Computer Networks 54 (2010) 1490–1501

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

A model based approach for improving router geolocation

S. Laki *, P. Mátray, P. Hága, I. Csabai, G. Vattay

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

ARTICLE INFO

Article history:
Received 17 March 2009
Received in revised form 1 December 2009
Accepted 4 December 2009
Available online 22 December 2009
Responsible Editor: A. Popescu

Keywords:
IP geolocation
Network topology
Active measurements
Path-latency model
One-way delay

ABSTRACT

This study outlines two novel techniques which can be used in the area of IP geolocation. First we introduce a detailed path-latency model to be able to determine the overall propagation delays along the network paths more accurately. The knowledge of accurate propagation delay values then leads to more precise geographic distance estimation between hosts and measurement nodes. Besides these distance values the evaluation process also takes into account the discovered topology between the measurement points, and end-to-end latency measurements as well. In addition to the application of the detailed path-latency model, we describe a method which utilizes high-precision one-way delay measurements to further increase the accuracy of router geolocation techniques. The precise one-way delay values are used as a "path-constraint" to limit the overall geographic distance between the measurement nodes. This approach can be used to localize all the network routers along the network path between the measurement nodes, but cannot be applied to end-host localization. The techniques introduced in this paper are validated in wide range of experiments performed in the ETOMIC measurement infrastructure and in PlanetLab.

1. Introduction

New location aware applications like e-commerce, censorships, web site content and advertisements have been appearing since the last years, which directs the attention to IP geolocation. The localization of IP addresses became important in scientific areas as well, e.g. to visualize the results of Internet measurements. Besides the visualization it could also help to investigate the geographic aspects of further interesting scientific problems like network topology, routing protocols, policy routing and pricing strategies. Nevertheless, determining geographical location of Internet hosts and routers by a single IP address poses many challenges, since there is no direct relationship between the IP address of a host and its geographic location.

Many of the existing geolocation services are based on databases which store organizational information assigned to IP domains, or try to infer location information from DNS names. Usually the accuracy of these services is insufficient due to the lack of reliable information. To overstep the uncertainty of techniques based on geolinguistic approach, recently active geolocation techniques have emerged. These techniques make an attempt to approximate geographical distances from delay measurements.

This paper outlines novel techniques which can be used in the area of geolocation and can be combined with other existing methods. First we introduce a novel detailed path-latency model to identify the contribution of various phenomena to the packet delay. Among these contributions we determine the ones which are related to the geographical distance that the packet traveled. Our model, instead of handling the packet delay as an irreducible unit, decomposes the overall path-wise packet delay to link-wise components like the processing delay, propagation delay and even ICMP Echo Reply generation time. In this way we are able to approximate the overall propagation delay

* Corresponding author. Tel.: +36 70 374 2646.
E-mail addresses: laki@complex.elte.hu (S. Laki), matray@complex.elte.hu (P. Mátray), haga@complex.elte.hu (P. Hága), csabai@complex.elte.hu (I. Csabai), vattay@complex.elte.hu (G. Vattay).

1389-1286/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2009.12.004

along the measurement path. The knowledge of accurate propagation delay values then leads to more precise geographic distance estimations between measurement nodes (also called *landmarks*) and target nodes. By using the basic multilateration technique the estimated distance values can be applied to infer the geographic location of Internet hosts.

It is possible to refine the path-latency model for a specific subset of IP addresses by using one-way delay values as additional constraints for the sum of the link-wise propagation delays. Taking one-way delays into account we can significantly improve the localization accuracy of routers in between the landmark nodes. Nowadays, just a few infrastructures provide this service, but we believe that in the future Internet the network nodes will be capable of measuring such important network metrics as one-way delay. Among the current infrastructural capabilities one-way delays can only be measured between the landmark nodes. These values can be used to limit the physical length of the overall network paths between any two landmark nodes.

Our study is purely based on active measurements and does not aim to include all the existing methods and possibilities to improve the accuracy of the location estimation. We do not rely on any passive methods like geolinguistic or location aware databases (e.g. WhoIS, DNS LOC records).

This paper aims to present the efficiency and usefulness of applying a detailed path-latency model in geolocation techniques. We also show how the extra information provided by precise one-way delay measurements can improve the accuracy of location approximation. Finally, the performance of the presented methods are investigated in real world experiments performed in the ETOMIC measurement infrastructure [1], and then the measurement results are validated in the GEANT2 research network [2].

The rest of the paper is organized as follows: in Section 2 we briefly overview the prior geolocation methods including constraint and topology based methods. Section 3 describes the detailed path-latency model which constitutes the basis of our geolocation technique. The application of this model for geographical distance approximation is presented in Section 4. Based on this distance approximation method we introduce several geographic constraints including the novel one-way delay based "path-constraint" in Section 5. The details of the data collection process, including the description of the applied tools and the performed experiments can be found in Section 6. In Section 7 we briefly describe how the geolocation problem can be solved by using global optimization techniques. In Section 8 we present a complete performance analysis focusing on the accuracy and consistency of the location estimation. The final section summarizes our results.

2. Related work

During the recent years several geolocation techniques have emerged, all of them aim to give an accurate approximation of the location of network hosts which are not known *a-priori*. Many of these techniques use WHOIS databases [3] and DNS names [4] to determine the location of a given router. From WHOIS databases one can retrieve the name and street address of the organization which registered the address block, however for a large ISP or a geographically dispersed organization the registered street address usually differs from the real location of its routers. A similar problem arises in the use of router DNS names, since the names can be both useful or misleading due to the naming conventions of the ISP [5].

In general, we can say that sometimes these passive methods are very accurate, but in some cases their errors are very large. To illustrate this problem we have localized some GEANT routers using *MaxMind*, which is one of the leading passive IP-to-location database. *MaxMind* provided "fake" locations for them. For example, for the Budapest-located GEANT router rt1.bud.hu.geant2.net *MaxMind* provides a location in Cambridge, UK, where the GEANT operator is registered. This mislocation is due to the fact that passive IP-location systems use previously registered DNS/WHOIS data, which leads to unreliable results in many cases (like as it is with GEANT routers).

Some proposals try to overcome these limitations. IP2-Geo [4] contains a measurement-based approach Geoping, which tries to approximate the geographical distance of network hosts on the basis of the packet delay measurements. A more mature approach is the simultaneous application of several delay constraints to infer the location of a network host. This is done by *constraint-based-geolocation* (CBG) techniques [6]. CBG introduces a triangulation-like method to combine the distance estimates from all landmarks. To estimate delay-distance relation, each landmark measures the delay from itself to all the others. From these inter-landmark measurements CBG can determine the *bestline* by data fitting [6]. In general, each delay measurement defines a circle around the landmark from where the delay was measured. The possible locations of the target node are determined by intersecting all of these circles. Most of the time this intersection produces a region in which the target node must be located.

In Fig. 1 the basic principles of CBG can be seen. In the ideal case the intersection is not empty and contains the real location of the target node. The case of valid distance estimation can be seen on Fig. 1a. If we underestimate the distances then it is possible that there are no points contained by all the circles. Therefore, the equation system is inconsistent and has no solution. This situation is shown by Fig. 1b. The effect of overestimation is not dangerous so much, but it can decrease the accuracy of the location estimates as it can be seen on Fig. 1c. In this case the region of intersection is larger which leads to much more solutions. On the other hand, this triangulation method might provide unreliable results for IP addresses that are located outside the convex hull of the landmarks. In this case the above intersection might be much larger then in the cases illustrated by Fig. 1.

