

Leveraging Buffering Delay Estimation for Geolocation of Internet Hosts

Bamba Gueye¹, Steve Uhlig^{2,*}, Artur Ziviani³, and Serge Fdida¹

¹ Université Pierre et Marie Curie,
Laboratoire d'Informatique de Paris 6 (LIP6)
{gueye, fdida}@rp.lip6.fr

² Université Catholique de Louvain,
Department of Computing Science and Engineering
suh@info.ucl.ac.be

³ National Laboratory for Scientific Computing (LNCC)
ziviani@lncc.br

Abstract. Geolocation techniques aim at determining the geographic location of an Internet host based on its IP address. Currently, measurement-based geolocation techniques disregard the buffering delays that may be introduced at each hop along the path taken by probe packets. To fill this gap, we propose the *GeoBuD* (**Ge**olocation using **B**uffering **D**elay estimation) approach. Although the network delay and the geographic distance between two Internet hosts have been shown to be related to some extent, leveraging buffering delay estimation at each hop for geolocation purposes is challenging for two reasons. First, correctly estimating the buffering delay at intermediate hops along a traceroute path for geolocation purposes depends on the accurate estimation of the geolocation of the intermediate routers. Second, even given an *a priori* knowledge of the location of the routers, estimating the buffering delays is difficult due to the coarse-grained information provided by delay measurements. Relying on traceroute measurements, we show that leveraging buffering delay estimation improves accuracy in the measurement-based geolocation of Internet hosts as well as the confidence that the geolocation service associates to each estimation.

Keywords: geolocation, buffering delay estimation, traceroute, multilateration.

1 Introduction

Geographically locating an Internet host from its IP address enables a diversified class of location-aware applications [1, 2, 3]. Examples of such applications comprise targeted advertising on web pages, displaying local events and regional weather, automatic selection of a language to first display the content of web pages, restricted content delivery following regional policies, authorization of transactions only when performed from pre-established locations, or locating pedo-criminality. Each application may have a different requirement on the resolution of the location estimation. Nevertheless, as IP

* Steve Uhlig is Postdoctoral fellow of the Belgian National Fund for Scientific Research (F.N.R.S).

addresses are in general allocated in an arbitrary fashion, there is no inherent relation between an IP address and the physical location of the corresponding physical interface. Therefore, inferring the geographic location of Internet hosts is a challenging problem.

Previous work on measurement-based geographic location of Internet hosts [4, 5, 6] relies on delay measurements between *landmarks*, *i.e.* hosts with well-known geographic location, to provide the position of a target host. In GeoPing [4], the positions of landmarks are used as the possible location estimates for a given target host. This leads to a discrete space of answers that may limit location accuracy because of the system's dependence on the number and placement of landmarks [5]. The *Constraint-Based Geolocation* (CBG) approach proposed by Gueye *et al.* [6] transforms delay measurements into distance constraints and then uses *multilateration* to estimate the geographic location of a given target host. Multilateration refers to the process of estimating a position using a sufficient number of distances to some fixed points, thus establishing a continuous space of answers instead of a discrete one. This multilateration with distance constraints provides an overestimation of the distance from each landmark to the target host to be located, thus determining a region that hopefully encloses the location of the target host. The centroid of this region is the location estimation provided by CBG. Further, the area size of this region is a confidence measure CBG associates with each given location estimation; the smaller the region, the more confident the system is in the provided estimation. Although showing relatively accurate results in most cases, these measurement-based approaches may have their accuracy disturbed by many sources of distortion that affect delay measurements. For example, delay distortion may be introduced by the circuitous Internet paths that tend to unnecessarily inflate the end-to-end delay [7, 8, 9]. Another source of distortion is the unpredictable buffering delay that packets face in queues at the intermediate routers along the end-to-end path. For an accurate geolocation of Internet hosts based on delay measurements, it is crucial to estimate and remove as much of the additional delay as possible.

This paper investigates the distortion introduced in delay measurements due to buffering delays and possible counter-measures to it in order to improve geolocation techniques. We present GeoBuD, a novel way of geolocating Internet hosts. We rely on traceroute measurements to estimate the buffering delay introduced along the path from each landmark to the target. Based on traceroute information about the successive RTTs at each intermediate hop, we estimate the buffering delay introduced by each of these hops. Our results show that the estimation of buffering delays introduced along the path allows the improvement of the geolocation estimation given by CBG. This is so because the additional delay distortions caused by buffering delay are removed from the overestimations of distance constraints that define the region enclosing the target host in CBG, thus allowing tighter overestimations that result in a smaller region. Smaller regions that still enclose the target host provide more accurate location estimation in CBG.

