An Empirical Evaluation of Landmark Placement on Internet Coordinate Schemes

Sridhar Srinivasan Ellen Zegura
Networking and Telecommunications Group
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332, USA
Email: {sridhar,ewz}@cc.gatech.edu

***Abstract*— Many applications can benefit from knowledge of latency between hosts. Recently, there have been proposals such as Global Network Positioning (GNP) [1] and Internet Coordinate System (ICS) [2] that attempt to provide this for hosts on the Internet. The schemes work by using a small set of nodes called "landmarks" to compute the coordinates of other hosts. In this paper, we investigate the effect of placement of landmarks on the performance of these proposals. We show that with random placement, the estimation error of these schemes varies widely. We propose heuristics for placing landmarks that show good performance. We also investigate the behaviour of these schemes with measurement errors and show that GNP is less affected by measurement errors than ICS.**

I. INTRODUCTION

Many current applications such as peer-to-peer file-sharing and content-delivery can be better served if some knowledge of latency between hosts is available. This is especially true in cases where the content is large or delay-dependent. For peer-to-peer applications, peer locations and latency information can be used to optimize the construction of the overlay (e.g. to avoid making neighbor links between hosts on opposite sides of the globe), which in turn can help in locating content "closer" to the peer. This "location-aware" construction can improve search latency, by avoiding long delay links unless necessary, and in selecting the "best" peer to get the content from when the search query is completed. In the case of content-delivery, knowledge of the client latencies to servers can help in directing clients to the closest server. Currently, some Content Delivery Networks provide this as a service while in most cases the client is simply offered a list of servers to choose from.

In the last few years, there have been several proposals for obtaining location and latency information about hosts. Broadly, they can be classified according to whether they are *measurement-based* or *coordinate-based*. In the category of coordinate-based schemes are Global Network Positioning (GNP) [1], and Internet Coordinate System (ICS) [2] which fix the location or coordinates of any host on the Internet in a Euclidean space using a small set of round-trip time (RTT) measurements. Using these coordinates, nodes can estimate their latency to other nodes without any other measurements.

In this paper, we restrict our discussion to the coordinate-based latency estimation schemes because of their relative ease of deployment as compared to the measurement-based schemes. The coordinate-based schemes use a small set of hosts as special nodes (called "landmarks" in this paper). Each landmark measures its latency to all other landmarks. These measurements are used to assign coordinates to the landmarks. A host measures its latency to these landmarks and performs some computation to obtain its coordinates relative to these nodes.

In both GNP and ICS, the number of landmarks and their relative locations determine the effectiveness of the latency estimates that they provide. Given that the set of potential landmarks in the public Internet can be huge, there is not much knowledge of how a particular placement of landmarks affects these schemes. In this paper, we investigate the *effect of placement of landmarks on the performance of the latency estimation proposals*. The authors of [1] showed that using random placement for GNP, the error in estimation can vary widely. They evaluated three heuristics for placing landmarks based on the inter-landmark separation and medians of clusters formed by landmarks. In [2] and [3], placement of landmarks in ICS was left as an open problem and the evaluation was performed using random sets of landmarks. In a follow up paper [4], the authors evaluate heuristics for ICS that are similar to the GNP heuristics. The authors conclude that a k-means clustering heuristic is best for a small number of landmarks while for larger sets, random selection is good enough. In our experiments with public `traceroute` servers scattered across 23 countries, we observed a significant variation in the quality of estimation even using large sets of landmarks. Given the potentially large number of available locations where landmarks can be placed, this conclusion seems surprising. There is also no work on analyzing the effect of errors in the latency measurements on the quality of estimation.

We propose a graph-theoretic method of classifying potential sets of landmarks according to the center and diameter of the sets with respect to a representation of the Internet topology. We then propose heuristics for landmark selection based on easily obtainable topological information.Specifically, we examine the placement of the landmarks based on topological information such as AS size and inter-landmark latency. The four heuristics we evaluate to select landmarks are the size of the ASes in which they are located, the minimum and maximum inter-landmark distances and a clustering method based on Singular Value Decomposition (SVD). We show that using the SVD and AS size heuristics, GNP gives good estimates and for ICS, the inter-landmark heuristics give good results. We also study the effects of errors in measurements on both schemes and find that GNP seems to be less affected by measurement errors than ICS.

