

Big-Bang Simulation for Embedding Network Distances in Euclidean Space

Yuval Shavitt, *Senior Member, IEEE*, and Tomer Tankel, *Student Member, IEEE*

Abstract—Embedding of a graph metric in Euclidean space efficiently and accurately is an important problem in general with applications in topology aggregation, closest mirror selection, and application level routing. We propose a new graph embedding scheme called Big-Bang Simulation (BBS), which simulates an explosion of particles under a force field derived from embedding error. BBS is shown to be significantly more accurate compared to all other embedding methods, including GNP. We report an extensive simulation study of BBS compared with several known embedding schemes and show its advantage for distance estimation (as in the IDMaps project), mirror selection, and topology aggregation.

I. INTRODUCTION

KNOWLEDGE of the distances between all pairs of a group of nodes can improve the performance of many practical networking problems, such as routing through a subnetwork and selecting the closest mirror server. However, measuring and, to a greater extent, disseminating this information becomes impractical even for a few tens of nodes, since the number of node pairs is quadratic in the number of nodes. Thus, researchers sought ways to reduce all pair-distance representation while preserving the distance in the reduced representation as close as possible to the original ones. Next we briefly describe two example networking problems, where all pair-distance information is required.

Routing through a subnetwork. When routing through an ATM subnetwork, the distances between all pairs of border nodes, i.e., nodes that are connecting the subnetwork to other subnetwork, are used to compute the shortest (or cheapest) path through the cloud [1]. For this end, each network advertises its distance matrix in a compressed manner, and it is recommended that the matrix representation be smaller than $3b$, where b is the number of border nodes [1]. The best compression technique that was suggested in the past [2] will be presented later.

Selecting the closest mirror. Recently, there was a large interest in using distance maps of the Internet to aid in tasks such as closest mirror selection and application layer multicast [3]–[6]. In the IDMaps project, it was identified that the number of possible nodes which represent the distance map granularity is in the thousands, which makes accurate distance dissemination impractical. Due to the practicality of the measurement

process and to reduce the representation, IDMaps suggests using a smaller number t of measurement points, *Tracers*, that measure distances among themselves and then using them as a reference distance map to the other network regions.

A relatively new approach to represent a network distance matrix is to map network nodes into points in a real Euclidean space. Such a mapping is designed to preserve the distance between any pair of network nodes close to the Euclidean distance between their geometric images. Such a mapping is called an embedding, and the embedding dimension d is the dimension of the Euclidean space. Ideally, graph edge lengths are exactly embedded in the geometric edges. However, it can be easily shown that an exact embedding is not always possible, e.g., in the case of a tree, and, in fact, in most cases embedding introduces some distortion. In a “good” embedding, the average and maximum distance distortions over all pairs of nodes are relatively low. The distance distortion is defined for each pair as the maximum of the ratio between the original and Euclidean distance and its inverse.

Outside the networking community, embedding has been used for quite a long time in many diverse research areas. Multidimensional Scaling (MDS) is widely used in areas of statistics and vision, in which its simplicity and efficiency is an advantage over more complicated methods. Classical metric MDS [7] develops the metric as a symmetric bilinear form and calculates the leading d eigen values of the corresponding matrix.¹ Recently, computer graphics researchers [8] suggested using MDS for mapping flat textures over curved surfaces with minimum distortion. Embedding is used extensively in bioinformatics and specifically for classification of protein sequences into similarity families [9].

Embedding of graph edges in the Euclidean plane is also achieved by the *spring embedder* algorithm for graph drawing applications. This algorithm is a heuristic based on a physical model which was first described in [10]. The heuristics which followed this direction, such as [11]–[13], are reviewed in a recent survey [14].

Theoretical bounds on the **maximal** pair distortion and the dimension of the target space were derived by Linial *et al.* [15] for any discrete metric space.² Perhaps the first use of graph embedding techniques in networking is due to Ng and Zhang [5], who suggested estimating Internet host distances by embedding distances among Tracer nodes and other network regions. Their embedding technique sought the minimum of the total square

Manuscript received December 28, 2002; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor R. Govindan. This work was supported in part by a grant from the United States–Israel Binational Science Foundation (BSF), Jerusalem, Israel, the Israel Science Foundation (ISF) Center of Excellence Program under Grant 8008/03, and the European Union 6th FP EVERGROW.

The authors are with the Department of Electrical Engineering, Tel-Aviv University, Tel Aviv 69978, Israel (e-mail: shavitt@eng.tau.ac.il; tankel@eng.tau.ac.il).

Digital Object Identifier 10.1109/TNET.2004.838597

¹Finding the leading eigenvalues of the matrix of a symmetric bilinear form is equivalent to Singular Value Decomposition (SVD).

²Though [15] and more recent publications discuss mapping to l_p , we will concentrate only on embedding in l_2 , which is the most popular in applications.

embedding errors over all nodes pairs, which is proportional to the **average pair** distortion.

In this paper, we present a new scheme for embedding, based on a novel idea, utilizing notions from Newtonian mechanics. We shall compare our scheme, *Big-Bang Simulation (BBS)*, to other embedding methods and show that it produces the best embedding over various parameter choices with reasonable complexity. This is not to say that, for special cases, e.g., when the number of embedded nodes is very small, it will outperform all other schemes, but it will be better than any other scheme when all cases are considered.

A. Big-Bang Simulation

BBS models the network nodes as a set of particles; each is the “geometric image” of a node, that is, the position of this particle in Euclidean space. The particles are traveling in that space under the effect of a potential force field. The force field reduces the potential energy of the particles that is related to the total embedding error of all particle pairs. Each pair of particles is pulled or repulsed by the field force induced between them depending on their pair embedding error, that is, the embedding error of the distance between them. As a particle accelerates under the effect of the force field, it is also attenuated by simulated friction force.

The advantage of the BBS scheme over conventional gradient minimization schemes, such as steepest decent and downhill simplex (DHS), is that the kinetic energy accumulated by the moving particles enables them to escape the local minima. Moreover, DHS, which was used by Ng and Zhang [5], is very sensitive to the initial vertices coordinates. The key idea which inspired the name “Big-Bang Simulation” is that all particles are **initially placed at the same point**, the origin. Starting with such initial condition, BBS obtains “good” embedding, that is, **low average** distortion and **quite low maximal** distortion, in several hundred simulation iterations for input graph sizes in the range $30 < n < 750$. The above good performance is equally achieved for a wide range of system friction coefficients.

Next we discuss in more detail one of the applications mentioned above for which we have applied our embedding and compare it to previous results.

B. Internet Distance Maps Application

While it was shown that IDMaps performs quite well, e.g., it correctly points a client to the closest mirror server in about 85% of the cases [4], it is far from being ideal.

Fig. 1 illustrates the main reason for inaccuracies in IDMaps: in the example, client H_1 estimates the network distance Δ_l to mirrors H_l , $l = 2, 3$, as a sum of several measured segments, $s_1 + s_t + s_l$. The inaccuracy of IDMaps estimation is larger when a nearest Tracer T_1 is shared by the two nodes. For instance, in the case depicted in Fig. 1(a), the shortest network distance is $\Delta_3 \ll \Delta_2$ although the smallest sum of segments is $s_1 + s_2 < s_1 + s_3$. On the other hand, in the case depicted in Fig. 1(b), the client is located near the Tracer T_2 , and both mirror hosts are located near another Tracer T_1 . The distance s_t between T_1 and T_2 is larger than two times the distances from the client and mirror hosts to their nearest Tracer. Thus, the distances Δ'_l ; $l =$




Fig. 1. Problem statement.

2, 3, are both dominated by s_t and are approximately equal to the estimated sum $s_1 + s_t + s_l$.

Having realized that IDMaps is blind to position, Ng and Zhang [5] suggested using coordinate-based mechanisms to predict Internet network distances: Tracers are embedded in a Euclidean space and each network region finds its image coordinates according to distances measured from the region to some Tracers. The estimated distance between two regions is then given by the distance between their images in Euclidean space.

However, due to the Internet AS structure that has a core in the middle and many tendrils connected to it [16], embedding distances between nodes located in different tendrils will result in a large embedding error. Thus, we suggest a **threshold criterion** to select either the Euclidean distance or the additive estimation of IDMaps. Our criterion performs well for multiple types of input graphs and different Tracers selection algorithms and outperforms both the additive IDMaps and GNP.

