

COMPLEX GEOGRAPHY OF THE INTERNET NETWORK*

ZSÓFIA KALLUS, PÉTER HÁGA, PÉTER MÁTRAY, GÁBOR VATTAY

Department of Physics of Complex Systems, Eötvös Loránd University
Budapest, Hungary

SÁNDOR LAKI

Department of Information Systems, Eötvös Loránd University
Budapest, Hungary

(Received April 12, 2011)

The geographic layout of the physical Internet inherently determines important network properties. In this paper, we analyze the spatial properties of the Internet topology. In particular, the distribution of the lengths of Internet links is presented — which was possible through spatial embedding of a representative set of IP addresses by applying a novel IP geolocalization service, called **Spotter**. The dataset is a result of a geographically dispersed topological discovery campaign. After showing the spatial likelihood of Internet nodes we present two approaches to describe the length distribution of the links. The resulting characterization reveals that the distribution can be separated into three characteristic distance ranges which can be mapped to the regional, transcontinental and intercontinental connections. These regimes follow a power-law function with different exponents.

DOI:10.5506/APhysPolB.42.1057

PACS numbers: 89.75.Da, 89.75.-k, 89.20.Hh

1. Introduction

One of the fundamental problems of Internet research is to elaborate our knowledge on the network's topology [1]. Besides the theoretical modeling of the network structure, a considerable amount of effort has been devoted to the construction and analysis of Internet maps [2, 3]. Despite the versatile

* Presented at the XXIII Marian Smoluchowski Symposium on Statistical Physics, "Random Matrices, Statistical Physics and Information Theory", Kraków, Poland, September 26–30, 2010.

and significant contribution of the related studies, some structural aspects of this global network still remained hidden. One of the reasons for that is that most of the existing work handle the underlying network as a graph and analyze its properties with the tools of graph theory. In this sense, these analyses were carried out in a “topological” space.

When geographic properties of network elements are not neglected anymore, new aspects of the Internet can be revealed [4, 5]. This, in return, opens the door for a wide range of research and application scenarios that could benefit from the deeper understanding of the spatial properties of the topology and traffic.

Despite its crucial role, until recently, there have been only a handful of quantitative studies dealing with the spatial properties of the physical Internet. New studies in this field are primarily limited by the challenges of obtaining reliable geographic location information for Internet resources.

In this work we apply **Spotter**, a newly emerged *measurement based* location estimation service to embed the Internet topology into the geographic space. This approach enables the investigation of the geographical properties of the Internet, such as the characterization of the lengths of Internet links. We extend the few existing empirical results that describe the length distribution of Internet links by determining the characteristic distance ranges of regional, transcontinental and intercontinental connections and presenting the corresponding power-law exponents.

2. Geographic embedding of the network

Spatially embedded networks have emerged and attracted attention in different domains of science during the last decades. Besides theoretically motivated analyses of spatial graphs a wide range of empirical studies have dealt with the quantitative characterization and the theoretical modeling of spatially embedded networks, such as airline, power grid, social or biological networks¹. However, in the case of the Internet, obtaining reliable geographic location information for network nodes is not a trivial task.

We apply a measurement based method to determine geographic location of target nodes. By doing so, we expect to prepare a reliable geographic embedding of the Internet and thus derive results that reflect a correct picture of the actual network.

In this section we describe the measurement method used for obtaining geographic location of network nodes, and the data set collected to represent the Internet for the link length analysis.

¹ For a recent review of spatial networks, please consult [6].

2.1. Theoretical overview of **Spotter**

During the recent years several geolocation techniques have emerged, all of them aim to give an accurate approximation of the location of network hosts which are not known *a priori*.

As presented in this section, the applied, publicly available active geolocation service, **Spotter**, is built on a probabilistic *measurement based approach*.

In the following, we give a brief summary of its basic methodology and its theoretical background based on [7].

