

On the Network Geography of the Internet

Péter Mátray, Péter Hága, Sándor Laki, István Csabai, Gábor Vattay
Dept. of Physics of Complex Systems, Eötvös Loránd University, Hungary
Email: {matray,haga,laki,csabai,vattay}@complex.elte.hu

Abstract—The geographic layout of the physical Internet inherently determines important network properties and traffic characteristics. To give insight into the geography of the Internet, we examine the spatial properties of the topology and routing. To represent the network we conducted a geographically dispersed traceroute campaign, and embedded the extracted topology into the geographic space by applying a novel IP geolocalization service, called *Spotter*. In this paper we present the frequency analysis of link lengths, quantify path circuitousness and explore the symmetry of end-to-end Internet routes.

I. INTRODUCTION

One of the fundamental problems of Internet research is to elaborate our knowledge on the network’s topology. Besides the theoretical modeling of the network structure, considerable amount of effort has been devoted to the construction and analysis of Internet maps [1], [2]. Despite the versatile and significant contribution of the related studies, some structural aspects of the Internet still remained hidden. One of the reasons for that is that most works describe the underlying network as a graph and analyze its properties with the tools of graph theory. In this sense, most analyses are carried out in a “topological” space where crucial details about the real world network are neglected. As a consequence, in these models important research questions might become infeasible to answer. Recently, a number of studies have attempted to go beyond the pure topological approach to unfold some previously overlooked facets of network connectivity. For instance, Boguña *et al.* [3] showed that by embedding the AS level topology into a hyperbolic space crucial findings on the network’s nature can be uncovered.

Similarly, by taking into account the geographic relations among network elements, we can uncover fundamental properties of the Internet [4], [5]. Several research and application scenarios could benefit from the deeper understanding of the spatial properties of the topology and traffic. For instance, different AS routing policies might cause major geographic asymmetries of forward and backward paths, and thus can highly affect end-to-end network behavior and potentially downgrade user experience. The large scale analysis of the geographic properties of routing could help identifying such routing anomalies. On the other hand, a detailed geographic view on the Internet would assist the verification of newly emerging Internet topology generators and would be highly beneficial in the development of efficient routing protocols, annotating existing Internet maps and also for the visualization of the results of network measurements.

Despite its inevitable importance, until recently there has been only a handful of quantitative studies dealing with the

spatial properties of the physical Internet. In their pioneering work Subramanian *et al.* [5] used a geolinguistic approach to infer router locations from DNS names. The authors quantified geographical network characteristics related to IP and AS level routing and also investigated geographic fault tolerance properties. The seminal paper of Lakhina *et al.* [4] discusses the geographic location of Internet resources and shows strong, power law-like relationship between population and interface count for economic regions. The paper also investigates the geographic extent of ASes and the differences between inter and intradomain links. The study utilizes two different traceroute-like topology discovery techniques and applies Whois and DNS based mapping services to assign geographic location to the discovered interfaces.

All related works that we are aware of determine the geographic position of Internet resources by means of *registry based* localization techniques (see Section II for details). In contrast, as a novelty in the field we apply a *measurement based* method called *Spotter* [6] to embed the topology into the geographic space. By this, we expect that we achieve a more reliable geographic embedding and thus our derived results reflect a more precise picture of the actual network.

The investigations presented in the paper include the analysis of link level geographic properties, the quantification of circuitousness and asymmetry of Internet routes. For a deeper analysis on the spatial properties of Internet routes please refer to [7]. To enable further analysis from other perspectives we share the data sets used in this study with the research community [8], [9].

II. PINNING DOWN THE TOPOLOGY

To represent the structure of the Internet we use the topology constructed from traceroutes measured between PlanetLab nodes. Obtaining consistent end-to-end routes by means of traceroute is not always a trivial task, as it is discussed by Augustin *et al.* [10]. To ensure that we capture a precise picture of the actual routes that packets follow we collected the traceroutes with fixed source and destination port numbers. The effect of load balanced routing is minimized this way, but the generality of our results is not violated.