C. Dynamic Distance Map

Due to the dynamic character of real networks, the calculated distance map has to be updated to track network distance changes. As with IDMaps [4], we expect only part of the Tracer pair distances to be measured at every measurement cycle. In the following, we discuss the map recalculation due to changes. We distinguish between two scenarios: slow variation, where less than 20% of the Tracer pair distances changes, and large transient, where more than 20% of the Tracer pair distances are changed between consecutive measurements.

Labovitz *et al.* [17] studied BGP route changes and concluded that, in reality, Internet routing is stable. They found that over

80% of the routes change at a frequency lower than once a day. Similarly, Paxson [18] found that a very high percentage of the routes are stable. Given that our measurements cycle is 15–20 min long, we can expect the slow variation scenario to be the common one, where large transient would happen rarely, due to major failures or attacks.

1) *Slow Variations*: Such update cycles recur constantly. In this case, one does not need to recalculate the coordinates from scratch and can use the previous coordinates as the initial condition for the embedding calculation. We verified that BBS embedding is insensitive to small changes in the input graph, in case the pairwise field force is given by the difference between the Euclidean and the network pair distances [see (7)]. Indeed, even after modifying as many as 20% of the input edges, and resuming simulation from the previous particles positions, our simulation CPU time is only 0.5% (for 150-node Waxman topology) and 5% (for 15-node BA topology) of the CPU time of the full calculation.

2) *Large Transient*: When network topology undergoes major change, e.g., failure of a major AS, we are required to recalculate all of the coordinates. In IDMaps, the number of measurements is $t^2 + rN$, where t is the number of Tracers, N the number of Address Prefixes (APs), and r is the number of measurement per AP. Due to the added accuracy of our embedding, we can use fewer Tracers than IDMaps, but we need more measurements per AP. That is, we use smaller t values and pay with higher r values, $r = 5$ compared to $r = 2$ in IDMaps. This gives us a factor of three more measurements, which we feel is worth the added accuracy.

The remainder of the paper is organized as follows. The BBS simulation, initial condition, friction, and an example embedding in the Euclidean plane are discussed in Section II. In Section III, BBS is compared to four other methods, using simulated graphs created according to both Waxman [19] and Barabási–Albert (BA) [20] methods. In Section IV, we apply BBS in two practical applications, topology aggregation and Internet distance maps, comparing it with previous results from [2], [4], and [5].

II. BIG-BANG SIMULATION

In this section, we shall discuss the embedding of network distances using BBS. We develop the potential field force equation and discuss the initial conditions for them and the effect of friction. Finally, an example of simulated metric is given, which is embedded perfectly in the $2d$ Euclidean space, that is, the linear plane.

A. Model

The vertices of the graph, the network nodes, are modeled as a set of particles, traveling in the Euclidean space under the affect of a potential force field. Each particle is the geometric image of a vertex. The field force is derived from potential energy which is equal to the total embedding error

$$E_T(v_1, v_2, \dots, v_n) = \sum_{\substack{i,j=1 \\ i>j}}^n E_{ij}(v_i, v_j) \quad (1)$$

where $v_k, k = 1, \dots, n$, are vectors designating the coordinates of the n network nodes in the target Euclidean space \mathcal{R}^d . The distance embedding error of a pair of particles, the “pair embedding error,” is denoted by E_{ij} . The field force induced between the two particles either pulls or repulses the particle pair, aiming to reduce their pair embedding error.

The rationale behind our approach is similar to locking an adaptive tracking loop using an increasing feedback constant to increase sensitivity or decrease the tracking bandwidth. Our method thus comprises several calculation phases performed sequentially. We start with an insensitive but less optimal potential function, and as we move along to the next phases we use an increasingly sensitive and adequate potential function. The sensitivity of the pair error functions in each phase increases with respect to the directional relative error (16). Numerical stability is maintained, because earlier phases are normally capable of substantially reducing the maximum pair embedding error.

In the first phase, the induced pair field force is equal to the difference between the Euclidean and network pair distances. The resulting field force can be realized by attaching an ideal spring with a fixed elastic coefficient to each pair of particles. The rest length of the spring in each pair is equal to the network pair distance. The objective of this phase is to find an approximate minimal energy configuration of the system of particles. The objective of the rest of the phases is to reduce the distortion of large relative error edges at the price of slightly increasing the average relative error. During each calculation phase, the particles are traveling in trajectories which tends to reduce the potential energy of the entire system. At the end of each phase, the system approximately achieves an equilibrium point where the potential energy is minimized.

At the beginning of the first phase, particles are placed at the origin. The field forces for the first iteration are chosen randomly. The initial position of particles in the next phase is at the point where the potential energy had achieved its global minimum in the previous phase. That point need not be the final position of the previous phase because the particles trajectories are stopped near but not precisely at equilibrium.

The potential field force in each calculation phase p is derived from a phase-specific potential energy $E_T^{(p)}$. The phase pair-embedding error function denoted by $E_{ij}^{(p)}$ assumes the form

$$E_{ij}^{(p)}(v_i, v_j) = \mathcal{F}(\|v_i - v_j\|, \Delta_{ij}), \quad \text{for } i \neq j \text{ and } v_i \neq v_j \quad (2)$$

where $\|\mathbf{x}\|$ is the Euclidean size of vector $\mathbf{x} \in \mathcal{R}^d$, i.e. $\|\mathbf{x}\| = \sqrt{\sum_{t=1}^d x_t^2}$, and $\Delta = (\Delta_{ij})$ is the distance matrix among network node pairs (i, j) . The potential function of the first phase $E_T^{(1)}$, (6), is equal to the simple squared error used in [5].

A calculation phase consists of several iterations which move the particles in discrete time intervals. The particles positions and velocities calculated in the current iteration are the input to the next iteration. An iteration begins with calculating the field force \vec{F}_i on each particle at the current particles’ positions. In all iterations, except the first iteration of phase 1, the field forces are derived from the potential energy (1). Next, assuming constant field forces in time interval $(t, t + \delta t)$, apply (27) and (28) to

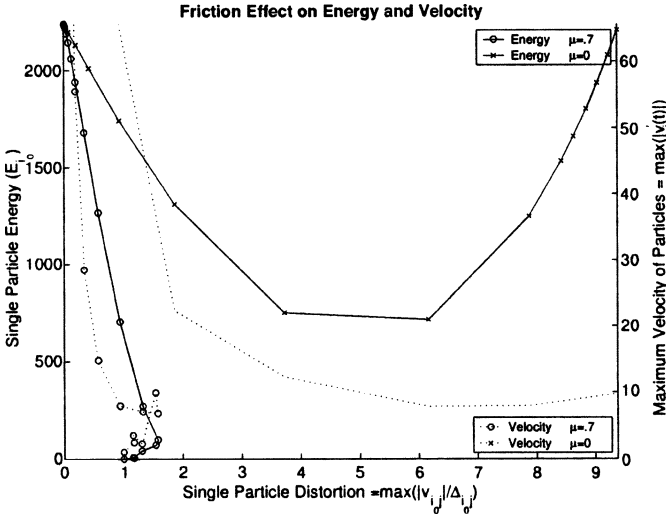


Fig. 2. Friction effect on energy and velocity.

calculate the positions and velocities at time $t + \delta t$. At the end of an iteration, the new potential energy (1) and other statistics are calculated at the new particle positions. They include the maximum particle velocity $\max_i \|\dot{v}_i\|$ and maximum symmetric pair distortion d_{ij} (9).

There are several numerical tradeoffs discussed in Appendix A: short versus long simulation time for each phase, in order to obtain efficiency while insuring good convergence, and large versus small simulation time-step, in order to obtain efficiency yet insuring particles are attracted to their minimum energy position.

In each iteration, we move all n particles, where $n \equiv |V|$ is the number of graph vertices being embedded. The force elements affecting each particle are induced by all other $(n - 1)$ particles. A single iteration thus has a complexity of $O(n^2)$. Because the total number of iterations I is independent of n , the overall complexity of the method is $O(In^2)$, that is, linear in the input size.

B. Friction Versus Kinetic Energy

As a particle accelerates under the affect of the force field, it is also attenuated by a simulated friction force. The friction force slows the particles down, so they can be drawn into wells of the potential energy. This effect is illustrated by Fig. 2, comparing two runs with an identical simulated metric input (see Section II-E) consisting of 20 nodes.