The remainder of the paper is structured as follows. Section 2 discusses the related work. Section 3 describes the CBG approach to estimate the geographic location of a given target host. Section 4 explains our methodology for performing the traceroutes and estimating the buffering delays along the traceroute measurements. Section 5 compares the results of GeoBuD and those of the CBG approach. Finally, Section 6 concludes our paper and discusses future work.

2 Related work

A DNS-based approach to provide a geographic location service of Internet hosts is proposed in RFC 1876 [10]. Nevertheless, the adoption of the DNS-based approach has been limited since it requires changes in the DNS records and administrators have little motivation to register new location records. Tools such as IP2LL [11] and NetGeo [12] query Whois databases in order to obtain the location information recorded therein to infer the geographic location of a host. This information, however, may be inaccurate or stale. Moreover, if a large and geographically dispersed block of IP addresses [13] is allocated to a single entity, the Whois databases may contain just a single entry for the entire block.

There are also some geolocation services based on an exhaustive tabulation between IP addresses ranges and their corresponding locations. Examples of such services are GeoURL [14], the Net World Map project [15], and several commercial tools [1, 2, 3].

Padmanabhan and Subramanian [4] investigate three different 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. This technique is the basis of GeoTrack [4], Visual-Route [16], GTrace [17], and the SarangWorld Traceroute project [18]. Quite often network operators assign names to routers that have some geographic meaning, presumably for administrative convenience. Nevertheless, not all names contain an indication of location. Since there is no standard, operators commonly develop their own rules for naming their routers even if the names are geographically meaningful. Therefore, the parsing rules to recognize a location from a node name must be specific to each operator. The creation and management of such rules is a challenging task as there is no standard to follow.
- 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. Knowing the location of some hosts in the cluster and assuming they are in agreement, the technique infers the location of the entire cluster. An example of such a technique is GeoCluster [4]. This technique, however, relies on information that is partial and possibly inaccurate. The information is partial because it comprises location information for a relatively small subset of the IP address space. Moreover, such information may be inaccurate because the databases rely on data provided by users, which may be unreliable.
- The third technique, GeoPing [4], is based on exploiting a possible correlation between geographic distance and network delay. 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. This assumption is similar to the one exploited by wireless positioning systems such as RADAR [19] concerning the relationship between signal strength and distance. Therefore, given a set of landmarks with a well-known geographic location, the location estimation for a target host is the location of the landmark presenting the most similar delay pattern to the one observed for the target host. In GeoPing, the number of possible location estimates is limited to the number of adopted landmarks, characterizing a discrete space of

answers. As a consequence, the accuracy of this discrete space system is directly related to the number and placement of the adopted landmarks [5].

To overcome the limitation of using a discrete space of answers, the Constraint-Based Geolocation [6] approach uses multilateration to yield a continuous space of answers. In the next section, we provide a brief background on how the CBG methodology operates as the goal in this paper is to investigate the impact that the leveraging of buffering delay can have on the geolocation of Internet hosts based on multilateration.

3 Background on the CBG Approach

In this section, we present a brief background on how CBG provides geolocation estimation for target hosts based on delay measurements.

3.1 Multilateration with Geographic Distance Constraints

The physical position of a given point can be estimated using a sufficient number of distances or angle measurements to some fixed points whose positions are known. When dealing with distances, this process is called multilateration.

Consider a set $\mathcal{L} = \{L_1, L_2, \dots, L_K\}$ of K landmarks. Landmarks are reference hosts with a well-known geographic location. For the location of Internet hosts using multilateration, CBG [6] tackles the problem of estimating the geographic distance from these landmarks towards the target host to be located, given the delay measurements from the landmarks. From a measurement viewpoint, the end-to-end delay over a fixed path can be split into two components: a deterministic (or fixed) delay and a stochastic delay [20]. The deterministic delay is composed by the minimum processing time at each router, the transmission delay, and the propagation delay. This deterministic delay is fixed for any given path. The stochastic delay comprises the queuing delay at the intermediate routers and the variable processing time at each router that exceeds the minimum processing time. Besides the stochastic delay, the conversion from delay measurements to geographic distance is also distorted by other sources as well, such as circuitous routing and the presence of redundant data. Anyway, it should be noted that no matter the source of distortion, this delay distortion is always additive with respect to the minimum delay of an idealized direct great-circle path.