The remainder of the paper is organized as follows. In the next section, we describe the two techniques that we consider. In Section III, we discuss the heuristics we use for our evaluation. We describe our data collection process and experimental methodology in Section IV. In the Section V, we describe our results and we conclude in Section VI.

II. DESCRIPTION OF COORDINATE SCHEMES

The first scheme to propose embedding hosts in a Euclidean space by assigning coordinates was Global Network Positioning (GNP) by Ng and Zhang [1]. Initially a small set of hosts, called landmarks, measure the latency to each other. Using these measurements they compute the coordinates of all the landmarks in a chosen coordinate space by minimizing the error between the measured distances and the computed distances. For all i, j, if d_{ij} represents the measured distance between the landmarks i and j and $\hat{d}_{ij}$ the Euclidean distance between the two nodes coordinates, the minimization is performed over

$$\Sigma_{i,j} Error(d_{ij}, \hat{d}_{ij}),$$

where $Error$ is defined as the squared error between the two hosts i and j. Any host that participates in this scheme measures its latency to all landmarks. Using these measurements, together with the coordinates of the landmarks, the host computes its own coordinates by minimizing the error between its measured and computed distances to the landmarks, that is, if A is the host, A tries to minimize the function $\Sigma_i Error(d_{Ai}, \hat{d}_{Ai})$.

In two closely related papers, Lim et al. [2] and Tang and Crovella [3], propose using Principal Component Analysis to assign coordinates to nodes. We describe the method used by Lim et al. called Internet Coordinate System (ICS) but the techniques are the same in both papers. First a matrix of inter-landmark distances is created using a small set of nodes called beacons or landmarks. Each of the n landmarks measures the distance to the other landmarks and creates an n-dimensional vector. These vectors are aggregated into the distance matrix D. This matrix D is then decomposed using Singular Value Decomposition [5] into 3 matrices,

$$D = UWV^T$$

The matrices U and V are column and row orthogonal and W is a diagonal matrix with the diagonal values w_i in decreasing order. These w_i are called the singular values of D. U consists of the eigenvectors of the matrix DD^T. Each element of D can be written in the form

$$D_{ij} = \Sigma_{k=1}^{n} w_k U_{ik} V_{jk}$$

If the initial w_is are much larger than the rest, it is possible to approximate the matrix D by

$$D'_{ij} = \Sigma_{k=1}^{r} w_k U_{ik} V_{jk}$$

where $r < n$. Thus, the n-dimensional coordinate of a landmark can be represented without much error by an r-dimensional coordinate. The first r columns of the matrix U become the orthogonal basis vectors of the new subspace. This subspace is shown empirically to be of a much smaller dimensionality than the original matrix D. A host A joining this network measures its distance to the landmarks and creates its n-dimensional vector d_A. Using the principal components U_r, it computes its new coordinate $c = U_r d_A$ in the subspace. Using these coordinates, the latency between any two nodes is calculated as the Euclidean distance between their coordinates.

In this paper we focus our attention on GNP and ICS as we believe they represent the most decentralized and practical approaches to distance estimation.

III. FACTORS AFFECTING DISTANCE ESTIMATION

In this section, we discuss the factors that affect the estimations provided by GNP and ICS. We then design heuristics for landmark selection based on these factors.

The initial coordinates assigned by these schemes are based on the measurement of the inter-landmark distances. These initial coordinates, in effect, decide the frame of reference in which all other client coordinates are placed. For GNP, all client coordinates are relative to the landmark coordinates and in ICS the inter-landmark matrix decides the subspace in which all the vectors are projected.

The location of the landmarks in the network topology also affects the estimates that a scheme produces. For an extreme example, consider a set of landmarks that are located in the same subnet. The distance from any client to all landmarks would be similar, leading to poor distance estimates.

For the practical deployment of a distance estimation scheme, the number of landmarks cannot be too large, as the overhead of measuring the latency to a large set of hosts can outweigh the usefulness of obtaining distance estimates. If the number of landmarks is too few, the quality of distance estimation can be affected. The failure of only a few of the landmarks can render the scheme ineffective. We investigate the performance of the schemes for various sizes of landmark sets in Section. V.

