Two-Tier Geographic Location of Internet Hosts

Bamba Gueye[1], Artur Ziviani[1], Serge Fdida[1],
José F. de Rezende[2], and Otto Carlos M. B. Duarte[2]

[1] Laboratoire d'Informatique de Paris 6 (LIP6)
Université Pierre et Marie Curie (Paris 6)
Paris, France
{Bamba.Gueye,Artur.Ziviani,Serge.Fdida}@lip6.fr

[2] Grupo de Teleinformática e Automação (GTA)
COPPE/Poli – Universidade Federal do Rio de Janeiro
Rio de Janeiro, Brazil
{rezende,otto}@gta.ufrj.br

Abstract. Multimedia delivery systems, such as Content Distribution Networks (CDNs), improve by knowing the geographic location of their clients. Therefore, we focus on a measurement-based geographic location service of Internet hosts. Such a service infers host locations from delay measurements taken from probe machines toward landmarks, which are hosts with a known geographic location, and the host to be located. We aim at mitigating the number of measurements generated in the network. We propose a two-tier hierarchical structure of landmarks to geographically locate an Internet host as opposed to a flat structure. In our two-tier structure, the upper level mitigates long distance measurements and the lower level keeps measurements within restricted areas. As a consequence, the two-tier geographic location structure significantly reduces the number of measurements and thereby it favors scalability.

1 Introduction

Geographically locating an Internet host from its IP address enables a diversified and interesting new class of location-aware applications. Examples of such Internet applications are targeted advertising on web pages, announce of local events and regional weather, or authorization of transactions only when performed from pre-established locations. In peer-to-peer networks, location-aware construction of overlay networks can avoid unnecessary high latency hops, thus improving routing performance [1]. Multimedia delivery systems, such as Content Distribution Networks (CDNs), can also benefit from knowing the location of their clients [2]. For example, benefits include the indication of nearby servers to clients or the location-based adaptation of the multimedia content.

A DNS-based approach to provide a geographic location service of Internet hosts is proposed in RFC 1876 [3]. This proposition, however, is not widely adopted since it requires changes in DNS structure and administrators have no motivation to register new location records. Different techniques [4] infer the geographic location of an Internet host from DNS names, from clustering the IP address space with BGP prefix information, or from delay measurements.

We focus on inferring the geographic location of an Internet host from delay measurements. The measurement-based location estimation of a host is based on the observation that hosts sharing similar delays to some fixed probe machines tend to be near each other geographically. Given a set of landmarks, which are hosts with a known geographic location, the location estimation for a target host is the location of the landmark with the most similar delay pattern to the host. Accuracy basically depends on the placement of landmarks and probe machines [5]. Previous works, such as Geoping in [4], use a flat structure of landmarks to locate Internet hosts from delay measurements. Nevertheless, this flat structure imposes a great measurement load in the network [6].

In this paper, we propose a two-tier hierarchical structure of landmarks to mitigate the number of measurements generated in the network to perform a location estimation. Although the use of hierarchies is a known technique in general, it has never to the best of our knowledge been applied in the context of measurement-based geographic location of Internet hosts. In our two-tier proposed approach, landmarks are disposed in two levels. At the upper level, we place few landmarks to determine a coarse-grained location estimation of hosts. At the lower level, we adopt subsets of landmarks to cover restricted areas indicated by the upper level. We use the lower level to possibly improve the accuracy of the location estimation achieved by the upper level. Results show that the proposed two-tier structure significantly reduces the number of measurements with respect to the flat structure. This favors the scalability of the two-tier proposition.

This paper is organized as follows. Section 2 presents the previous flat structure used to estimate the location of hosts from delay measurements. In Section 3, we describe the design of the two-tier geographic location of Internet hosts. We evaluate the proposed two-tier structure as opposed to the flat structure in Section 4. In Section 5, we present our conclusions.

2 Flat Measurement-Based Host Location

We formalize the flat measurement-based host location as follows. Consider a set $\mathcal{L} = \{L_1, L_2, \ldots, L_K\}$ of K landmarks. Landmarks are any reference hosts able to echo `ping` messages and 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 using a flat structure, which are detailed along this section. The N probe machines perform delay measurements toward the flat structure of K landmarks (Fig. 1(a)). Each probe machine P_x keeps a delay vector $\mathbf{d}_x = [d_{1x}, d_{2x}, \ldots, d_{Kx}]^T$ to the set of K landmarks. Given a host T to be located, the location server asks each probe machine to measure the delay to host T (Fig. 1(b)). Each probe machine then returns to the location server a new delay vector $\mathbf{d}'_x = [d_{1x}, d_{2x}, \ldots, d_{Kx}, d_{Tx}]^T$, which is composed by the delay vector $\mathbf{d}_x$ and the just measured delay to host T (Fig. 1(c)).