Figure 1 illustrates the multilateration in CBG using the set of landmarks $\mathcal{L} = \{L_1, L_2, L_3\}$ in the presence of some additive distance distortion due to imperfect measurements. Each landmark L_i intends to infer its geographic distance constraint to a target host τ with unknown geographic location. Nevertheless, the inferred geographic distance constraint is actually given by $\hat{g}_{i\tau} = g_{i\tau} + \gamma_{i\tau}$, *i.e.* the real geographic distance $g_{i\tau}$ plus an additive geographic distance distortion represented by $\gamma_{i\tau}$. This purely additive distance distortion $\gamma_{i\tau}$ results from the possible presence of some additive delay distortion. As a consequence of having additive distance distortion, the location estimation of the target host τ should lie somewhere within the gray area (*cf.* Figure 1) that corresponds to the intersection of the overestimated geographic distance constraints from the landmarks to the target host.

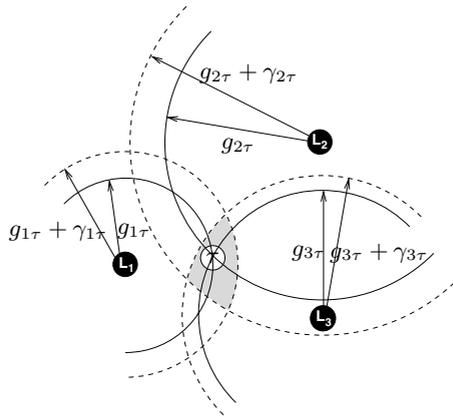


Fig. 1. Multilateration with geographic distance constraints

3.2 From Delay Measurements to Distance Constraints

Recent work [21, 4, 22] has investigated the correlation between geographic distance and network delay. Figure 2 provides an example of the relation between the distance and the delay for one of the landmarks we used in our measurements towards the remaining landmarks of our dataset (further details on the experimental data used are found in Section 5). The *bestline* shown in Figure 2 for a given landmark L_i is defined as the line that is closest to, but below all data points (x, y) , where x expresses the actual great-circle geographic distance between this given landmark and all the other landmarks in the set, while y represents the measured RTT between the same pairs. The equation of the bestline is defined as

$$y = m_i x + b_i. \tag{1}$$

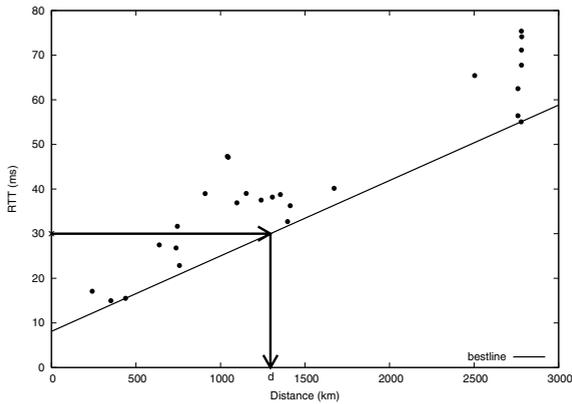


Fig. 2. Sample scatter plot of geographic distance and network delay

It should be noted that each landmark finds its slope m_i and its positive intercept b_i based only on delay measurements between the available landmarks. For further details about the computation of b_i and m_i , we refer the reader to [6]. The presence of a positive intercept b_i in the bestline reflects the presence of some localized delay. Each landmark uses its own bestline to convert the delay measurement towards the target host into a geographic distance constraint. A delay measurement from the considered landmark of Figure 2 towards a particular target host τ is transformed into a distance constraint by projecting the measured delay on the distance axis using the computed bestline of this landmark. For example, if the measured delay is 30 ms, the distance constraint is d , as illustrated by the thick arrow in Figure 2. This estimated geographic distance constraint $\hat{g}_{i\tau}$ between a landmark L_i and a target host τ is derived from the delay $d_{i\tau}$ using the bestline of the landmark as follows:

$$\hat{g}_{i\tau} = \frac{d_{i\tau} - b_i}{m_i}. \quad (2)$$

Each landmark L_i localizes a given destination τ inside a circle whose radius is the obtained distance constraint $\hat{g}_{i\tau}$. The region formed by the intersection of all these circles from the set of landmarks is called in CBG the *confidence region*. CBG provides the centroid of this confidence region as the location estimation for the target host.