The figure shows the movement of one of the particles during the first calculation phase of each run. The horizontal axis depicts the maximum ratio between Euclidean distance and original network distance for all node pairs attached to that one particle i_0 . As our particle converges to its perfect embedding, its horizontal position approaches the distance ratio 1. The solid lines depicts the potential energy of the particle i_0 , given by $E_{i_0} = \sum_{j=1, j \neq i_0}^n E_{i_0 j}(v_{i_0}, v_j)$. The dashed line depicts the maximum velocity of all particles, $\max_{i=1}^n (\|\dot{v}_i\|)$, which is closely related to the total kinetic energy of the system. The lines marked with circles (o) depict the case where the dynamic friction coefficient is $\mu_k = 0.7$, whereas the lines marked with a cross (x) depict no friction, that is, $\mu_k = 0$.

Each marked point represents 10 ms of CPU time. The total CPU time per phase 1 is thus 200 ms with friction and 160 ms without friction. When comparing the solid and dashed lines, we should recall that initially our particles are placed at the origin, where the distance ratio equals 0. The effect of friction is to attenuate the particles that are moving away from each other, so that immediately as $\|v_{i_0 j}\|/\Delta_{i_0 j} > 1$, the field force attracts the particles back to each other and they are stopped at equilibrium of $\|v_{i_0 j}\|/\Delta_{i_0 j} = 1$. Without friction, however, the particles are moving away too fast and the attracting field force just cause them to oscillate forever with constant velocity.

The friction force attenuating each particle depends on the normal force, which is particle-dependent, and on the constant friction coefficients of the system, defined in Appendix E. Different friction coefficients, denoted μ_s, μ_k , are accounted for static and moving particles, respectively.

Fig. 2 might suggest that BBS is sensitive to the system friction coefficients. Thus, we checked the sensitivity to system friction coefficients of the CPU time, indicating calculation iterations count, and the maximum and average of symmetric pair distortion d_{ij} . Due to space limitations, we present sensitivity of results for a fixed graph size, that is $n = 50$. Note that in the case that the graph size is fixed, a calculation iteration has constant complexity, so that the CPU time is proportional to total number of iterations. Fig. 3 illustrates results for $n = 50$ tracers in Waxman and BA topologies. The figure indicates that all three values are almost indifferent to friction in the wide range $0.003 \leq \mu_k \leq 3$. As for larger and smaller graph sizes, in limited tests we've found similar indifference to friction coefficients.

C. Potential Field Force

We shall now calculate the field force $\vec{F}_{i_0}^*$, derived from the potential energy (1), which affects one particle i_0 . This force is given by the opposite of the potential energy gradient with respect to the position of its particle v_{i_0} as follows:

$$\vec{F}_{i_0}^* = -\nabla_{v_{i_0}} E_T(v_1, \dots) = -\sum_{i,j=1, i>j}^n \nabla_{v_{i_0}} E_{ij} \quad (3)$$

where $\nabla_x f(\cdot)$ denotes the gradient with respect to vector $x = (x_1, x_2, \dots, x_d)$ defined as

$$\nabla_x f \equiv \left(\frac{\partial}{\partial x_1} f \quad \frac{\partial}{\partial x_2} f \quad \dots \quad \frac{\partial}{\partial x_d} f \right).$$

Let F_{ij} denote the magnitude of the field force element induced on particle i_0 by particle j , along the course of vector $v_{j i_0}$. Then the force affecting particle i_0 equals the sum of induced forces by the other particles

$$\vec{F}_{i_0}^* = \sum_{\substack{j=1 \\ j \neq i_0}}^n F_{i_0 j} v_{j i_0}. \quad (4)$$

The sign of F_{ij} determines whether the induced force pulls or repulse the two particles i and j . If $F_{ij} > 0$, the induced force pulls the two particles together, whereas if $F_{ij} < 0$ the induced

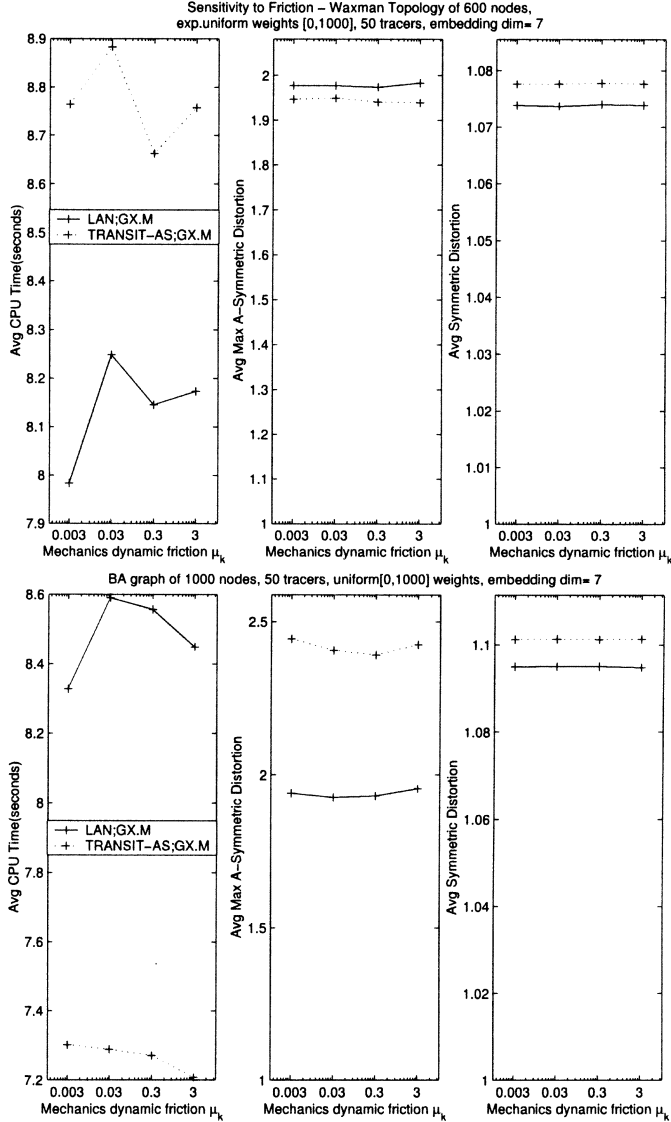


Fig. 3. Sensitivity to friction coefficients.

force repulse them apart. In Appendix C, we calculate the pair field force and show that it is given by

$$F_{i_0j} = \frac{d}{dx} \mathcal{F}(x, \Delta_{i_0j}) \Big|_{x=\|v_{ji_0}\|}. \quad (5)$$

Namely the field force induced on a particle by another particle is given by the derivative of the pair embedding error with respect to the Euclidean distance between the particles.

D. Error Functions of Phases

At the first phase the pair embedding error, $E_{ij}^{(1)}$ is given by

$$\begin{aligned} E_{ij}^{(1)}(v_i, v_j) &= \mathcal{F}(\|v_i - v_j\|, \Delta_{ij}) \\ &= (\|v_i - v_j\| - \Delta_{ij})^2 \end{aligned} \quad (6)$$

that is the squared pair distance error. The field force induced between two particles is given by (5), that is

$$\begin{aligned} F_{ij}^{(1)} &= \mathcal{F}_x(\|v_i - v_j\|, \Delta_{ij}) \\ &= 2(\|v_i - v_j\| - \Delta_{ij}). \end{aligned} \quad (7)$$

Namely, in the first phase, the force induced between each particle pair is equal to twice the distance error of the two particles. Substituting (30) for the field force at $t = 0$, we find

$$\vec{F}_i(t)|_{t=0} = 2 \sum_{\substack{j=1 \\ j \neq i_0}}^n (\|u(i, j)\| - \Delta_{ij}) \hat{u}(i, j). \quad (8)$$

The error functions we chose below for the rest calculation phases are increasingly sensitive to directional relative error (16). An equally good output should result from choosing any similar series of increasingly sensitive error functions.

The symmetric pair distortion d_{ij} is defined as the maximum between the expansion and the contraction ratios

$$d_{ij} = \max\left(\frac{\|v_i - v_j\|}{\Delta_{ij}}, \frac{\Delta_{ij}}{\|v_i - v_j\|}\right). \quad (9)$$

The pair embedding error function of the second phase is selected as follows:

$$E_{ij}^{(2)} = (d_{ij} - 1)^2. \quad (10)$$

Substituting (6) into the above, we find

$$E_{ij}^{(2)} = \frac{E_{ij}^{(1)}}{\min(\|v_i - v_j\|, \Delta_{ij})^2}. \quad (11)$$

Thus, the second phase function is the square of pair distance error, divided by the square of minimum between embedding and network pair distance.