Another technique is where the topology information and latency measurements are used in the location estimation. This method type is called *topology based geolocation* (TBG) [7]. TBG localizes all the intermediate routers be-

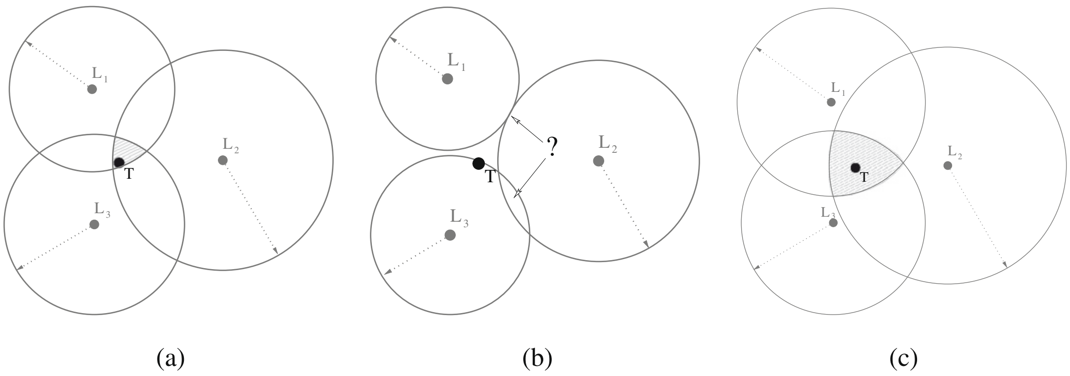

Fig. 1. The constraint-based geolocation use a triangulation-like method to determine a region in which the target is assumed to lie (a) shows correct distance estimates where the target node lies in a small intersection. The effect of and overestimation can be seen on (b) and (c).

tween the landmarks and the target node. This approach is based on link-latency estimations and on precise topology discovery. The basic tools of this method are `traceroute` and interface clustering applications.

In the following we introduce a method that combines `CBG` and `TBG` techniques. Opposite to the prior works, we use a detailed network model to determine geographical distances from the landmarks and we introduce new types of geographical constraints using one-way delay measurements.

3. Path-latency model

The delay experienced by a packet as it passes through the network is a sum of contributions from various phenomena. Based on [8,9] and [10] these phenomena can be divided into the following groups:

- *Routing delay* is the time that a packet spends inside a router. The routing delay can be modeled by the sum of the following three parts:
  - *Processing delay*, denoted by D_{pc}, represents the time that is needed for the router to determine where to forward the packet and to move the packet to the corresponding output port. Usually this delay can be modeled by the sum of a few router dependent constants and a time value that is proportional to the packet size.
  - *Queuing delay*, denoted by D_q, is a time that a packet waits until the preceding packets are transmitted.
  - *Other delays* which can appear in routers due to some occasional events which are not due to queuing and packet processing. This is usually neglected in modeling.
- *Transmission delay* is a time needed to place a packet onto a link. For a packet of size p and a link with capacity C the transmission delay can be expressed by $D_{tr} = p/C$.
- *Propagation delay* represents the time that is needed for a packet to reach from one end of a link to the other end. This is mostly related to the length of the link cable (e.g. the geographical distance) between the source and destination nodes. For a link with a physical length s in which the signal propagate with speed v the propagation delay can be expressed by $D_{pg} = s/v$.

The packet delay on hop i is a sum of the contributions listed above:

$$d^i = D^i_{pc} + D^i_q + D^i_{tr} + D_{pg}(\mathbf{n_{i-1}}, \mathbf{n_i}). \tag{1}$$

Since the packet delay is an additive metric, the per-hop delays can be summed up for each hop along the path. By this, the overall packet delay of a given packet can be written as:

$$d(\mathbf{s},\mathbf{d}) = \sum_{i=1}^{H} d^i = \sum_{i=1}^{H} \left(D^i_{pc} + D^i_q + D^i_{tr} + D_{pg}(\mathbf{n_{i-1}}, \mathbf{n_i}) \right), \tag{2}$$

where the number of hops is denoted by H, and the measurement path is $\mathbf{s} = \mathbf{n_0}, \mathbf{n_1}, \ldots, \mathbf{n_H} = \mathbf{d}$. From the above expression the only contribution that is related to the geographical distance is the propagation delay.

Usually the overall packet delay $d(\mathbf{s},\mathbf{d})$ is used to estimate the delay parameters, since in general it is not possible to directly measure the propagation delay. To determine the value of $D_{pg}(\mathbf{s},\mathbf{d}) = \sum_{i=1}^{H} D_{pg}(\mathbf{n_{i-1}}, \mathbf{n_i})$ overall propagation delay, we have to subtract the different kind of delay quantities from the overall packet delay

$$D_{pg}(\mathbf{s},\mathbf{d}) = d(\mathbf{s},\mathbf{d}) - \sum_{i=1}^{H} \left(D^i_{pc} + D^i_q + D^i_{tr} \right). \tag{3}$$

In case of no queuing the contribution of the queuing delay is neglected $D^i_q = 0$, while for a given probe packet size the contribution of the transmission delay is constant for all the probe packets at a given hop. Since we have no detailed information about the delay values at individual routers we treat the processing and transmission delays with a constant value $D^i_{pc} = D_{pc}$ and $D^i_{tr} = D_{tr}$ for each $i = 1, \ldots, H$ hop. In this case the overall propagation delay between source $\mathbf{s}$ and destination $\mathbf{d}$ can be written as

$$\begin{aligned} D_{pg}(\mathbf{s},\mathbf{d}) &= d(\mathbf{s},\mathbf{d}) - \sum_{i=1}^{H} \left(D^i_{pc} + D^i_{tr} \right) \\ &= d(\mathbf{s},\mathbf{d}) - H \cdot (D_{pc} + D_{tr}) = d(\mathbf{s},\mathbf{d}) - H \cdot d_h, \end{aligned} \tag{4}$$

which shows that the propagation delay can be determined from the overall packet delay $d(\mathbf{s}, \mathbf{d})$, the number of hops along the measurement path H and the constant per-hop delay value d_h.

In the following subsections we show how the components of the d_h per-hop delay can be approximated.

3.1. Transmission delay of a hop

Usually the transmission delay can be neglected, since its contribution to the total delay is very small. For example, if we use small ($p = 56$ bytes) IP packets over a $C = 1$ Gbps capacity link the contribution of the transmission delay is only $D_{tr} = p/C = 0.448\ \mu s$. Even the low capacity edge links with 100 Mbps cause an error of around 1 km. On the other hand, for routers with higher physical capacities the transmission delay could be even smaller. Hence we neglect the transmission delay:

$$D_{tr} \approx 0\ \mu s. \tag{5}$$

3.2. Processing delay of a hop

The processing delay represents the time needed to process an incoming packet and the time needed to prepare the packet for further transmission on the next link. This delay depends on several software and hardware factors, like the network protocol, the computational power of the router and also the efficiency of the network interface cards. Besides these variable factors we treat the processing delay as a constant value for each router along the network path: $D^i_{pc} = D_{pc}$ for $i = 1, \ldots, H$.