4 Buffering Delay Estimation Via Traceroutes

CBG builds its distance constraints on a per-landmark basis. In contrast, based on a per-destination, GeoBuD transforms delay measurements into distance constraints. In practice, paths for different destinations may suffer from different distortions. To take that into account, we first replace the linear model of CBG by decomposing it in a per-hop basis. So in the GeoBuD approach, for each landmark L_i and each target host τ , we model the delay $y_{i\tau}$ as

$$y_{i\tau} = m_i x_{i\tau} + b_{i\tau}, \quad (3)$$

where m_i represents the propagation speed of data along the path computed only between the landmarks, $x_{i\tau}$ represents the geographic distance constraint between landmark L_i and destination τ , and $b_{i\tau}$ represents the total buffering delay along the path from L_i to a target host τ . In our measurements, the value of m_i actually represents 2 times the propagation speed of light in fiber, as m_i captures both the signal propagation aspect and the fact that the delay on which we are relying is a RTT hence contains both the forward and the return path. To estimate the total buffering delay $b_{i\tau}$, we estimate the buffering delay at each hop along the path from L_i to target τ based on traceroute between these nodes. The output of the traceroute measurements hopefully provides the different intermediate nodes that compose the path, as well as the delay between each pair of consecutive intermediate router.

For example, suppose that we perform a traceroute from landmark L_i towards some target host τ . The traceroute is composed of n intermediate hops, the last hop being the one that arrives at the target node. For each hop k of the traceroute that answers with an ICMP message `TIME exceeded`, we have an RTT measurement. If by any reason an intermediate router along the traceroute path does not answer with the ICMP message

TIME exceeded, we disconsider this hop as we lack a delay measurement for this particular node. To estimate $b_{i\tau}$ in Equation (3), we actually estimate its components b_k along the traceroute path using

$$\Delta RTT_{k+1} = RTT_{k+1} - RTT_k = m_i \times dist(k, k + 1) + b_{k+1}, \quad (4)$$

where k represents the k^{th} intermediate router on the traceroute path for which we were able to have a delay measurement and geographical location. The term RTT_k denotes the minimum RTT value out of the 3 RTTs measurements obtained for a given hop¹ k and $dist(k, k+1)$ represents the geographic distance between nodes k and $k + 1$. Note that m_i is the same as the one in Equation (1). The sum of the $dist(k, k + 1)$ for each k from 0 to $n - 1$ gives the estimation of the geographic path length followed by the traceroute. Thus, we estimate the buffering delay b_k at each hop k in a straightforward way from Equation (4) as

$$b_k = \Delta RTT_k - m_i \times dist(k - 1, k). \quad (5)$$

It is clear from Equation (5) that we need to estimate the geographic distance between each pair of consecutive intermediate routers along the traceroute path in order to be able to estimate b_k . This implies knowing the geographic location of these routers. It is unlikely to have an *a priori* knowledge of the geographic location of all possible intermediate routers along a traceroute path. This would actually amount to being able to geolocate any Internet node, *i.e.* the actual intent of the geolocation service under investigation. Therefore, the estimation of buffering delay along the path demands successive use of the geolocation service on each node identified along the traceroute path until reaching the target host.

5 GeoBuD Evaluation

In order to estimate the buffering delay b_k and the $dist(k, k + 1)$ at each hop k along a path between each landmark and each target host τ , we have considered two datasets: First, we considered nodes located in the U.S. We have used 29 PlanetLab nodes [24] as landmarks and 87 AMP nodes [25] as targets. The dataset we consider is composed by traceroute measurements performed on October 17th 2005 from our landmarks towards targets hosts. For the second dataset we performed on November 21st 2005 traceroutes from 27 PlanetLab nodes located in Western Europe towards 57 RIPE nodes [26], also located in Western Europe. In CBG methodology, landmarks perform ping measurements towards a given target host to locate it. Traceroute measurements from the same landmarks towards the same targets were performed simultaneously with ping measurements, in order to have similar network conditions for both CBG and traceroute measurements.