A. Characterizing a set of landmarks

The inter-landmark latencies and the location of the landmark set in the network topology are important factors that affect the quality of the estimates returned by both GNP and ICS. To compare different sets of landmarks based on these factors, we first formalize the notions of inter-landmark distances and location. We use the following concepts from graph theory for this purpose. For a graph $G = (V, E)$, let $\delta(u, v)$ be the length of a shortest path from a vertex u to a vertex v. The *eccentricity* of a vertex v is defined as the longest of the shortest paths from v to any vertex, i.e.,

$e(v) = max\{\delta(v, w) : w \in V\}$. The *radius* of G is defined as $radius(G) = min\{e(v) : v \in V\}$. The *center* of G is defined as the set of vertices that have eccentricity equal to the radius of G. Also, the *diameter* of G is defined as the maximum eccentricity of any vertex in G.

The above definitions formalize the center and diameter of the entire graph but for characterizing landmark sets, we need similar definitions for subsets of vertices from the graph. To define the center of an arbitrary subset of vertices $S \subset V$, we begin by defining the eccentricity of a vertex with respect to an arbitrary set of vertices. Let $es(u, S) = max\{\delta(u, v) : v \in S\}$. We can define the subgraph induced by this set of vertics S as $G' = (V', E')$, where $V' = \{u : \forall v \in S, \delta(v, u) \leq es(v, S)\}$ and $E' = \{(u, v) : u, v \in V', (u, v) \in E\}$. Intuitively, the center of this set lies "inside" the collection of vertices forming this set.

Based on these definitions and by representing the Internet in the form of an AS-level graph, we can compute the center of the Internet graph and the centers of various set of landmarks.

To fix the location of a set of landmarks, we require a notion of distance from the center of the Internet to the center of the set of landmarks. For this work, we define the distance between the Internet center and a set of landmarks as the average length of the shortest paths from the nodes in the center of the Internet to the nodes comprising the center of the set of landmarks.

$$Distance(G, G') = avg\{\delta(u, v) | u \in center(G), v \in center(G')\}$$

This characterization allows us to explore the effect of selecting sets of landmarks that are in different locations in the network topology with different inter-landmark distances.

B. Heuristics for landmark selection

For a heuristic to select landmarks based on location in the topology, we use the degree of the ASes in which the landmarks are located in. The size of an AS (as indicated by its degree in the Internet AS graph) has strong correlation with its location in the Internet topology [6]. Selecting the landmarks to be those located in the largest ASes of the set of potential landmarks, we attempt to pick those hosts that are the closest to the core of the Internet topology. We call this heuristic **AS-Large**.

As a measure of the size of the set of landmarks, instead of the diameter, we consider the maximum latency between the hosts. The **Max-Dist** heuristic tries to place the landmarks as far apart from each other based on the measured latency. An iterative procedure picks the server that maximizes the distance to the remaining candidate servers in the set. The objective of this heuristic is to determine the effect of separation between the landmarks. Intuitively, we expect the maximum separation should perform well, but it was reported in [1] that a similar criterion works poorly. For comparison, we also select a set of landmarks based on the minimum distance between hosts called **Min-Dist** which, at each step selects the server that minimizes the distance to the existing landmarks as the next landmark to be chosen.

Our last heuristic, **SVD**, is an attempt to group the servers into clusters and picking servers from each cluster. The clustering is based on the SVD of the inter-server distances. From the SVD of the distance matrix D, we consider U as the eigenvectors of the affinity matrix and use the technique of segmentation using eigenvectors [7] to compute the clusters of servers. From each cluster, we select a server to be a landmark.

In our experiments, we considered other heuristics based on the smallest AS sizes and the minimum and maximum variants of the Dist heuristic, but due to lack of space we only report on the four mentioned above. The performance of the other heuristics was generally worse than the four shown here.

IV. EXPERIMENTAL METHODOLOGY

To evaluate landmark placement for the distance estimation schemes, our main requirement is a large set of servers to act as potential landmarks. Previous experiments for GNP and ICS were restricted in the number of landmarks that were used. GNP used a pool of 19 servers as potential landmarks whereas the evaluation of the ICS method relied on existing measurement projects for the data.

