Network Distance Estimation With Dynamic Landmark Triangles *

Zhihua Wen
Case Western Reserve University
Electrical Engineering & Computer Science
10900 Euclid Avenue
Cleveland, OH 44106-7071
zhihua.wen@case.edu

Michael Rabinovich
Case Western Reserve University
Electrical Engineering & Computer Science
10900 Euclid Avenue
Cleveland, OH 44106-7071
misha@eecs.case.edu

ABSTRACT

This paper describes an efficient and accurate approach to estimate the network distance between arbitrary Internet hosts. We use three landmark hosts forming a triangle in two-dimensional space to estimate the distance between arbitrary hosts with simple trigonometrical calculations. To improve the accuracy of estimation, we dynamically choose the "best" triangle for a given pair of hosts using a heuristic algorithm. Our experiments show that this approach achieves both lower computational and network probing cost over the classic landmarks-based approach [3] while producing more accurate estimates.

Categories and Subject Descriptors

C.2.m [**Computer Systems Organization**]: COMPUTER-COMMUNICATION NETWORKSMiscellaneous

General Terms

Algorithms, Experimentation, Measurement

Keywords

network distance estimation

1. INTRODUCTION

The modern Internet is experiencing a rapid growth in large-scale distributed applications utilizing overlay and peer-to-peer networks. The performance and scaling properties of these applications crucially depend on forming overlay topologies and exercising communication paths according to the network distances between hosts. Unfortunately, the scale of these systems also makes direct on-demand distance measurements between host pairs impractical.

A number of approaches have been proposed to handle this problem by predicting the distance rather than measuring it directly. In particular, the Global Network Positioning (GNP) approach has proved influential due its ability to produce high quality distance estimation [3]. However, its high quality estimates come at high computational and probing costs. A number of approaches were subsequently proposed to reduce either of these overheads while claiming comparable accuracy. However, some of these proposals reduced only the computational costs but not probing overhead. Other proposals required global coordination between the entire set of hosts being measured.

This paper presents a novel approach for distance estimation using dynamically selected landmark triangles, which achieves *both* higher prediction accuracy *and* drastically lower computational and probing overheads than GNP.

2. ALGORITHMS AND METHODOLOGY

Similar to GNP and subsequent approaches, we estimate the distance between two hosts from the measured distance between each host and specially deployed *landmark hosts*. Unlike GNP, we actually use many more landmarks but carefully select only a small number of them for any given prediction. Our basic algorithm employees only three landmarks, forming a triangle in a two-dimensional Euclidean space, to estimate the distance between two arbitrary end hosts using simple and efficient trigonometrical calculations. Our goal is to dynamically select a specific triangle for a particular pair of hosts among a large set of landmarks that is likely to produce a high-accuracy estimate. We first analyze the relationship between landmark triangles and their accuracy of distance estimates using a development data set with 15 selected landmarks and 455 potential triangles. We used each possible triangle to estimate the distance between all pairs of end hosts and compared the quality of each triangle. The results lead to the following observations.

A big triangle is more likely to produce decent accuracy for end-hosts that are far apart as shown in Figure 1(a) and 1(c). A small triangle generates poor accuracy when it is far removed from both end-hosts as shown in Figure 1(a). But it generates high accuracy when it is close to one of the end-hosts, regardless of the distance between the end hosts, as shown in figures 1(b) and 1(c). If we wanted to use one static triangle for all host pairs, a big triangle is likely to be a better choice for more host pairs because any particular small triangle works well only for a relatively small number of hosts in its vicinity. However, if we can select different triangles for different host pairs, an appropriate small triangle can produce a better estimate.

Following the above observations, we have designed a heuristic algorithm for dynamic triangle landmark selection. The basic idea is that for two given hosts, we try to find a small triangle close to either of them and use it for estimation. If

*This material is based upon work supported by the National Science Foundation under Grant No. CNS-0721890.

SIGMETRICS'08, June 2–6, 2008, Annapolis, Maryland, USA.
ACM 978-1-60558-005-0/08/06.