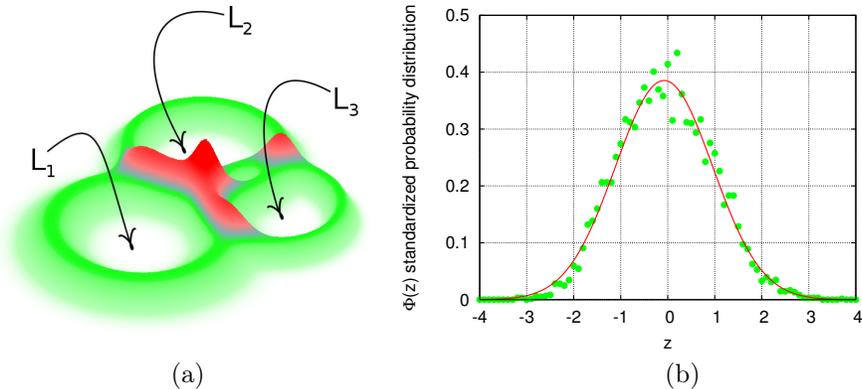


Fig. 1. Target localization from three landmarks. (a) The sum of the probability density surfaces for three landmarks illustrates the ring-like structure of the individual spatial densities. The locations denoted by light gray (green) color are less probable than the dark gray (red) ones where the three rings meet. (b) The probability distribution of the standardized delay-distance data indicates generic behavior and describes the radial profile of the rings.

In the typical scenario of measurement based geolocation we have *landmark* nodes with known geographic location and a *target* node without known position. To approximate the location of the target we measure propagation delays from the landmarks to the target, and then convert the delays into geographic distances based on a delay-distance model. The resulting set of distance constraints is used to determine the target's estimated location with a triangulation-like method [8,9]. The heart of this process is the delay-distance model.

Through a large scale calibration process it became possible to determine a delay-distance probability function reflecting the overall probabilistic characteristics of the Internet as a whole network. This means that by measuring a round-trip delay time between any two of its nodes a probability density function profile, characterizing the distance between these two nodes, (see

figure 1(b)), is identified. With the help of a series of landmark nodes with known geographic coordinates, this probabilistic relation is used to determine a spatial probability density function measured for any of the nodes previously only identified with their IP addresses, but with an unknown geolocation.

Compared to other methods, **Spotter** claims to give localization results with higher accuracy than state-of-the-art measurement techniques [7]. This has also been verified through different performance analyses that confirm the city-level precision of this geolocation service.

2.2. Geographic representation of the Internet

To represent the geographic structure of the Internet we followed a twofold approach. First, we conducted a geographically dispersed topology discovery campaign by utilizing the well-known `traceroute` command. Then we applied the **Spotter** geolocation service (see Section 2.1) to embed the identified routers into the geographic space. To represent the structure of the Internet we use a data set constructed from the previously collected topology data measured between 700 nodes of the PlanetLab research network. We used all the PlanetLab nodes in source and destination roles resulting a collection of more than 400,000 traceroute experiments and identified 16,065 distinct IP addresses. **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).

2.3. Spatial likelihood of Internet nodes

When visualizing the resulting data set, the collective results of multiple probabilistic measurements give rise to different options. The thematic measurement campaigns, where a series of related network interfaces are localized, can easily result in several thousands of individual spatial distributions. Since, at this scale, there is a significant overlap between these spatial distributions, aggregating them into one common geographic map in order to show the arising geographic structure inevitably reduces the amount of information carried by each of the original, individual probability distributions — such as their individual spatial extent, maximum places or weighted centers. This only means that what we can read from a resulting collective map will greatly depend on the method of aggregation, and different methods will be suitable to give the answer to different questions regarding the results of the campaign.

Our intention is to preserve the distribution-like property of the highly accurate geolocation method, so that our results remain correct. For this reason our choice is not to reduce the individual distributions to the mapping of the center points of the probability distributions, so that the resulting maps do not give the false impression of infinite-precision localization results.

Besides the trivial solution of simply assigning the local maximum probability values to each of the cells — that would result in a great amount of information loss — one can choose between several other possibilities.

Hereby we present two different approaches that preserve different information about the results of large scale localization campaigns. Both of them show the regions covered by the localization results. The first kind of map also depicts the “weights” of the determined concentration centers compared to each other — these centers are almost inevitably formed, since targets are mostly organized following the structure of populated places: the urban areas.

On the other hand, the resulting maps of the second method emphasize the locations of target centers — regardless of them being the location of multiple targets — and thus possibly dominating the entire data set — or only hosting a small number of IP addresses, and disappearing beside the highly targeted locations. The high accuracy city-level resolution, that is attainable with the **Spotter** measurement service, enables us to create these kinds of maps, and to reveal the underlying information without loss of the distribution-like characteristics of the individual targets.