In our traceroute experiments, for the U.S. dataset, we have been able to geolocate 1153 distinct intermediate routers excluding the AMP hosts, out of a total of 1408 traversed routers, thus leading to geolocating 82% of the intermediate routers. In the W.E

¹ Each step of a traceroute consists in sending 3 consecutive UDP packets towards the destination using an increasing TTL value. In our measurements we rely on native traceroutes [23].

dataset we have located 1235 routers among the 1328 routers that we have encountered. So 93% of W.E. routers are located. It should be noted that most of undiscovered routers are typically located in the vicinity of the source or the destination, so the resulting error in the traceroute path length estimation is due to be small. For each of these located routers, we relied on the CBG-based GeoLIM project [27] to find out their geographic location. We cross-checked the results obtained with the GeoLIM project with *rockettrace* provided by the scriptroute tool suite [28]. In addition, some hops along the traceroutes underwent congestion by the time the measurements are carried. These traffic conditions, however, may not last for the whole time of the traceroute measurements. In such a case, only a few intermediate hops along the traceroute path exhibit a very large RTT value. If any of the intermediate hops along a traceroute exhibits a RTT value larger than any RTT of its succeeding hops along the traceroute path, we disconsider this particular hop. If we were to take such inflated RTTs into account we would overestimate the buffering delay for that hop.

After performing the geolocation of the intermediate routers, we compute the set of b_k values along each of different traceroute paths for the located intermediate nodes using Equation (4). In some cases, the estimated b_k are negative, in which case they were not considered. We had 21% of negative b_k 's corresponding to 4043 among 19172 b_k 's that we have computed for AMP hosts. For RIPE hosts the percentage of negative b_k 's is 14% for 11908 b_k 's found. Most cases where the b_k were negative correspond to situations where ΔRTT_{k+1} is very small or negative. This is due to variations in the network conditions along the path of the traceroute during the experiments. Hence, to consider a particular b_k , we require that $\Delta RTT_{k+1} > 0$.

For each landmark L_i and target τ , we have then a corresponding $b_{i\tau} = \sum_{k=1}^{n-1} b_k$, where n denotes the number of intermediate hops along the traceroute path from landmark L_i and target τ . To transform delay measurements into distance constraints, we can use the following equation derived from Equation (3):

$$x_{i\tau} = \frac{y_{i\tau} - b_{i\tau}}{m_i}. \quad (6)$$

GeoBuD uses the distance constraints given by Equation (6) to localize a given target host. The distance constraints obtained by GeoBuD are expected to be tighter than those provided by the CBG method. Using these new tighter distance constraints, in spite of the number of negative b_k 's, the confidence region shrinks, thus increasing both the accuracy of the location estimation and the system's confidence on these location estimation as shown in Section 5.1.

5.1 Shrinking the Confidence Region

Figure 3 compares the cumulative probability distributions of the confidence region of GeoBuD and the CBG approach. On the x -axis, we have the surface area of the confidence region for different target hosts. On the y -axis, we show the probability that the location estimation for the target hosts have a confidence region smaller than x .

One can observe on Figure 3 the improvement due to buffering delay estimation for areas smaller than 10^7 km². With CBG, 72% of the target hosts located in the U.S. have a confidence region smaller than 10^6 km². For GeoBuD and the same confidence region

surface, we have about 86% of the target hosts. With CBG, 49% of the target hosts have a confidence region smaller than 10^5 km^2 , whereas for GeoBuD 63% of the target hosts are within such a confidence region. For hosts located in W.E., GeoBuD localizes 10% of the target hosts with a confidence region inferior to 10^2 km^2 . For reference, a surface area of 10^5 km^2 is slightly larger than Portugal or the U.S. state of Indiana.