For last two phases, we selected the following two functions:

$$E_{ij}^{(3)} = \exp(E_{ij}^{(2)})^{\frac{3}{4}} - 1 = \exp^{(d_{ij}-1)^{\frac{3}{2}}} - 1 \quad (12)$$

$$E_{ij}^{(4)} = \exp^{E_{ij}^{(2)}} - 1 = \exp^{(d_{ij}-1)^2} - 1. \quad (13)$$

The derivation of the induced forces in the other phases following the first phase is given in Appendix D.

E. Example of Perfect Embedding

We take as a simple example the metric representing the distances among fixed points in a Euclidean space. Note that the Euclidean distances between any set of fixed points is a metric because they satisfy the triangle inequality. Moreover, this metric can be embedded perfectly, i.e., with distortion 1 for all pairs, by choosing the coordinates of those fixed points as geometric images for its vertices. For ease of presentation, we stick to the simplest case of two-dimensional space, that is, the XY plane. Next we calculate the input metric, that is, the $(n/2)(n-1)$ Euclidean distances among these n points, and run our BBS procedure with $d = 2$ on this metric.

If the embedding is accurate, we would expect that Euclidean distances would match exactly, up to the rounding error, to the

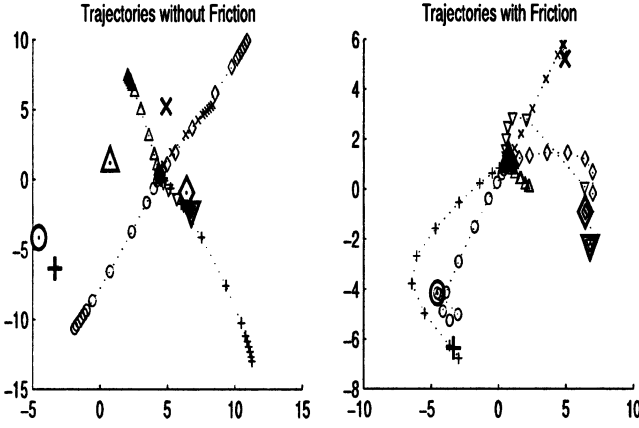


Fig. 4. Friction effect on simulated metric trajectories.

simulated metric distances. We experimented with BBS embedding with and without friction. The results are illustrated in Fig. 4. We scattered $n = 20$ random points uniformly in a 16×16 rectangle, $(x_i, y_i) \in (-8, +8) \times (-8, +8)$. For presentation clarity, we have drawn only the trajectories of the first six particles out of $n = 20$. The trajectories on the right side are from a run with zero friction, $\mu_k = \mu_s = 0$, whereas on the left side they are with friction coefficient of $\mu_k = 0.7$ and $\mu_s = 0.9$. Trajectories were matched to the input points using only the shift, rotation, and reflection transformations. We placed an enlarged marker at each of the input points and marked the trajectory of the corresponding particle with a smaller marker of the same type. Indeed, all trajectories on the right, with friction, converged to their corresponding input points. Especially interesting is one of the right trajectories, marked with the triangle pointing up, which left the origin and then returned back home to its point at the origin. To the contrary, none of the trajectories with zero friction, except for the trajectory matched by the transformation, have even come close to their corresponding input points.

III. EMBEDDING METHOD COMPARISON

In this section, we compare our BBS method to three other embedding methods, a topology aggregation method [21], and a force-directed graph drawing method [11].

A. Other Embedding Methods

We begin by briefly describing the four methods.

1) *Semi-Definite Programming*: (---)³ SDP is a well-known optimization technique that has drawn attention from diverse areas. The use of SDP for graph embedding was introduced in the seminal work of Linial *et al.* [15]. The SDP problem in primal standard form is

$$\max \text{tr}(CX) \quad \text{s.t.} \quad A(X) = \begin{bmatrix} \text{tr}(A_1 X) \\ \text{tr}(A_2 X) \\ \dots \\ \text{tr}(A_k X) \end{bmatrix} = a \quad (14)$$

where all of the matrices A_i , X , and C are symmetric, and a is a vector. The matrices A_i and C and the vector a depend on the

³The graphs in this section use a unique marker for each method which is given in the heading of the paragraph describing each method. Our embedding method, Big-Bang Simulation (BBS), is marked with “×.”

metric Δ , and X is the variable. A Cholesky decomposition of the solution to the SDP problem, $X = MM^t$, yields the embedding coordinates matrix $M_{n \times n}$. The embedding dimension, that is, the number of columns of matrix M , is reduced by Johnson and Lindenstrauss’ [22] random projection from \mathcal{R}^n into \mathcal{R}^d . The Matlab code we used is due to [23]. It executes the Brian Borchers SDP solver [24].

2) *Multidimensional-Scaling* (---): The classical metric MDS [7] has a simple and low-complexity implementation. The Matlab code we used is described in [8, Sec. 4].

3) *Down-Hill Simplex* (---): This minimization method has been known for a very long time [25] and was recently used as an embedding method in [5]. For Euclidean embedding, the minimization target function is either equal to the total embedding error from phase 1, $E_T^{(1)}$, (6); or the normalized pair error used in GNP [5] as follows:

$$E_{ij}^{(GNP)} = \left(\frac{\|v_i - v_j\| - \Delta_{ij}}{\Delta_{ij}} \right)^2. \quad (15)$$

We experimented with both unnormalized and normalized minimization target functions, and both yielded similar results. The results for DHS throughout this paper were calculated with the normalized target function (15).

4) *MSTxRST* (---): The MSTxRST topology aggregation procedure, discussed in [2], represents a network by an aggregated graph of at most $(x + 1)n$ edges. The edges in the aggregated graph are the union of a minimum spanning tree (MST) of a clique that is built from the shortest paths of the original graph and x or more random spanning trees (RSTs) of this clique. For instance, in MST2RST, the aggregated graph has $3n$ edges. This method had the lowest distance distortion among the aggregation methods considered in [2], and our comparison was done using the C-code used in that work.

5) *KK89 Force Directed Graph Drawing* (---): This algorithm was used for a different purpose, that is, to create a straight-line drawing of a graph with as much symmetry as possible. The mechanical model of this algorithm, called *spring embedder*, resembles our model. The drawing process simulates mechanical systems, where vertices are replaced by rings and edges are replaced by springs. Many attempts were made to draw graphs utilizing the spring model. However, in most of these attempts, springs connect only between neighbor vertices, and the spring original length is set to the corresponding edge length. We did limited testing of BBS in which forces were considered only between neighboring nodes and verified that it *did not converge well* for simulated metric graphs of medium sizes, that is 50–100 nodes. Thus, all of the embedding methods compared here take as their input the clique built from the shortest paths of the input graph.

The first attempt to model such a clique with springs is due to Kamada and Kawai [11]. We decided to compare Kamada and Kawai, because their KK89 algorithm is simple, and its embedding error function is same as ours in phase 2, except that Δ_{ij}^2 replaces the denominator in (11). We set the algorithm constants as follows:

$$\begin{aligned} L_0 &= 10 \\ K &= 1 \\ \epsilon &= 0.00001. \end{aligned}$$

The second-order Newton–Raphson (NR) method of KK89 is very sensitive even with a simulated metric graph of small size, that is, 15–20 nodes. Thus, the following changes were made to the original algorithm.

- outer The outer loop of the KK89 algorithm selects a particle with maximal gradient size $\|\nabla_{v_m} E_T\|$ and executes the inner loop with that particle. Originally, this loop continues as long as the maximal gradient size is larger than ϵ . However, we stop after 1000 times, leaving particles at their last position.
- inner The inner loop of the KK89 algorithm finds a local minimum of $E_T(\dots, v_m, \dots)$ as a function of the position v_m , using second-order NR. Originally, this loop continues until $\|\nabla_{v_m} E_T\| \leq \epsilon$. However, we stop after 500 iterations and check if the gradient size $\|\nabla_{v_m} E_T\|$ is smaller than its size at the beginning of the loop. If it is, the outer loop continues and selects the next particle to be moved, but otherwise we abort the outer loop.
- region We check after each iteration of second-order NR that its particle remains inside the ball with radius L_0 , that is, twice the size of the drawing. In case the result point is outside the ball, we place the particle at random point in a 0.001-neighborhood of the origin.

