements and the actual trail of the underlying physical links.

As a demonstrative example, we drew two geo-paths that are grouped in the same cluster in Fig. 1. Those routes differ at the IP-level and their geo-paths are also different, but since the distance (marked with dashed blue line) of the intermediate node of one route from the geo-path of the other route is smaller than 50 km, our clustering algorithm rules the geographic difference between these two routes negligible.



Fig. 1: Route merging example

IV. CALCULATION OF GEOGRAPHIC ROUTE DIVERSITY

Our final step to capture the geographic route diversity is to define and evaluate a Geographic Diversity Index (GDI). We require GDI to produce values to a given route set between a source and destination pair such that multiple geo-paths spanning over large geographical areas get a higher GDI.

A. Requirements

Before describing the computation of GDI in details we highlight the key attributes which make a route set "more" diverse in a geographical sense against an other. Let us assume a source (S) and a target (T) node as endpoints and two "cover" routes S-A-B-C-T (R1) and S-F-G-H-I-T (R3) as shown in Fig. 2¹. Let us assume that nodes are positioned according to their geographic positions. Let us also assume that the GDI for this setting is \hat{r} . Our first goal is to reward higher route count. Therefore a route set with two routes must achieve lower GDI than the same route set extended with another arbitrary route. That is if we add R2 to our theoretical example and obtain \hat{r}_2 as the GDI for this amended route set, then we expect $\hat{r} < \hat{r}_2$. Our second goal is to reward higher geographic distance between routes. Therefore if we modify the route set by reducing the distance between the routes, then the GDI of the new route set must be lower: by replacing R2 with R4 (new GDI is \hat{r}_4), we




Fig. 2: Geo-diversity example

¹Lines are curved to distinguish different routes using the same link.

require the GDI to fall. It follows that geo-paths with varying distances (from other routes) increase more the GDI (e.g. R5 and \hat{r}_5) than parallel geo-paths (e.g. R4 and \hat{r}_4) closer than R2. Therefore we require the order between GDI values as:

$$\hat{r} < \hat{r}_4 < \hat{r}_5 < \hat{r}_2 \quad (1)$$

B. Geographic Diversity Index (GDI)

We model geo-paths of Internet routes as collections of sections between their consecutive nodes. Let us assume two routes defined by their nodes: $P = \{a, b, c, d\}$ and $L = \{e, f, g\}$. Naturally, the distance $\delta(a, L)$ between node a and route L translates to the distance between a and the closest point (not necessarily a node) of route L to a . Let $\Delta(P, L) = \{\forall u \in P | \delta(u, L), \forall u \in L | \delta(u, P)\}$ be a vector containing all possible node distances between P and L . In the toy example of Fig. 3 $\Delta(P, L) = \{\delta(a, L), \delta(b, L), \delta(c, L), \delta(d, L), \delta(e, P), \delta(f, P), \delta(g, P)\}$.

In order to satisfy the requirements listed in Sec. IV-A we define GDI for a given set of routes as follows. First, we define a diversity index between two routes as:

$$d(P, L) = \left(1 - \text{Var}'(\Delta(P, L))\right) \text{Mean}(\Delta(P, L)), \quad (2)$$

where Var' denotes the variance of its arguments normalized to the interval $[0..1]$. Second, we define the diversity between a single route P and a set of routes \mathbb{V} as:

$$\mathcal{D}(P, \mathbb{V}) = \min_{L \in \mathbb{V}} d(P, L) \quad (3)$$

Finally, we quantify the overall GDI for a given set of routes. In order to calculate this, we use a step-by-step method. Let us assume a set \mathbb{V} containing the routes. The process starts from an empty set $\mathbb{U} = \{\}$. In the 0. step we search in set \mathbb{V} for the two routes having the highest diversity score $d_0 = \max_{P, L \in \mathbb{V}} (d(P, L))$ and move paths $\text{argmax}_{P, L \in \mathbb{V}} (d(P, L))$ to \mathbb{U} . In the i -th step we compute $\mathcal{D}_i = \max_{P \in \mathbb{V}} (\mathcal{D}(P, \mathbb{U}))$ and move the path $\text{argmax}_{P \in \mathbb{V}} (\mathcal{D}(P, \mathbb{U}))$ to \mathbb{U} . The process terminates when \mathbb{V} is empty. Finally we compute GDI as:

$$\text{GDI} = d_0 + \sum \mathcal{D}_i. \quad (4)$$

To demonstrate that the proposed method of GDI calculation satisfies the requirements we set, we show the GDIs of the route sets defined in Sec. IV-A (with a grid cell size is ~ 84 km in Fig. 2) in Table I. Note, that the produced GDI values fulfill Eq. 1.