For both methods, as an input, there are a number of N individual spatial probability distributions $f(\text{lng}, \text{lat})$ that **Spotter** has returned for N different IP-address localizations. For the K^{th} distribution that means, that

$$\int f_K(\text{lng}, \text{lat})dA = 1,$$

where the integral goes over the surface of the Earth, and the (lng, lat) values represent the geographic coordinates of a point. If the surface of the Earth is divided into a number of M disjoint cells — as in the case of the “hierarchical triangular mesh” (htm) method [8, 9] used by **Spotter** — in the discrete equation this integral becomes a sum of the local integrals over the cells, *i.e.* the sum of “cell probability” values:

$$\sum_{i=1}^M \int_{i^{\text{th}}\text{cell}} f_K(\text{lng}, \text{lat})dA = \sum_i p_{K,i} = 1.$$

The joint probability distribution map

The first method comprises of the formation of a 2-dimensional, histogram-like, aggregated spatial probability distribution. The value assigned to the i^{th} cell is calculated by the simple local summation of the cell probability values assigned to it by each of the overlapping individual distributions. In return, integrating these locally aggregated cell probability values over the entire globe results N , the number of different IP addresses on the map:

$$\sum_{i,K} p_{K,i} = N.$$

This graphic representation can be seen as the joint spatial probability density function of the targets, and by assigning a corresponding color scale to these values, it instantly shows the concentration of well localized target areas. Nevertheless, when dealing with high scale campaigns, this method can still hide the under-represented individual target locations, for example when several partially overlapping distributions hide the original boundaries.

Localizing areas of target centers

This second approach can be useful when one is interested in showing the structure of well localized target centers. In this case we proceed as follows: a transformation of the individual probability distributions will take place prior to the aggregation, and a suitable form of collective distribution will be derived by means of local weighted averages of the transformed distribution values.

Our aim with the transformation is to characterize the territory covered by the original function in a manner emphasizing that the weighted center of the distribution is the *strongest* point of all covered cells — but without discarding the *weaker* ones.

One way to do that is to consider the expected value of error, measured in distance, when choosing a cell over the others — given the cell probability values assigned to each of the cells by the measurement. It results in a function that has a minimum value and is concentricly increasing with the distance from its minimum place. This makes it suitable for our purpose.

This minimum value equals to the average absolute deviation of the spatial distribution, and is assigned to the point with coordinates corresponding to the expected value derived from the spatial probability distribution (the place with the geographic coordinates corresponding to the weighted average position). By minimizing this expected error function at least the cell that is closest to the center will have the minimum value assigned². Thus

² There are cases where the center does not fall into the area, but this does not change the use of the transformation.

this function may not only be regarded as a cost function that one wishes to minimize when choosing between cells covered by the distribution, but it also can localize the center of the distribution and give a measure of how “far away” the possible location cell — with non-zero cell probability — is from the center. Also, the distinction between different distributions of different extents becomes possible: the greater the minimal value, the greater the average extent of the distribution³. This means that for the K^{th} distribution and the i^{th} cell — if it is in the covered area

$$\langle \text{error}_i \rangle_K = \sum_{j \in K} p_{K,j} d(i, j),$$

where $d(i, j)$ stands for the average geographic distance between the i^{th} and the j^{th} cells, and is calculated as the great-circle distance between their centers.

Following this transformation the aggregation method is closed by calculating the local weighted average for each cell, taking into account the relative weights of probabilities assigned to the cell by the different overlapping distributions. For the i^{th} cell

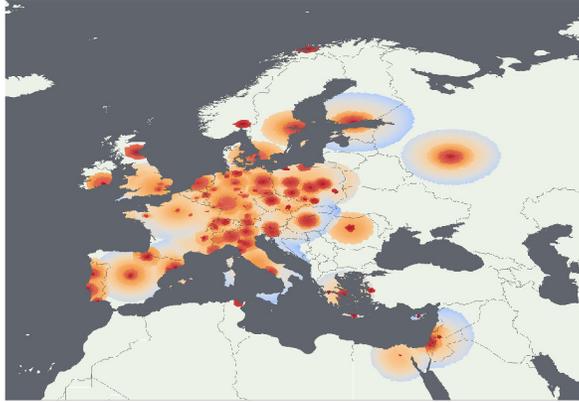
$$\langle \langle \text{error}_i \rangle \rangle = \sum_K \frac{p_{K,i}}{\sum_K p_{K,i}} \langle \text{error}_i \rangle_K.$$

