Improving the Accuracy of Measurement-Based Geographic Location of Internet Hosts

Artur Ziviani [a,*], Serge Fdida [b], José F. de Rezende [c], and Otto Carlos M. B. Duarte [c]

[a] *Laboratório Nacional de Computação Científica (LNCC)*
Av. Getúlio Vargas, 333 - 25651-075, Petrópolis, Brazil

[b] *Laboratoire d'Informatique de Paris 6 (LIP6)*
Université Pierre et Marie Curie (Paris 6)
8, rue du Capitaine Scott - 75015, Paris, France

[c] *Grupo de Teleinformática e Automação (GTA)*
COPPE/EE – Universidade Federal do Rio de Janeiro
P.O. Box 68504 - 21945-970, Rio de Janeiro, Brazil

Abstract

Location-aware applications take into account from where the users are accessing and thereby can offer novel functionalities in the Internet. This paper focuses on improving the accuracy of a geographic location service that relies on delay measurements to locate Internet hosts. Host locations are inferred by comparing delay patterns of geographically distributed landmarks, which are hosts with a known geographic location, with the delay pattern of the target host to be located. We deal with two problems that influence the accuracy of the resulting location estimation: (i) the placement of the landmarks and the probe machines that perform the delay measurements; and (ii) how to best measure the similarity between the delay patterns of the landmarks and the one observed for the target host. For the landmark placement problem, we propose a demographic approach to improve the representativeness of each landmark with respect to the hosts to be located. Given a limited number of landmarks, results show that a demographic placement provides closer landmarks and more accurate location estimations for most hosts. Concerning the placement of probe machines, we show that they have to be sparsely placed to avoid gathering redundant data. Furthermore, we define and evaluate three similarity models. Experiments show that other similarity models outperform the commonly adopted Euclidean distance, resulting then in a more accurate geographic location of Internet hosts.

Key words:
Internet host location, measurement, demographic placement, similarity models

Preprint submitted to Elsevier Science *8 September 2004*

1 Introduction

Knowing the geographic location of an Internet host from an identification of that host, such as a name or IP address, enables a whole new class of location-aware applications. These applications take into account where the users are accessing from. Examples of novel location-aware applications are: local advertising on web pages, automatic selection of a language to first display the content, accounting the incoming users based on their positions, restricted content delivery following regional policies, and authorization of transactions only when performed from pre-established locations. In the current Internet, however, there is no direct relationship between host identification and the host physical location. The novel location-aware applications then require the deployment of a geographic location service for Internet hosts.

A DNS-based approach to provide a geographic location service of Internet hosts is proposed in RFC 1876 [1]. Nevertheless, the adoption of the DNS-based approach is restricted since it requires changes in the DNS records and administrators have no motivation to register new location records. Tools such as IP2LL [2] and NetGeo [3] query Whois databases in order to obtain the location information recorded therein and then infer the geographic location of a host. However, if a large and geographically dispersed block of IP addresses is allocated to a single entity, the Whois databases may contain just a single entry for the entire block.

Padmanabhan and Subramanian [4] investigate three important techniques to infer the geographic location of an Internet host. The first technique infers the location of a host based on the DNS name of the host or another nearby node obtained using `traceroute`. For example, the name `bcr1-so-2-0-0.Paris.cw.net` indicates a router located in Paris, France. This technique is the base for GeoTrack [4], VisualRoute [5], and GTrace [6]. The creation and management of parsing rules is, however, a challenging task as there is no standard to follow. As the position of the last recognizable router in the path toward the host to be located is used to estimate the position of such a host, a lack of accuracy is also expected. The second technique splits the IP address space into clusters such that all hosts with an IP address within a cluster are likely to be co-located. An example of such a technique is GeoCluster [4]. Nevertheless, this technique is based on information that may be inaccurate because the databases rely on data provided by users, which may be unreliable to provide correct location information. The third technique is based on delay measurements and the exploitation of a possible correlation between geographic

* Corresponding author, tel.: +55 24 2233-6199, fax: +55 24 2233-6198.
Email addresses: `ziviani@lncc.br` (Artur Ziviani), `Serge.Fdida@lip6.fr` (Serge Fdida), `rezende@gta.ufrj.br` (José F. de Rezende), `otto@gta.ufrj.br` (Otto Carlos M. B. Duarte).

distance and network delay. Such a technique is the base for GeoPing [4]. The location estimation of a host is based on the assumption that hosts with similar network delays to some fixed probe machines tend to be located near each other. Given a set of landmarks with a well-known geographic location, the location estimation for a target host to be located is the location of the landmark presenting the most similar delay pattern to the one observed for the target host.

In this paper, we focus on improving the accuracy of the geographic location estimation of Internet hosts inferred from delay measurements. Hence, we investigate the correlation between geographic distance and network delay. This correlation is weak to moderate if considered worldwide, whereas we show it is stronger in regions with richer connectivity. We use the term rich, or poor, connectivity to represent the variety of connectivity and transit options found in a certain region at both router and autonomous system levels. An environment with rich connectivity is expected to be able to find more geographically straightforward paths from source to destination. Moreover, we identify two key points that influence the accuracy of the Internet host location from delay measurements. The accuracy basically depends on the placement of landmarks and probe machines as well as on how efficiently the similarity between delay patterns is evaluated. Therefore, we aim at improving the accuracy of the host location estimation by: (i) strategically placing landmarks and probe machines [7]; and (ii) selecting models to best measure the similarity between the delay pattern of each landmark and the one of the target host [8].

Landmarks are expected to reflect where most users and hosts are. The number and position of the landmarks are key points for the accuracy of the location estimation and for the impact on network load due to measurements. The problem of finding the best location to place Internet resources, to reduce both the traffic load on networks and the delay perceived by end users, was addressed by Krishnan *et al.* [9] for caches, Qiu *et al.* [10], Radoslavov *et al.* [11], Cronin *et al.* [12], and Bassali *et al.* [13] for mirrors. We use, in this paper, similar concepts to find the best placement of landmarks and probe machines to provide accurate geographic location estimations for the majority of Internet hosts. We propose and evaluate a demographic placement approach that considers the geographic distribution of users, and consequently of hosts to be located, to place landmarks and probe machines. Results show that the demographic placement provides a relatively small number of landmarks able to represent a large portion of users (hosts) within a limited coverage distance. Adopting fewer landmarks to locate a host implies in a lower amount of measurement traffic. We also verify that for a limited number of landmarks, the demographic placement improves the representativeness of each landmark. This improvement results in closer landmarks and more accurate location estimations for the most part of hosts to be located. We also apply our demographic approach to place the probe machines. Probe machines are placed on sites likely to have enough network infrastructure to make their deployment feasible and in a fashion to avoid gathering redundant measurement data.

Another key issue concerning the location estimation accuracy is the evaluation of how similar a landmark and the target host delay patterns are. Measuring the similarity between different items is fundamental for accuracy of recommender systems [14] and pattern analysis [15]. In our case, a similarity model compares the delay patterns and determines the landmark with the most similar delay pattern with respect to the one of the target host. Hence, we investigate different similarity models and evaluate their accuracy for the measurement-based geographic location of an Internet host. Padmanabhan and Subramanian [4] adopt the Euclidean distance as a way to assess the similarity of the observed delay patterns. One of our similarity models considers the previously adopted Euclidean distance and we use this distance as a reference for evaluation. As is further detailed in Section 5.2, Euclidean distance tends to be less robust to violations of the triangle inequality that are present in some parts of the Internet [16–18]. Results show that the city-block distance outperforms the Euclidean distance, thus providing more accurate location estimations of the target host.

This paper is organized as follows. A formalization of the host location inference based on delay measurements is introduced in Section 2. In Section 3, we study the correlation between geographic distance and network delay. Section 4 presents and evaluates the demographic placement proposition. Section 5 defines and compares the similarity models we evaluate. In Section 6, we present our conclusions.

2 Inferring Host Geographic Locations from Delay Measurements

GeoPing [4] infers a host geographic location from delay measurements. In general, a moderate correlation between distance and delay prevents the capture of such a relationship under a mathematical model. Therefore, GeoPing adopts an empirical approach based on the observation that hosts sharing similar delays to other fixed hosts tend to be near each other geographically.

We formalize the problem of inferring a host location from delay measurements as follows. Consider a set $\mathcal{L} = \{L_1, L_2, \ldots, L_K\}$ of K landmarks. Landmarks are reference hosts with a well-known geographic location. Consider a set $\mathcal{P} = \{P_1, P_2, \ldots, P_N\}$ of N probe machines. Fig. 1 illustrates the steps in inferring a host location from delay measurements, which are detailed along this section. Dotted lines represent the measurements taken by the probe machines while the solid lines indicate information exchange. The probe machines periodically

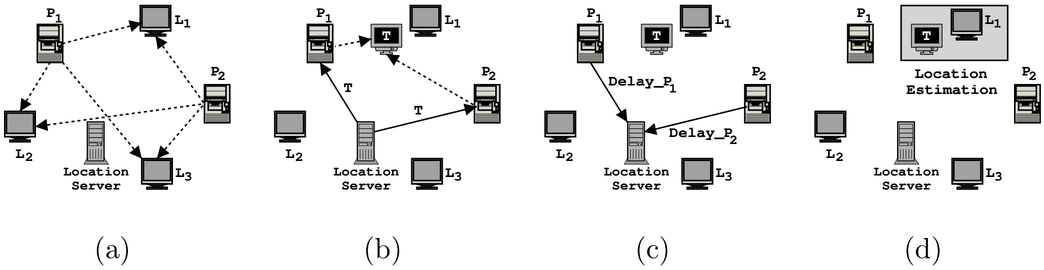

Fig. 1. Inferring a host location from delay measurements.

determine the network delay, which is actually the minimum delay of several measurements, to each landmark (Fig. 1(a)). Therefore, each probe machine P_x, $1 \leq x \leq N$, keeps a delay vector $\mathbf{d}_x = (d_{1x}, d_{2x}, \ldots, d_{Kx})$, where d_{ix} is the delay between the probe machine P_x and the landmark $L_i \in \mathcal{L}$. Suppose one wants to determine the location of a given target host T. A location server that knows the landmark set $\mathcal{L}$ and the probe machine set $\mathcal{P}$ is then contacted. The location server asks the N probe machines to measure the delay to host T (Fig. 1(b)). Each probe machine P_x, $1 \leq x \leq N$, returns a delay vector $\mathbf{d}'_x = (d_{1x}, d_{2x}, \ldots, d_{Kx}, d_{Tx})$, i.e., the delay vector $\mathbf{d}_x$ plus the just measured delay to host T (Fig. 1(c)). After receiving the delay vectors from the N probe machines, the location server is able to construct the delay matrix $\mathbf{D}$ with dimensions $(K+1) \times N$:

$$\mathbf{D} = \begin{bmatrix} d_{11} & d_{12} & \ldots & d_{1N} \\ d_{21} & d_{22} & \ldots & d_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ d_{K1} & d_{K2} & \ldots & d_{KN} \\ d_{T1} & d_{T2} & \ldots & d_{TN} \end{bmatrix} \tag{1}$$

The delay vectors gathered by the demanding location server from the probe machines correspond to the columns of the delay matrix $\mathbf{D}$. The location server then compares the lines of the delay matrix $\mathbf{D}$ to estimate the location of host T. To infer the location of host T, the landmark L presenting the most similar delay pattern with respect to the delay pattern of host T is determined. The corresponding location of the landmark L is the location estimation of host T (Fig. 1(d)). The delay matrix $\mathbf{D}$ combined with the knowledge of the location of the landmarks of the set $\mathcal{L}$ compose a delay map recording the relationship between network delay and geographic location.