r	2182
\hat{r}_4	3299
\hat{r}_5	3467
\hat{r}_2	3729

TABLE I: GDI values for the geo-diversity example

V. RESULTS

In this section we show the geodiversity results we achieved from the measurement data set. First we show through an illustrative example the difference between the raw `traceroute` outputs and the geographically equivalent routes achievable after applying our route clustering algorithm. Second, we present the compression rates that we were able to attain on the whole measurement set. Finally, we show the GDI results that we calculated for the already clustered route set.




Fig. 3: Distance example

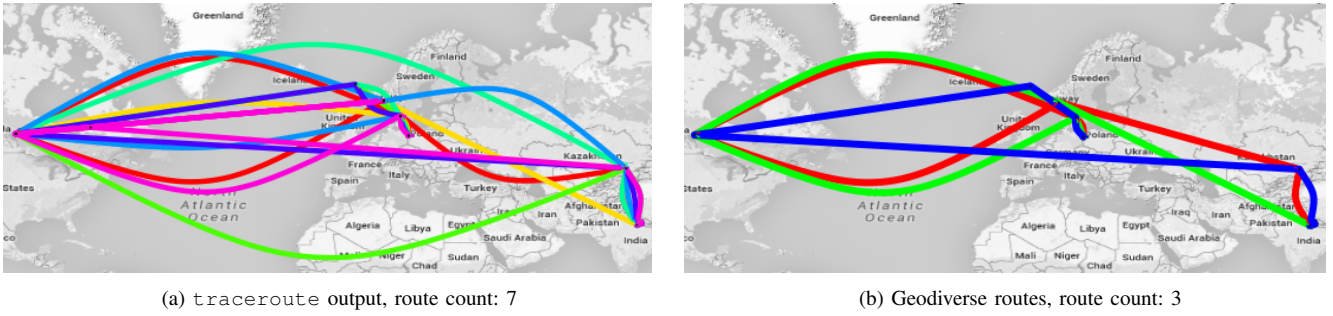


Fig. 4: Route count comparison

A. An example for route clustering

An example of the results of our route clustering algorithm is shown in Fig. 4. Arcs represent hops between the localized IP nodes obtained from `traceroute` output. Note, that arcs do not indicate a real link trajectories, merely distinguish routes on the same intermediate links. On the left-hand side one can observe all the routes that the `traceroute` measurements yielded between a source node located in Poland and a destination node located in India (Fig. 4a). Between these two hosts there exist 7 different routes on the IP-level, but only 3 geographically different routes, i.e., geo-paths (Fig. 4b).

B. Compression ratio of route counts

Stepping up from one example, here we show the overall results in terms of route clustering on the whole measurement data set. We call the ratio of the number of original routes and the resulting clusters as the geo-compression ratio. In Fig. 5 we plot the empirical distribution of this geo-compression ratio for all source and destination host pairs in our data set.




Fig. 5: ECDF of compression ratios for all source-destination pairs

The fraction of source-destination pairs for which the route clustering does not decrease the number of geographically equivalent routes is large (around 80%). In these cases the main reason for the poor compression performance is the fact that only one route is yielded by `traceroute`. Another reason, in less significant number of cases, is that the multiple recorded routes run through nodes at exactly the same locations (Sec. III-B). On the remaining 20% of source-destination pairs we applied our clustering algorithm with remarkable results: this is shown in Fig. 5. In more than 4% of the all cases we could obtain a geo-compression ratio higher than 2.

C. Geo-diversity results

Using the GDI metric that we defined in Sec. IV to characterize the geo-diversity of routes between a source and a destination pairs we show how the calculated values compare to the theoretical maximum of the same metric, taking only the number of geographically equivalent routes and the length of the longest one into account (not their actual trajectories). For this hypothetical maximal value, denoted as MGDI, we place a number of routes forming triangle shapes, the longest one reaching to the highest, so that their GDI would be the largest. In Fig. 6 we plot the distribution of the ratios of GDI over MGDI for those source-destination pairs between which we found at least 2 geographically different routes (20% of all pairs, as mentioned above). The results show that for 80% of these cases the GDI of routes is less than 10% of their MGDI, i.e., the theoretical maximum diversity given the number of routes and the length of the longest route.