Examples of different regions

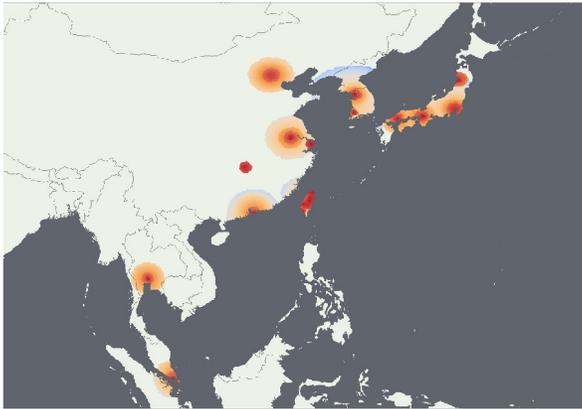
In figures 2 and 3 we illustrate results of the embedding with the center localization process in three different regions of the world. The expected error values are represented by a diverging color map [10]. (The htm cell level resolution of the illustrations is realized by the use of the Quantum GIS software, an Open Source Geographic Information System [11]. The different geographical and cultural aspects of the maps were created with the use of *Natural Earth* vector maps [12].)

We find that the spatial distribution of the PlanetLab router set is strongly correlated with the structure of the urban areas — which is in accordance with the anticipated similarity of geographic distribution of populated places and the Internet. This can be observed in figure 3 where urban areas are indicated for the East-Coast cities of the North of the USA Other urban areas with high target localization density show a qualitatively similar picture.

³ If the transformation was to calculate the expected value of the square of the error distances, this minimum value would be assigned to the center as well, but it would be equal to the sigma of the original spatial distribution.



(a) Inter-PlanetLab routers in Europe.

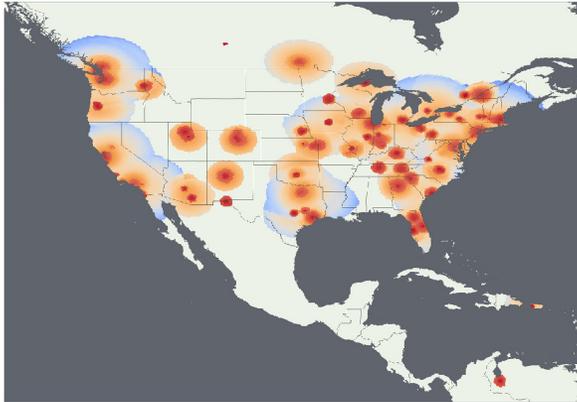


(b) Inter-PlanetLab routers in Asia.

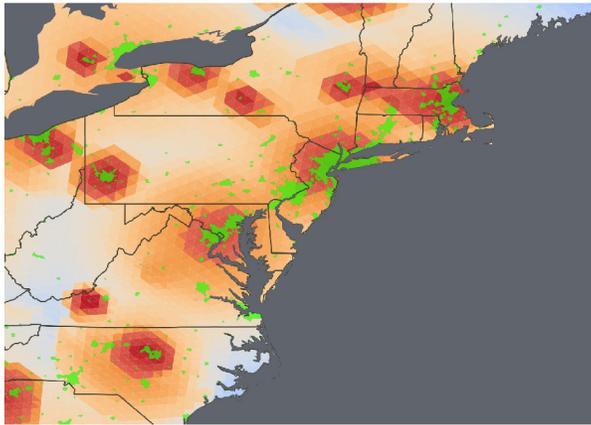
Fig. 2. Likelihood of router positions is shown with well localized centers being colored to darker red shades. On these maps over 11,000 embedded nodes are represented along network paths between PlanetLab node pairs.

3. Modeling the link length distribution

In order to gain insight into the spatial character of the Internet we address the following question: *do the distances between interconnected routers follow specific rules?* To answer this question we introduce the notation r_i for the i^{th} node in our topology and the notation $l_{i,j}$ to 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`, as discussed in the previous sections. By calculating the great-circle distance between neighboring routers we can approximate the length of the IP level links in the topology.



(a) Inter-PlanetLab routers in the USA



(b) Inter-PlanetLab routers in the USA with urban areas indicated.

Fig. 3. Likelihood of router positions is shown with well localized centers being colored to darker red shades. On these maps over 11,000 embedded nodes are represented along network paths between PlanetLab node pairs.

Although insight into the actual distribution of link lengths would provide important input for the refinement of topology models, there is only a very limited number of available papers presenting such results for the Internet. Such knowledge would help the validation of existing models and enable their annotation with link latency values. In the following we analyze the distribution of link lengths from a qualitative perspective to expand our knowledge on the spatial structure of the network.