The specific value of the processing delay can be determined by measuring router transmission times at very low traffic intensities. The authors of [9] performed thorough investigation of the processing delay of UDP and ICMP probes in different kind of Cisco routers. They found that the processing delay is independent from the packet size, and the delay values for UDP and ICMP packets are almost the same: $D^{UDP}_{pc} \approx 97.9\ \mu s$ and $D^{ICMP}_{pc} \approx 101\ \mu s$. They also found some outliers due to some busy periods of the routers, which we ignore. In correspondence with these and other experimental results of [9,11] we use the

$$D_{pc} \approx 100\ \mu s, \tag{6}$$

approximation for both UDP and ICMP packets. Based on the observations (5) and (6) in this paper we use

$$d_h = D_{pc} + D_{tr} = 100\ \mu s, \tag{7}$$

as a constant per-hop delay for all routers along the measurement path. However, this per-hop value is treated as a parameter of the model, and can be easily fine-tuned based on newer results of router processing time experiments.

3.3. ICMP echo reply generation time

The delay estimation of various other geolocation techniques utilizes ICMP Echo (i.e. `ping`) measurements to approximate the overall delay between the target and the landmark nodes. To apply the previously described path-latency model for a `ping` measurement, a minor but important extension is necessary, which aims to include the generation time of the Echo Reply packets as an extra delay value. During a measurement the ICMP Echo Request packet takes up the previously described delays at each hop. After the probe packet reaches the destination node it is terminated and a newly generated ICMP Echo Reply packet is sent back. To model this process we need an additional term in the delay expression which describes the generation time of the Echo Reply packet.

In general the round-trip delay is meant as the sum of the delays collected on the forward and backward directions:

$$d_{rt} = d_{fw} + d_{bw}. \tag{8}$$

Assuming that the very same probe packet is going around from the source to the destination and then back to the source node, the ICMP based round-trip delay should be written as

$$d(\mathbf{s}, \mathbf{d}, \mathbf{s}) = d_{fw} + D_g + d_{bw}, \tag{9}$$

where D_g denotes the extra time elapsed due to generating the ICMP Echo Reply packet, while d_{fw} and d_{bw} represent the overall delay on the forward and backward directions, respectively.

In the case of symmetric routing, the D_g packet generation time can be determined by subtracting the overall round-trip time from the sum of the link-wise round-trip times. This special experiment scenario could not be performed in general, but the GEANT2 [2] research network provides services via its public Looking Glass interface that enabled us to collect the required path-and link-wise round-trip delays between the core routers. In Fig. 2 the schematic view of this measurement can be seen, where (a), (b) and (c) represent the link-wise round-trip measurements, while (d) represents the measurement of the overall round-trip delay of the path. The terms of (9) can be identified with the forward and backward arrows, while the D_g extra due to the Echo Reply packet generation is symbolized by a small clock in each row. To estimate D_g, the overall round-trip delay has to be subtracted from

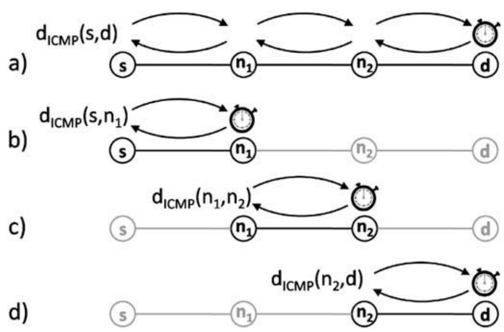

Fig. 2. Schematic view of the experiments performed in GÉANT2 network to determine the D_g ICMP Echo Reply packet generation time in case of a symmetric network path (a) shows the measurement of the overall round-trip delay of the path including the ICMP Echo Reply packet generation, while (b), (c) and (d) represent the link-wise round-trip measurements.

the sum of the link-wise round-trip delays, and the result should be divided by $H - 1$.

The packet generation times and the sum of propagation delays can be separated in the sum of link-wise round-trip delays in the following way:

$$\sum_{i=1}^{H} d(\mathbf{n_{i-1}}, \mathbf{n_i}, \mathbf{n_{i-1}}) = \sum_{i=1}^{H} (d_{\text{fw}}(\mathbf{n_{i-1}}, \mathbf{n_i}) + D_g + d_{\text{bw}}(\mathbf{n_{i-1}}, \mathbf{n_i})) = H \cdot D_g + \sum_{i=1}^{H} (d_{\text{fw}}(\mathbf{n_{i-1}}, \mathbf{n_i}) + d_{\text{bw}}(\mathbf{n_{i-1}}, \mathbf{n_i})), \quad (10)$$

where $\mathbf{n_0} = \mathbf{s}$ and $\mathbf{n_H} = \mathbf{d}$ denote the source and destination nodes. The overall round-trip delay can also be separated to the sum of the per-link one-way delays and a term describing the extra ICMP Echo Reply generation time:

$$d(\mathbf{s}, \mathbf{d}, \mathbf{s}) = D_g + \sum_{i=1}^{H} (d_{\text{fw}}(\mathbf{n_{i-1}}, \mathbf{n_i}) + d_{\text{bw}}(\mathbf{n_{i-1}}, \mathbf{n_i})). \quad (11)$$

From (10) and (11) the ICMP Echo Reply generation time is determined as

$$D_g = \frac{1}{H-1} \left(\sum_{i=1}^{H} d(\mathbf{n_{i-1}}, \mathbf{n_i}, \mathbf{n_{i-1}}) - d(\mathbf{s}, \mathbf{d}, \mathbf{s}) \right). \quad (12)$$

In these formulas we assume that the value of the D_g packet generation time is the same for all the routers along the path.

To determine the specific value of ICMP Echo Reply generation time we performed several experiments with the `GEANT2` Looking Glass service [2]. `GEANT2` connects 34 countries through 30 national research and education networks and provides basic services at user level. The Looking Glass service provides a `ping` and `traceroute` interface, which makes it possible to run experiments from all of their routers. Based on the `GEANT2` topological information we defined a large number of symmetric network paths between the core routers, where all the routers belonging to these paths are accessible via the Looking Glass interface. By performing `ping` measurements along these paths we were able to collect all the terms appearing in (12) that are needed to estimate the D_g value for the backbone routers. The data collection method is identical to the previously described measurement scenario shown in Fig. 2.

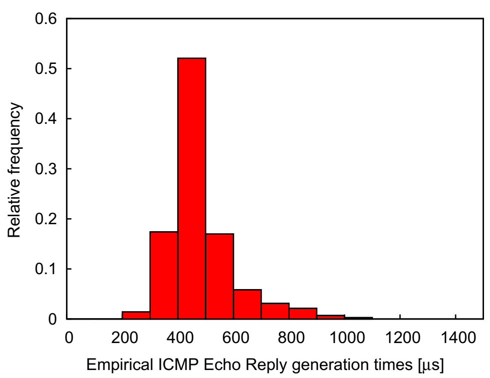

Fig. 3. Histogram of ICMP Echo Reply generation times based on RTT measurements. Each bin is 100 µs wide.

The histogram of the obtained empirical D_g packet generation times is shown in Fig. 3. The histogram has one peak around 500 µs, and the vast majority of the measured values are between 300 µs and 1000 µs. In the sense of our optimization method described in Section 7, the consistency of the estimation might be ensured by choosing the minimal

$$D_g = 300\ \mu s \quad (13)$$

value in the path-latency model. These measurement results are in accordance with the literature [12,13].

4. Distance approximation

In this section we discuss the relation between the overall propagation delay and the geographical distance. The real geographical distance s between the a source node $\mathbf{s}$ and a destination node $\mathbf{d}$ can be approximated by s^*, which can be written as

$$s(\mathbf{s}, \mathbf{d}) \leqslant s^* = c \cdot \frac{d(\mathbf{s}, \mathbf{d})}{2}, \quad (14)$$

where c denotes the speed of light in vacuum and $d(\mathbf{s}, \mathbf{d})$ represents the overall round-trip delay along the path. In Section 3 we assumed a symmetric network path and divided the overall round-trip delay into equal portions for the forward and backward directions. Since this approximation is too permissive, we need to elaborate firm constraints for the $d(\mathbf{s}, \mathbf{d})$ delay value and for the c signal propagation speed parameter used in the above approximation.