In order to cover a dispersed part of the Internet we utilized more than 700 PlanetLab nodes located all over the world. From the collected more than 400,000 traceroute experiments we identified 16,065 distinct IP addresses. As different IP addresses might belong to unique interfaces of the same router the number of individual routers might be less. Based on these data we built the topology of the underlying network, where

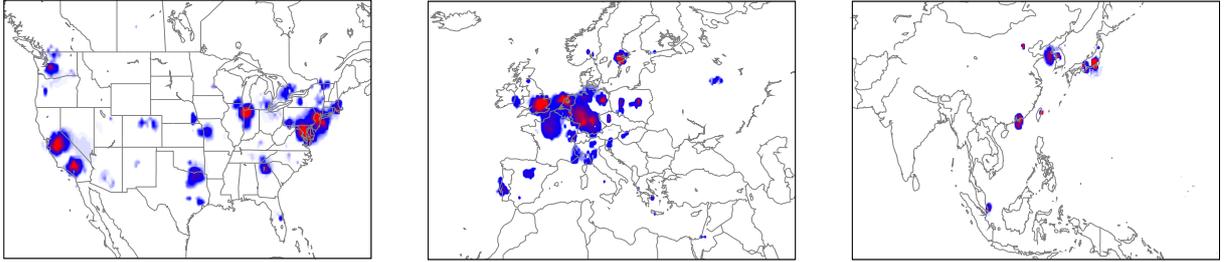


Figure 1. Router density visualization for the collected data set in U.S., Europe and East Asia.

the vertices represent the discovered IP addresses and the edges symbolize the neighborhood relation in the traceroute measurements.

To pin down the network topology onto the globe we need to determine the geographic coordinates of its IP addresses. There are two basic concepts to realize this task. One option is to query *registry based* services like Maxmind [11] or other similar IP geolocation databases. These services usually gain location information from user entered positions or previously registered Whois and DNS data. Whois stores location information only on organization level, which provides unreliable results for geographically dispersed organizations, where the registered street address usually differs from the real location of the hosts. A similar problem arises in the use of DNS names, since the names can be misleading due to the naming conventions of the ISP [12], [13].

In general, registry based methods can provide very accurate location approximations, but due to their inherent limitations they are not truly appropriate for the embedding of non-eyeball IP addresses. In many cases their errors are very large for entire blocks of IP addresses, as it is discussed in [6], [13], [14], and thus the provided resolution might be insufficient for router level analysis.

To overcome the above issues, in this study we use *Spotter* [6], an active geolocation service to map the investigated topology onto the globe. *Spotter* localizes IP addresses individually by measuring delays from PlanetLab landmarks and applying a probabilistic delay-distance model. Consequently, the obtained embedding is solely affected by measured physical properties and do not rely on uncontrolled registry data. According to [6] the median accuracy for router localization is ≈ 30 km which provides sufficient city level precision. We believe that due to these features *Spotter* is more appropriate for large scale router localization.

Spotter successfully provided location estimates for 15,339 distinct IP addresses. The achieved 95% success rate can be considered as high. This is probably due to the speciality of our IP set (i.e. the set contains only those routers that have already replied to traceroute packets). The results presented in this paper strongly depend on the reliability and the accuracy of the location mapping. To minimize the effect of mislocalizations we rerun *Spotter* five times and created a strict test to identify IP addresses without consistent estimation on the city level. We discarded two thousand addresses

that did not pass the consistency test. The remaining 13,120 addresses are considered well positioned. Between these we identified 44,116 links in our extracted topology for which we have reliable location information for both endpoints.

In Figure 1 we illustrate the result of the embedding in three regions of the world. For each target *Spotter* returns a spatial probability density of the possible locations and its expected value as the location estimate. For the sake of clear demonstration the maps show the aggregated probability density for the well positioned routers. The probability values are represented by colors, where red indicates the highest joint probabilities. Due to the aggregated nature of the visualization the map can be interpreted as the router density map of our topology. The colored regions can be easily associated with cities and the red peaks substantiate the intuition that a vast number of routers are hosted in major urban areas. In North America the largest hubs are located in the Northeast Megalopolis, Chicago and California, while in Europe London, Amsterdam, Stockholm and the Frankfurt area attract the largest connectivity. Both maps contain several other regions with significant number of hits. In East Asia the picture is slightly different since we can identify only a handful of important spots including Tokyo, Seoul, Beijing, Hong-Kong and Singapore.

III. CHARACTERIZATION OF NETWORK LINKS

In this section, we investigate the geographic properties of network links. Let r_i denote the i^{th} router in our topology and the link $l_{i,j}$ represent a directed connection between r_i and r_j . The spatial embedding of r_i is given by the latitude and longitude coordinates returned by *Spotter*. We assign a length to $l_{i,j}$ by calculating the $S(r_i, r_j)$ great circle distance between r_i and r_j .¹

We determined S for all the neighboring routers to extract the basic distance properties of the embedded network. Figure 2(a) presents the relative frequencies of different link lengths with 50 km bin size. 50.17% of the 44,000 links fall into the first bin. This indicates that half of the links span urban distances. The other half is not evenly distributed, rather peaks are appearing at certain values. Links having

¹Here, we note that our data set is collected on the IP level, thus hops in the lower network layers might remain hidden and our link notion cannot be directly corresponded to the physical cable connections running between routers.