3.1. Large scale properties of the probability distribution

To provide an overall view on link lengths we plotted their complementary cumulative distribution in figure 4(a). The linear decaying on the log-linear plot indicates logarithmic relationship, implying that the probability distribution can be approximated as

$$P(d) \sim 1/d. \quad (1)$$

This approximation gives a general view on the large-scale geographic structure of the Internet over four orders of magnitude. Our observation is in accordance with the conclusions of [2] where the authors present similar results for a narrower distance range⁴.

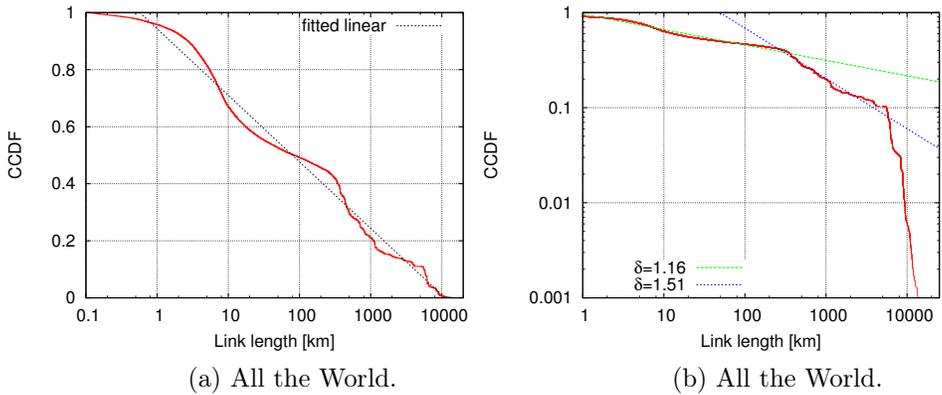


Fig. 4. CCDFs of the link lengths. In figure (a) the log-linear plot indicates a general, logarithmic relationship between the link lengths and their cumulative probability values. Figure (b) depicts the power law behavior on log-log scales.

3.2. Characteristic distance ranges

With an alternative representation of the same data a slightly different character of the distribution can be exposed. Now, in figure 4(b) the data is plotted on log-log scales. The sectional linear decaying on the logarithmic plot suggests that the distribution has three different power-law regimes with different exponents. Accordingly, the probability density can be written as

$$P(d) \sim d^{-\delta}, \quad (2)$$

⁴ The authors of [2] omit distances below ≈ 200 km from their analysis. We suspect that the reason for that is the inability of their applied registry based embedding to resolve shorter distances.

where δ is a regime-dependent parameter. The first breakpoint appears at ≈ 350 km, which can be interpreted as the typical size of an economic region, while the second breakpoint appears at $\approx 5,300$ km which matches well with the typical size of a continent. Consequently, the emerging regimes can be corresponded to *regional* connections (below 350 km), *transcontinental* connections (between 350 and 5,300 km) and *intercontinental* connections (above 5,300 km). The cumulative shares of these ranges are 60%, 30% and 10%, respectively.

We note that a similar separation between the regional and transcontinental scales appears in [4]. The authors investigate the distance-sensitivity of router connectivity and find significantly different behavior below and above ≈ 350 km. We suspect that this phenomenon is not an accidental coincidence, but rather suggestive of some hidden driving forces that shape the network's evolution. Further research could confirm that well-studied graph-theoretic metrics (*e.g.* the degree distribution or the density of the network) also expose distinguishable statistical behavior over the identified regimes. As our traceroute-based dataset is not appropriate for such graph-theoretic investigations, addressing these questions is beyond the scope of this paper.

One would expect that certain factors (*e.g.* geographic, geopolitic or economic constraints), which vary from continent to continent, have a traceable effect on the distribution of link lengths. To verify this we investigate the intracontinental data for USA and Europe, where a sufficient number of links are available for statistical inference. We depict the resulting length distributions in figures 5(a) and 5(b). Here, the same regimes can be observed as