Regarding the c signal propagation speed, we might use an effective velocity considering the physical properties of all the links along the path, since their material constants are well known. Beyond these constants we must consider some kind of effective link curvature as well, since the network cables are not running straight from the source to their destination, due to practical reasons. To take these properties into account it is sufficient to introduce only one new parameter, since we cannot distinguish the effects of the varying signal propagation speeds and the varying cable lengths. Both have the very same symptom: the increased delay of probe packets. In this approach the s geographical distance can be written as

$$s(\mathbf{s}, \mathbf{d}) \leqslant s^* = c \cdot r \cdot \sum_{i=\mathbf{1}}^{H} D_{pg}(\mathbf{n_{i-1}}, \mathbf{n_i}) = c \cdot r \cdot D_{pg}(\mathbf{s}, \mathbf{d}), \quad (15)$$

where the one-way propagation delay is used instead of the half of the overall round-trip delay, and a new parameter r, the *geographic signal propagation rate* is introduced to describe the conversion rate between the real geographical distance and the measured signal propagation delay in c units. This parameter describes curvatures in all of the 3 dimensions including level differences as well. For obvious reasons this rate must be $0 < r < 1$. While (15) provides more precise approximation for the geographical distance than (14) does, for numerical distance estimations we need to know the value of r. An approximation of r is presented in the next subsection.

4.1. Velocity of signal propagation

The conversion between propagation delay and geographic distance is a key point in active geolocation techniques. To determine the relation between them, we need a conversion ratio to describe the velocity of signal propagation in the network cable considering that the cable is not running as the crow flies. This new parameter r is the ratio of a signal's transmission speed compared to the speed of light in vacuum.

Henceforward we use $c = c_{\text{vacuum}}$, since the introduced r parameter can both express the rate of signal propagation in any transport medium and the link curvatures. To determine the numerical value of r, we have collected a wide range of experimental results in the GEANT2 research network, using its Looking Glass service. From the known router positions we were able to determine their real geographic distances, while the latency values were also directly measured between them. Based on these data we approximated the r conversion rate according to the path-latency model. For neighboring nodes **a** and **b** the conversion rate is estimated as

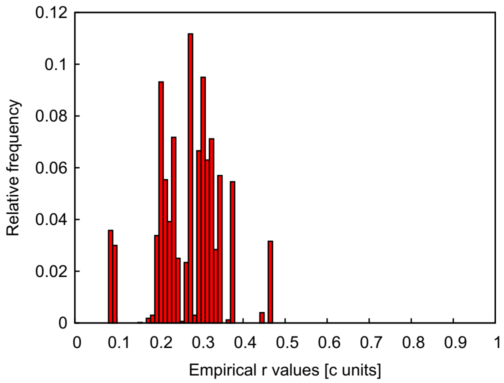

Fig. 4. Histogram of the empirical velocities of signal propagation in network. Each bin is 0.01 units wide.

$$r^*(\mathbf{a},\mathbf{b}) = \frac{s(\mathbf{a},\mathbf{b})}{c \cdot \frac{D_{pg}(\mathbf{a},\mathbf{b}) + D_{pg}(\mathbf{b},\mathbf{a})}{2}}, \tag{16}$$

where $D_{pg}(\mathbf{a},\mathbf{b})$ and $D_{pg}(\mathbf{b},\mathbf{a})$ are the link propagation delays from **a** to **b** and vice versa. The empirical minimum, average and maximum r values are 0.08, 0.27 and 0.47 respectively, as can be seen on Fig. 4. These observations are in accordance with the measurement results in [7].

Considering the evaluation method described in Section 7, the distance s^* in (15) should be an upper limit for the real distance s, otherwise the evaluation might lead to an inconsistent equation system in (28)–(31). Fig. 5 shows the estimated and the real geographical distances between GEANT-neighbors. It can be seen that the usage of the minimum and average ratios can lead to significant distance underestimation. On the other hand, by using the maximum velocity the probability of underestimation is small enough to avoid inconsistency in the evaluation. However, in this case the method usually overestimates the geographical distance, which of course decreases the accuracy of the geolocation estimation. As a conclusion of the above reasoning, we use the

$$r = 0.47, \tag{17}$$

value for the signal propagation in the network. We note that this value was measured on wired network paths since the connections between the GEANT core routers are terrestrials. If we examine the effects of satellite connections then it will be found that our model still remains consistent. It is a well known fact, that satellite connections generate very large latencies since the signal travels to the satellite and then back to the surface. Thus in this case the effective signal propagation velocities are much smaller, while the delay values are much larger than as it is with terrestrial connections. Although this phenomena might lead to distance over estimations, the consistency of our approach is not worsened, since it never causes distance underestimation.

5. Geographic constraints

The main goal in geolocation is to determine the unknown coordinates of network nodes. Geographic con-

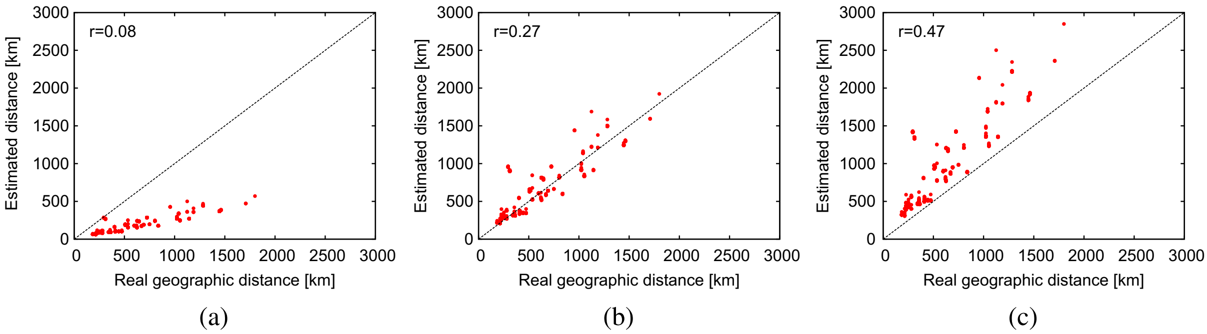

Fig. 5. The effect of using different signal propagation velocities in distance estimations. (a) shows the effect of using the minimum of the observed velocity values, which leads to distance underestimation. (b) represents the application of the average velocity value. In this case both under and overestimations occur. In (c) we use the maximum of the measured velocities, which decreases the possibility of distance underestimations.

straints can mark out areas in the world map in which the target nodes are located with high probability. These constraints are obtained via latency and topology measurements from the landmarks. In this section we overview the different type of measurements that can be used to define geographic constraints. First, we extend the well known `ping` based round-trip time measurement with the path-latency model to limit the distance between a landmark and a target node. Next, we describe a technique to estimate per-link latencies, and finally it is shown how one-way delay measurements can be used to define geographic constraints. In this paper we limit ourselves to using latency measurements in constraint definitions, and do not apply any passive methods.

5.1. Round-trip time based method

The `ping` tool is widely used in network measurements due to its simplicity and because it does not depend on having control over the target node. Although its inaccuracy is well known, the measured round-trip delays can give a constraint which defines an upper limit for the geographical distance between the landmark and the target node.