The placement problem we deal with involves the number of needed landmarks and where to place such landmarks and probe machines. We are thus interested in where to place a finite number of landmarks to maximize the representativeness of each placed landmark. Fewer landmarks imply a lower amount of measurement traffic injected in the network. Note that probe machines may perform the measurements toward the set of landmarks in a unsynchronized way, avoiding a scalability problem of measuring distances to all hosts at the same time. Furthermore, the initiative of performing measurements is kept at the probe machines to allow the use of oblivious hosts as landmarks. The amount of measurements may be evaluated as follows. Let Δ denote the time interval adopted by the probe machines to periodically gather the delay from the landmarks of the set $\mathcal{L}$. The total number of measurements $\mathcal{M}$ to estimate the location of h hosts in a time interval τ is

$$\mathcal{M}(h,\tau) = 2N\left(\left\lceil\frac{\tau}{\Delta}\right\rceil K + h\right). \tag{2}$$

It should be noted that each measurement may consist of one to several delay samples, but only the minimum value is considered to not take into account delays due to network congestion. In the case we send p `ping` packets to estimate the minimum RTT between a probe machine and a landmark, the amount of measurement traffic injected in the network is actually given by $p\mathcal{M}$.

3 Correlation between Geographic Distance and Network Delay

In order to study the correlation between geographic distance and network delay, we adopt two datasets:

- LibWeb – delay measurements performed from the LIP6 laboratory located in Paris, France, to 135 target hosts with well-known locations all over the world in June 2002. The set of target hosts is mainly composed of university sites extracted from library web (LibWeb) servers around the world [19]. The geographic distribution of the target hosts is as follows: 56 in North America, 44 in Western Europe, 11 in Asia, seven in Eastern Europe, 7 in Latin America, 4 in the Middle East, 3 in Africa, and 3 in Oceania. From the 135 original target hosts, 109 hosts have answered the `ping` requests. The considered delay is the minimum of several measurements to not take into account delays due to congestion in intermediate routers.
- RIPE – data collected from the Test Traffic Measurements (TTM) project of the RIPE network [20]. The dataset we consider is composed by the 2.5 percentile of the delay observed from each RIPE host to each other host in the set during a period of 10 weeks from early December 2002 until February 2003. All 55 hosts on the RIPE network are equipped with a GPS card, thus allowing their exact geographic position to be known. The hosts in the RIPE network are geographically distributed as follows: 42 in Western Europe, 5 in the US, 3 in Eastern Europe, 2 in the Middle East, 2 in Oceania, and 1 in Asia.

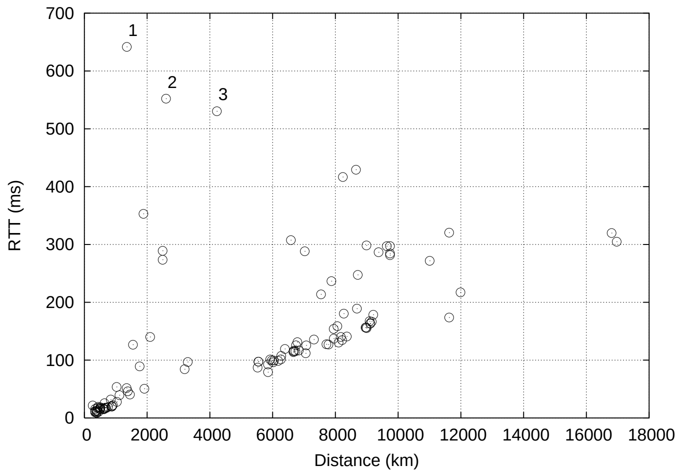

Fig. 2. Scatter plot of geographic distance and network delay (LibWeb).

The geographic locations of the hosts composing both the Libweb and the RIPE datasets are well known. Given the latitude and the longitude of two points, the geographic distance between them is derived using Vincenty's formulae [21].

3.1 Experimental Results for the LibWeb Data Set

Fig. 2 presents the scatter plot of the geographic distance and the minimum delay between our probe machine and each target host for the LibWeb dataset. A weak correlation is observed between geographic distance and network delay worldwide, resulting in a coefficient of correlation $R = 0.2971$. Some points are significantly away from the others. For about the same geographic distance, the observed delay may be greater by one order of magnitude. For example, points 1, 2, and 3 correspond to hosts located in Algeria, Turkey, and Iran, respectively. The `ping` packets from Paris to Algeria (point 1), for instance, actually make their way through routers located in New York. Geographic properties of Internet routing are studied in further detail in [22]. Routes toward some locations take directions far from the straightforward physical direction. Spring *et al.* [23] show that interconnection policies between ISPs directly contribute to end-to-end paths being significantly longer than necessary. This inconvenience is strongly reinforced by the poor connectivity in certain regions.

Poor connectivity weakens the correlation between geographic distance and network delay. From the set of 109 answering hosts, the 80 hosts located in North America and Western Europe have been identified. These regions have the richest connectivity linking their hosts [24–26]. Fig. 3 shows the correlation between geographic distance and network delay considering only the hosts located in North America and Western Europe. The regions that

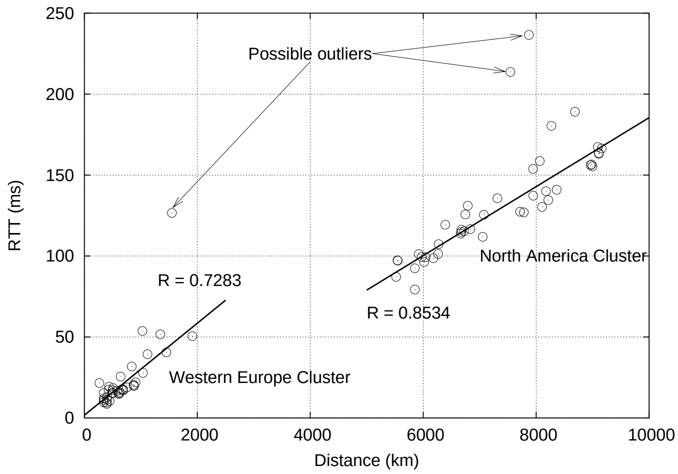

Fig. 3. Correlation between distance and delay (richer connectivity in LibWeb).

have a richer connectivity have a much stronger correlation between geographic distance and network delay. Even if there are a few isolated points away from the other points, two main clusters are observed in Fig. 3. These clusters correspond to hosts located in Western Europe and North America. Excluding the points that clearly remain out of the pack for both clusters (one outlier for the Western Europe cluster and two outliers for the North America cluster) results in the following coefficients of correlation: $R = 0.7283$ for the Western Europe cluster and $R = 0.8534$ for the North America cluster. These outliers result from hosts in regions that are likely to have poorer connectivity than the average of each group. For example, the outliers in North America are located in Hawaii and Alaska in contrast to the remaining hosts in this region that are located in the continental U.S. and Canada.

3.2 Experimental Results for the RIPE Data Set

In contrast with the LibWeb dataset used in the experiment of Section 3.1, where one probe machine gathers delay measurements from 135 landmarks, the RIPE dataset gives us the possibility of considering multiple measurement points. As we have the measured delay from each single host to any other host in the RIPE network, any host in the set may be viewed as a probe machine, as a landmark, or both. It should be noted that RIPE hosts measure the one-way delay between them. This is possible because the RIPE hosts are equipped with GPS cards, allowing them to be enough synchronized. As a consequence, the delays observed between each pair of RIPE hosts are asymmetric, presumably due to asymmetries in routing and in concurrent traffic load. Therefore, we consider individually the viewpoint of each host toward the other hosts in the set. The correlation between geographic distance and network delay within the whole RIPE dataset is shown in Fig. 4. For the whole RIPE dataset, there

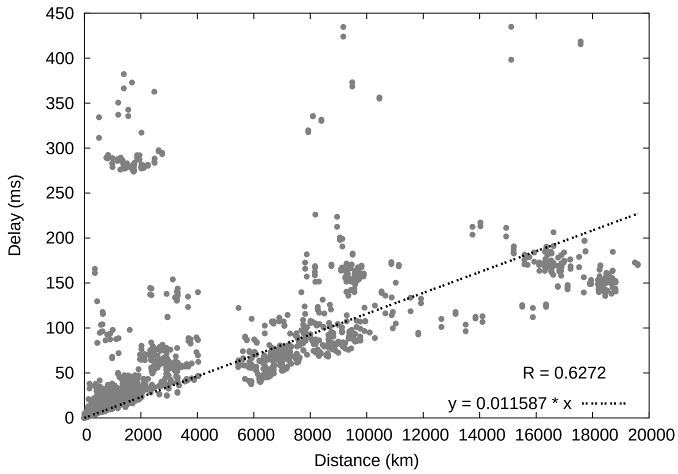

Fig. 4. Correlation between distance and delay for the RIPE dataset.