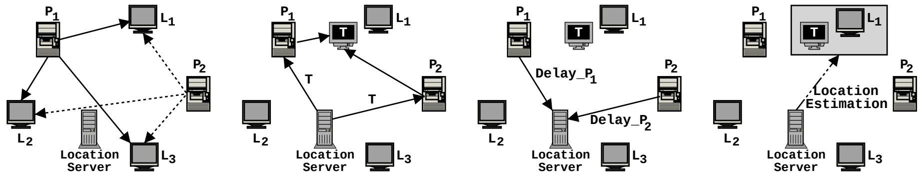

(a) (b) (c) (d)

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

After receiving the delay vectors from the N probe machines, the location server is able to construct the delay matrix $\mathbf{D}$ with dimensions $(K+1) \times N$:

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

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. Geoping [4] uses Euclidean distance as a means to find the landmark with the most similar delay pattern with respect to the one of the host to be located. In other words, the landmark L with the smallest Euclidean distance $e_{LT} = \sqrt{(d_{y1} - d_{T1})^2 + (d_{y2} - d_{T2})^2 + \ldots + (d_{yN} - d_{TN})^2}$ from host T, where $y = 1, \ldots, K$, is the nearest landmark with respect to T. The location of the landmark L is used as the location estimation of the host T (Fig. 1(d)).

The accuracy of the location estimation in the flat structure basically depends on the location and on the number of the landmarks. Better location estimations for most hosts need a large number of landmarks. Nevertheless, augmenting the number of landmarks considerably increases the number of measurements generated in the network. To model this, let Δ denote the time interval adopted by the probe machines to periodically gather the delay from the landmarks of the set $\mathcal{L}$. For the flat structure, the total number of measurement messages $\mathcal{M}_{\text{flat}}$ to estimate the location of t hosts in a time interval τ is

$$\mathcal{M}_{\text{flat}}(\tau) = 2N\left(\left\lceil \frac{\tau}{\Delta} \right\rceil K + t\right). \tag{2}$$

The multiplicative factor 2 when evaluating the number of measurement messages is used to take into account that each `ping` measurement is actually composed by two messages, i.e. the ECHO_REQUEST and the ECHO_RESPONSE messages. It should also be noted that each measurement may consist of one to several delay samples, but only the minimum value is considered to avoid delays due to network congestion. In the case we send w 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 $w \times \mathcal{M}_{\text{flat}}$. In our analysis, we consider $w = 1$ for simplicity without affecting the relative performance of the flat and the two-tier structures.

The minimum number of measurement messages $\mathcal{M}_{\text{flat}}^{\min}$ to estimate the location of t hosts in a flat structure is

$$\mathcal{M}_{\text{flat}}^{\min}(t) = 2N(K + t). \tag{3}$$

3 Two-Tier Geographic Host Location

The two-tier hierarchical structure we propose aims at mitigating the number of measurements generated in the network. At the upper level of the hierarchical structure we place few landmarks that are sparsely distributed to reduce long distance measurements. Each landmark covers a large coverage distance, leading to a coarse-grained location estimation. The subsets of landmarks at the lower level are obtained by decreasing the coverage distance of the landmarks located in the upper level. Thereby, we obtain a more accurate location estimation at the lower level. The set of probe machines remain unchanged with respect to the flat structure. The main user agglomerations worldwide are considered as candidate sites to place the landmarks (further details are given in Section 4.1). We adopt the notation described in Table 1 to determine the distribution of landmarks through these agglomerations in a general hierarchical structure with q levels. This hierarchical structure leads to the subsets $\mathcal{L}_s^q$ of the set of landmarks $\mathcal{L}$ as given by

$$\mathcal{L} \supseteq \bigcup_{s,q} \mathcal{L}_s^q. \tag{4}$$

In the two-tier structure, we have a unique set $\mathcal{L}_1^1$ of K_1^1 landmarks at the upper level ($q = 1$). The lower level ($q = 2$) contains the sets $\mathcal{L}_s^2$ with the respective K_s^2 landmarks, for all s. We place landmarks in the two-tier structure according to the user population distribution, extending the demographic placement policy proposed in [5] from a flat to a two-tier structure. The basic idea behind the demographic placement is to place landmarks following the user distribution to reflect where most hosts to be located are. Recent findings [7] indicate a strong correlation between population and router density in economically developed countries. Using the notation from Table 1, this approach is formulated by the objective function

$$\max \sum_q \sum_s \sum_i h_{is}^q Z_i^q. \tag{5}$$

Table 1. Notation for a general hierarchical structure of q levels.