As we mentioned before, the round-trip time contains not only the delays on the forward, but also on the backward direction along the measured network path. If we generalize (9) then the following equation is obtained.

$$d(\mathbf{s},\mathbf{d},\mathbf{s}) = \sum_{i=1}^{H_{\text{fw}}} d^i + D_g + \sum_{i=H_{\text{fw}}+1}^{H_{\text{fw}}+H_{\text{bw}}} d^i = D_g + \sum_{i=1}^{H_{\text{fw}}+H_{\text{bw}}} \left(D^i_{pc} + D^i_q + D^i_{tr} + D_{pg}(\mathbf{n_{i-1}},\mathbf{n_i})\right), \quad (18)$$

where H_{fw} and H_{bw} denotes the number of hops on the forward and backward directions respectively. This delay formula is built up by the very same link-wise delay terms as (2), but for a longer network path including the possibly asymmetric backward direction as well. Here the delay terms belong to two different probe packets: the ICMP Echo Request packet on the forward direction, and the ICMP Echo Reply packet on the backward direction. We assume that the size of these two ICMP packets are the same, and thus their D_{tr} transmission and D_{pg} propagation delays can be treated in the same way along both directions. At the destination node an extra delay term appears due to the termination of the ICMP Echo Request packet and the generation of the ICMP Echo Reply packet. Taking this into account, the $D_{pg}(\mathbf{s},\mathbf{d},\mathbf{s})$ round-trip propagation delay is written as

$$D_{pg}(\mathbf{s},\mathbf{d},\mathbf{s}) = d(\mathbf{s},\mathbf{d},\mathbf{s}) - \sum_{i=1}^{H_{\text{fw}}+H_{\text{bw}}} \left(D^i_{pc} + D^i_{tr}\right) - D_g = d(\mathbf{s},\mathbf{d},\mathbf{s}) - (H_{\text{fw}} + H_{\text{bw}}) \cdot d_h - D_g, \quad (19)$$

where we assume equal processing and transmission delays, and negligible queuing delays at each hop. Based on (15) and the above expression we can limit the geographical distance between the source and destination nodes:

$$s(\mathbf{s},\mathbf{d}) \leqslant s^* = c \cdot r \cdot \min\{D_{pg}(\mathbf{s},\mathbf{d}), D_{pg}(\mathbf{d},\mathbf{s})\} \leqslant \frac{1}{2} \cdot c \cdot r \cdot D_{pg}(\mathbf{s},\mathbf{d},\mathbf{s}) = \frac{1}{2} \cdot c \cdot r \cdot (d(\mathbf{s},\mathbf{d},\mathbf{s}) - (H_{\text{fw}} + H_{\text{bw}}) \cdot d_h - D_g), \quad (20)$$

when the path is symmetric, its propagation delay is the same for both directions. In this case we can take the half of the round-trip time to approximate the one-way delay. It can be easily seen that in case of asymmetric paths, this approximation provides an upper bound for the minimal latency between the two endpoints. This is because either the uplink or the downlink delay is overestimated by this value. In this case, the real geographical distance could be even less than the above s^* limit.

5.2. Per-link delay estimation

Assuming that the network topology is known, we introduce a method to estimate per-link latencies, which then lead to the approximation of per-link distances. These information can improve the accuracy of the location estimation, since they give additional constraints to the geolocation. For all $\mathbf{n_i}$ and $\mathbf{n_j}$ nodes on the network path between the $\mathbf{s}$ source and $\mathbf{d}$ destination, the distance can be written as the difference of their distance from the common source node:

$$s(\mathbf{n_i},\mathbf{n_j}) = |s(\mathbf{s},\mathbf{n_i}) - s(\mathbf{s},\mathbf{n_j})|. \quad (21)$$

With the use of (20) this distance can be written as

$$s(\mathbf{n_i},\mathbf{n_j}) = \frac{1}{2} \cdot c \cdot r \cdot (d(\mathbf{s},\mathbf{n_j},\mathbf{s}) - d(\mathbf{s},\mathbf{n_i},\mathbf{s}) - (H(\mathbf{s},\mathbf{n_j},\mathbf{s}) - H(\mathbf{s},\mathbf{n_i},\mathbf{s})) \cdot d_h), \quad (22)$$

where $H(\mathbf{s},\mathbf{n_i},\mathbf{s})$ is the sum of the number of hops on the forward (H_{fw}) and backward (H_{bw}) directions. For consecutive $\mathbf{n_{i-1}}, \mathbf{n_i}$ nodes the distance only depends on the difference of their overall delay. Assuming that the round-trip delay to $\mathbf{n_i}$ is the greater than it is to $\mathbf{n_{i-1}}$, the absolute value can be dropped:

$$s(\mathbf{n_{i-1}},\mathbf{n_i}) = s(\mathbf{s},\mathbf{n_i}) - s(\mathbf{s},\mathbf{n_{i-1}}) = \frac{1}{2} \cdot c \cdot r \cdot (d(\mathbf{s},\mathbf{n_i},\mathbf{s}) - d(\mathbf{s},\mathbf{n_{i-1}},\mathbf{s}) - 2 \cdot d_h. \quad (23)$$

5.3. One-way delay as geographic constraint

If we can measure precise one-way delays between all the landmark nodes and the paths between them are also available, then the one-way delays provide the value of $d(\mathbf{s},\mathbf{d})$ in (2). Using (4) and (15) we can approximate the geographical distance along the full path as

$$\sum_{i=1}^{H-1} s(\mathbf{n_i},\mathbf{n_{i+1}}) \leqslant s^* = \cdot(d(\mathbf{s},\mathbf{d}) - H \cdot d_h). \quad (24)$$

This equation gives constraints for the coordinate variables of the $\mathbf{n}_i (i = 1,\ldots,H-1)$ nodes. In practice, these $\mathbf{n}_i$ nodes correspond with the hops along the measurement path between s and d. Since precise one-way delays can be measured only between landmarks, the constraint provided by the above equation can be applied to localize inter-landmark routers, while cannot be used for end-host targets. In Section 8 we show that using the one-way delay constraints we can improve the localization accuracy of network routers.

6. Data collection

In our geolocation system the data collection and the evaluation are separated into two independent phases. As a first step in determining the location of network nodes we collect delay data from round-trip and one-way delay experiments, while topological information are gathered by means of `traceroute` measurements. When all the data are available, we build up constraints and solve the given equation system as a global optimization problem.

Both the delay and topology data were collected in the `ETOMIC` system(European Traffic Observatory Measurement InfrastruCture) using its nodes as landmarks [1]. `ETOMIC` contains 18 PC based high-precision, GPS synchronized active probing nodes which are deployed at selected European locations. Every node is equipped with an Endace DAG card which open the door to high-precision end-to-end measurements between `ETOMIC` nodes (e.g. one-way delay measurements). These cards provide a way to time-stamp in-coming and out-going packages with 10 μs precision. All the collected measurement data are stored in the `ETOMIC`'s Network Measurement Virtual Observatory [14].

We measured both round-trip delays to every target node and one-way delays between landmark pairs. In our data collection a single round-trip delay measurement session contained 25 ICMP probes with 56 bytes packet size and the sessions were repeated 5 times in time-dispersed way. To measure one-way delay between `ETOMIC` nodes we used UDP packets with the same, 56 bytes packet size. To decrease the effect of the D_q queuing components in (4), every one-way delay measurement session contained more than 100 time-dispersed UDP probes. The UDP probes were sent and received by high-precision DAG cards to obtain accurate one-way delay values. For both type of delay values we determine the base values, i.e. the minimal round-trip and minimal one-way delay values. We assume that the effect of queuing delays are neglected in this way.