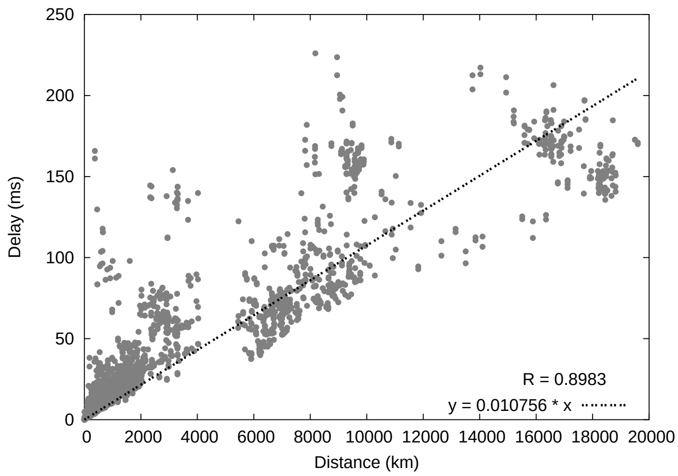

Fig. 5. Correlation disregarding the outlier host in the RIPE dataset.

is a moderate correlation between distance and delay ($R = 0.6272$). Nevertheless, we have observed that one single host in the RIPE network significantly contributes to weaken such a correlation. Fig. 5 presents the correlation for the RIPE dataset without the outlier host. A significantly stronger correlation between distance and delay ($R = 0.8983$) is observed disregarding the outlier host. Performing a `traceroute` toward this outlier host, one single link on the path strongly contributes to increase the end-to-end delay. In accordance with the experiments presented in Section 3.1, we observe once more that a richer connectivity strengths the correlation between geographic distance and network delay.

From our study we observe a weak to moderate correlation between geographic distance and network delay worldwide. Similar results are presented in [27]. Nevertheless, we observe a much stronger correlation on regions with richer connectivity, indicating that the correlation becomes stronger as connectivity becomes richer. Furthermore, within regions with richer connectivity, we observe that hosts located in nearby areas present similar delays to fixed measurement points. Therefore, we verify the assumption that hosts with similar network delays to some fixed probe machines tend to be located near each other. Recent findings [24,25] indicate a strong correlation between population and router density in economically developed countries. Moreover, most users, and consequently most hosts to be located, are likely to be in regions presenting richer connectivity, whereby a stronger correlation between distance and delay holds.

Two main aspects contribute to the robustness of the host inference from delay measurements against factors that may weaken the correlation between distance and delay. First, delay is measured from multiple geographically distributed probe machines rather than from one single location. Second, the minimum delay, among several delay samples, is considered rather than an individual delay sample.

4 Demographic Placement of Landmarks and Probe Machines

A landmark is a reference to be used as location estimation of a certain quantity of hosts supposed to be nearly located. As a landmark is a reference for the geographic position of a set of nearby hosts, it should ideally represent the position of as many hosts as possible. Areas with high host density with a co-located landmark provide location estimations that reflect more accurately the positions of a large number of hosts. Therefore, landmarks are expected to indicate where most users (and hosts) are. A strong correlation between population and router density is found in economically developed countries and in urban areas characterized by population density peaks [24,25]. Most users, and consequently most hosts to be located, are likely to be in regions with rich connectivity. We propose a demographic placement approach to address the issue of placing landmarks and probe machines. In this approach, we place landmarks and probe machines according to the user (host) population distribution.

The main urban agglomerations spread worldwide are considered since they offer the highest concentration of users (hosts to be located). We consider all

Table 1
Top 10 urban agglomerations worldwide.

Rank	Urban Agglomeration	Country	Total Population
1	Tokyo	Japan	34,900,000
2	New York	USA	21,600,000
3	Seoul	South Korea	21,150,000
4	Mexico City	Mexico	20,750,000
5	São Paulo	Brazil	20,250,000
6	Bombay	India	18,150,000
7	Osaka	Japan	18,000,000
8	Delhi	India	17,150,000
9	Los Angeles	USA	16,800,000
10	Jakarta	Indonesia	15,850,000

urban agglomerations with more than one million inhabitants [28], totalizing 407 locations. For illustration, Table 1 shows the top 10 urban agglomerations worldwide. It is known that the Internet infrastructure varies dramatically across different regions throughout the world. Therefore, we weight the populations of the different agglomerations with the number of Internet users in the country the agglomeration belongs to over the total population of the country. In applying such a weight, we estimate the main user agglomerations worldwide to be covered by the demographic placement. Table 2 presents the top 10 user agglomerations worldwide out of the total 407 considered agglomerations. Data on estimations of the Internet users and total population of each country are available in [29]. Denoting the set of user agglomerations as $\mathcal{A}$, the landmark placement problem can now be formally defined. The set $\mathcal{A}$ represents the 407 user agglomerations, totalizing 173,696,253 estimated users [28,29]. The cumulative user distribution over $\mathcal{A}$ is presented in Fig. 6(a). The mean distance between each pair of elements in $\mathcal{A}$ is 8167 km. The cumulative distribution of such a distance is shown in Fig. 6(b). The landmark placement problem is to find a set of landmarks $\mathcal{L} \subseteq \mathcal{A}$ with K landmarks subject to an optimization condition $f(\mathcal{L})$. The possible locations of landmarks are limited to the locations of the agglomerations ($\mathcal{L} \subseteq \mathcal{A}$). From the viewpoint of the optimization condition $f(\mathcal{L})$, the solution to the placement problem under this restriction may be inferior to a solution that allows landmarks to be placed anywhere, i.e. not necessarily on the locations of the agglomerations. Allowing landmarks to be placed anywhere might provide geographically closer sites to place landmarks with respect to the fixed set of agglomerations $\mathcal{A}$. The resulting locations, however, might include sites with no network infrastructure such as rural areas, forests, oceans, and alike. Limiting the placement of landmarks to the locations of the agglomerations is thus justified as it is suitable to place

Table 2
Top 10 user agglomerations worldwide.

Rank	User Agglomeration	Country	User Population
1	New York	USA	11,496,837
2	Los Angeles	USA	8,941,984
3	Tokyo	Japan	7,449,579
4	Seoul	South Korea	6,755,013
5	Chicago	USA	5,003,253
6	Washington	USA	4,178,248
7	London	UK	3,868,029
8	San Francisco	USA	3,858,892
9	Osaka	Japan	3,842,190
10	Philadelphia	USA	3,353,244

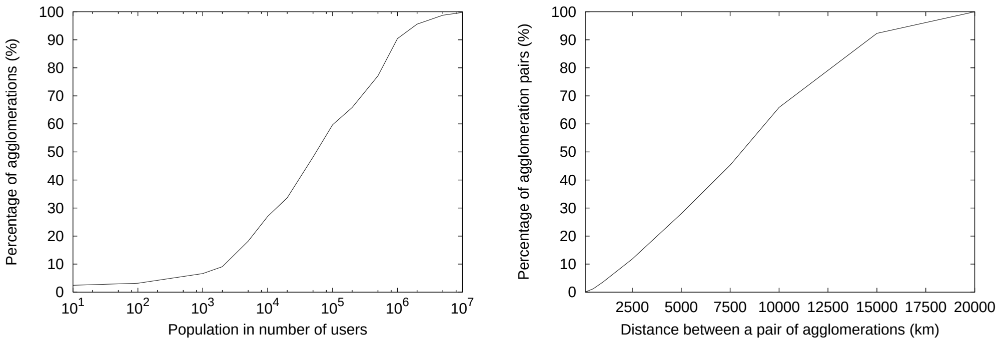

(a) User distribution (b) Distance distribution

Fig. 6. Cumulative user and distance distributions over the set $\mathcal{A}$.

landmarks where is likely to exist a previous network infrastructure. Furthermore, agglomerations with a high level of network infrastructure provide a large set of hosts able to be used as landmarks.

Eligible landmarks are any host able to echo `ping` messages and known to be located within the indicated agglomerations, thus resulting in a potentially large number of eligible landmarks in large user agglomerations. Such landmarks may even be unsuspecting participants in the procedure. As a consequence, for each considered urban agglomeration, a set of oblivious hosts may be selected to constitute a pool of eligible landmarks. As probe machines need to use one landmark from a certain urban agglomeration, such probe machines pick one of the eligible landmarks located in the urban agglomeration. The association between a geographic location and the pool of eligible landmarks located therein may be carried out using DNS, performing load balancing among co-located landmarks. Each pool of landmarks might be named `new-york.netlandmark.info`, `london.netlandmark.info`, and so forth rather than adding IP addresses into the code of probe machines, thus enhancing robustness against individual landmark failures.

We consider two complementary approaches to determine the placement of the set of landmarks $\mathcal{L}$ of size K taking into account the concentration of users in $\mathcal{A}$. Such approaches have been used to determine the placement of different kinds of facilities like fire stations, hospitals, or police departments [30]. In the first approach, given the maximum distance from an agglomeration to a landmark, we want to determine the minimum number of landmarks needed to cover all agglomerations. If the condition of covering all agglomerations is relaxed, a smaller number of landmarks may cover a large portion of the considered space of users as the user distribution is unequal throughout the agglomerations. Thus, given the maximum distance between an agglomeration and a landmark, we want to know how many landmarks are needed to cover at least a certain portion of the considered users, if not all. This model is known as the maximum covering location model [31]. In the second approach, after fixing the number of landmarks (K) to be located, we minimize the maximum distance from any agglomeration to its nearest landmark. This problem is known as the K-center problem [31].

The placement of probe machines is also investigated. We apply the demographic approach to place probe machines on sites with enough network infrastructure to make their deployment feasible and in a fashion to avoid gathering redundant measurement data. Table 3 presents the adopted notation to model the problem of placing landmarks and probe machines.