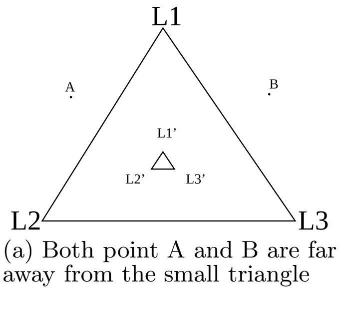

Figure 1: The effect of the landmark triangle position and size on the accuracy of inter-host distance estimation.

we cannot find such a triangle, we resort to the general big triangle.

In order to save the network probing traffic, our estimation involves two rounds of probes. In the first round, we use a general, big triangle to obtain a rough position of each end host, with accuracy comparable to GNP using only three landmarks in a two-dimensional space. We then use these rough position estimates as the basis for selecting another landmark triangle using our algorithm, which is likely to obtain a high accuracy distance estimation for the current end hosts. This second triangle produces the final distance estimation in the second round of probing.

3. EXPERIMENTS AND RESULTS

Having developed our approach using the development data set, we evaluated its accuracy using a different, much bigger assessment data set. Both data sets were obtained using DipZoom measurement infrastructure [1], which provides programmatic access to a number of measuring points (MPs) around the world from a locally run Java application. We used over 200 DipZoom MPs to collect the assessment data set. We selected 100 MPs as the candidate landmarks and predicted the distance among the remaining MPs and also between the remaining MPs and around 1900 Gnutella peers selected from a list compiled at the University of Oregon [2]. We then compared predicted distances with the direct ping measurements from the MPs to the target hosts.

Following GNP as well as a number of subsequent landmark-based studies, we used relative error to quantify the accuracy of our distance estimates:

$$\frac{|predicteddistance - measureddistance|}{min(measureddistance, predicteddistance)} \qquad (1)$$

Figure 2 shows the cumulative distribution functions of the relative error of distance predictions. The figure shows the overall error for all host pairs (the curve labeled "dynamic triangles") as well as separately for different groups of host pairs. The group denoted "small triangles (partial)" includes the host pairs that used a small nearby triangle for distance estimation as the result of dynamic triangle selection; the group "general triangle (partial)" represents host pairs that utilized the general triangle.

Figure 2 indicates that the quality of the dynamic triangles estimates is considerably higher than GNP using the default setting of 15 landmarks and 8 dimensions. The median and average relative errors of our estimates are 4.1% and 12.9% vs., respectively, 9.64% and 19.15% for GNP.

Figure 2 also shows that the accuracy suffers significantly when no suitable small triangle is found, and the estimates must be produced from the general triangle. Fortunately, with 100 candidate landmarks, only 2760 host-pairs out of nearly a quarter million total had to resort to the general triangle, so the overall accuracy was not affected.

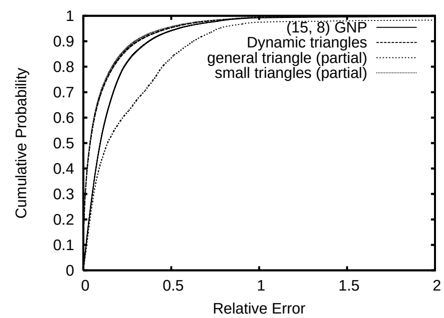

Figure 2: Accuracy of distance estimates using dynamic landmark triangles.

Besides accuracy, a key advantage of dynamic triangles is their low probing and especially computational overhead. Our approach requires only 3-6 probes per host in a host-pair. In terms of computational overhead, because we use simple trigonometrical calculations instead of searching for a non-linear optimization solution as GNP, dynamic triangles have drastically lower computational overhead. Overall, the dynamic triangles approach finishes the distance estimation for all quarter million host pairs in 15 seconds, compared to over 9 minutes for GNP.

4. REFERENCES

[1] DipZoom. http://dipzoom.case.edu.

[2] Gnutella Topology Snapshots. http://mirage.cs.uoregon.edu/P2P/info.cgi.

[3] T. S. E. Ng and H. Zhang. Predicting internet network distance with coordinates-based approaches. In *INFOCOM*, 2002.