Besides collecting the delay values we also discover the topology among the measurement nodes. Therefore, a large number of `traceroute` experiments were collected between each `ETOMIC` node pairs. For the UDP probe packets of the `traceroute` experiments we used fixed sender and receiver ports to avoid multipath routing effects of load-balance routers [15]. The discovered inter-`ETOMIC` paths define a directed graph which contains network interfaces and connections between them according to the traceroute paths. However, several traceroute experiments can discover the same router, the assigned topological entities will be different, since the traceroute experiments arriving from different directions discover different interfaces of the same router. If the interfaces belong to the same router, we can cluster them into a single entity, since their geographical position is the same. By doing so we decrease the number of unknowns in the equation array to be optimized and we increase the precision and the stability of the geolocation optimization.

A recent, reliable technique the `Mercator` tool [16] was used to cluster the router interfaces and to reduce the number of nodes in the network graph. `Mercator` sends UDP probes to different ports of network interfaces to force the target router to reply with an ICMP Port-unreachable message. The source IP address of the error message is compared to the original destination address of the sent UDP probe. If they differ, then we found two interfaces that belong to the very same router, since they are aliases of each other. Fig. 6 represents how Mercator resolves aliases. We applied `Mercator` on a large interface-set which was collected by `traceroute` discovery between all the `ETOMIC` measurement node pairs. The original interface-set contained 1192 IP addresses, from which `Mercator` created 160 different clusters containing 584 interfaces. The remaining interfaces are also handled as clusters with a single interface. In this way we were able to identify 768 clustered nodes.

The applied clustering method significantly reduces the number of nodes in the topology and hence the number of unknowns in the evaluation method.

7. Solving geolocation as global optimization

We can handle the geolocation problem as a graph optimization task. The coordinates of the unknown routers represent the variables in the optimization problem. The

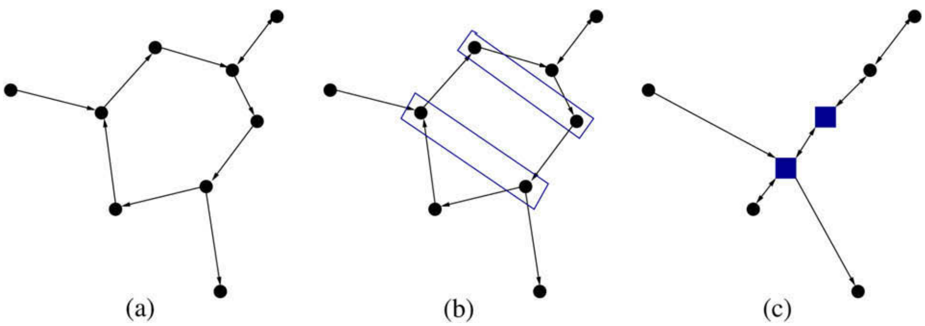

Fig. 6. Clustering interfaces (a) shows the original topology graph where the black nodes represent network interfaces. It can be seen in (b) that Mercator found two nontrivial clusters which contain interfaces belonging to the same router (c) contains the reduced topology graph where every node represents a single entity (e.g. a router).

goal is to determine the values of these variables according to the applied delay constraints.

Let us introduce the following notations: L is the set of landmarks, LAT and LON are the sets of possible latitude and longitude coordinates, respectively. Let X be a set of unknown router positions. In case we have m constraints related to one-way delay and round-trip time measurements, they can be written in a general form as:

$$f_i(L,X) \leqslant C_i^f (i = 1, \ldots, m). \quad (25)$$

Let's suppose, we have additional n constraints related to per-link delay experiments. These constraints can be written in a similar form:

$$g_j(L,X) = C_j^g (j = 1, \ldots, n). \quad (26)$$

In these formulas the C_i^f and C_j^g components are non-negative constants coming from the delay measurements directly. The left sides of these equations are non-linear functions of the coordinates to be localized. They contain the following great-circle distance in the constraint definitions:

$$s(\mathbf{a},\mathbf{b}) = \rho \cdot \arccos(\sin(\mathbf{a}_{lat}) \cdot \sin(\mathbf{b}_{lat}) + \cos(\mathbf{a}_{lat}) \cdot \cos(\mathbf{b}_{lat}) \cdot \cos(\mathbf{b}_{lon} - \mathbf{a}_{lon})),$$

where $\rho \approx 6371$ km is the Earth's radius.

As we mentioned before, the per-link delay estimations can be unreliable due to the inaccuracy of round-trip time measurements. To describe these estimation errors we introduce new $0 \leqslant E_j \in E$ variables, where E is the set of error variables. By this, (26) can be reformulated as follows:

$$g_j(L,X) = C_j^g + E_j (j = 1, \ldots, n). \quad (27)$$

We use these E_j variables to define the cost function of our optimization problem. The equation system can be solved by minimizing the overall tension in the network graph, i.e. minimizing the sum of the E_i^2 values. Hence, our geolocation problem can be formulated as the following global optimization task:

$$\min \sum_{E_j \in E} E_j^2, \quad (28)$$

$$f_i(L,X) \leqslant C_i^f (i = 1, \ldots, m), \quad (29)$$

$$g_j(L,X) = C_j^g + E_j (j = 1, \ldots, n), \quad (30)$$

$$x \in X : x_{lon} \in LON, \quad x_{lat} \in LAT. \quad (31)$$

There are well documented techniques [17] that can be applied to solve the above non-convex optimization problem. In this study we used a modified gradient method with adaptive step size to determine router locations by minimizing the objective function within the parameter space.

8. Performance analysis

Next we present the performance analysis of the evaluation, focusing on the role that different constraints played in the geolocation process. In our study, five different scenarios are investigated. The first three can be applied to localize not only routers but also end-hosts. These scenarios are the following. First, we analyze a simple case, called `Geo-R`, where only round-trip time constrains are considered. Next, we also take into account the hop-counts of the measurement paths (`Geo-Rh`). Next, also the link-latency constraints are used (`Geo-RhL`).

Besides the above constraints we introduce two more settings that use one-way delay constraints as well. As we mentioned in Section 5.3 this special constraint provides higher accuracy for the geolocalization of inter-landmark routers. However, this refinement cannot be generally applied for end-host localization. In the fourth scenario we use both round-trip and one-way delay constraints (`Geo-RhO`), and in the last setting all the introduced constraints are applied (`Geo-RhOL`). A short explanation of these notations can be seen in Tables 1 and 2.

By means of `ETOMIC` experiments we collected two reference node sets that contain backbone routers with known geographic positions (`Ref-1` and `Ref-2`). Locating a target node that is nearby a landmark, for example when they are in the same city, yields small estimation errors due to small measured delay values. To avoid misleading conclusions we selected nodes that are far from the landmarks. We performed several experiments to investigate the accuracy provided by our geolocation scenarios. Although the evaluation technique gives location estimates for all the nodes in the paths between the landmarks, the estimation errors are calculated only for the reference sets.

The `Ref-1` dataset contains 41 different interfaces, including `GEANT2` routers and several other ones. We define `Ref-2` to increase the reliability of the applied methods. This is a subset of `Ref-1` with 20 elements, where all the nodes are in the convex hull spanned by the `ETOMIC` landmarks. The presented approach is slightly different from the usual geolocation techniques, where any IP addresses can be localized. Although our methods can provide more accurate location estimation for nodes that

Table 1
Short explanation of notations used in the paper.