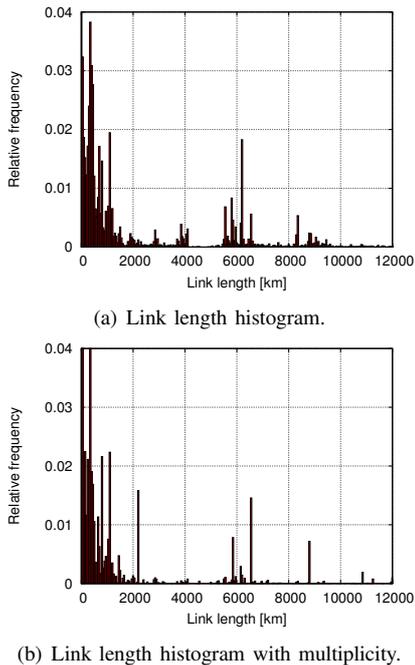


Figure 2. Relative frequencies of link lengths. In Figure (a) each unique link is represented once, while in Figure (b) links are represented with a weight proportional to their prevalence in the traceroute data set.

length between 50 and 1000 km cover the 30.17% of the set. It is not possible to assign these points to a few links, since there are plenty of location pairs in this characteristic range, including Chicago, IL–Washington, D.C., Frankfurt, DE–Copenhagen, DK. Between 1700 and 4100 km we can identify three distinct groups, most of which belong to North American links (e.g. Houston, TX–Los Angeles, CA and Seattle, WA–New York, NY). The gap between 4100 and 5500 km indicates the absence of IP hops in the Atlantic Ocean. Transatlantic connections between US and Europe appear in the 5500 – 6500 km range, while Transpacific links between Tokyo and US West Coast appear above a similar gap from 6500 to 8300 km.

To reveal the significance of different connections in traffic forwarding a weighted variant of the topology is introduced. To assign weights we use edge betweenness, which quantifies the number of shortest paths passing through a link. To approximate the actual betweenness values we extract the number of occurrences of a link in the full mesh traceroute set. The aggregated weight of the topology sums up to 7.6 million, which equals to the number of discovered links with multiplicity. The plot in Figure 2(b) presents the histogram of link lengths weighted up with their multiplicities, and thus places emphasis on the frequent links while playing down the rare ones. Under 2200 km there is no essential change compared to Figure 2(a). However, at 2200 km 121,000 links give rise to a high peak, while up to 4100 km the rest of the links are overshadowed. This indicates that a few backbone links carry the majority of the Transamerican traffic. The longest of these backbone links constitute the peak at

2200 km. A closer look at the data set shows that 99.2% of this peak correspond to connections between Los Angeles, CA and Houston, TX. This nexus plays a significant role for interconnecting the East and West Coast of the U.S., transferring the 1.5% of the overall network traffic in our data set. This significant share is built up from only 39 unique directed links and is almost invisible in Figure 2(a).

A similar phenomenon can be observed around 6000 km, where Amsterdam–New York and Frankfurt–Washington, D.C. pairs dominate the Transatlantic relations with 70 and 251 distinct directed links respectively, pointing out considerable parallelism between the cities. The corresponding total link occurrences are 55,000 and 111,000 demonstrating that network traffic between U.S. and Europe is almost exclusively served by these cities. Since the applied data set is incomplete in its nature, we do not argue that these two city level connections carry all Transatlantic traffic in general. However, we believe that a more complete sampling of Internet routes would yield a similar picture, with only a handful of gateway cities.

Around 8500 km the same shift in details can be recognized. From Figure 2(a) to Figure 2(b) the San Francisco–Tokyo relation at 8350 km technically disappears, which means that its weight is negligible, even though large number of parallel connections exist. Simultaneously, the Los Angeles–Tokyo linkage at 8800 km gains importance with 55,000 occurrences in our traceroute data set. At the periphery of the plot two new peaks appear, corresponding to the New York–Tokyo and the Seattle–Auckland liaisons.