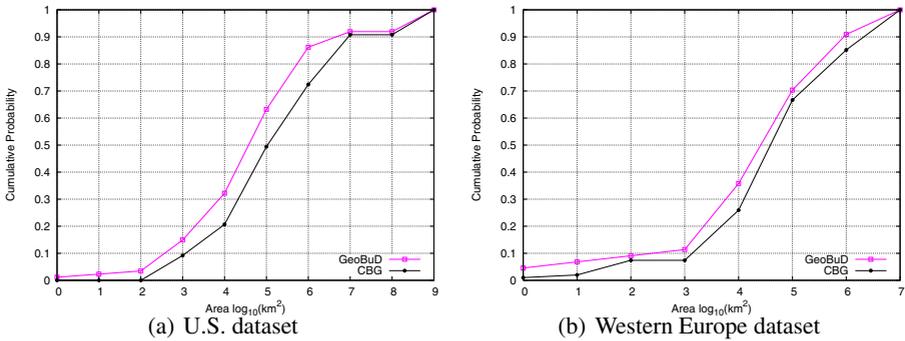


Fig. 3. Confidence regions provided by GeoBuD and CBG in km^2

5.2 Location Estimation Error

We might expect that reducing the surface area of the confidence region by estimating the buffering delay would also reduce the error observed in location estimation. In Figure 4, we show the cumulative probability of the error observed in the obtained location estimation. The estimation error for a given target host is the difference between its actual geographic location and its location estimate. The performance gap between GeoBuD and CBG is larger in the U.S. dataset. For the U.S. dataset, 80% of the target hosts, the estimation error is smaller using GeoBuD compared to CBG. The median of the location estimation error is of 144 km for GeoBuD, while of 228 km for CBG. In the W.E. dataset, it is 100 km and 137 km for GeoBuD and CBG respectively.

5.3 Upper Bounds on Distance Constraints

The fundamental idea of CBG relies on the controlled distance overestimation provided by the distance constraints. The goal is to overestimate in a controlled way so that distance constraints provide the smallest possible region that still encloses the target host. In practice, however, it is important to verify that distance constraints inferred between each landmark and the target hosts effectively provide an upper bound on the actual geographic distance between them. To evaluate whether CBG provides upper bounds on the actual distance, Figure 5 provides the cumulative distribution of the distances for each landmark-target pair.

Figure 5 compares the cumulative distribution of the estimated distances for each landmark-target pair using CBG, the estimation of the traceroute path length, and GeoBuD. The traceroute path length was computed by adding the geographic distances between the intermediate nodes along the traceroute which we were able to geolocate. For estimated distances larger than 1000 km, Figure 5 shows that CBG indeed provides

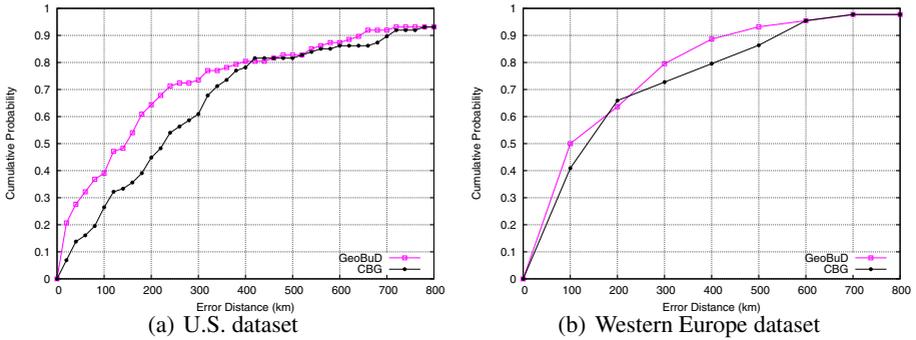


Fig. 4. Location estimation error for GeoBuD and CBG

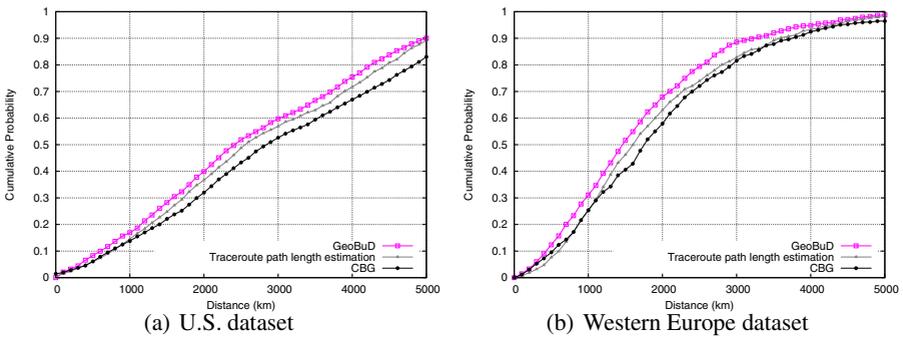


Fig. 5. Comparison of distance constraints and path length

an upper bound on the actual geographic distance, as even the estimated geographic length of the traceroute path is smaller. For distances smaller than 1000 km, CBG sometimes is close or below the estimated length of the traceroute path, as the estimated traceroute path length is also an upper bound on the actual geographic distance. Concerning GeoBuD, we observe in Figure 5 that it is a stricter upper bound on the distance than CBG or the estimated traceroute path length.