$\mathcal{L}_s^q$	subset s of landmarks at level q
h_{is}^q	number of users at agglomeration i covered by subset s of level q
U_s^q	number of covered users by the subset s of level q; $U_s^q = \sum_j h_{sj}^q$
K_s^q	number of landmarks at subset s of level q
a_{ij}^q	$\begin{cases} 1 & \text{if agglomeration } i \text{ can cover agglomeration } j \text{ in level } q \\ 0 & \text{if not.} \end{cases}$
X_i^q	$\begin{cases} 1 & \text{if one places a landmark on agglomeration } i \text{ in level } q, \\ 0 & \text{if not.} \end{cases}$
Z_i^q	$\begin{cases} 1 & \text{if agglomeration } i \text{ is covered in level } q, \\ 0 & \text{if not.} \end{cases}$
c_s	$\begin{cases} 1 & \text{if } \mathcal{L}_s^q \subset \mathcal{L}_j^{q-1}\ \forall j, \\ 0 & \text{if not.} \end{cases}$

The objective function (5) maximizes the number of users nearby a placed landmark in the level q and is subject to the following constraints:

$$Z_i^q \leq \sum_j a_{ij}^q X_j^q \qquad \forall i, q \tag{6}$$

$$\sum_j X_j^q = K_s^q \qquad \forall s, q \tag{7}$$

$$Z_i^q = 0, 1 \qquad \forall i, q \tag{8}$$

$$X_i^q = 0, 1 \qquad \forall i, q \tag{9}$$

The constraint (6) states that users at the agglomeration i are covered at level q if at least one site that covers agglomeration i is selected to host a landmark of level q. The constraint (7) stipulates that we locate no more than K_s^q landmarks at each subset s of level q. The constraints (8) and (9) are the constraints of integrality for the decision variables Z^q and X^q.

In a hierarchical structure of landmarks, for a given target host, the number of measurements to locate it depends on the user (host) distribution and on the required accuracy of the location estimation. If some agglomerations have more users (hosts) than others, it is likely that measurements toward the landmarks covering these agglomerations would be more frequent. Supposing that each host has an equal probability of being located, the probability p_s^q, given by

$$p_s^q = \frac{\sum_{j=1}^{K_s^q} h_{js}^q}{\sum_{j=1}^{K_s^{q-1}} h_{js}^{q-1}}, \quad \forall s, q \tag{10}$$

expresses the likelihood of falling into a specific subset s of landmarks located at level q. In the two-tier structure, since $q = 2$, we have the probabilities p_s^2 for each of the K_1^1 subsets composing the lower level. For example, the probability p_1^2 reflects the probability of using the first subset ($s = 1$) of the lower level ($q = 2$) to refine the location estimation. Therefore, for the two-tier structure, where we have two levels ($q = 2$) and just one subset of landmarks at the upper level, we formulate the probability p_s^2 as

$$p_s^2 = \frac{\sum_{j=1}^{K_s^2} h_{js}^2}{\sum_{j=1}^{K_1^1} h_{js}^1} = \frac{U_s^2}{U_1^1}, \quad \forall s. \tag{11}$$

Knowing the probability of falling into each subset of landmarks of the lower level allows us to evaluate the measurement load of the two-tier structure. The expected average minimum number of measurement messages in the two-tier structure to estimate the location of t hosts is

$$\mathcal{M}_{2\text{tier}}^{\min}(t) = 2N\left[(K_1^1 + t) + \sum_{s=1}^{K_1^1}(K_s^2 - c_s)p_s^2\right], \quad \forall s. \tag{12}$$

It should be noted that the subsets of landmarks $\mathcal{L}_s^q$ are not necessarily disjoints. In the two-tier case, nothing prevents a landmark L to represent a large region at the upper level, i.e. being an element of $\mathcal{L}_1^1$, and to also be comprised in some subset $\mathcal{L}_s^2$ at the lower level to serve as a possible location estimation. In this case, we have then $\mathcal{L}_1^1 \cap \mathcal{L}_s^2 = L$. If the same landmark is present at the two levels, it is unnecessary to perform two measurements toward this landmark. The parameter c_s controls this case to avoid considering twice a measurement that has been only performed once.