B. Environment Details

The network graphs in our comparison were created according to both Waxman [19] and Barabási–Albert (BA) [20] methods. Only a subset of the graph nodes was selected to be embedded according to the characteristics of the two applications. For IDMaps [4], the subset was selected using two Tracer placement methods, TRANSIT or LAN (stub), which select the Tracers among the highest or lowest degree nodes, respectively. For topology aggregation [2], the subset includes all of the border nodes which are located on the edges of the Waxman topology rectangle area.

To increase the confidence, each experiment was conducted on ten networks using five sets of random weights per network. Namely, each point in the comparison graph results from 50 embedding experiments.⁴ For Waxman networks, the edge weights were taken as $\lfloor 10^{\text{exp}} + .5 \rfloor$, where exp are i.i.d. uniformly distributed in the interval (0,3]. For BA networks, the edge weights are i.i.d. uniformly distributed in the interval [1,1000].

C. Performance Metrics

The *symmetric pair distortion*, d_{ij} , is defined in (9). The *directional relative error*, E_{rel} , was defined by [5, eq. 4] as

$$\frac{\text{predicted distance-measured distance}}{\min(\text{measured distance; predicted distance})}$$

and for Euclidean embedding is given by

$$E_{rel} \equiv \frac{\|v_i - v_j\| - \Delta_{ij}}{\min(\|v_i - v_j\|, \Delta_{ij})^2}. \quad (16)$$

⁴However, for $n = 450$ tracers, we performed only 5×1 embeddings, due to their large space and time requirements.

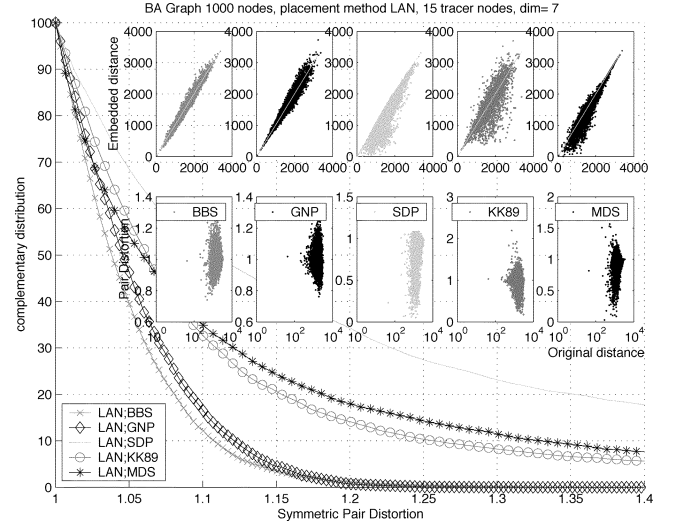


Fig. 5. BA symmetric pair distortion.

Comparing with (11), we find that $|E_{rel}| = \sqrt{E_{ij}^{(2)}}$, and, substituting (10), we thus have $d_{ij} = 1 + |E_{rel}|$. The *average symmetric pair distortion* is given by $\bar{d}_{ij} = 1 + |E_{rel}|$.

As a measure of the worst-case distortion, we use the *two-sided embedding distortion* defined as

$$\min \left(c_1 c_2 : c_1 \Delta_{ij} \geq \|v_i - v_j\| \geq \frac{1}{c_2} \Delta_{ij} \right) \quad (17)$$

over all node pairs $i, j = 1, \dots, n; i \neq j$.

An alternative measure, defined by Linial *et al.* [15], is the *one-sided distortion*

$$\min \left(c : \Delta_{ij} \geq \|v_i - v_j\| \geq \frac{1}{c} \Delta_{ij} \right) \quad (18)$$

over all node pairs $i, j = 1, \dots, n; i \neq j$. The above measures are comparable since, from linear contraction of our coordinates, we obtain $c = c_1 c_2$. However, this contraction increases the other metric \bar{d}_{ij} .

D. Comparison Results

A different marker depicts each of the embedding methods compared in following figures. Lines with no marker depict SDP method, lines marked with “*” depict the MDS method, “+” depict the DHS method, “•” depict the MSTxRST method, and “o” depict the KK89 method. Some figures uses different line styles to distinguish between the several tracer placement methods shown on the same picture.

1) *Symmetric Pair Distortion*: We compare the accuracy of the five methods using the complementary symmetric distortion distribution over all pairs of the 50 embedding experiments. Fig. 5 compares all methods, except from topology aggregation (MSTxRST), for the LAN placement methods of $n = 15$ tracers in the BA topology, which is typical to the IDMaps application discussed here. BBS is more accurate than DHS, having smaller complementary distribution along the entire range of the distortion. For example, the probability $P(d_{ij} > 1.1)$, is 0.11 for BBS versus 0.15 for DHS, i.e., a 35% increase. SDP is much worse than both. The insets in the top right depict the nonsymmetric

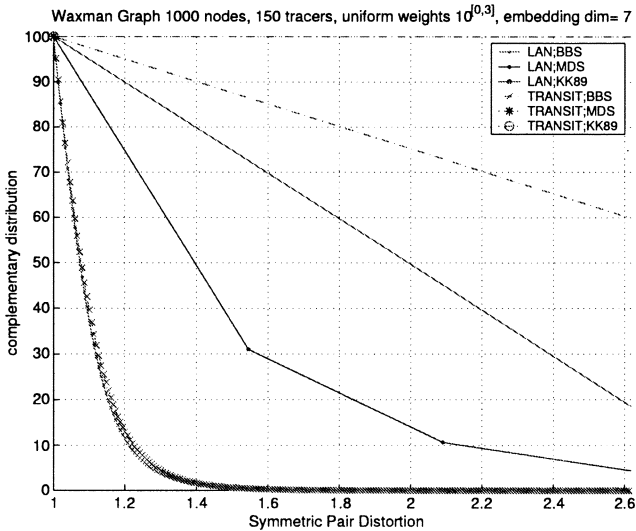


Fig. 6. Waxman symmetric pair distortion.




Fig. 8. Performance metrics versus BA embedding dimension.




Fig. 7. Performance metrics versus BA graph size.

pair distortion in the lower graphs and the embedded distance in the upper graphs, both versus original pair distances. The SDP embedding contract nearly all edges whereas DHS and BBS have similar masses of contracted and expanded edges. Fig. 6 compare MDS, KK89, and BBS for the LAN and TRANSIT placement methods of $n = 150$ tracers in the Waxman topology. We rule out SDP and DHS as practical embedding methods for large n due to their long running times and sensitivity, respectively. The accuracy of BBS is much better than MDS, e.g., $P(d_{ij} > 1.3)$, is 0.05 for LAN-BBS and TRANSIT-BBS versus 0.6 for LAN-MDS and 0.999 995 for TRANSIT-MDS.

2) *Embedded Graph Size*: Fig. 7 depicts the performance of the different methods as a function of the number of embedded nodes. For $n > 30$ or 70, the long running times of SDP or KK89 and high sensitivity of DHS exclude them from the comparison. For BBS, we depicted in dotted lines a linear fit of the

average CPU time graph, which indicates that BBS has complexity Cn^2 , $\forall n \in [10450]$. The value of C calculated for a Pentium-IV 2.0-GHz processor is 6×10^{-4} . The BBS embedding distortion is the lowest for all graph sizes except for $n \geq 150$, where MST6RST distortion is smaller, e.g., for $n = 450$, the distortions are 4 versus 15.5, respectively. BBS has the lowest average symmetric pair distortion in all graph sizes. The symmetric pair distortion of MST6RST and SDP are not comparable with the rest of the methods, since all MST6RST edges are expanded and most SDP edges are contracted.

3) *Embedding Dimension*: Fig. 8 illustrates the effect of the embedding dimension on the BA topology with $n = 15$ tracers. Naturally, the performance of all methods improves when the embedding dimension increases. For the BA topology, the knee point, where the improvement diminishes, is at $d = 7$, as was found by Ng and Zhang [5].