4.1 Maximum Covering Location Model

The maximum covering location model is obtained when the number of covered hosts to be located is maximized, considering a limited number of landmarks and a fixed coverage distance of each landmark. A landmark covering a certain region is a location estimation for the hosts within that region. The demographic placement of landmarks takes into account the concentration of users within the agglomerations. Using the notation from Table 3, the maximum covering location model is expressed by the objective function

$$f_1(\mathcal{L}) = \max \sum_i h_i Z_i \tag{3}$$

Table 3
Adopted notation for placing landmarks and probe machines.

g_{ij}	geographic distance between agglomerations i and j
h_i	number of users at agglomeration i
G	geographic coverage distance
W	maximum distance from any agglomeration to the nearest landmark
M	minimum distance between any pair of placed probe machines
a_{ij}	$\begin{cases} 1 & \text{if agglomeration } i \text{ can cover demands at agglomeration } j, \\ 0 & \text{if not.} \end{cases}$
X_i	$\begin{cases} 1 & \text{if there is a landmark on agglomeration } i, \\ 0 & \text{if not.} \end{cases}$
Z_i	$\begin{cases} 1 & \text{if agglomeration } i \text{ is covered,} \\ 0 & \text{if not.} \end{cases}$
Y_{ij}	$\begin{cases} 1 & \text{if agglomeration } i \text{ is assigned to a landmark at site } j, \\ 0 & \text{if not.} \end{cases}$
Q_i	$\begin{cases} 1 & \text{if there is a probe machine on agglomeration } i, \\ 0 & \text{if not.} \end{cases}$

with the following constraints:

$$Z_i \leq \sum_j a_{ij} X_j \qquad \forall i \tag{4}$$

$$\sum_j X_j \leq K \tag{5}$$

$$X_j = 0, 1 \qquad \forall j \tag{6}$$

$$Z_i = 0, 1 \qquad \forall i \tag{7}$$

The constraint (4) states that users at the agglomeration i are covered if at least one site that covers agglomeration i is selected to host a landmark. The constraint (5) stipulates that we locate no more than K landmarks. The constraints (6) and (7) are the constraints of integrality for the decision variables X and Z.

We adopt a standard greedy approach [31] to obtain the maximum coverage with time complexity $O(|\mathcal{A}|^2 K)$. After placing the first landmark to cover the most uncovered demand, the algorithm greedily looks for the best location for the next landmark until K landmarks are placed.

Fig. 7 shows the percentage of covered users achieved for three policies of landmark placement and coverage distances of 10 and 500 km. The used policies were: random placement, geographic placement, and demographic placement. Under random placement, the landmarks are randomly selected, disregarding

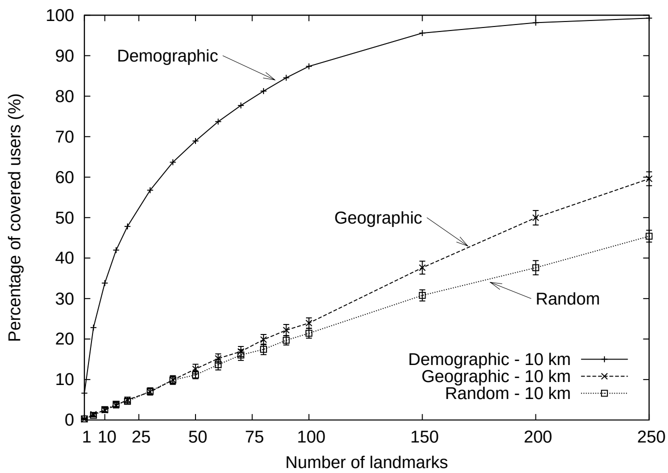

(a) $G = 10$ km

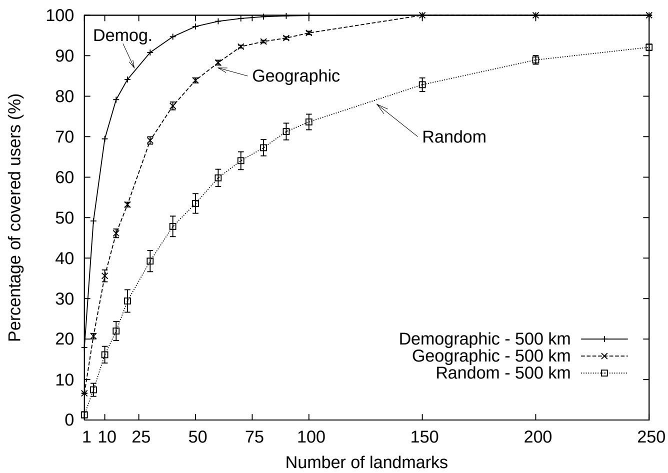

(b) $G = 500$ km

Fig. 7. Covered users in demographic, geographic, and random placement.

both the concentration of users and the proximity between agglomerations. Geographic placement disregards the user concentration ($h_i = 1, \forall i$) within the candidate agglomerations to locate landmarks. The number of agglomerations within the range given by the coverage distance G around an agglomeration determines the weight of the agglomeration. Among agglomerations with equal weight, the elected agglomeration to locate the landmark is randomly chosen. Error bars in the results from the random and geographic placement represent the 99% confidence interval. The demographic placement of landmarks significantly improves the representativeness of the chosen landmarks in terms of hosts to be eventually located when compared with the random and the geographic placement policies.

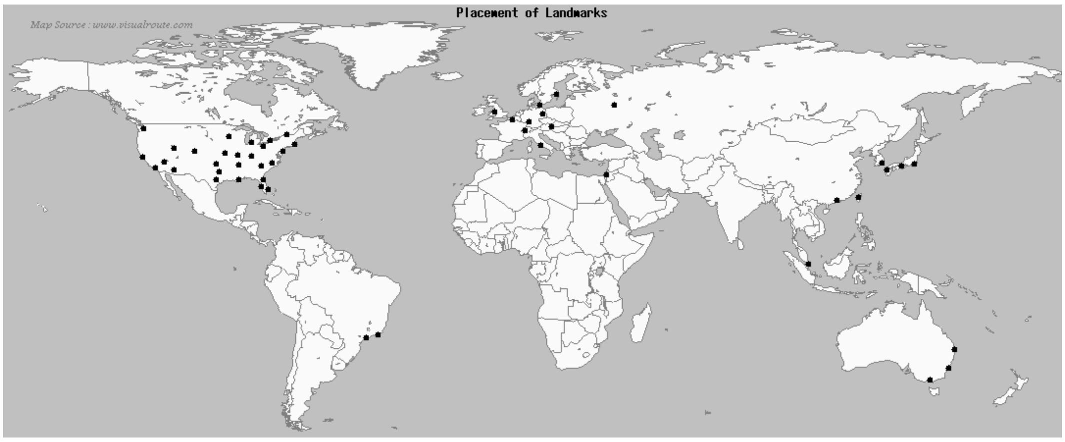

Fig. 8. Placement of 50 landmarks for a coverage distance of 250 km.

The proposed methodology provides not only the number of needed landmarks to cover a certain percentage of users, but also the specific geographic locations of such landmarks. The geographic distribution of the landmarks indicated by the demographic placement approach is visualized using the GeoPlot tool [32]. Fig. 8 illustrates the demographic placement of 50 landmarks to cover almost 90% of the considered number of users for a coverage distance of 250 km. In Fig. 8, landmarks are concentrated in three main regions: USA, Western Europe, and Japan. Such results are consistent with recent findings [24,25] that show a high density of Internet infrastructure on such regions. The high density of placed landmarks on these three regions reflects the concentration of users therein. There are some areas of relatively high user concentration elsewhere, as in Brazil and Australia. Results indicate where to place landmarks to improve the representativeness of each landmark given the constraints on the accuracy and the number of landmarks to be placed.

4.2 K-center Problem

We now tackle the problem of minimizing the maximum distance between an agglomeration and the nearest landmark while considering a fixed number of landmarks. One can notice that even if an agglomeration i is within the coverage distance of a landmark j ($a_{ij} = 1$), it may be assigned to another closer landmark k ($Y_{ij} = 0$ and $Y_{ik} = 1$). Therefore, $a_{ij} \geq Y_{ij}$ because there is no reason to assign an agglomeration to a landmark other than the closest one. This minimization problem is known as the K-center problem [31] and may be formulated using the notation defined in Table 3 by the objective function

$$f_2(\mathcal{L}) = \min W \tag{8}$$

and it is subject to the following constraints:

$$\sum_j X_j = K \tag{9}$$

$$Y_{ij} \leq X_j \qquad \forall i, j \tag{10}$$

$$W \geq \sum_j g_{ij} Y_{ij} \qquad \forall i \tag{11}$$

$$X_j = 0, 1 \qquad \forall j \tag{12}$$

$$Y_{ij} = 0, 1 \qquad \forall i, j \tag{13}$$

The constraint (9) stipulates that K landmarks are to be located. The constraint (10) states that agglomeration i can only be covered by site j if site j has a landmark. The constraint (11) defines the lower bound on the maximum distance, which is being minimized. In other words, the maximum distance (W) must be greater or equal than the distance between any agglomeration i and the landmark j to which the agglomeration i is assigned. The constraint (12) is the integrality constraint on the decision variable X. The constraint (13) requires the agglomeration i to be assigned to only one landmark j.

In the demographic approach, in order to consider the concentration of users in each agglomeration, constraint (11) should be replaced by $W \geq h_i \sum_j g_{ij} Y_{ij}, \forall i$. Otherwise, the minimization problem leads to a geographic placement of landmarks. One solution to the problem is to enumerate each possible subset of size K out of the $|\mathcal{A}|$ candidate locations and then verify which one provides the minimum value of W. Nevertheless, even for moderate values of $|\mathcal{A}|$ and K such an enumeration is not realistic, as the K-center problem is known to be NP-Complete [33]. An alternative binary search algorithm [31] to approximately solve the K-center problem may be outlined as follows:

$G_{\max} \leftarrow \max_{i,j}\{g_{ij}\}$
$G_L \leftarrow 0$; $G_H \leftarrow G_{\max}$
while $(G_H \neq G_L)$ **do**
    $G \leftarrow \lfloor (G_H + G_L)/2 \rfloor$
    Compute the smallest set of landmarks $\mathcal{L}(G)$ that covers the entire set of agglomerations $\mathcal{A}$ for a coverage distance of G.
    if $(|\mathcal{L}(G)| \leq K)$ **then**
        $G_H \leftarrow G$
    else
        $G_L \leftarrow G + 1$
    end if
end while
G_L is the solution to the objective function and $\mathcal{L}(G_L)$ provides the locations of the K landmarks for the solution.