4 Evaluation

4.1 Experimental Setting

For our experiments, the main urban agglomerations spread worldwide are considered since they are likely to offer the highest concentration of users. We consider all urban agglomerations with more than one million inhabitants in a total of 407 agglomerations [8]. It is known that the Internet infrastructure varies dramatically across different regions throughout the world. Therefore, we weight the populations of the different urban agglomerations with the number of Internet users in the country the agglomeration belongs to over the total population of the country. In applying this weight, we estimate the main user agglomerations worldwide to be used in our experiments. We denote as $\mathcal{A}$ the set with these user agglomerations. The set $\mathcal{A}$ represents the 407 main user agglomerations, totalizing 173,696,253 estimated users. The mean distance between each pair of elements in $\mathcal{A}$ is 8167 km. Data on estimations of the Internet users and total population of each country are available in [9].

Algorithm 1 Greedy approach to the Maximum Covering Location Model

$\mathcal{L} \leftarrow \emptyset$; $\mathcal{A}' \leftarrow \mathcal{A}$
while ($|\mathcal{L}| < K$) **do**
  Find $A \in \mathcal{A}'$ that covers the most uncovered demand
  Set $\mathcal{C} \subseteq \mathcal{A}'$ as the set of agglomerations covered by A
  $\mathcal{L} \leftarrow \mathcal{L} \cup A$; $\mathcal{A}' \leftarrow \mathcal{A}' - \mathcal{C}$
end while
$\mathcal{L}$ is the set of K landmarks

4.2 Building the Two-Tier Structure

At the upper level of the two-tier structure, we maximize the number of covered users (hosts) for a limited number of landmarks. This results in a Maximum Covering Location Model (MCLM) [10]. We adopt the greedy approach outlined in Algorithm 1 to solve the MCLM problem with time complexity $O(|\mathcal{A}|^2 K)$. The algorithm greedily places landmarks to cover the most uncovered demand until K landmarks are placed. Thereby, we determine the landmark set $\mathcal{L}_1^1$ at the upper level containing K_1^1 landmarks.

At the lower level of the two-tier structure, we perform a set coverage [10] within each of the large areas covered by the upper level. In other words, each landmark placed at the upper level considers a large coverage distance to cover a large number of hosts. For each area covered by a landmark of the upper level, we place extra landmarks to cover with a smaller coverage distance all agglomerations within this restricted area. We therefore determine the landmark subsets $\mathcal{L}_s^q$ that compose the lower level.

4.3 Results

First we compare the minimum number of measurement messages in the flat and in the two-tier hierarchical structure following Equations (3) and (12). The presented results are for 10 probe machines locating one target host. We consider G the target coverage distance for the final location estimation. Therefore, given a base value for G, we build the upper level of the two-tier structure with a value $\alpha \times G$. We use the multiplicative factor α to control the relationship between the upper and the lower levels in the two-tier hierarchical structure. For the upper level, we solve the Maximum Covering Location Model with the value $\alpha \times G$ as the coverage distance in Algorithm 1. For the lower level, we perform a set coverage using G as the coverage distance over each one of the large areas associated with landmarks from the upper level. In other words, we establish landmarks at the lower level that cover areas with a coverage distance G within the areas delimited by $\alpha \times G$ indicated by the landmarks at the upper level. In order to achieve comparability between the flat and the two-tier structures, we build the flat structure with the same number of resulting landmarks at the lower level of the two-tier structure with a coverage distance of G.