To understand why GeoBuD outperforms the original CBG, we need to recall how CBG transforms the delay into a distance constraint. For a given landmark, CBG overestimates the actual geographic distance by calibrating its transformation of the delay into a distance constraint by defining the value of b (see Equation (1)) based on the targets having the lowest delay measurement. This approach has the advantage of providing a conservative upper bound on the distance, as showed in Figure 5. However, its drawback compared to GeoBuD is of obtaining larger confidence regions.

In Figure 6 we plot the cumulative probability of the distance ratio over all landmark-target pairs, *i.e.* the ratio of the estimated distance (with CBG and GeoBuD) to the actual geographic distance. The purpose of Figure 6 is to study how each approach overestimates the actual distance. For instance, if we were to know that all estimated distances have a distance ratio larger than some value, then we could re-calibrate the

estimated distance by dividing all distance estimates by this factor. Unfortunately, we can see on Figure 6 that there is a small fraction of the estimated distances that do not overestimate the actual distance (ratio = 1). In fact, 3% of the landmark-target pairs do not overestimate the actual distance for CBG, and 13% for GeoBuD (see Figure 6(a)). In Figure 6(b) we have 5% and 7% for CBG and GeoBuD respectively. If we were to perform this re-calibration for the hosts having a distance ratio of 1, these distance constraints would underestimate the actual distance, potentially leading to an empty confidence region. From Figure 6, we can also see that for a distance ratio smaller than 4, GeoBuD provides a tighter overestimation of the distance than CBG. This is another illustration of the improvement of GeoBuD compared to the CBG approach.

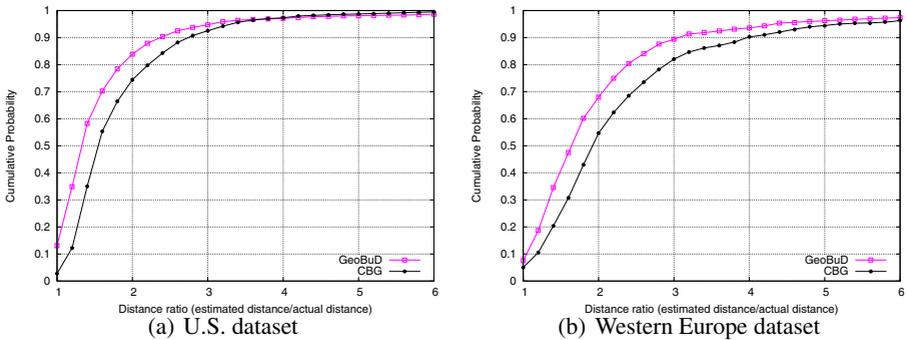


Fig. 6. Cumulative probability of distance ratios

6 Conclusion

In this paper we have shown that estimating the buffering delays at intermediate hops along the traceroute between a landmark and a target host enables to improve the accuracy of the geolocation of Internet hosts. Based on traceroute measurements, we estimated the buffering delays at intermediate hops. By combining these buffering delay estimation with a multilateration technique (CBG [6]), we were able to shrink the confidence region where the target host is located. Results show that, with GeoBuD we obtain more accurate location estimation as well. As further work, we see the implementation of our approach as an on-line tool like GeoLIM [27]. We also aim at converging towards a confidence region as small as possible to provide better location estimation. We might also refine our estimation of the buffering delay by considering the potential existence of a bottleneck link on the path.

References

1. Qwerks, Inc., *WhereIsIP*, <http://www.jufsoft.com/whereisip/>.
2. MaxMind LLC, *GeoIP*, <http://www.maxmind.com/geoip/>.
3. Quova Inc., *GeoPoint*, <http://www.quova.com/>.
4. Venkata N. Padmanabhan and Lakshminarayanan Subramanian, "An investigation of geographic mapping techniques for Internet hosts," in *Proc. of the ACM SIGCOMM'2001*, San Diego, CA, USA, Aug. 2001.