Notation	Short explanation
d	Overall delay along the network path
d_{fw}, d_{bw}	Delay suffered on forward And backward directions
d_h	Overall delay suffered at router h
$d(\mathbf{a},\mathbf{b})$	Overall delay from node **a** to node **b**
$d(\mathbf{a},\mathbf{b},\mathbf{a})$	Overall round-trip delay, Measured by ICMP packets
D_{pc}	Processing delay
D_q	Queuing delay
D_{tr}	Transmission delay
D_{pg}	Propagation delay
$D_{pg}(\mathbf{a},\mathbf{b})$	Propagation delay From node **a** to node **b**
$D_{pg}(\mathbf{a},\mathbf{b},\mathbf{a})$	Round-trip propagation delay From node **a** to node **b** and back to node **a**
D_g	Extra delay appearing due To packet generation At the destination node
$s(\mathbf{a},\mathbf{b})$	Geographical (great-circle) distance Between node **a** and node **b**
$r(\mathbf{a},\mathbf{b})$	Velocity of signal propagation(in c units) Between node **a** and node **b**

Table 2
Analyzed case studies.

Settings	Applied constraints
`Geo-R`	Special round-trip time constraints (the D_g ICMP Echo Reply generation time and the d_h processing and transmission delay components are eliminated from Eq. (20))
`Geo-Rh`	Round-trip time
`Geo-RhL`	Round-trip time and link-latency
`Geo-RhO`	Both round-trip time and one-way delay
`Geo-RhOL`	All type of introduced constraints

belong to the topology spanned by the landmarks. Hence, the direct comparison of our results to the prior works is not feasible.

In Fig. 7 and 8 we plotted the CDF curves of the location estimation errors for the two reference datasets. Both figures present all the five measurement scenarios. The main parameters of the observed error distributions can be seen in Table 3 and Table 4 for the reference sets.

By comparing `Geo-R` and `Geo-Rh` scenarios, one can observe that the median and mean errors are decreased by using the path-latency model, while on the maximum error the model has only a slight influence. In case of taking into account one-way delay constraints also the maximum error can be reduced, as it is shown in the table for `Geo-RhO` and `Geo-RhOL` scenarios. It is evident that the round-trip time based geolocation is inaccurate, since the area of possible locations could be wide due to the distance overestimation effect of inaccurate delay measurements. Using one-way delay constraints the range of possible locations of target nodes can be reduced significantly. In case of both reference datasets the observed estimation errors were much smaller than in the previous scenarios.

However, in case of both reference set high mean and standard deviation in the errors can be noticed. The main reason of this high values is the limited number of landmarks that were involved in the measurement: during the experimentation only 11–12 `ETOMIC` nodes were online. To give a more accurate reasoning for the higher error values we have investigated the worst localized reference nodes and found that the routes from the landmarks to these targets were very similar. We feel that for these nodes the measurement paths and the measurements themselves were not independent enough to provide more accurate geolocalization. We plan to deploy more `ETOMIC` nodes in the near future which will solve this problem and will improve the accuracy as well.

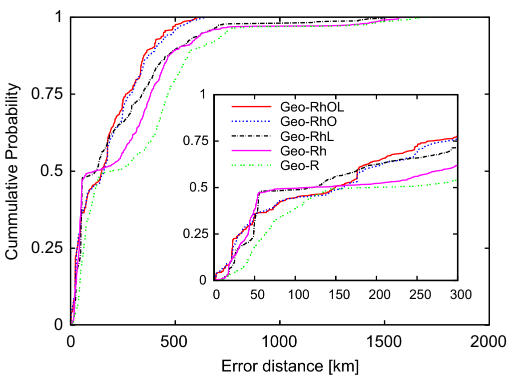

Fig. 7. Estimation errors of different scenarios for `Ref-1` reference data.

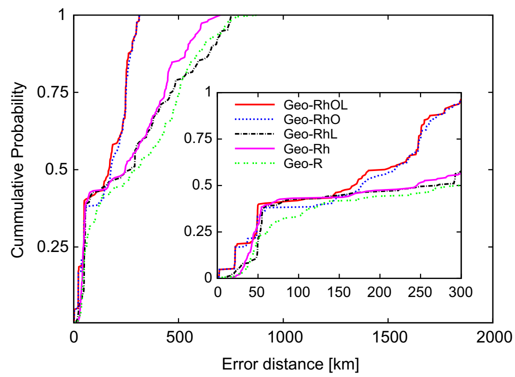

Fig. 8. Estimation errors of different scenarios for `Ref-2` reference data.

Table 3
Accuracy with different constraints on Ref-1 dataset (km).

Settings	Mean error	Maximum error	Standard deviation
`Geo-R`	304	1708	308
`Geo-Rh`	246	1602	288
`Geo-RhL`	213	1554	249
`Geo-RhO`	177	645	157
`Geo-RhOL`	169	609	149

Table 4
Accuracy with different constraints on Ref-2 dataset (km).

Settings	Mean error	Maximum error	Standard deviation
`Geo-R`	305	878	236
`Geo-Rh`	251	699	205
`Geo-RhL`	281	751	241
`Geo-RhO`	156	313	104
`Geo-RhOL`	149	312	104

In Fig. 9 we illustrate how the real and the estimated locations for some selected routers are related to each other in the worst (`Geo-R`) and the best (`Geo-RhOL`) cases.

8.1. *End-host localization using* `PlanetLab`

As we previously mentioned, the path-latency model and the round-trip time constraint (`Geo-Rh`) can be used to localize not only network routers but also end-hosts with public IP address. To localize end-hosts there are many existing methods like `GeoPing`, `CBG`, `TBG` or passive IP-to-location databases which were discussed in Section 2. Since `Geo-Rh` is an active measurement-based technique we choose `GeoPing` as a reference end-host geolocator to introduce the capabilities of our model based approach.

Next we present a world wide experiment which involves 151 different `PlanetLab` nodes. During the experi-

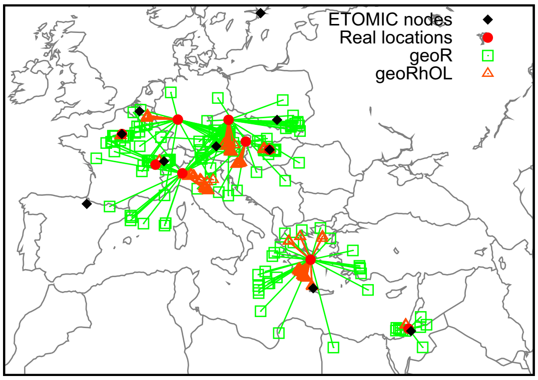

Fig. 9. Deviations between the estimated coordinates and the real locations for selected routers of REF-2 for two scenarios. One can see that the symbols representing the locations estimated based on the detailed path-latency model(Geo-RhOL) are much closer to the real locations than the symbols representing the estimations based on only the round-trip delay measurements (Geo-R).

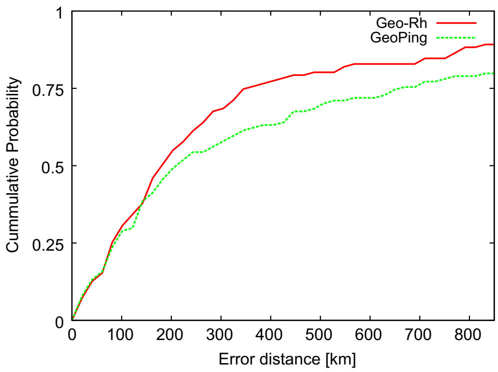

Fig. 10. Estimation errors of end-host localization for PlanetLab dataset.

mentation these nodes were used both as targets and landmarks. We localized each host using Geo-Rh and GeoPing methods to compare their accuracy. We selected a target node from the PlanetLab nodeset and considered the remaining hosts as landmarks to perform the location estimation. This procedure was repeated to evaluate the resulting location estimation for each host in the PlanetLab dataset.