Fig. 2 compares the volume of measurement messages injected in the network by the flat and two-tier structures. We consider coverage distances G of 100 km

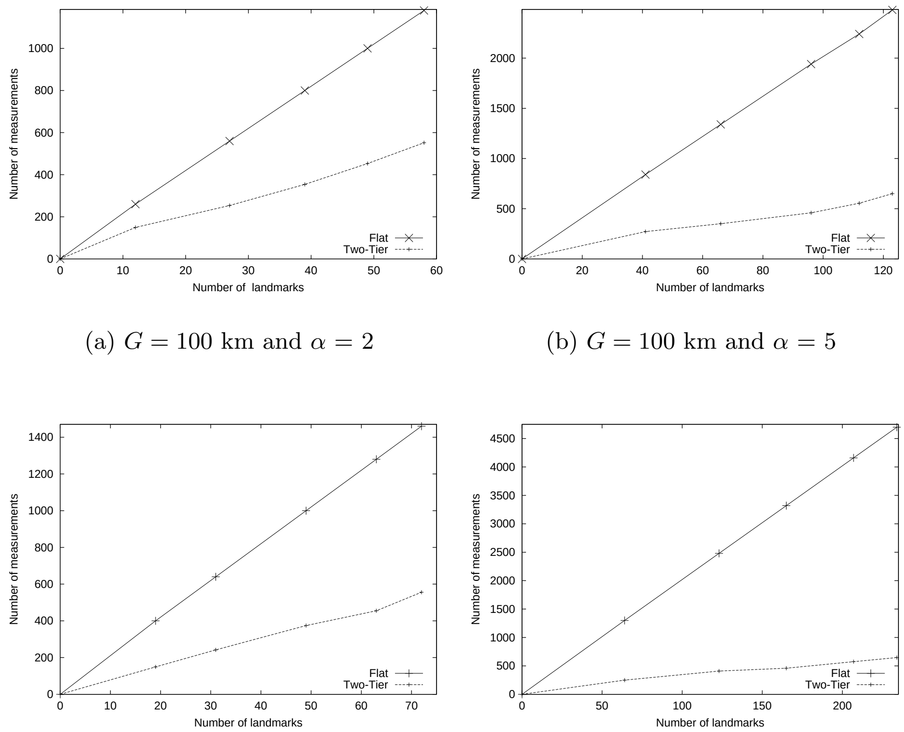

(a) $G = 100$ km and $\alpha = 2$ (b) $G = 100$ km and $\alpha = 5$

(c) $G = 250$ km and $\alpha = 2$ (d) $G = 250$ km and $\alpha = 5$

Fig. 2. Impact of measurement traffic in both flat and two-tier structures.

and 250 km with multiplicative factors α of 2 and 5. To obtain the data shown in Fig. 2, we fix the base coverage distance G. For a given α, we place at the upper level 5, 10, 15, 20, and 25 landmarks and then we show the respective number of landmarks at the lower level of the two-tier structure. With $\sum_i K_i^2$ landmarks at the lower level, we use this same number in the flat structure and then we compare the resulting volume of measurement messages.

The two-tier structure significantly mitigates the number of measurements generated in the network. The flat structure has a linear increase in the number of measurements as a function of the number of landmarks. Meanwhile, results for the two-tier structure present a much slower increase in the volume of measurement traffic for the same number of landmarks. Therefore, the two-tier structure performs a much smaller amount of measurements to achieve a comparable number of possible location estimations. This result is a consequence of few measurements at the upper level to achieve a coarse-grained location estimation

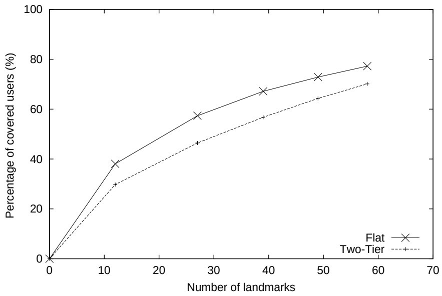

(a) $G = 100$ km and $\alpha = 2$

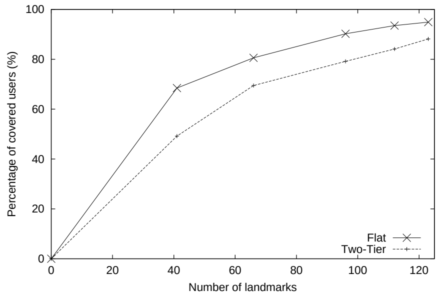

(b) $G = 100$ km and $\alpha = 5$

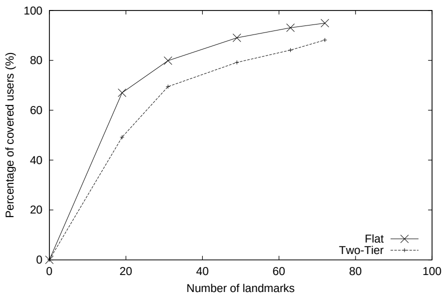

(c) $G = 250$ km and $\alpha = 2$

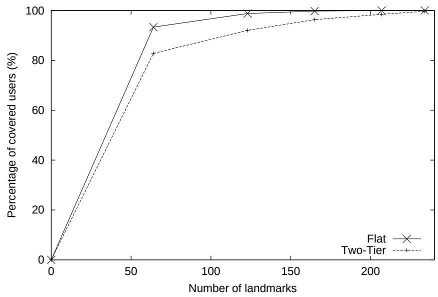

(d) $G = 250$ km and $\alpha = 5$

Fig. 3. Percentage of covered users in both flat and two-tier structures.