Fig. 6: Geo-diversity results given by GDI/MGDI

VI. CONCLUSION

Our goal was to find out where on Earth the packets travel exactly when traffic is carried over the Internet. We discovered that between two given points packet flows are not so scattered as the diversity of *traceroute* results suggest. We actually saw very few host pairs between which more than 1 geographically different routes exist. We showed that 80% of endpoints with more than one route has less than a 10% MGDI value which indicates low diversity in terms of geographic distance. The knowledge we gained from this study about the geographical diversity of Internet routes is useful for several applications. In this section we give a few examples of these, and we discuss the weaknesses of our method.

A. Applications

Estimating bufferbloat: In order to measure Internet delay correctly one must fight many sources of inaccuracy: if one-way latency measurement is possible between two hosts, their clocks must be synchronized, if not, several issues come up: misleading *traceroute* results due to load balancing and MPLS tunnels, different return path of the ICMP_REPLY when using *ping*, etc. Indeed, even if the topology is correctly discovered, many aspects of the actual operation of the network equipment pieces affect the measured delay. In order to somehow infer the impact of bufferbloat from the total delay, a very hot topic nowadays, it is important to have an idea about the propagation delay of the packets. Since the propagation delay is closely related to the traveled geographic distance, the sets of geographically equivalent paths, discovered by our method, provide important input to the analysis of the bufferbloat phenomenon.

Network resilience: Network resilience, in its classical sense, is a well-studied research domain [30]. When network links are not going down individually, but instead are affected en masse due to a regional catastrophe, let it be a natural disaster, a power blackout or an EMP attack, then the geographic topology of the Internet suffers correlated link failures. In order to be ready for this, planning geographic redundancy of Internet paths can use the results of our method as an input.

B. Discussion

After discussing the potential role of our method in various use cases, we turn to the weaknesses of it that we are aware of. First, by applying measurements created by the relatively simple tool *traceroute*, we do not tackle IP-aliasing when building the paths before clustering nodes close-by to each other. One could argue that performing an already documented merging method targeting IP aliases might yield the same result as the geomerging we do. Second, one might question the accuracy of MaxMind, the tool we use to position the nodes. However its accuracy is explored in details in [28] and in long run measurements, we plan to use active measurement based geolocation tools, such as Spotter [31]. Third, it can be argued that estimating geo-paths using straight lines between IP-level nodes may be misleading. However, our city-sized threshold ensures that as long as there are IP-level nodes in close proximity of fiber ducts' ends, this is avoided with high probability. Finally, the proposed GDI and MGDI metrics might seem simplistic, but we argue that to define a diversity metric between routes, many key attributes must be taken into account, and a trade-off between various features and scenarios must be accepted.

REFERENCES

- [1] M. Pathan *et al.*, *A taxonomy of CDNs*. Springer, 2008.
- [2] M. Castro *et al.*, "Topology-aware routing in structured peer-to-peer overlay networks," in *Future directions in distributed computing*. Springer, 2003.
- [3] E.-K. Lua *et al.*, "A survey and comparison of peer-to-peer overlay network schemes." *IEEE Communications Surveys and Tutorials*, vol. 7, no. 1-4, pp. 72–93, 2005.
- [4] A. Greenberg *et al.*, "The cost of a cloud: research problems in data center networks," *ACM SIGCOMM CCR*, vol. 39, no. 1, pp. 68–73, 2008.