For our experiments, we used a set of public traceroute servers[1] to act as our landmarks. These servers are scattered across the world and seem to be mostly located in commercial ISPs. We use the round-trip times reported by the traceroute servers as the latency measure from the client to the server. We initially picked a set of 157 traceroute servers to perform our probes from. Out of these servers, we identified a set of 64 servers, located in 23 different countries as reported by the NetGeo tool, which successfully performed traceroutes to most clients. We call this set of servers LS.

To obtain the set of clients, we used BGP tables available from the RouteViews project [8] to divide the Intenet address space into the advertised routable prefixes. For each prefix, we then probed for an authoritative DNS nameserver that was located in the same AS as the prefix. Of these nameservers, a significant proportion do not respond to traceroutes which we presume is due to the presence of firewalls. After processing, we were left with 194 clients which we probed.

We also used the nodes comprising the Planet-Lab network for a second client set PL. This consists of 192 nodes distributed among universities in more than 20 countries. The inter-host latencies of these nodes is available and this gives us an accurate estimate of the latencies between them.

The traceroute servers we used are a public resource and so we wanted to ensure that no traceroute server was queried too quickly.Queries to a server were spaced approximately five minutes apart. The queries were issued sequentially to all traceroute servers over a period of approximately 2 hours for each client. Thus, the measurements from the set of traceroute servers to the clients are not concurrent.We believe that this reflects the way in which a measurement service would function with clients joining at different times and hence, we do not believe this affects our results.

---

[1] A list of such servers is available at `http://www.traceroute.org`.

For all the data-sets collected, we select a set of servers to be our Landmarks, then run the GNP and ICS algorithms to compute the coordinates of the clients. We also compute the coordinates of the remaining servers. To compute the estimation accuracy of each method, we compare the measured latency between each of the remaining servers and the set of clients with the predicted latency. The metric used is the relative error defined [1] as

$$\left|\frac{predicted\ distance\ -\ measured\ distance}{min(predicted\ distance, measured\ distance)}\right|.$$

Ideally, the relative error should be zero indicating that the predicted latency exactly matches the measured latency. The relative error metric emphasizes the difference between the predicted and measured distances, providing a better estimate of the performance of the different schemes.

V. RESULTS

We begin by considering the effect of selecting landmarks randomly on the estimation error. We then investigate various heuristics to select landmarks from a set of servers. The experiments in the next two sections were performed from the set LS to the 194 nameservers which were used as clients.

A. *Random selection of Landmarks*

To find out if there is a benefit to selecting landmarks using any method other than choosing them at random, we perform a set of experiments in which we pick the landmarks randomly. For these experiments, each run was set up by first selecting a set of servers to represent landmarks from the LS set and then running the distance estimation scheme to the client set. The results of 30 such experiments for each landmark set size are plotted as boxplots for GNP and ICS in Figures 1 and 2. The boxplot provides a simple graphical summary of the data with lines at the lower quartile, median and upper quartile values. The whiskers extend from the ends of the box to the most extreme data value within 1.5 times the range of the middle quartiles. Data outliers are explicitly shown beyond the whiskers. As a measure of the estimation quality, we plot the 90th percentile value of the relative error as the number of landmarks is varied. The smaller the value of the relative error, the better the estimation of the latencies between hosts.

We see from the figures that in both schemes estimation accuracy varies depending on the particular landmark set that is used. The variation persists when evaluated over different-sized sets of landmarks. This indicates that the placement of landmarks is important to both schemes irrespective of the number of landmarks used and also reaffirms the observations in [1] about the diminishing improvements.

For GNP we note that increasing the number of landmarks from 12 to 20 improves the estimation as shown by the decreasing value of the median. For set sizes from 20 to 48, the data show no clear trend but there are significant outliers for the 48 landmark sets. This suggest that a range of landmarks between 20 – 32 provides the best results.

For ICS, the benefit of increasing the number of landmarks appears to be in reducing the range of the estimation error as shown by the size of the whiskers. The number of dimensions of the landmark data is a small number (in our experiments, it is usually between 4-6), implying that increasing the number of landmarks doesn't improve estimation, but helps in reducing the noise in the data, thereby providing more consistent results. This can be seen in the reduction of the whiskers of the boxplots as the number of landmarks increases.