and limited measurements at the lower level within restricted areas. Hence, the two-tier structure contributes to enhance the scalability of a measurement-based geographic location service of Internet hosts.

Fig. 3 presents the percentage of covered users achieved by the flat and two-tier structures. This metric represents how much of the user space each structure is able to cover within a given coverage distance. Our findings show that, for the same number of landmarks, the percentage of covered users in a flat structure is greater than in the two-tier structure. Nevertheless, this gap decreases for larger values of α (Fig. 3(d)). From Fig. 3, we can also estimate how many additional landmarks would be needed to the two-tier structure to cover the same number of users as the flat structure. Although the percentage of covered users in the flat structure is greater than in a two-tier structure, the number of measurements generated in the network for a flat structure is very large. Therefore, we significantly mitigate the impact of measurement load at the expense of a small reduction on the number of covered users. Comparing Fig. 2 and Fig. 3, we observe this trade-off.

5 Conclusion

This paper has focused on a measurement-based geographic location service of Internet hosts. We aimed at reducing the number of measurements generated in the network by probe machines toward landmarks and hosts to be located. To achieve this goal, we have proposed and evaluated a two-tier hierarchical structure as opposed to the previously adopted flat structure. Results show that the two-tier structure significantly mitigates the impact of measurements in the measurement-based geographic location of Internet hosts. This comes at the expense of a small reduction on the number of covered users. Nevertheless, our evaluation demonstrates that the reduction in the measurement load by far compensates the slight reduction in covered users. The accuracy can be further tuned by adding more lower layer, i.e. localized, measurements. Mitigating measurement impact also favors the scalability of the two-tier proposition, contributing to a more scalable measurement-based geographic location service of Internet hosts. Such a service can be viewed as an underlying infrastructure for the deployment of novel location-aware applications and location-based multimedia services in the Internet.

Acknowledgment

This work is supported by CAPES/COFECUB, FUJB, and CNPq.

References

1. Ratnasamy, S., Handley, M., Karp, R., Shenker, S.: Topologically-aware overlay construction and server selection. In: Proc. of the IEEE INFOCOM'2002, New York, NY, USA (2002)
2. Sripanidkulchai, K., Maggs, B., Zhang, H.: Efficient content location using interest-based locality in peer-to-peer systems. In: Proc. of the IEEE INFOCOM'2003, San Francisco, CA, USA (2003)
3. Davis, C., Vixie, P., Goowin, T., Dickinson, I.: A means for expressing location information in the domain name system. Internet RFC 1876 (1996)
4. Padmanabhan, V.N., Subramanian, L.: An investigation of geographic mapping techniques for Internet hosts. In: Proc. of the ACM SIGCOMM'2001, San Diego, CA, USA (2001)
5. Ziviani, A., Fdida, S., de Rezende, J.F., Duarte, O.C.M.B.: Demographic placement for Internet host location. In: Proc. of the IEEE GLOBECOM'2003, San Francisco, CA, USA (2003)
6. Ziviani, A., Fdida, S., de Rezende, J.F., Duarte, O.C.M.B.: Toward a measurement-based geographic location service. In: Proc. of the Passive and Active Measurement Workshop - PAM'2004. Lecture Notes in Computer Science (LNCS), Antibes Juan-les-Pins, France (2004)

7. Lakhina, A., Byers, J.W., Crovella, M., Matta, I.: On the geographic location of Internet resources. IEEE Journal on Selected Areas in Communications **21** (2003) 934–948
8. Brinkhoff, T.: (City Population) http://www.citypopulation.de.
9. Central Intelligence Agency (CIA): The World Factbook 2002. (2002) http://www.cia.gov/cia/publications/factbook.
10. Daskin, M.S.: Network and Discrete Location. John Wiley & Sons, New York, NY (1995)