Although both histograms display connections below 12,000 km, there are several links spanning even longer distances. The most significant among these is the Chicago–Taiwan nexus (just above 12,000 km), while the most extreme relation is between Los Angeles and Singapore, constituted from 1493 occurrences of 7 unique links of length 14,100 km. For a more detailed discussion on the distribution of link lengths please refer to [7].

IV. CIRCUITOUSNESS

The routing of Internet traffic is shaped by the routing policies applied by individual ASes on the way. As a consequence, the realized paths typically do not coincide with the shortest paths neither in geographical nor in graph-theoretical sense. The extent of the deviation might highly affect important end-to-end traffic properties such as the delay or jitter.

To quantify the geographic deviation of a P Internet path from the great circle course, in the following we introduce the $C(P)$ relative circuitousness. Let P denote the path from source s to destination d . We represent P with the p_i endpoint routers of its consecutive links: $P = (s, p_1, \dots, p_N, d)$. Here, each p_i corresponds to an r_j router from our extracted topology. The aggregated length of P is defined as the sum of the lengths of its individual links: $L(P) = \sum_{i=0}^N S(p_i, p_{i+1})$, while $G(P)$ denotes the $S(s, d)$ great circle distance between the endpoints of P . By this, $C(P)$ can be formulated as the ratio of the $L(P)$ aggregated length and the $G(P)$ great circle distance. $C(P)$ takes the value of 1 in the theoretical case

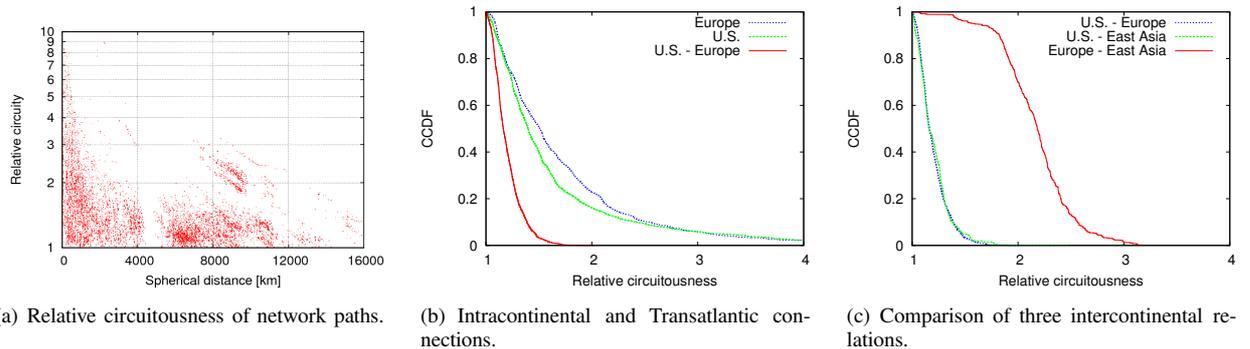


Figure 3. Different aspects of circuitousness. In (a) the relative circuitousness values are presented against the corresponding spherical distances. Figure (b) shows the similarity between the two intracontinental regions and their deviation from the Transatlantic liaison. Figure (c) presents that the Europe–Asia relation has significantly higher deviation values compared to the other two intercontinental liaisons.

when the links of the path exactly follow the great circle course. In practical cases $C > 1$ and its magnitude reflects the extent of deviation.

The analysis of an end-to-end network path from the original traceroute data set is feasible only in case we have location information for all of its hops. Consequently, we filtered out those end-to-end paths that contain even a single unpositioned router, yielding 51,000 fully localized paths. We present three aspects of circuitousness in Figure 3 for the sanitized data set. In Figure 3(a) the derived $C(P)$ values are displayed versus the $G(P)$ spherical distance of the endpoints on a semi-logarithmic scatter plot. Figure 3(b) and 3(c) show the complements cumulative distributions of the relative circuitousness for different geographic regions.

In Figure 3(a) a significant dispersion can be observed for geographic path lengths below 2000 km, while for longer distances the vast majority of the points exhibit moderate relative circuitousness. The observation is also confirmed by Figure 3(b) where a clear separation appears between the two intracontinental (U.S.–U.S., Europe–Europe) and the Transatlantic (U.S.–Europe) relations. Only 3.3% of the Transatlantic routes show relative circuitousness above 1.5, while for the intracontinental routes these fractions are 41.5% and 51.1%. It can be also recognized that deviations in the U.S. and Europe are slightly different. Although we do not have a straightforward explanation for this finding, one possible cause for the increased deviations in Europe might be its different political-administrative structure.