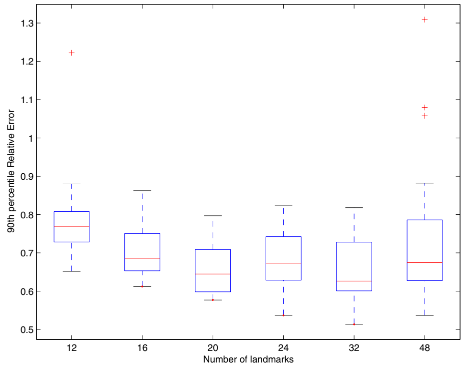

Fig. 1. Estimation error of GNP using random sets of landmarks

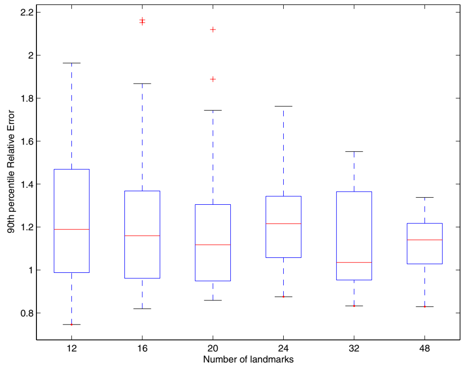

Fig. 2. Estimation error of ICS using random sets of landmarks

B. *Using heuristics to select landmarks*

We next examine the heuristics for selecting the landmarks. From our set of available servers, we select the servers located in the the largest ASes to form the set AS-Large. The degrees of the ASes in AS-Large range from 10 to 520. Table I shows the center and diameters of the sets generated using these heuristics. The values show that AS-Large is close to the Internet center with the landmarks relatively close to each other with a diameter of 4, while Max-Dist, although its center is closer, actually has a much larger diameter indicating that the landmarks are spread out farther apart. The results of these experiments for ICS and GNP are plotted in Figures 3 and 4 respectively. From the graphs, we note that for GNP, the best performing heuristics are the SVD and AS-Large while for ICS, the best performing are Min-Dist and Max-Dist.

To test the heuristics and our conclusions, we apply the heuristics to a different data-set. In this experiment, we use a subset of the LS set of servers to traceroute to the PL client

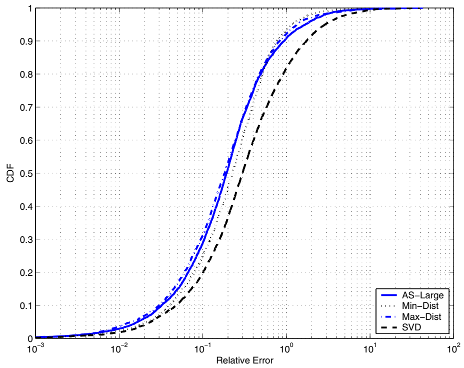

Fig. 3. Estimation error of ICS using the heuristics

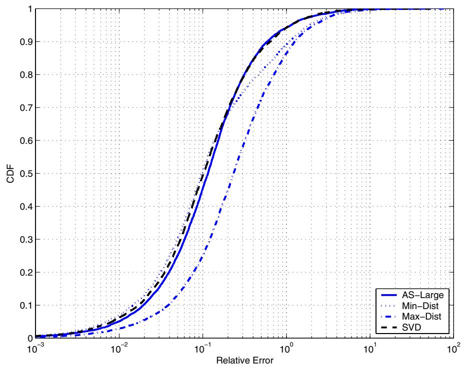

Fig. 4. Estimation error of GNP using the heuristics

set (the set of clients consisting of Planet-Lab nodes). The experiment was conducted using 61 servers from the LS set to 112 Planet-Lab nodes. The results are plotted in Figures 5 and 6. We observe that the best performing heuristics for GNP are the SVD and AS-Large heuristics, while for ICS the Max-Dist and Min-Dist heuristics are the best.