For $d < 7$, the performance gap between all other methods and ours is significant, as can be seen by the two right graphs in Fig. 8. The difference is larger for the two-sided embedding distortion. The larger performance gap in embedding distortion is explained by the improvement of large distortion pairs in the last two phases of our calculation. Although the main objective of SDP is to reduce the embedding distortion, with LAN placement it has a larger embedding distortion compared to both DHS and BBS for all dimensions. However, with TRANSIT placement for $d \geq 7$, the SDP distortion is the smallest. The results for the Waxman topology with $n = 70$ and 150 tracers are similar and are thus omitted. However, for $n > 30$, the sensitivity of DHS rules it out as a viable method.

4) *Input Graph Sensitivity*: We tested embedding of the BA and Waxman topologies with 15 TRANSIT and 150 BORDER nodes, respectively. We either increased or decreased the weights of 15% of the input edges by 10% and disconnected or reconnected an additional 5% of the edges. The input graph is embedded first with regular initial conditions, and the modified graph is embedded with initial positions of the unmodified graph embedding. The test is repeated three times where the

initial positions in each embedding are the output positions of previous embedding. We compare the accuracy and complexity of the three embedding by embedding the three modified graphs again with regular initial conditions. With unnormalized square embedding error, (6), the sensitivity of BBS is **low**. The calculations from previous graph position achieves comparable performance metrics with 13 iterations and 0.01–0.15 CPU seconds for BA or Waxman topology, respectively, compared to 200–1000 iterations and 0.15–10 CPU seconds, with regular initial conditions.

IV. APPLICATIONS

A. Topology Aggregation

Topology aggregation is used in hierarchical networks to compactly represent the cost of traversing a network between every possible entrance and exit points. If the aggregation decreases the cost of an edge, it means that a routing application that is using the aggregated view to find a route in the network may select a path with a higher cost than it expects. It is generally considered undesirable to receive such bad news, since, e.g., if the traversal cost means delay, we may choose a route which violates the application bound. Thus, for topology aggregation, we seek an embedding that favors length increase over length decrease.

The MSTxRST aggregation procedure (Section III-A4) insures that an edge length in the aggregation is never smaller than the original. However, in the embedding methods we have discussed so far, some distances are contracted while others are expanded. A transformed embedding in which no edges are contracted (or expanded) was discussed in Section III-C. The problem with the transformed embedding is that it increases the average pair distortion which is, of course, important for aggregation performance.

An alternative way to favor expansion in our embedding is an introduction of a **price factor** denoted \mathbf{P} , $\mathbf{P} \gg 1$, in the definition of pair embedding error functions

$$\hat{E}_{ij}^{(p)}(v_i, v_j) = \begin{cases} E_{ij}^{(p)}(v_i, v_j), & \text{if } \frac{\|v_i - v_j\|}{\Delta_{ij}} \geq 1 \\ \mathbf{P}E_{ij}^{(p)}(v_i, v_j), & \text{otherwise.} \end{cases} \quad (19)$$

Thus, the weight of a contracting pair in the total embedding error will be larger than the weight of an expanding pair. Such a price factor can be directly incorporated into the DHS method, but should pose inherent difficulty for MDS since it is not linear. We introduce the price factor into our calculation at the end of the first calculation phase. Particles are placed at the best position of the first phase, and then moved by a modified field force incorporating the price factor \mathbf{P}_1 ; $\mathbf{P} \gg \mathbf{P}_1 > 1$. As particles near an equilibrium of the modified field force, the price factor is increased again, and particles continue from the previous equilibrium point to the next equilibrium. This procedure is repeated until the price factor is increased to the final value \mathbf{P} . In the rest of the phases, the calculation continues with the field force which directly incorporates \mathbf{P} .

Fig. 9 illustrates the effect of the price factor $\mathbf{P} = 512$ on our embedding method compared to embedding without it, i.e., $\mathbf{P} = 1$. The middle graph shows that the price factor decreases

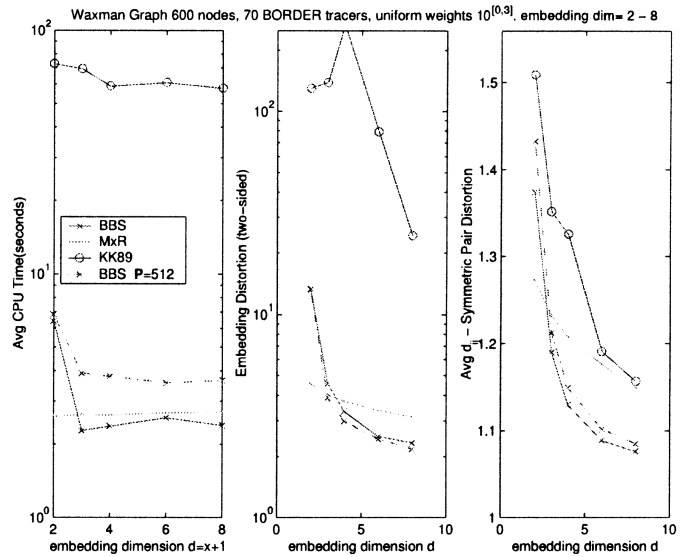


Fig. 9. Waxman symmetric versus asymmetric aggregation.

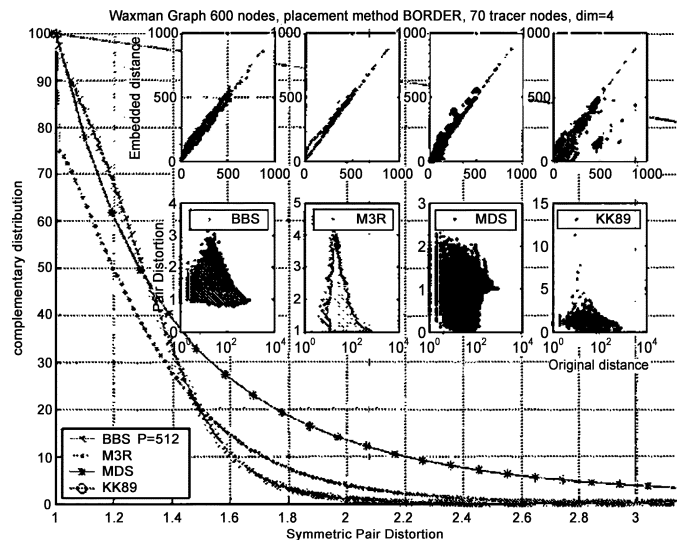


Fig. 10. Aggregation distortion histograms.

our two-sided embedding distortion by approximately half at $d = 2$ and, for $d > 2$, its effect is only modest. However, with the price factor, the average symmetric pair distortion increases approximately by 2%, which translates to a 10% increase of the absolute relative error. Here the knee point where the improvement diminishes is at $d = 6$. Although at $d = 2$ (or $x = 1$) our performance is worth less than MSTxRST, at $d = x + 1 \geq 3$ it supersedes MSTxRST.

Fig. 10 compares the symmetric pair distortion histograms of MST3RST and BBS with $d = 4$ and $\mathbf{P} = 512$. The insets in the top right illustrate nicely the effect of the price factor which is to force all pairs to expand rather than contract. Almost all of the pairs of BBS are above the $y = x$ line in the upper inset and the $y = 1$ line in the lower inset.

B. Internet Distance Estimation

IDMaps [4] is a project that aims to build a global architecture for Internet host distance estimation and distribution. The archi-




Fig. 11. Directional relative error (stand alone).

ecture is based on Tracers, which are instrumentation boxes, that are placed in the Internet. Each Tracer measures distances to other Tracers and to APs that are close to it. These measurements are multicast to topology servers which combine them to an estimated map of the Internet.

Euclidean embedding yields smaller relative distance errors compared to IDMaps, especially when both nodes are sharing a single nearest Tracer [see Fig. 1(a)]. The main problem with Euclidean distance estimation is underestimation of large measured distances. This is an inherent problem when embedding the BA topology, as all shortest paths between distant nodes must go through a small number of core nodes [16]. Indeed, due to the triangle inequality, the Euclidean distance between distant nodes is bound to be smaller than their shortest path which goes through the core nodes. On the other hand, IDMaps sum of segments accurately estimate the longer paths going through the core nodes, as in Fig. 1(b).