In Figure 3(a) a gap corresponding to the Atlantic Ocean is visible between 4100 and 5500 km. Around 6000 km a compact dense region reflects the fact that every Transatlantic connection incorporates a single link with length comparable with the aggregated length of the path. In this range the effect of AS level routing policies fades away due to the shortage of possible Transatlantic liaisons. A group of paths appear in the 8–10,000 km range containing 2800 end-to-end routes with $C > 1.5$. These paths can be considered significantly deviated. A thorough analysis reveals that mainly connections between Europe and East Asia are responsible for this set. As most of their traffic is routed through the U.S., the distance

that packets traverse is extremely stretched. The reason behind this routing routine is the poor direct connectivity through the continental Asia. Figure 3(c) verifies the argument from a different perspective, where three intercontinental liaisons are represented by their complements cumulative distribution functions of C . The plot shows that the vast majority of the routes between Europe and Asia are highly affected by circuitousness. 96% of these connections have $C > 1.5$, while for the Transatlantic and Transpacific routes these fractions are 3.3% and 4.5%, respectively.

V. ANALYSIS OF GEOGRAPHIC SYMMETRY

The symmetry of Internet routing is in the forefront of interest since Paxson’s seminal paper from 1996 [15]. One of the motivations behind the research is that the presence of asymmetries might highly affect end-to-end network behavior and potentially downgrade user experience. Although numerous papers have analyzed routing asymmetries on both AS and router levels (e.g. [16], [17]) the spatial aspects of this phenomenon have not yet been investigated systematically.

To reveal how frequently the asymmetries manifest on the scale of geographic terms we analyze the symmetry of network path circuitousness. To quantify this we define the function $Y(P, P') = L(P)/L(P')$, where P and P' represent the forward and backward routes between the same endpoints. Without restriction of generality we fix the order of the two endpoints so that $L(P) < L(P')$ is ensured. Hence, we eliminate the duplicate representation of route pairs and obtain an unambiguous description. As a consequence, the values of Y fall between 0 and 1, where the latter indicates a perfect match between the aggregated length of the two directions. Smaller values refer to asymmetric end-to-end paths. The calculation of Y values requires route pairs fully localized in both directions. We filtered out those end-to-end route pairs that did not satisfy this condition and obtained 12,194 forward-backward pairs. The corresponding Y values are presented in Figure 4(a) as a function of the spherical distance of the endpoints. Figure 4(b) shows the complements cumulative distribution of the Y values in a geographic break up. The scatter plot suggests that there are significant asymmetries, especially in

the intracontinental range (0 – 4000 km). Nevertheless, if we examine the prevalence of serious asymmetries in the cumulative distributions, it turns out that around 80% of the route pairs have $Y > 0.9$ and thus these can be considered symmetric. For Transatlantic connections we observe even stronger matches, here 93% of the pairs exhibit high symmetry. This result is due to the limited number of gateways, and it is in accordance with the results of the Section IV.

In case of the connections between East Asia and U.S. or Europe, only 79% of the route pairs have Y values above 0.9. The corresponding cumulative curve suggests that the intercontinental connections of Asia display similar symmetry characteristics as the intracontinental ones in U.S. and Europe. A detailed analysis shows that the higher asymmetries for Asia are caused by a surprising routing behavior. We examined three outlier groups marked as A, B and C on Figure 4(a). These sets can be clearly mapped to a few hundred UK–Hong Kong, California–Hong Kong and California–Singapore routes, respectively. The major asymmetries corresponding to A, B and C are all caused by the very same symptom, namely that forward and backward routes completely differ in their geographic layout. While inbound traffic to Hong Kong and Singapore is routed through Europe and the continental Asia, outbound traffic is headed through Transpacific links both to US and Europe. An interesting consequence of this phenomenon is that in the above cases ICMP ping packets might actually take a full eastward trip around the globe.²

VI. CONCLUSION

In our study we presented an analysis of the geographic properties of Internet links and end-to-end routes. The location estimates are provided by the *Spotter* geolocation service. From the embedded data set we identified the most frequently appearing network links and presented the importance of the few intercontinental gateways. We quantified end-to-end route circuitousness in a geographic break up and found that intracontinental routes exhibit higher geographic path deviations than intercontinental ones. We extended the usual notion of routing symmetry by spatial aspects and showed that the majority of intercontinental routes exhibit moderate asymmetries, since their traffic is served by a handful of gateway cities. We hope that our analyses present important steps to a more accurate and comprehensive representation of the global Internet.