To eliminate the effects of using only one set of potential landmarks, we performed an experiment in which we "invert" the previous measurements. Using the inter-node distances for the set of Planet-lab nodes, we performed an experiment in which the Planet-lab nodes are the set of potential landmarks and the `traceroute` servers are the clients. The results of this are plotted in Figures 7 and 8 for ICS and GNP respectively.

For GNP, the SVD heuristic performs the best again, but the AS-Large performs worse than the Max-Dist heuristic. We explain this behaviour based on the set of hosts in the PL data-set. There are several sets of hosts that are on the same network, i.e., their IP addresses differ only in the fourth octet. While selecting hosts for the AS-Large heuristic, if two or more hosts are on the same network in a large AS, all these hosts would be selected by the heuristic. Clients measuring latency to such hosts would get very similar values and this reduces the effectiveness of the estimation. For ICS, the results are similar to the previous graphs except that the SVD and AS-Large heuristic perform better. The best heuristic for ICS is still the Min-Dist heuristic.

TABLE I

MEASURE OF LANDMARK SETS

Heuristic	Dist. to Internet center	Diameter
AS-Large	1.125	4
Min Dist	1.2757	5
Max Dist	1.0714	6
SVD clustering	1.3281	6

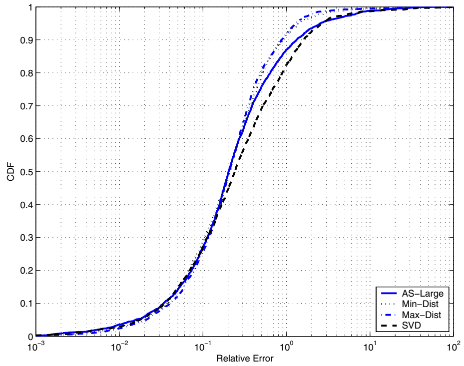

Fig. 5. Estimation error to the set of Planet-Lab nodes using ICS with all heuristics

From the three sets of experiments, we conclude that the SVD heuristic is the best for GNP while for ICS, the Min-Dist performs the best.

C. Performance with errors in measurement

In a realistic environment the measurements made by the hosts are subject to errors. Errors can occur when the landmarks are measuring the inter-landmark distance or when clients are measuring to landmarks. In the latter case, it is clear that the client will calculate erroneous coordinates for its position but this doesn't affect the performance of the system as a whole. When the landmarks make errors in measurement, the potential for affecting the entire system is higher as discussed in Section III. To investigate this, we consider the

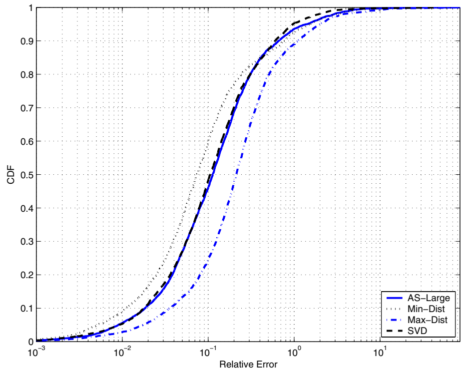

Fig. 6. Estimation error to the set of Planet-Lab nodes using GNP with all heuristics

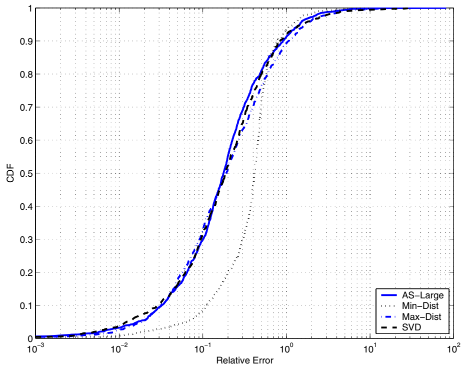

Fig. 7. Estimation error to the set of `traceroute` servers from Planet-Lab nodes using ICS with all heuristics