Fig. 10 shows the CDF curves of location estimation errors for the PlanetLab nodes performing the above two methods. The mean error distances of Geo-Rh and GeoPing evaluations were 437 km and 713 km, respectively. However, in case of Geo-Rh 66% of the target nodes were localized with less than 280 km error, while in case of GeoPing this error was 460 km.

It can be seen that our basic method provides apparently better accuracy than GeoPing, while we can assume that the Geo-RhO and Geo-RhOL methods would perform even better. Unfortunately these improved methods could not be tested in this wider verification area, since PlanetLab today does not provide high-precision one-way delay measurements between its node pairs.

9. Conclusion

In this paper we investigated a model based approach of geolocation to demonstrate how an accurate approximation of the propagation delay can improve the accuracy of distance estimation. In contrast to the prior works, our path-latency model enables the separation of the propagation and per-hop delays in the overall packet latency. Besides the model based approximation of propagation delays, high-precision one-way delay measurements are used to define novel geographic constraints. This type of constraint yields additional information into the geolocation process by limiting the overall physical length of a given measurement path. This paper demonstrate that both the detailed path-latency model and the novel one-way delay constraints can significantly increase the accuracy of location estimates. The investigated techniques are tested and validated in a wide range of experiments performed in the ETOMIC and PlanetLab measurement infrastructures and in the GEANT2 research network. In the future we will able to extend the number of landmark nodes and routers used for reference node set with the precise active measurement infrastructure of the OneLab2 project [18]. The introduced method can be included in existing geolocation frameworks to improve their accuracy.

Acknowledgements

The authors thank the partial support of the National Office for Research and Technology (NAP 2005/KCKHA005) and the EU ICT OneLab2 Integrated Project (Grant agreement No. 224263).

References

[1] D. Morato, E. Magana, M. Izal, J. Aracil, F. Naranjo, F. Astiz, U. Alonso, I. Csabai, P. Hága, G. Simon, J. Steger, G. Vattay, ETOMIC: a testbed for universal active and passive measurements, in: Proceedings of the IEEE TRIDENTCOM 2005, Best Testbed Award, Trento, Italy, 23–25 February 2005, pp. 283–289.
[2] GÉANT network <http://www.geant2.net/> and Looking Glass service <http://www.stats.geant2.net/lg/>.
[3] K. Harrenstien, M. Stahl, E. Feinler, RFCc954- NICNAME/WHOIS, IETF, October 1985.
[4] Venkata N. Padmanabhan, Lakshminarayanan Subramanian, An investigation of geographic mapping techniques for internet hosts, in: Proceedings of the ACM SIGCOMM, San Diego, CA, USA, August 2001, pp. 173–185.
[5] M. Zhang, Y. Ruan, V. Pai, J. Rexford, How DNS misnaming distorts internet topology mapping, in: Proceedings of the USENIX Annual Technical Conference, 2006.
[6] Bamba Gueye, Artur Ziviani, Mark Crovella, Serge Fdida, Constraint-based geolocation of internet hosts, IEEE/ACM Transactions on Networking 14 (6) (2006) 1219–1232.
[7] E. Katz-Bassett, J. John, A. Krishnamurthy, D. Wetherall, T. Anderson, Y. Chawathe, Towards IP geolocation using delay and topology measurements, in: Proceeding of the AMC IMC 2006, Rio de Janeiro, Brazil, 2006, pp. 71–84.
[8] Mark Crovella, Balachander Krishnamurthy, Internet Measurement: Infrastructure, Traffic and Applications, John Wiley & Sons Inc., 2006.
[9] D. Constantinescu, A. Popescu, Modeling of one-way transit time in IP routers, in: IEEE Advanced International Conference on Telecommunications (AICT '06), Guadeloupe, French Caribbean, February 2006.
[10] Ana Hernandez, Eduardo Magana, One-way delay measurement and characterization, in: IEEE ICNS '07, Washington, DC, USA, p. 114.

[11] K. Mochalski, J. Micheel, S. Donnelly, Packet delay and loss at the auckland internet access path, in: Proceedings of the PAM 2002, Fort Collins, Colorado, March 2002, pp. 46–55.
[12] Fabien Viger, Active probing with ICMP packets, Technical Report, 2003.
[13] R. Govindan, V. Paxson, Estimating router ICMP generation delays, in: Proceedings of the Passive and Active Measurements Workshop(PAM), 2002.
[14] P. Mátray, I. Csabai, P. Hága, J. Stéger, L. Dobos, G. Vattay, Building a prototype for network measurement virtual observatory, in: Proceedings of the ACM SIGMETRICS-MineNet 2007, San Diego, CA, USA, 12 June 2007.
[15] Weiguang Shi, M.H. MacGregor, Pawel Gburzynski, Load balancing for parallel forwarding, IEEE/ACM Transaction Network 13(4) 790–801.
[16] R. Govindan, H. Tangmunarunkit, Heuristics for internet map discovery, in: IEEE INFOCOM 2000, Tel Aviv, Israel, March 2000, pp. 1371–1380.
[17] L. Vandenberghe, S. Boyd, Semidefinite programming, SIAM Review 38 (1) (1996).
[18] OneLab Future Internet Test Beds project, <http://www.onelab.eu/>.

Sandor Laki received the M.Sc. degree in computer science from the Eotvos Lorand University, Budapest, Hungary, in 2007. Currently he is working towards his Ph.D. in computer science at the Department of Physics of Complex Systems.
His primal research interest includes Internet measurement techniques, adaptive protocols and network modelling, especially IP geolocation and bandwidth estimation. Since 2008 he is developing techniques to improve the accuracy of measurement-based IP geolocation and participating in numerous national and EU research projects.

Peter Matray received his M.Sc. degree in mathematics and computer science from Eotvos Lorand University, Budapest, Hungary in 2005.
Currently he is doing his Ph.D. studies at the Department of Physics of Complex Systems. His attention is mainly focused on the database aspects of network measurements. Besides that, he is partly involved in active probing Internet measurements (especially network tomography) and visualization.

Peter Haga received his Ph.D. degree in physics from the Eotvos Lorand University, Budapest, Hungary in 2008. From 1999 he is a permanent member of the Communication Networks Laboratory at the Eotvos Lorand University.
He was one of the founders of the ETOMIC measurement infrastructure. Currently he is a teaching assistant at the Department of Physics of Complex Systems. His research interest includes network measurement techniques, traffic modelling and adaptive protocols.

Istvan Csabai (Associate Professor, Department of Physics of Complex Systems, Eotvos University, Budapest) received the Masters Degree (1989) and Ph.D. (1992) in physics from Eotvos University.
His research interests covers wide areas of data mining, large scientific databases, modelling and measuring communication networks, simulation of complex systems. He was among the first authors (1994), who have identified and analysed the self-similar nature of internet traffic. Leader and participant of several national and international research projects.

Gabor Vattay (Full Professor, Eotvos University, Budapest) received his Ph.D. in physics from the Hungarian Academy of Sciences in 1994. Founder and Director of the Communication Networks Laboratory at Eotvos University since 2000. Leader of the EU supported Internet measurement effort European Traffic Observatory Measurement Infrastructure (http://www.etomic.org) at Collegium Budapest, Institute for Advanced Study.
He published about 40 papers in physics and networking and edited the book Complex Dynamics in Communication Networks. His interest includes modeling, measuring and design of communication networks.