Using the binary search approach, we are able to solve the K-center problem with time complexity $O(|\mathcal{A}|K^2 \log G_{\max})$. This approach defines a lower bound G_L and an upper bound G_H on the maximum distance between an agglomeration and the nearest landmark. The approach then successively narrows the range between such bounds until they converge into the smallest coverage distance that allows a set of K landmarks to cover all agglomerations. Such a coverage distance is the smallest maximum distance between an agglomeration and the nearest landmark for K placed landmarks, thus being the approximate solution of the objective function (8). The set of K landmarks that covers all agglomerations for such a coverage distance is the solution to the K-center problem. The binary search algorithm considers unweighted distances, resulting in a geographic placement of landmarks. In the demographic strategy, to consider demand-weighted distances, step 1 in the binary search algorithm should state $G_{\max} \leftarrow \left\lceil \max_{i,j}\{g_{ij}\}\right\rceil \left\lceil \max_i\{h_i\}\right\rceil$. Furthermore, in computing the smallest set of landmarks that covers the entire set of agglomerations (step 5 in the binary search algorithm), an agglomeration j is able to cover an agglomeration i if $g_{ij}h_i \leq G$.

We compare the results from the random, geographic, and demographic placement policies in Fig. 9. The geographic placement performs the best, providing the lowest maximum geographic distances between agglomerations and the nearest landmark. The demographic approach provides higher maximum and average distances, but such results mask the concentration of users within the agglomerations. Adding more landmarks may even keep the maximum distance observed under the demographic placement policy unchanged. A remote agglomeration with low user concentration may be kept far from the nearest landmark as additional landmarks are used to decrease the distance between more user-populated agglomerations to the nearest landmark. One observes such a situation in Fig. 9(a) between 50 and 150 placed landmarks as the maximum distance keeps leveled off. In spite of that, as shown in Fig. 9(b), the average distance to the nearest landmark keeps decreasing in this same range as the density of landmarks in denser user areas increases.

The demographic placement considers the user concentrations within the different agglomerations. As a consequence, the demographic placement pushes the worst-case distances toward the farthest and least user concentrated agglomerations. High user concentrated agglomerations are assigned to closer landmarks and the agglomerations with the highest user concentrations host the landmarks. These results are presented in Fig. 10 using 50 and 100 placed landmarks. The demographic strategy provides smaller distances between the agglomerations and their nearest landmark for the most part of users at the expense of leaving farther landmarks for the remote and least user concentrated agglomerations.

Fig. 11 illustrates the difference in adopting either unweighted (geographic

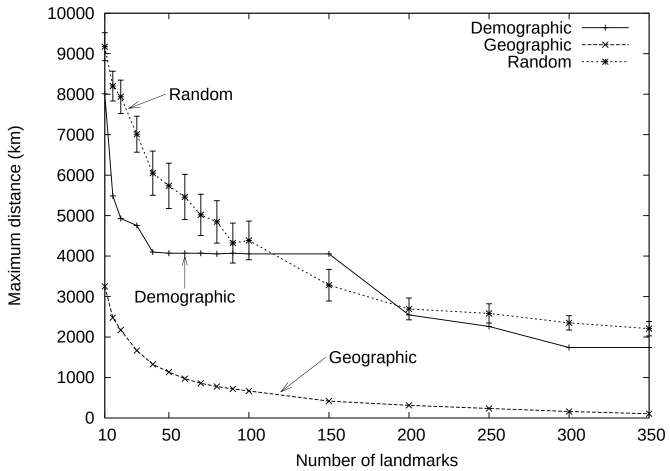

(a) Maximum geographic distance

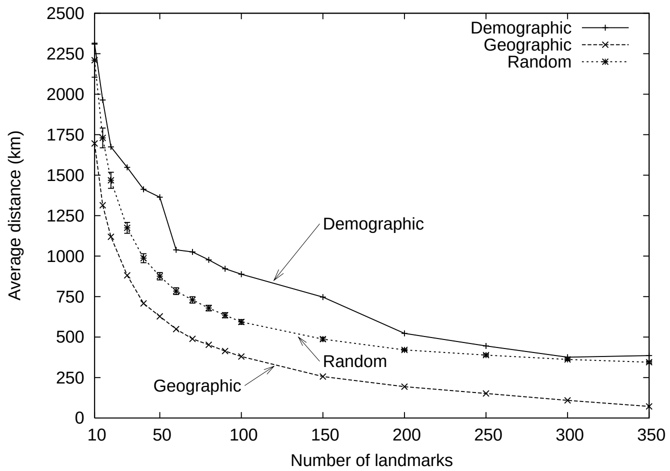

(b) Average geographic distance

Fig. 9. Distance from an agglomeration to the nearest landmark (km).

placement) or weighted (demographic placement) distances to be minimized in the K-center problem. Fig. 11(a) presents the placement of 50 landmarks considering the geographic placement, thus disregarding the concentration of users. Such a disregard of the concentration of users leads to a more uniform geographic distribution of the 50 landmarks worldwide and to a lower maximum unweighted or purely geographic distance W between any agglomeration and its nearest landmark. In the other hand, Fig. 11(b) presents the placement of 50 landmarks considering weighted distances under the demographic placement. The maximum (unweighted or purely geographic) distance W between an agglomeration and the nearest landmark is higher, but the worst cases are pushed to remote agglomerations with the lowest concentrations of users. The

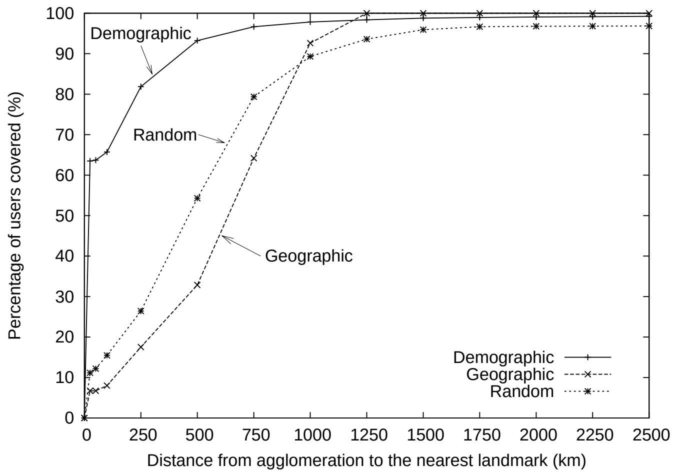

(a) $K = 50$

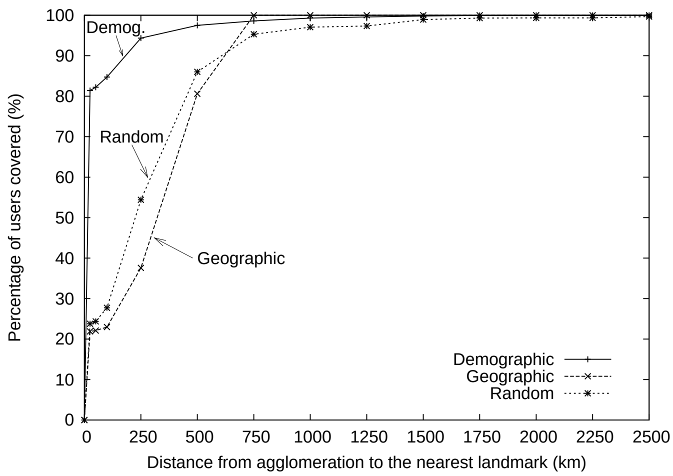

(b) $K = 100$

Fig. 10. Covered users as a function of the distance to the nearest landmark.

result is a dense placement of landmarks in the regions presenting the highest concentrations of users like the USA, Western Europe, and Japan.

4.3 Placement of Probe Machines

In this subsection, we address the issue of where to place the probe machines. Probe machines measure the delay to a target host and regularly gather the delays to landmarks. Probe machines that are near each other may share common paths to some remote landmarks or target hosts, thus providing redundant information to the location estimation decision. Therefore, our first

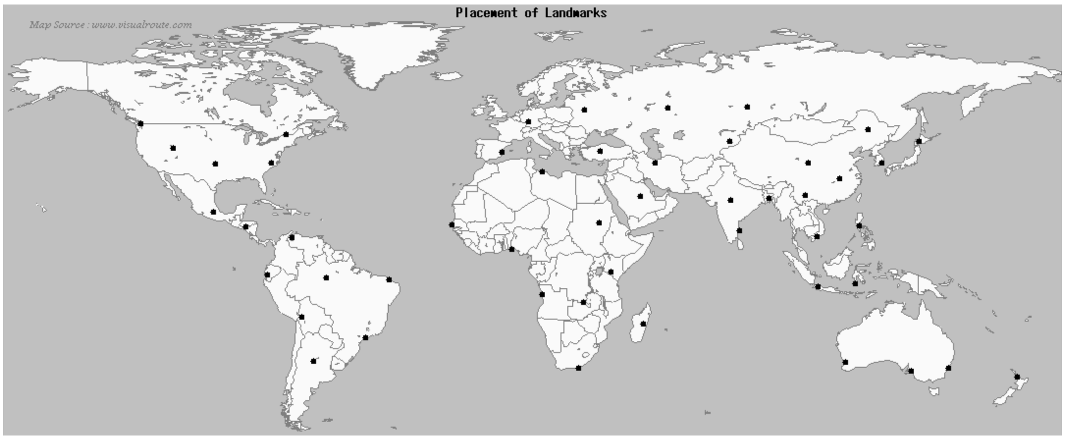

(a) Geographic placement (disregards the concentration of users)

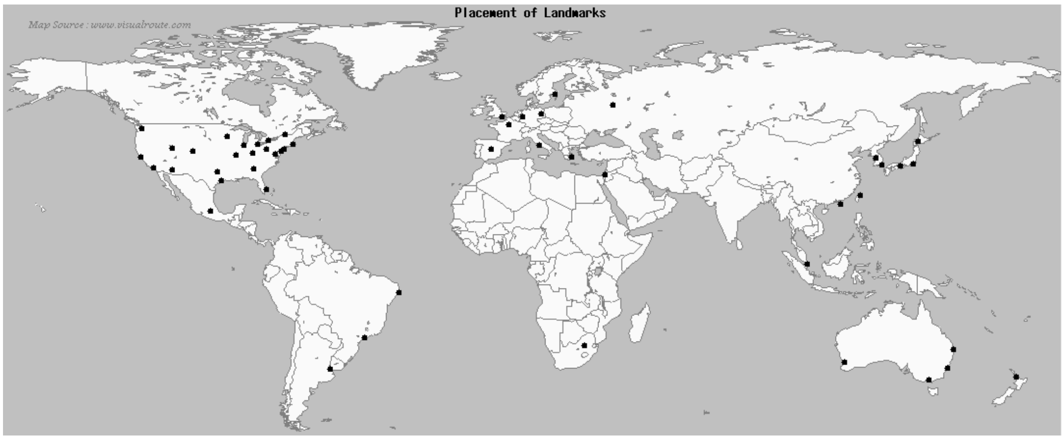

(b) Demographic placement (considers the concentration of users)

Fig. 11. K-center solution for placing 50 landmarks.

goal is to make a geographically sparse distribution of probe machines to avoid shared paths. This sparse placement of probe machines avoids a possible correlation between the delay patterns gathered from the landmarks. In order to make the deployment of probe machines feasible, they are placed on agglomerations with better network infrastructure. The second goal is thus to maximize the number of users on the agglomerations selected to host probe machines. We adopt the number of users (hosts) as a means to reflect the level of network infrastructure on a given agglomeration.