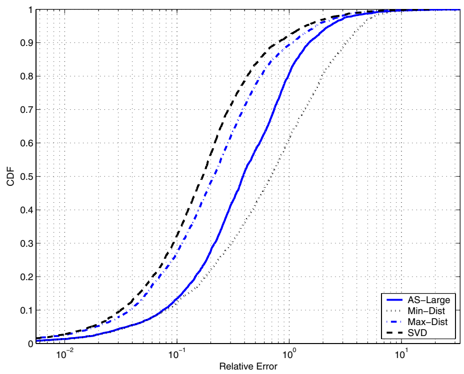

Fig. 8. Estimation error to the set of `traceroute` servers from Planet-Lab nodes using GNP with all heuristics

performance of GNP and ICS with one landmark making erroneous measurements while measuring latency to the other landmarks. For this we added 250ms, 500ms and 1000ms to a single randomly selected landmark's measurements to simulate the error that might result from a routing error or congestion at a network through which the landmark's traffic passes. The results using landmarks from the LS set are shown in Figures 9 and 10. We also plot the results obtained from the unaltered data for comparison.

We observe that, in general, GNP seems to offer better performance with the SVD and AS-Large heuristics being relatively unaffected by the errors. For ICS the heuristic Max-Dist is relatively unaffected, although the magnitude of the change is larger compared to GNP as can be seen by the difference in the scale of the y-axis of the figures.

VI. CONCLUSIONS

Given the recent interest in distance estimation schemes and their utility in many applications, there have been several proposals for coordinate-based distance estimation schemes. In this paper, we evaluate two of these schemes, Global Network Positioning (GNP) and Internet Coordinate System (ICS), with a focus on the role of landmark selection on performance.

We show that picking a set of landmarks at random is generally not a good approach as the results can vary widely depending on the random selection. We evaluated several heuristics for selecting landmarks over different data-sets and find that for GNP, selecting servers from clusters based on the SVD heuristic gives good results while for ICS, picking the landmarks that are farthest apart in the available set using the Max-Dist heuristic works well.

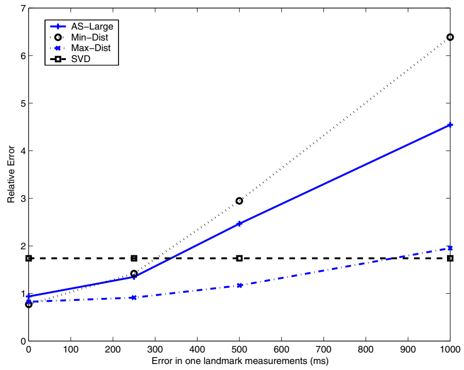

Fig. 9. Estimation error of ICS from TR servers with error in one landmark

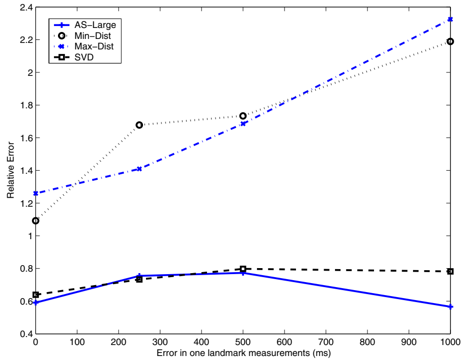

Fig. 10. Estimation error of GNP from TR servers with error in one landmark

We also examined the performance of the schemes with respect to errors during the measurement process. In our experiments, the GNP heuristics were less affected by errors in measurements between landmarks than the ICS heuristics.

REFERENCES

[1] T. S. E. Ng and H. Zhang, "Predicting Internet network distance with coordinates-based approaches," in *INFOCOM*, 2002.
[2] H. Lim, J. Hou, and C.-H. Choi, "Constructing Internet Coordinate System based on delay measurement," in *Internet Measurement Conference*, 2003.
[3] L. Tang and M. Crovella, "Virtual landmarks for the Internet," in *Internet Measurement Conference*, 2003.
[4] ——, "Geometric exploration of the landmark selection problem," in *Passive and Active Measurement Workshop*, 2004.
[5] G. Strang, *Linear Algebra and its Applications*, Third ed. Thomson, 1998.
[6] L. Subramanian, S. Agarwal, J. Rexford, and R. Katz, "Characterizing the Internet Hierarchy from multiple vantage points," in *Infocom*, 2002.
[7] P. Perona and W. Freeman, "A factorization approach to grouping," *Lecture Notes in Computer Science*, vol. 1406, pp. 655–??, 1998.
[8] http://antc.uoregon.edu/route views/, "Route views project."