- [5] D. Awduche *et al.*, “Overview and principles of internet traffic engineering,” RFC 3272, Tech. Rep., 2002.
- [6] D. Clark *et al.*, “Interconnection in the internet: the policy challenge,” in *Research Conference on Communication, Information and Internet Policy*, 2011.
- [7] J. P. Sterbenz, D. Hutchison, E. K. Çetinkaya, A. Jabbar, J. P. Rohrer, M. Schöller, and P. Smith, “Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines,” *Computer Networks*, vol. 54, no. 8, pp. 1245–1265, 2010.
- [8] J. P. Sterbenz, E. K. Cetinkaya, M. A. Hameed, A. Jabbar, S. Qian, and J. P. Rohrer, “Evaluation of network resilience, survivability, and disruption tolerance: analysis, topology generation, simulation, and experimentation,” *Telecommunication systems*, vol. 52, no. 2, pp. 705–736, 2013.
- [9] J. Rak, *Resilient Routing in Communication Networks*. Springer, Computer Communications and Networks series, 2015.
- [10] C. Kaufman, R. Perlman, and M. Speciner, *Network security: private communication in a public world*. Prentice Hall Press, 2002.
- [11] W. Lou, W. Liu, and Y. Fang, “Spread: Improving network security by multipath routing,” in *Military Communications Conference, 2003. MILCOM’03. 2003 IEEE*, vol. 2. IEEE, 2003, pp. 808–813.
- [12] Y. Schwartz *et al.*, “On the diversity, stability and symmetry of end-to-end internet routes,” in *INFOCOM Workshops*. IEEE, 2010.
- [13] W. Mühlbauer *et al.*, “Building an AS-topology model that captures route diversity,” *ACM SIGCOMM CCR*, vol. 36, no. 4, pp. 195–206, 2006.
- [14] R. Teixeira *et al.*, “In search of path diversity in ISP networks,” in *SIGCOMM IMC*. ACM, 2003.
- [15] Y. Shavitt *et al.*, “Geographical internet pop level maps,” in *Traffic Monitoring and Analysis*. Springer, 2012, pp. 121–124.
- [16] L. Subramanian *et al.*, “Geographic properties of internet routing,” in *Annual Technical Conference*. USENIX, 2002.
- [17] M. Zhang, Y. Ruan, V. S. Pai, and J. Rexford, “How dns misnaming distorts internet topology mapping,” in *USENIX Annual Technical Conference, General Track*, 2006, pp. 369–374.
- [18] S. Laki *et al.*, “A detailed path-latency model for router geolocation,” in *TridentCom*. IEEE, 2009. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4976258
- [19] M. Boguná, F. Papadopoulos, and D. Krioukov, “Sustaining the internet with hyperbolic mapping,” *Nature Communications*, vol. 1, p. 62, 2010.
- [20] Y. Shavitt *et al.*, “A structural approach for PoP geo-location,” in *INFOCOM Workshops*. IEEE, 2010.
- [21] J. Han *et al.*, “An experimental study of internet path diversity,” *IEEE Transactions on Dependable and Secure Computing*, vol. 3, no. 4, pp. 273–288, 2006.
- [22] “Geolocation and online fraud prevention from MaxMind,” <http://www.maxmind.com/>, accessed: 2014-12-22.
- [23] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani, “iplane: An information plane for distributed services,” in *Proceedings of the 7th symposium on Operating systems design and implementation*. USENIX Association, 2006, pp. 367–380.
- [24] CAIDA, “Macroscopic Internet Topology Data Kit (ITDK),” <http://www.caida.org/data/internet-topology-data-kit>.
- [25] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman, “Planetlab: an overlay testbed for broad-coverage services,” *ACM SIGCOMM Computer Communication Review*, vol. 33, no. 3, pp. 3–12, 2003.
- [26] B. Augustin *et al.*, “Avoiding traceroute anomalies with Paris traceroute,” in *IMC*. ACM, 2006. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1177100>
- [27] S. Laki *et al.*, “Spotter: A model based active geolocation service,” in *INFOCOM*. IEEE, 2011. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5935165
- [28] B. Huffaker *et al.*, “Geocompare: a comparison of public and commercial geolocation databases - Technical Report ,” Cooperative Association for Internet Data Analysis (CAIDA), Tech. Rep., May 2011.
- [29] Y. Shavitt and N. Zilberman, “A study of geolocation databases,” *CoRR*, vol. abs/1005.5674, 2010. [Online]. Available: <http://arxiv.org/abs/1005.5674>
- [30] W. Najjar *et al.*, “Network resilience: a measure of network fault tolerance,” *IEEE Transactions on Computers*, vol. 39, no. 2, pp. 174–181, 1990.
- [31] P. Mátray *et al.*, “On the spatial properties of internet routes,” *Computer Networks*, vol. 56, no. 9, pp. 2237–2248, 2012. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S1389128612000989>