The problem of placing probe machines is to find a set $\mathcal{P} \subseteq \mathcal{A}$ with N probe machines that maximizes the minimum weighted distance between any pair of placed probe machines. We adopt the weighted distance between agglomerations to consider the user concentration at each agglomeration and thus place probe machines on agglomerations that have better network infrastructure.

Using the notation defined in Table 3, the problem of placing probe machines is given by the objective function

$$f_3(\mathcal{P}) = \max M \tag{14}$$

and is subject to the following constraints:

$$\sum_i Q_i = N \tag{15}$$

$$M \le h_i \sum_j g_{ij} Q_j \qquad \text{if } Q_i = 1, \forall i \tag{16}$$

$$Q_i = 0, 1 \qquad \forall i \tag{17}$$

The constraint (15) states that N probe machines are to be located. The constraint (16) defines the upper bound on the minimum distance between two placed probe machines, which is being maximized. The minimum weighted distance M must be lesser or equal than the distance between any pair of probe machines i and j. This constraint makes no sense if $Q_i = 0$ (no probe machine at agglomeration i) as M is the weighted distance between a pair of probe machines i and j. The constraint (17) is the integrality constraint on the decision variable Q.

We propose a greedy approach to solve the problem of placing probe machines with time complexity $O(|\mathcal{A}|^2 N)$. This greedy approach is as follows:

1: $\mathcal{A}' \leftarrow \mathcal{A}$
2: $\mathcal{P} \leftarrow A \in \mathcal{A}$ with $\max_i\{h_i\}$
3: **while** ($|\mathcal{P}| < N$) **do**
4: Find $A \in \mathcal{A}'$ that maximizes $D = \sum_i dist(A, P_i)$, $\forall P_i \in \mathcal{P}$
5: $\mathcal{P} \leftarrow \mathcal{P} \cup A$
6: $\mathcal{A}' \leftarrow \mathcal{A}' - A$
7: **end while**
8: $\mathcal{P}$ is the set of N probe machines

The set of probe machines is initialized with the agglomeration that presents the highest user concentration. Afterwards the algorithm greedily places a new probe machine to maximize the weighted distance between the new probe machine and the previously placed probe machines until N probe machines are placed. The first placed probe machine influences the provided solution. Nevertheless, the decision to place the first probe machine on the agglomeration with the highest user concentration is consistent with the argument of placing probe machines where there exists enough network infrastructure.

We compare the results provided to the problem of placing probe machines by the demographic, geographic, and random placements in Fig. 12. Under the geographic placement, to consider unweighted distances, the constraint (16) is replaced by $M \le \sum_j g_{ij} Q_j$, if $Q_i = 1$, for all i. In Fig. 12(a), we compare the

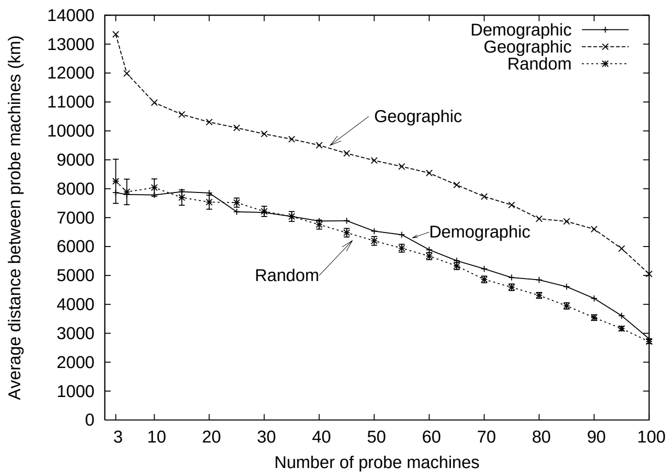

(a) Distance between probe machines

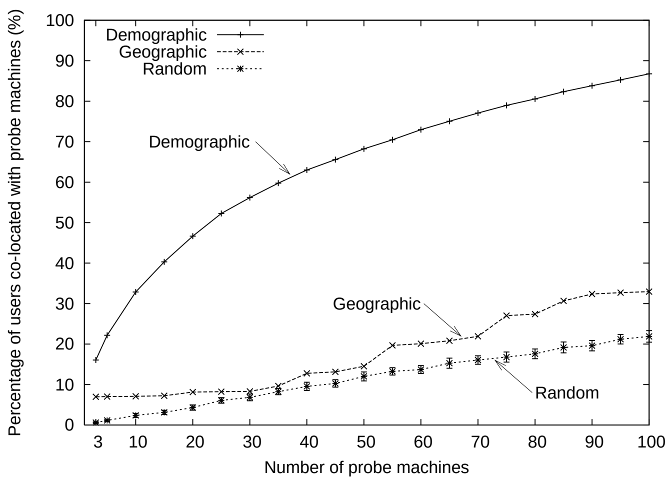

(b) Users co-located with probe machines

Fig. 12. Placing probe machines.

average distance observed between all pairs of the N placed probe machines for the three placement policies. The geographic placement performs the best in sparsely distributing the probe machines as it presents the largest average distances. Nevertheless, such results place probe machines on remote user agglomerations that do not necessarily have a good network infrastructure since the geographic placement disregards user concentration throughout the agglomerations.

Fig. 12(b) shows that the demographic placement locates the probe machines on agglomerations with high density of users, which is our adopted criterion to measure the level of network infrastructure. The demographic placement largely outperforms the geographic and random placements. In the demographic placement, probe machines are placed on agglomerations that have a sufficient network infrastructure to make the deployment of probe machines feasible. Meanwhile, the demographic placement is still able to provide a relatively sparse distribution of probe machines to avoid the shared paths toward landmarks and hosts to be located as shown in Fig. 12(a).

5 Measuring the Similarity between Delay Patterns

In this section, we investigate how to best measure the similarity between the delay pattern of each landmark and the one observed for the target host. The delay patterns result from the partial viewpoints gathered by the distributed probe machines. The landmark that presents the most similar delay pattern with respect to the one of the target host provides the location estimation of that host. Measuring the similarity of the concerned delay patterns is thus a key point for the accuracy of the host location from delay measurements.

5.1 Similarity Models

We define the similarity function $\mathcal{S}(\mathbf{x}, \mathbf{y}) : \mathbb{R}^N \rightarrow [0, 1]$ to measure the degree of similarity between two delay patterns $\mathbf{x}$ and $\mathbf{y}$ where N is the number of adopted probe machines. We adopt such a function to evaluate the degree of similarity between the delay patterns gathered by the probe machines from each landmark and from the target host to be located. In this subsection, we investigate the adoption of three known similarity models [34]: Cosine-based, correlation-based, and distance-based. Each similarity model provides its own manner of implementing the similarity function $\mathcal{S}(\mathbf{x}, \mathbf{y})$. The resulting similarity level between the delay patterns falls in the interval $0 \leq \mathcal{S}(\mathbf{x}, \mathbf{y}) \leq 1$. The closer the similarity level is to 1, more similar the delay patterns are.

5.1.1 Cosine-based Similarity

In the first similarity model, the two delay patterns are thought of as two vectors in a N-dimensional delay space. The similarity between them is measured by computing the cosine of the angle θ between these two vectors. The cosine-based similarity between the delay patterns $\mathbf{x}$ and $\mathbf{y}$, denoted by $\mathcal{S}_{\cos}(\mathbf{x}, \mathbf{y})$, is then given by

$$\mathcal{S}_{\cos}(\mathbf{x}, \mathbf{y}) = \cos\theta = \frac{\mathbf{x} \cdot \mathbf{y}}{\|\mathbf{x}\|\|\mathbf{y}\|} \tag{18}$$

where "$\cdot$" denotes the dot-product of the two vectors and $\|\mathbf{x}\|$ is the Euclidean size of vector $\mathbf{x} \in \mathbb{R}^N$, i.e. $\|\mathbf{x}\| = \sqrt{\sum_{i=1}^{N} x_i^2}$. It should be noted that the cosine of the angle θ is already in the range of $[0,1]$ because the delay patterns $\mathbf{x}$ and $\mathbf{y}$ are both positive vectors.