Throughout this section we have estimated distances in BA graphs, using three sets of random weights per each of five simulated BA graphs. Fig. 11 compares the additive estimation of IDMaps with Euclidean embedding by GNP and BBS methods using the *directional relative error*, (16). We only experimented with LAN and TRANSIT Tracer placement methods, illustrated in the top and bottom pictures, respectively, because LAN placement is the most easy to deploy and TRANSIT placement yields the best mirror selection performance [4]. We used 15 Tracers and measured distances from each host to all 15 Tracers, which matches the conditions of the similar figures in [5]. The groups of vertical lines 75 ms apart depict the distribution of relative errors for measured distances belonging to the 75-ms interval. The lines marked with square, upright triangle, and circles depict GNP, IDMaps, and BBS, respectively. The method marker

is placed at the average relative error point, and the stars depict the median. Each line has whiskers at the 5, 25, 75, and 95 percentiles.

For each placement and calculation method, we randomly picked 150 nodes out of the 1000 graph nodes in each of the 15 simulated BA graphs. We estimated the distance from all other graph nodes to each picked node, that is, a total of 2, 247, and 750 distance pairs per method. The thick lines depict the overall count of measured pair distances per interval.

IDMaps additive estimation has larger positive relative errors than Euclidean estimation for short distances, but is more accurate for longer distances. Therefore, the best distance estimate is achieved by selecting Euclidean estimation for short distances and IDMaps additive estimation for longer distances, based on a threshold of the Euclidean distance. Unfortunately, the optimum threshold point changes for different placement methods, graph topologies, and ranges of random edge weights.

An alternative for selecting between Euclidean and IDMaps additive estimations is using the ratio R between the two, given by

$$R = \frac{\text{Euclidean distance}}{\text{IDMaps additive}} \quad (20)$$

and the estimated distance is selected as follows:

$$\begin{cases} \text{Euclidean distance,} & \text{if } R < R_{Th} \\ \text{IDMaps additive,} & \text{otherwise.} \end{cases} \quad (21)$$

Fig. 12 illustrates the improvements in accuracy of our estimation compared to IDMaps additive, with threshold $R_{Th} = 0.45$. As in the previous figure, we used 15 Tracers and measured distances from each host to all 15 Tracers.

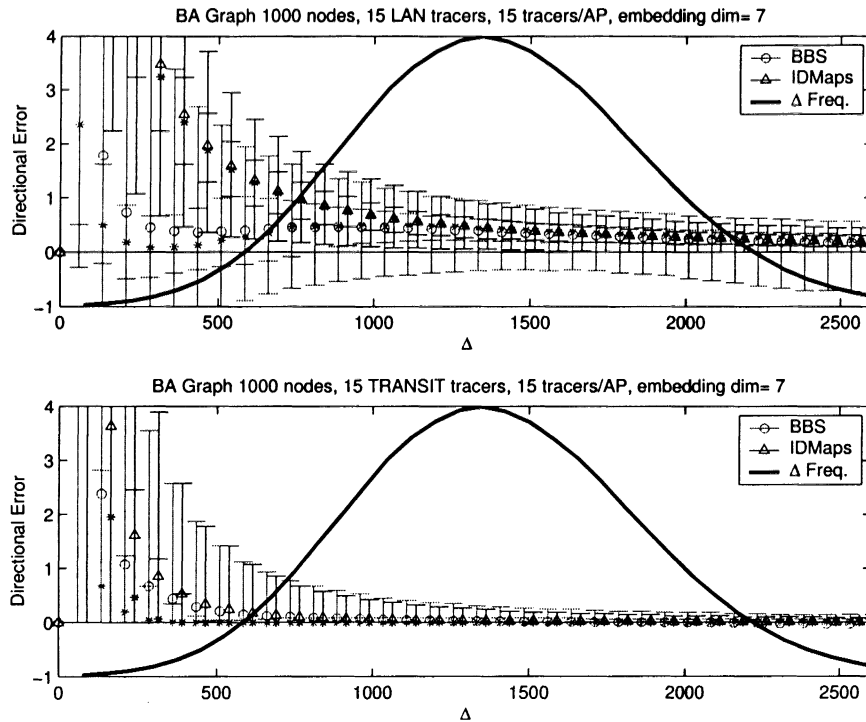


Fig. 12. Directional relative error (threshold selection).

One could have selected the closet mirror as the one with smallest distance estimated by (21). Such naive approach, however, does not maintain the ordering among estimated distances, as some were estimated by additive IDMaps and some as Euclidean distances. We calculate the ratio

$$R' = \frac{\min_k \{\text{Euclidean distance}\}}{\min_k \{\text{IDMaps additive}\}} \quad (22)$$

where the minimum among mirrors for either method is achieved by the closet mirror denoted by $k_{\min}^{(\cdot)}$. The closet mirror is then selected as

$$\begin{cases} k_{\min}^{(\text{Euclidean})}, & \text{if } R' < R'_{\text{Th}} \\ k_{\min}^{(\text{IDMaps additive})}, & \text{otherwise.} \end{cases} \quad (23)$$

We compared the mirror selection accuracy of IDMaps additive with BBS using the selection criterion of (23) with $R'_{\text{Th}} = .45$, i.e., the same threshold value used for distance estimation. Following [4], we randomly selected 10 mirror servers and estimated the closet mirror to each of the rest of the graph nodes acting as clients. The client decision is considered correct if it selects the mirror whose client–mirror distance is at most twice the optimal distance. For each mirror, group rank accuracy is defined as the percentage of correct client decisions. Fig. 13 illustrates the average cumulative distribution function (CDF) of rank accuracy. Each mark is the average of the CDFs from the 15 simulated graphs, where each CDF consists of 300 mirror group experiments performed on a single graph. The number of Tracer distance measurements per AP, is specified in the legend after the “×” mark and is depicted with increasing marker sizes.

The accuracy of additive IDMaps improves with three Tracer measurements compared to one. Using TRANSIT placement,

IDMaps does not improve with more than three Tracer measurements, and, using LAN Tracer placement, it becomes even less accurate with the additional measurements. The accuracy of our threshold selection, however, improves with each additional Tracer measurement. three measurements, it is more accurate than additive IDMaps for both LAN and TRANSIT placement. With 15 measurements for LAN placement, it is nearly as accurate as IDMaps for TRANSIT placement, pointing clients to the closet mirror server with confidence 0.95 in 94% of the cases, compared to 88% of the cases of additive IDMaps for LAN placement.

V. CONCLUDING REMARKS

We presented a novel scheme for embedding a graph metric in a d -dimensional Euclidean space and showed that, with one exception (SDP for very small networks with LAN placement and high dimension), BBS was always the most accurate embedding scheme. In addition, BBS execution time is second only to MDS, but MDS has a stability problem in large graphs and works well only with high d . In addition, MDS is not applicable to topology aggregation, as we have stated earlier. Finally, BBS is insensitive to its only arbitrary parameters, the dynamic and static friction coefficients, as demonstrated in Fig. 3, thus no fine tuning is required.

We demonstrated the efficiency of our scheme for important networking problems: topology aggregation, closet mirror selection, and distance estimation. We believe our method can be applied to other problems as well, such as routing in ad hoc networks and efficient building of peer-to-peer networks and application layer multicast.

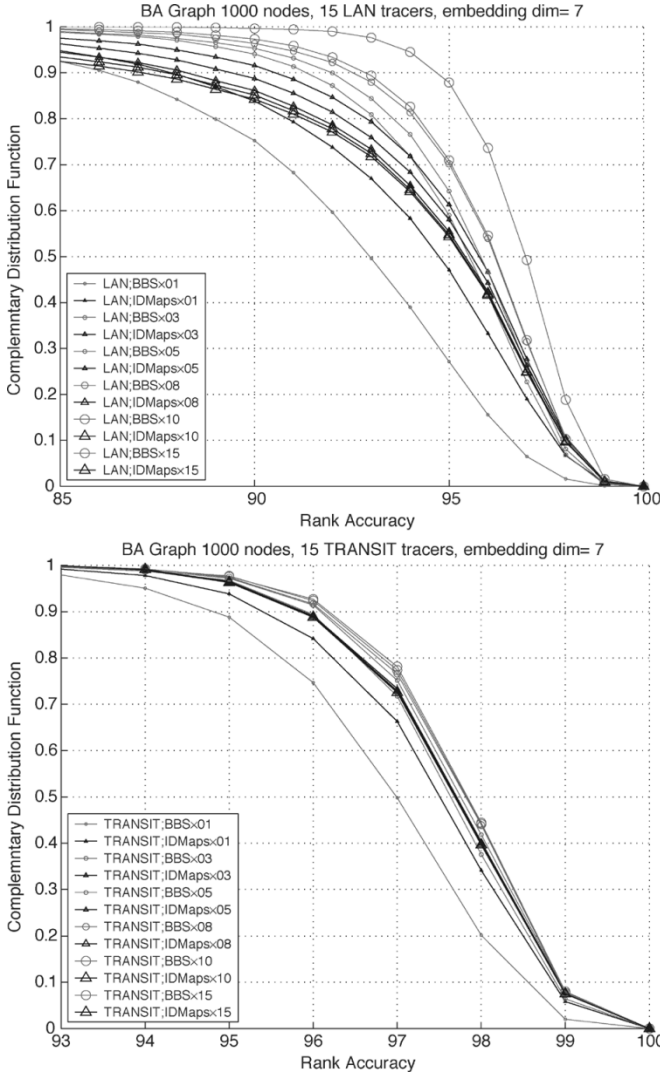


Fig. 13. Mirror selection accuracy (threshold selection).