ACKNOWLEDGMENT

The authors thank the partial support of the National Office for Research and Technology (NAP 2005/ KCKHA005), the National Science Foundation OTKA 7779 and 80177, the EU ICT OneLab2 (Grant No.224263) and the EU FIRE NOVI project (Grant No.257867).

²Just like Phileas Fogg did in [18]. While in 1873 the round trip took 80 days, today it takes ≈ 300 ms.

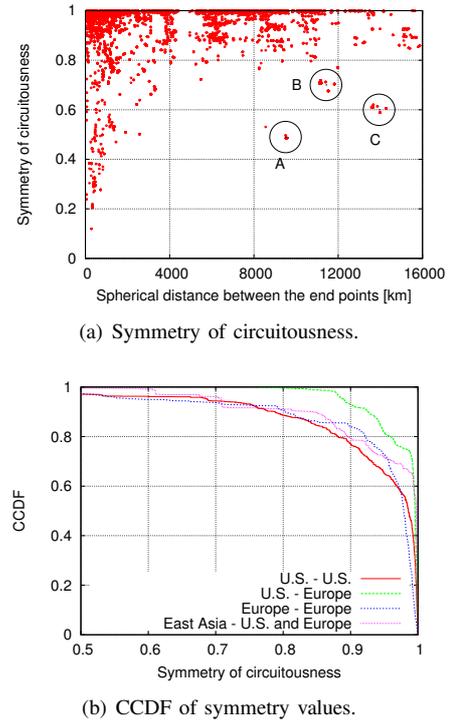


Figure 4. A scatter plot and a complementary cumulative distribution function of the symmetry of circuitousness values.

REFERENCES

- [1] R. Pastor-Satorras and A. Vespignani, "Evolution and Structure of the Internet: a Statistical Physics Approach", Cambridge Univ. Press, 2004.
- [2] CAIDA Internet Topology Research, www.caida.org/research/topology
- [3] M. Boguna, F. Papadopoulos and D. Krioukov, "Sustaining the internet with hyperbolic mapping", *Nature Communications*, 1, 2010.
- [4] A. Lakhina, J. W. Byers, M. Crovella and I. Matta, "On the Geographic Location of Internet Resources", *JSAC*, Vol. 21 Issue 6, p934-948, (2003).
- [5] L. Subramanian, V. N. Padmanabhan, R. H. Katz, "Geographic Properties of Internet Routing", *USENIX 2002 Annual Conference*, 2002.
- [6] S. Laki, P. Mátray, P. Haga, T. Sebok, I. Csabai and G. Vattay, "Spotter: A Model Based Active Geolocation Service", *IEEE INFOCOM 2011*, Shanghai, China, 2011.
- [7] P. Mátray, P. Haga, S. Laki, G. Vattay and I. Csabai, "On the Spatial Properties of Internet Routes", under submission.
- [8] P. Mátray, I. Csabai, P. Haga, J. Steger, L. Dobos, G. Vattay: "Building a Prototype for Network Measurement Virtual Observatory", *ACM SIGMETRICS - MineNet 2007*, San Diego, CA, USA, 2007.
- [9] Data Release, <http://spotter.etomic.org/geodataset>
- [10] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien and R. Teixeira, "Avoiding traceroute anomalies with Paris traceroute", *Internet Measurement Conference*, October 2006.
- [11] MaxMind, <http://www.maxmind.com>
- [12] M. Zhang, Y. Ruan, V. Pai, and J. Rexford, "How DNS misnaming distorts Internet topology mapping", *USENIX Annual Conference*, 2006.
- [13] B. Gueye, S. Uhlig and S. Fdida, "Investigating the Imprecision of IP Block-Based Geolocation", In *Proceedings of PAM 2007*, Louvain-la-neuve, Belgium, LNCS 4427, pp. 237-240, 2007.
- [14] Y. Shavitt and N. Zilberman, "A study of geolocation databases", *arXiv:1005.5674*, 2010.
- [15] V. Paxson, "End-to-End Routing Behavior in the Internet", *Proceedings of SIGCOMM '96*, p.25-38, 1996.
- [16] Y. He, M. Faloutsos, S. Krishnamurthy and B. Huffaker. "On routing asymmetry in the internet." In *GLOBECOM*, 2005.
- [17] Z. M. Mao, L. Qiu, J. Wang and Y. Zhang, "On AS-level Path Inference" in *ACM SIGMETRICS*, Banos, Alberta, Canada, 2005.
- [18] J. Verne, "Around the World in Eighty Days", 1873.