5.1.2 Correlation-based Similarity

In this model, similarity is measured by computing the coefficient of correlation between the two delay patterns $\mathbf{x}$ and $\mathbf{y}$. The coefficient of correlation is defined as

$$\mathrm{corr}(\mathbf{x},\mathbf{y}) = \frac{\mathrm{cov}(\mathbf{x},\mathbf{y})}{\sigma_{\mathbf{x}}\sigma_{\mathbf{y}}}, \tag{19}$$

where $\mathrm{cov}(\mathbf{x},\mathbf{y})$ denotes the covariance between delay patterns $\mathbf{x}$ and $\mathbf{y}$, and $\sigma_{\mathbf{x}}$ is the standard deviation of $\mathbf{x}$. The correlation-based similarity, denoted by $\mathcal{S}_{\mathrm{cor}}(\mathbf{x},\mathbf{y})$, is scaled to the interval $[0,1]$ and given by

$$\mathcal{S}_{\mathrm{cor}}(\mathbf{x},\mathbf{y}) = \frac{\mathrm{corr}(\mathbf{x},\mathbf{y})+1}{2}. \tag{20}$$

5.1.3 Distance-based Similarity

We represent the distance $u(\mathbf{x},\mathbf{y})$ between two delay patterns $\mathbf{x}$ and $\mathbf{y}$ by

$$u(\mathbf{x},\mathbf{y}) = L_p = \left[\sum_{i=1}^{N} |x_i - y_i|^p\right]^{\frac{1}{p}}, \quad p > 0. \tag{21}$$

The distance function $u(\mathbf{x},\mathbf{y})$ belongs to the L_p family of functions. When $p=1$, we have the Manhattan or city-block distance. In contrast, for $p=2$, we have the Euclidean distance. Furthermore, $0 < p < 1$ results in a non-metric distance function adequate to be used if distances do not satisfy the triangle inequality [35]. Shepard [36] argues in favor of an exponential decay function for a distance-based similarity model. A flexible distance-based similarity model, denoted by $\mathcal{S}_{\mathrm{dis}}(\mathbf{x},\mathbf{y})$, which includes the exponential decay function is given by

$$\mathcal{S}_{\mathrm{dis}}(\mathbf{x},\mathbf{y}) = \mathrm{e}^{-(u(\mathbf{x},\mathbf{y})/\beta)^{\alpha}}, \quad \beta > 1,\ \alpha > 0, \tag{22}$$

where β is a scaling factor defined as

$$\beta = \max_{i,j}\left(1, \quad \frac{1}{2}\left(\frac{\sum_i \mathbf{x}_i}{\|\mathbf{x}\|} + \frac{\sum_j \mathbf{y}_j}{\|\mathbf{y}\|}\right)\right). \tag{23}$$

Since $\mathcal{S}_{\mathrm{dis}}(\mathbf{x},\mathbf{y})$ should decrease as $u(\mathbf{x},\mathbf{y})$ increases, then $\alpha > 0$. The same kind of function is explored for measuring similarity in different contexts like pattern analysis and similarity theory [15,35,37].

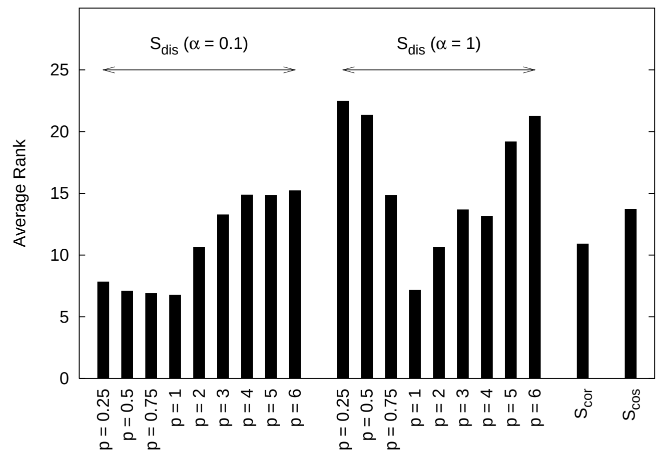

Fig. 13. Average rank of the ideal landmark.

5.2 Experimental Results for the Similarity Models

In this subsection, we analyze the performance of the different similarity models to provide a location estimation of a target host to be located. We adopt the RIPE dataset to evaluate the three similarity models. The RIPE hosts are equipped with GPS cards, allowing us to know their accurate geographic location. This allows us to compare the estimated locations with the real ones and, as a consequence, derive our performance results. The hosts in the RIPE dataset are considered the target host to be located one at a time. The remaining hosts in the set are then used as probe machines and landmarks to perform the location estimation of the target host. We repeat this procedure to evaluate the resulting location estimation of each host in the dataset.

5.2.1 Ranking Landmarks

As we know the geographic position of each host in the RIPE dataset due to the GPS cards, we are able to determine the *ideal* landmark to be chosen by the similarity models being evaluated. The ideal landmark is the geographically closest landmark to the target host. We analyze the performance of the similarity models by ranking the ideal landmark for each target host. Each model provides a list of landmarks in descending order of similarity presented by their delay patterns with respect to the delay pattern of the target host. We then observe the rank of the ideal landmark on the ordered list resulting from each evaluated similarity model. This rank is obtained for each host in the RIPE dataset being considered as a target at a time. Fig. 13 presents the average rank of the ideal landmark for each similarity model.

The best performance in ranking the landmarks is obtained by the distance-based similarity model when $p = 1$, equivalent to the city-block distance. The city-block distance also has the lowest sensitivity to the considered values of α since it presents a similar performance independent of the value of this parameter. Moreover, the city-block distance outperforms the Euclidean distance previously adopted in [4] as well as the cosine-based and correlation-based similarity models. Furthermore, the previously adopted Euclidean distance performs similarly to the correlation-based similarity model. For data following a standard normal distribution, i.e. zero mean and unit variance, the correlation is indeed a linear transformation of the squared Euclidean distance. This relationship is given by $\text{corr}(\mathbf{x}, \mathbf{y}) = 1 - L_2^2(\mathbf{x}, \mathbf{y})/2N$, where $L_2(\mathbf{x}, \mathbf{y})$ is the Euclidean distance between $\mathbf{x}$ and $\mathbf{y}$. Even if in our experiments the delay patterns $\mathbf{x}$ and $\mathbf{y}$ clearly do not follow the standard normal distribution, our findings indicate a certain level of correlation between the correlation-based and the distance-based (with Euclidean distance) similarity models.

We note that the distance-based similarity model is more sensitive to the variation of the distance parameter p for larger values of α. The adopted values of p and α exert a great influence on the performance of the distance-based similarity model. The results obtained by the correlation-based and cosine-based similarity models are even comprised within the performance range of the distance-based model. For $\alpha = 0.1$, relatively good performances are achieved by the distance-based model with non-metric distances $(0 < p < 1)$. Some experiments show that the triangle inequality does not hold over all parts of the Internet [16–18]. These experiments indicate that the adoption of non-metric distances may be more adequate than metric distances on specific parts of the Internet. Euclidean distance as well as higher order distances tend to be less robust to violations of the triangle inequality that are present in some parts of the Internet.

The cumulative probability of ranking the ideal landmark for different similarity models is depicted in Fig. 14. We compare the distance-based similarity model (presented in the form $\mathcal{S}_{\text{dis}}$–p–α) for city-block and Euclidean distances with the correlation-based and cosine-based similarity models. As the Euclidean distance has been previously adopted, we use it as a performance reference for comparison. We observe that the distance-based similarity model with city-block distances outperforms the remaining ones. The correlation-based similarity model provides an equivalent performance to the distance-based similarity model with Euclidean distance.

5.2.2 Distance Accuracy

We now evaluate the distance accuracy of the location estimation of each considered similarity model. The location estimation corresponds to the location of the landmark chosen by each similarity model, i.e. the first ranked landmark in the resulting ordered list of each similarity model. Therefore, we

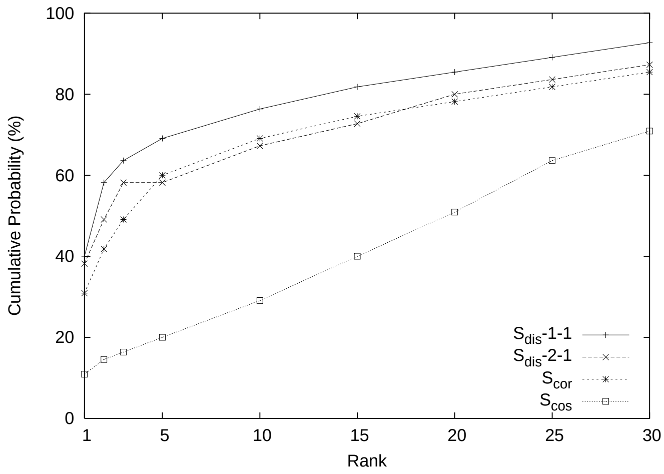

Fig. 14. Cumulative probability of the rank of the ideal landmark.

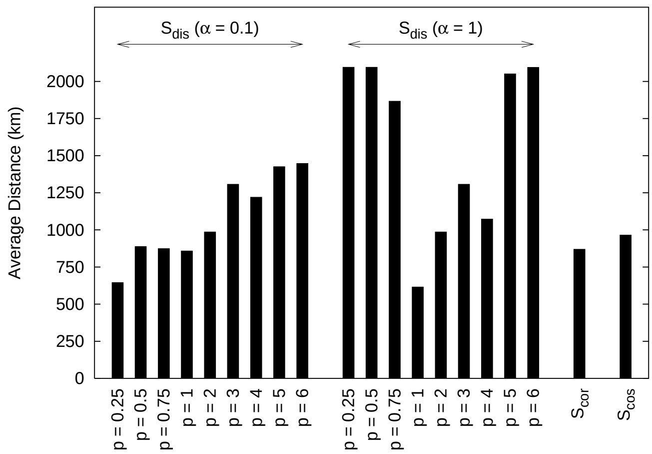

Fig. 15. Average distance between the target host and the location estimation.

evaluate the performance of the different similarity models by comparing the error distance from the selected landmark and the target host to be located.

Fig. 15 presents the average distance between the target host and the elected closest landmark. Such results are the average of each RIPE host being considered a target host one at a time. The remaining hosts are then used as landmarks and probe machines. As in the ranking landmarks evaluation, the best performance is achieved by the distance-based similarity model for the city-block distance ($p = 1$) and $\alpha = 1$. Nevertheless, the insensitivity to α is not the same for the distance accuracy as in the evaluation of ranking the ideal landmark. The correlation-based similarity model presents a similar performance to the non-metric distances with $\alpha = 0.1$, outperforming the distance-based similarity model with the reference Euclidean distance.