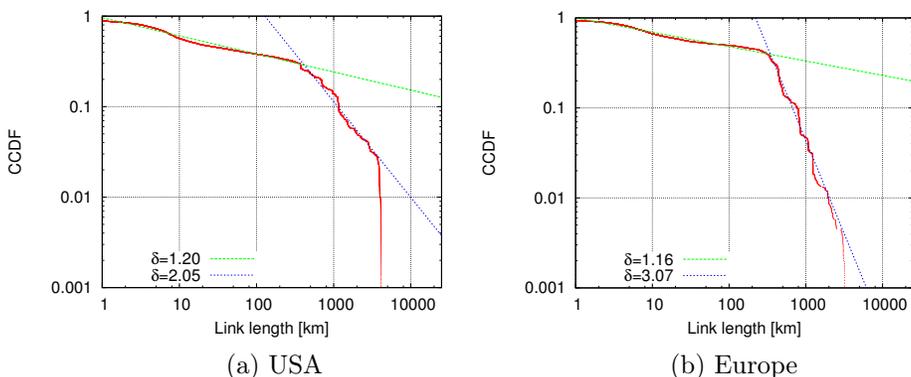


Fig. 5. CCDFs of the link lengths for different regions. Figures (a) and (b) depict the power law behavior on log–log scales for different geographic regions. The exponent of the probability distribution can be calculated by subtracting 1 from the fitted values.

before, certainly without the intercontinental section. In both decomposed distributions the breakpoint appears at ≈ 350 km, while the fitted power-law exponents slightly differ in the USA and Europe. In other words, the share between short and long links is shifted, reflecting structural differences between the two continents. The deeper analysis of this phenomenon is not trivial and requires further research. Table I lists the fitted δ power-law exponents for the regional and transcontinental regimes.

TABLE I

Power-law exponents for distance distributions.

	USA	Europe	All the World
Regional connections	1.20	1.16	1.16
Transcontinental connections	2.05	3.07	1.51

Finally, we emphasize that while (1) provides a good link length model for all scales, a χ^2 analysis shows that for particular scales (2) gives more accurate description of the link length distribution. These descriptions could refine the existing graph-theoretical models of the Internet with fundamental spatial properties comprising geographical, economical and political decisions as well.

4. Conclusion

In this study the analysis of the geographic properties of Internet were presented. Our findings are based on an experimental data collection and the spatial embedding of the representative set of Internet nodes. The location estimates are provided by **Spotter**, a measurement based geolocation service. The location estimates are visualized on geographic maps by their spatial likelihood function which enables the investigation of spatial correlation with urban areas. The length distribution of the Internet links has been investigated, and two approaches are presented for their description. The distribution can be separated into three characteristic distance ranges, identified as *regional*, *transcontinental* and *intercontinental* connections. On these regimes the distribution shows specific power-law function behavior with different exponents. Although the data set used in this study is incomplete, it is our belief that this analysis takes us one step closer to a more accurate representation of the global Internet, and opens the door for further studies of the geographic reality of this abstract network.

This work was partially supported by the National Office for Research and Technology (NAP 2005/KCKHA005), the National Science Foundation OTKA 7779, the National Development Agency (TAMOP 4.2.1/B-09/1/KMR-2010-0003) and the EU FIRE NOVI project (Grant No. 257867).

REFERENCES

- [1] R. Pastor-Satorras, A. Vespignani, *Evolution and Structure of the Internet: a Statistical Physics Approach*, Cambridge University Press, Cambridge 2004.
- [2] S.H. Yook, H. Jeong, A.L. Barabási, *PNAS* **99**, 13382 (2002).
- [3] M. Zhang, Y. Ruan, V. Pai, J. Rexford, How DNS Misnaming Distorts Internet Topology Mapping, in: Proc. of USENIX Annual Technical Conference, 2006.
- [4] A. Lakhina, J.W. Byers, M. Crovella, I. Matta, *IEEE J. Sel. Areas Commun.* **21**, 934 (2003).
- [5] L. Subramanian, V.N. Padmanabhan, R.H. Katz, Geographic Properties of Internet Routing, USENIX 2002 Annual Conference, 2002.
- [6] M. Barthelemy, *Phys. Rep.* **499**, 1 (2011) [arXiv:1010.0302v2 [cond-mat.stat-mech]].
- [7] S. Laki *et al.*, Spotter: A Model Based Active Geolocation Service, in: Proc. of IEEE INFOCOM 2011, April 10–15, 2011, Shanghai, China 2011.
- [8] A. Szalay *et al.*, Indexing the Sphere with the Hierarchical Triangular Mesh, Technical Report, MSR-TR-2005-123, Microsoft Research, 2005.
- [9] The HTM library, <http://skyserver.org/htm>
- [10] M. Kenneth, Diverging Color Maps for Scientific Visualization, Conference Paper, 5th International Symposium on Visual Computing, December 2009.
- [11] Quantum GIS Project, <http://www.qgis.org>
- [12] Natural Earth, naturalearthdata.com