5. Artur Ziviani, Serge Fdida, José Ferreira de Rezende, and Otto Carlos Muniz Bandeira Duarte, "Improving the accuracy of measurement-based geographic location of Internet hosts," *Computer Networks, Elsevier Science*, vol. 47, no. 4, pp. 503–523, Mar. 2005.
6. B. Gueye, A. Ziviani, M. Crovella, and S. Fdida, "Constraint-based geolocation of internet hosts," *IEEE/ACM Transactions on Networking*, 2006, to appear.
7. H. Tangmunarunkit, R. Govindan, S. Shenker, and D. Estrin, "The impact of routing policy on internet paths," in *Proc. of the IEEE INFOCOM'2001*, Anchorage, AK, USA, Apr. 2001.
8. Lakshminarayanan Subramanian, Venkata N. Padmanabhan, and Randy Katz, "Geographic properties of Internet routing," in *Proc. of USENIX 2002*, Monterey, CA, USA, June 2002.
9. H. Zheng, E. K. Lua, M. Pias, and T. Griffin, "Internet Routing Policies and Round-Trip-Times," in *Proc. of the Passive and Active Measurement Workshop – PAM'2005*, Boston, MA, USA, Apr. 2005.
10. Christopher Davis, Paul Vixie, Tim Goodwin, and Ian Dickinson, "A means for expressing location information in the domain name system," *Internet RFC 1876*, Jan. 1996.
11. University of Illinois at Urbana-Champaign, *IP Address to Latitude/Longitude*, <http://cello.cs.uiuc.edu/cgi-bin/slamm/ip211/>.
12. David Moore, Ram Periakaruppan, Jim Donohoe, and Kimberly Claffy, "Where in the world is netgeo.caida.org?," in *Proc. of the INET'2000*, Yokohama, Japan, July 2000.
13. M. Freedman, M. Vutukuru, N. Feamster, and H. Balakrishnan, "Geographic locality of IP prefixes," in *Proc. of ACM/SIGCOMM Internet Measurement Conference – IMC 2005*, Berkeley, CA, USA, Oct. 2005.
14. *GeoURL*, <http://www.geourl.org/>.
15. *Net World Map*, <http://www.networldmap.com/>.
16. Visualware Inc., *VisualRoute*, <http://www.visualware.com/visualroute/>.
17. CAIDA, *GTrace*, <http://www.caida.org/tools/visualization/gtrace/>.
18. *Sarangworld Traceroute Project*, 2003, <http://www.sarangworld.com/TRACEROUTE/>.
19. Paramvir Bahl and Venkata N. Padmanabhan, "RADAR: An in-building RF-based user location and tracking system," in *Proc. of the IEEE INFOCOM'2000*, Tel-Aviv, Israel, Mar. 2000.
20. C. J. Bovy, H. T. Mertodimedjo, G. Hooghiemstra, H. Uijterwaal, and Piet van Mieghem, "Analysis of end-to-end delay measurements in Internet," in *Proc. of the Passive and Active Measurement Workshop – PAM'2002*, Fort Collins, CO, USA, Mar. 2002.
21. Artur Ziviani, Serge Fdida, José Ferreira de Rezende, and Otto Carlos Muniz Bandeira Duarte, "Toward a measurement-based geographic location service," in *Proc. of the Passive and Active Measurement Workshop – PAM'2004*, Antibes Juan-les-Pins, France, Apr. 2004, Lecture Notes in Computer Science (LNCS) 3015, pp. 43–52.
22. Stijn van Langen, Xiaoming Zhou, and Piet van Mieghem, "On the estimation of Internet distances using landmarks," in *Proc. of the International Conference on Next Generation Teletraffic and Wired/Wireless Advanced Networking – NEW2AN'04*, St. Petersburg, Russia, Feb. 2004.
23. V. Jacobson, *Traceroute Software*, 1999, <ftp://ftp.ee.lbl.gov/traceroute.tar.Z>.
24. *PlanetLab: An open platform for developing, deploying, and accessing planetary-scale services*, 2002, <http://www.planet-lab.org>.
25. *NLANR Active Measurement Project*, 1998, <http://watt.nlanr.net/>.
26. *RIPE Test Traffic Measurements*, 2000, <http://www.ripe.net/ttm/>.
27. *GeoLIM Project*, <http://planetlab-01.ipv6.lip6.fr:10000/cbg.php/>.
28. Neil Spring, Ratul Mahajan, and Thomas Anderson, "Quantifying the causes of path inflation," in *Proc. of the ACM SIGCOMM'2003*, Karlsruhe, Germany, Aug. 2003.