The cumulative probability of the distance between the target host and the selected closest landmark is shown in Fig. 16. Again, we compare the correlation-

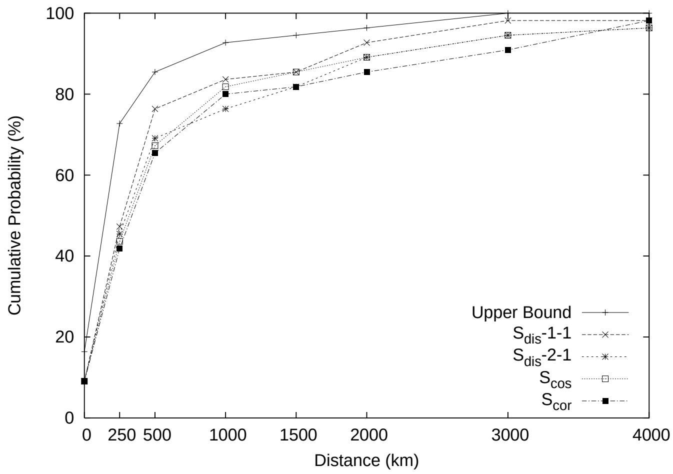

Fig. 16. Cumulative probability of the distance (km).

based and the cosine-based similarity models with the distance-based similarity model for city-block and Euclidean distances. The "Upper Bound" is the best possible performance for the RIPE dataset. This upper bound consists of the distances from the target hosts to their respective ideal landmarks. The city-block distance ($\mathcal{S}_{\mathrm{dis}}$–1–1) outperforms the remaining models, presenting the closest performance to the upper bound.

It should be noted that the RIPE dataset is relatively small. Using one host as the target host results in 54 landmarks to choose from to infer a location estimation. Even if the elected landmark is the geographically closest landmark to the target host, not necessarily it is nearby that target host. For example, the RIPE host in Tokyo is the only one in Asia and its ideal landmark is located in Finland, 7819 kilometers away. As a consequence of the limited number of elements in the RIPE dataset, the average distance from the ideal landmarks to their respective target hosts is 405 kilometers. This distance is the best average distance the similarity models can reach for the RIPE dataset.

6 Conclusion

This paper focuses on a geographic location service of Internet hosts based on a technique that infers host locations using delay measurements from geographically distributed landmarks. We aim at improving the accuracy of the geographic location estimation of the target host to be located.

The contributions of this paper are based on three key points. First, we study the correlation between geographic distance and network delay. Poor connectivity weakens such a correlation. We observe for two different datasets that the correlation between distance and delay becomes stronger as connectivity becomes richer. Thereafter, we identify two key issues that influence the accuracy of the resulting geographic location estimation of a target host: the placement of landmarks and probe machines, and the similarity model that compares the observed delay patterns. Thereby we address the problem of strategically placing landmarks and probe machines to improve the accuracy of the geographic location estimation. We propose and evaluate a demographic placement approach that considers the geographic distribution of users (hosts) to place landmarks and probe machines. Results show that the proposed demographic placement allows a relatively small number of landmarks to represent a large portion of users within a limited coverage distance. Fewer landmarks also imply a lower amount of measurement traffic in the network. These landmarks are placed on areas of high user density, thus providing closer landmarks and more accurate location estimations for most hosts to be located. Probe machines are placed on locations that have enough network infrastructure and sparsely distributed to avoid gathering redundant data. Furthermore, in order to address the second key issue on accuracy, we investigate three different similarity models and how accurate they are for Internet host location from delay measurements. This investigation is carried out using experimental data from the RIPE measurement infrastructure. The similarity models select the landmark that provides the location estimation of the target host to be located, thus being a key component for accuracy. Despite such a fact, Euclidean distance was the only similarity model that has been previously adopted. Our findings show that the distance-based similarity model with $\alpha = 1$ and city-block distance outperforms in accuracy the remaining similarity models. This result includes the Euclidean distance that has been used as a reference for comparison. Therefore, strategically placing the landmarks and probe machines as well as using a more adequate similarity model contribute to improve the accuracy of the measurement-based geographic location of Internet hosts.

Acknowledgment

This work is supported by CAPES/COFECUB, FUJB, and CNPq. During the development of this work, Artur Ziviani had a scholarship from CAPES/Brazil to support his Ph.D. studies under a cooperation between LIP6 and UFRJ. The authors are thankful to Timur Friedman and Tijani Chahed for the access to the RIPE network.

References

[1] C. Davis, P. Vixie, T. Goodwin, I. Dickinson, A means for expressing location information in the domain name system, Internet RFC 1876 .

[2] University of Illinois at Urbana-Champaign, IP Address to Latitude/Longitude, `http://cello.cs.uiuc.edu/cgi-bin/slamm/ip2ll/`.

[3] D. Moore, R. Periakaruppan, J. Donohoe, K. Claffy, Where in the world is netgeo.caida.org?, in: Proc. of the INET'2000, Yokohama, Japan, 2000.

[4] V. N. Padmanabhan, L. Subramanian, An investigation of geographic mapping techniques for Internet hosts, in: Proc. of the ACM SIGCOMM'2001, San Diego, CA, USA, 2001.

[5] Visualware Inc., VisualRoute, `http://www.visualware.com/visualroute/`.

[6] CAIDA, GTrace, `http://www.caida.org/tools/visualization/gtrace/`.

[7] A. Ziviani, S. Fdida, J. F. de Rezende, O. C. M. B. Duarte, Demographic placement for Internet host location, in: Proc. of the IEEE GLOBECOM'2003, San Francisco, CA, USA, 2003.

[8] A. Ziviani, S. Fdida, J. F. de Rezende, O. C. M. B. Duarte, Similarity models for Internet host location, in: Proc. of the IEEE International Conference on Networks - ICON'2003, Sydney, Australia, 2003, pp. 81–86.

[9] P. Krishnan, D. Raz, Y. Shavitt, The cache location problem, IEEE/ACM Transactions on Networking 8 (5) (2000) 568–582.

[10] L. Qiu, V. N. Padmanabhan, G. M. Voelker, On the placement of web server replicas, in: Proc. of the IEEE INFOCOM'2001, Anchorage, AK, USA, 2001.

[11] P. Radoslavov, R. Govindan, D. Estrin, Topology-informed Internet replica placement, Computer Communications 25 (4) (2002) 384–392.

[12] E. Cronin, S. Jamin, C. Jin, A. R. Kurc, D. Raz, Y. Shavitt, Constrained mirror placement on the Internet, IEEE Journal on Selected Areas in Communications 20 (7) (2002) 1369–1382.

[13] H. S. Bassali, K. M. Kamath, R. B. Hosamani, L. Gao, Hierarchy-aware algorithms for CDN proxy placement in the Internet, Computer Communications 26 (3) (2003) 251–263.

[14] B. Sarwar, G. Karypis, J. Konstan, J. Riedl, Item-based collaborative filtering recommendation algorithms, in: Proc. of the International World Wide Web Conference - WWW10, Hong Kong, 2001.

[15] S. Santini, R. Jain, Similarity measures, IEEE Transactions on Pattern Analysis and Machine Inteligence 21 (9) (1999) 871–883.

[16] P. Francis, S. Jamin, C. Jin, Y. Jin, D. Raz, Y. Shavitt, L. Zhang, IDMaps: A global Internet host distance estimation service, IEEE/ACM Transactions on Networking 9 (5) (2001) 525–540.

[17] S. Savage, A. Collins, E. Hoffman, J. Snell, T. Anderson, The end-to-end effects of Internet path selection, in: Proc. of the ACM SIGCOMM'99, Cambridge, MA, USA, 1999.

[18] S. Banerjee, T. G. Griffin, M. Pias, The interdomain connectivity of PlanetLab nodes, in: Proc. of the Passive and Active Measurement Workshop - PAM'2004, Lecture Notes in Computer Science (LNCS) 3015, Antibes Juan-les-Pins, France, 2004.

[19] LibWeb, `http://sunsite.berkeley.edu/Libweb`.

[20] RIPE Test Traffic Measurements, `http://www.ripe.net/ttm/`.

[21] T. Vincenty, Direct and inverse solutions of geodesics on the ellipsoid with application of nested equations, Survey Review 22 (176) (1975) 88–93.

[22] L. Subramanian, V. N. Padmanabhan, R. Katz, Geographic properties of Internet routing, in: Proc. of USENIX 2002, Monterey, CA, USA, 2002.

[23] N. Spring, R. Mahajan, T. Anderson, Quantifying the causes of path inflation, in: Proc. of the ACM SIGCOMM'2003, Karlsruhe, Germany, 2003.

[24] A. Lakhina, J. W. Byers, M. Crovella, I. Matta, On the geographic location of Internet resources, IEEE Journal on Selected Areas in Communications 21 (6) (2003) 934–948.

[25] S.-H. Yook, H. Jeong, A.-L. Barabási, Modeling the Internet's large-scale topology, Proc. of the National Academy of Sciences (PNAS) 99 (2002) 13382–13386.

[26] A. Ziviani, S. Fdida, J. F. de Rezende, O. C. M. B. Duarte, Toward a measurement-based geographic location service, in: Proc. of the Passive and Active Measurement Workshop - PAM'2004, Lecture Notes in Computer Science (LNCS) 3015, Antibes Juan-les-Pins, France, 2004, pp. 43–52.

[27] G. Ballintijn, M. van Steen, A. S. Tanenbaum, Characterizing Internet performance to support wide-area application development, Operating Systems Review 34 (4) (2000) 41–47.

[28] T. Brinkhoff, City Population, `http://www.citypopulation.de`.

[29] Central Intelligence Agency (CIA), The World Factbook 2002, `http://www.cia.gov/cia/publications/factbook` (Jan. 2002).

[30] C. Toregas, R. Swain, C. Revelle, L. Bergman, The location of emergency service facilities, Operations Research 19 (6) (1971) 1363–1373.

[31] M. S. Daskin, Network and Discrete Location, John Wiley & Sons, New York, NY, 1995.

[32] CAIDA, GeoPlot, http://www.caida.org/tools/visualization/geoplot/.

[33] M. R. Garey, D. S. Johnson, Computers and Intractability: A Guide to the Theory of NP-Completeness, W. H. Freeman & Co., New York, NY, 1979.

[34] A. D. Gordon, Classification: Methods for the Exploratory Analysis of Multivariate Data, Chapman and Hall, 1981.

[35] D. W. Jacobs, D. Weinshall, Y. Gdalyahu, Classification with non-metric distances: Image retrieval and class representation, IEEE Transactions on Pattern Analysis and Machine Inteligence 22 (6) (2000) 583–600.

[36] R. N. Shepard, Toward a universal law of generalization for psychological science, Science 237 (1987) 1317–1323.

[37] D. M. Ennis, J. J. Palen, K. Mullen, A multidimensional stochastic theory of similarity, Journal of Mathematical Psychology 32 (4) (1988) 449–465.