APPENDIX METHOD TECHNICALITIES

A. Numerical Tradeoffs

The phase of calculation is ended if one of the following conditions, concerning the total energy E_T and rest statistics, are met.

- 1) Particles are almost perfectly embedded; The distortion satisfies either $E_T < \epsilon$ or $\max_{i,j}(d_{ij}) < 1 + \epsilon'$.
- 2) Particles are almost at a halt; the maximum particle velocity decreases below threshold.
- 3) Particles are near an equilibrium; the difference between local minimum and maximum of potential energy decreases below threshold.
- 4) Slow convergence rate of energy; the reduction pace of potential energy is below threshold.
- 5) Divergence of particles; the maximum velocity grows above threshold.

However, except for the first condition, which is always checked, the other conditions are checked only after the physical system's time is greater than the phase's minimum period.

The minimum period for the first phase equals 15 s and the minimum period for the other phases is 10 s.

The time step δt for the next iteration is adjusted according to the total energy and rest statistics. In the beginning of each phase the time step is small, to prevent the system from oscillating. The particles acquire velocity and a definite direction, which reduces the potential energy. The time step is gradually increased as the energy continues to decrease. There is a tradeoff between increasing the time step for greater numerical efficiency and keeping it small to detect and attract particles to global minimum points of the potential energy function. This tradeoff is handled by backtracking the particles to their previous position if the calculated statistics indicates minima of potential energy was skipped.

B. Movement Equations

The instantaneous acceleration of an object is, according to the second movement law of Newton

$$\vec{a} \equiv \frac{d^2}{dt^2} \vec{x} = \frac{\vec{F}_C}{m} \quad (24)$$

where \vec{F}_C is the combined force affecting the object, \vec{x} is its position at time t ,⁵ and m is its mass. Assuming unit mass for all particles, $m = 1$, and neglecting the variation in the combined force \vec{F}_C within a small time interval $(t, t + \delta t)$, we have

$$x(t + \delta t) - x(t) = \dot{x}(t)\delta t + \frac{1}{2}\ddot{x}(t)(\delta t)^2 \quad (25)$$

and

$$\dot{x}(t + \delta t) - \dot{x}(t) = \ddot{x}(t)\delta t \quad (26)$$

where \dot{x} , and \ddot{x} denotes the first and second t -derivatives, that are the velocity and acceleration of the particle respectively. The combined force affecting particle i at position v_i is given by $\vec{F}_{C_i} = \vec{F}_i - \vec{F}r_i$, where \vec{F}_i is the field force affecting this particle, and $\vec{F}r$ is its simulated friction force, discussed below. Thus, we have⁶

$$\dot{v}_i(t + \delta t) - \dot{v}_i(t) = (\vec{F}_i - \vec{F}r_i)|_t \delta t \quad (27)$$

$$v_i(t + \delta t) - v_i(t) = \dot{v}_i(t)\delta t + \frac{1}{2}(\vec{F}_i - \vec{F}r_i)|_t (\delta t)^2. \quad (28)$$

1) *Initial Conditions:* The initial conditions for all phases except the first phase are: all particles are at rest at the point where the potential energy achieved its global minimum along the particles trajectories of the previous phase. The initial conditions for first phase are given by

$$\left\{ \begin{array}{l} v_i(t)|_0 = 0 \\ \dot{v}_i(t)|_0 = 0 \end{array} \right\} \quad \text{for } i = 1, 2, \dots, n \quad (29)$$

that is, all particles are at rest in the origin. At the origin, however, its impossible to decompose the field force into induced

⁵For notational convenience, we'll omit the vector notation from position x .

⁶Our particle positions are denoted by v , to emphasize they are images of vertices, although this letter normally denotes mechanics velocity.

forces between pairs since $v_{ji} \equiv v_j - v_i = 0$; for all i, j : At $t = 0$, in the first iteration, we thus set the field forces at $t = 0$ to

$$\begin{aligned} \vec{F}_i(t)|_0 &= \sum_{\substack{j=1 \\ j \neq i_0}}^n \mathcal{F}_x(x, \Delta_{ij})|_{x=\|u(i,j)\|} \hat{u}(i, j) \\ u(i, j) &= \begin{pmatrix} u_1^{ij} & u_2^{ij} & \dots & u_d^{ij} \end{pmatrix} \end{aligned} \quad (30)$$

where \hat{x} denotes the unit vector along the course of vector x . Here u_i^{ij} , that are coordinates of vectors $u(i, j)$, are independent random variables uniformly distributed across the interval $(0, 1]$. The pair embedding error function of phase 1 $\mathcal{F}^{(1)}$ and the derived field forces are given in Section II-D.

Substituting (30) and (29) into the approximate velocity and position (27) and (28), we find the position and velocity of all particles at time $t = \delta t$.

C. Pair Field Forces

In (2), we have assumed that the pair embedding error E_{ij} depends only on pair Euclidean distance, so, using the chain derivation rule, we obtain

$$\nabla_{v_{i_0}} E_{ij} = \frac{d}{dx} \mathcal{F}(x, \Delta_{ij}) \Big|_{x=\|v_{ij}\|} \cdot \nabla_{v_{i_0}} \|v_{ij}\| \quad (31)$$

where $v_{ij} \equiv v_i - v_j$ is the vector connecting between particle positions v_j and v_i . Taking the derivative of its Euclidean distance, we find

$$\nabla_{v_{i_0}} \|v_{ij}\| = \begin{cases} \frac{v_{ij}}{\|v_{ij}\|} = -\hat{v}_{ji}, & \text{if } i = i_0 \text{ and } j \neq i_0 \\ 0, & \text{otherwise.} \end{cases} \quad (32)$$

So for $i \neq i_0$ all derivatives are zero, and the other derivatives are the opposite of the unit vectors along the course of the vector v_{ji_0} connecting between particle positions v_{i_0} and v_j . Substituting (31), (32) in (3) and comparing with the force expression (4), we thus get (5).

One can easily check that signs are correct in (5). If $F_{ij} > 0$ then the derivative of pair embedding error is positive, so in order to decrease it the Euclidean distance should be decreased. The pair Euclidean distance is decreased by the induced field force because in case $F_{ii} > 0$, the two particles are pulled by the field force. The case of $F_{ij} < 0$ follows immediately from symmetry.

D. Field Forces of Rest Phases

We derive the expressions of field forces from error function of rest calculation phases. Following (5), we derivate the error function of phase 2 (10), and, using (7), we obtain

$$\begin{aligned} F_{ij}^{(2)} &= \frac{dE_{ij}^{(2)}}{d\|v_{ji}\|} \\ &= F_{ij}^{(1)} \begin{cases} \Delta_{ij}^{-2}, & \text{if } \|v_{ji}\| > \Delta_{ij} \\ \frac{1}{\|v_{ji}\|^2}, & \text{if } \|v_{ji}\| < \Delta_{ij} \end{cases} \\ &= 2(\|v_{ji}\| - \Delta_{ij}) \begin{cases} \frac{1}{\Delta_{ij}}, & \text{if } \|v_{ji}\| > \Delta_{ij} \\ \frac{\Delta_{ij}}{\|v_{ji}\|^3}, & \text{if } \|v_{ji}\| < \Delta_{ij} \end{cases}. \end{aligned} \quad (33)$$

Similarly, for phases 3 and 4, whose error functions are defined in (12) and (13), and substituting (33) we obtain

$$\begin{aligned} F_{ij}^{(3)} &= \frac{3}{4} \frac{F_{ij}^{(2)}}{\sqrt{d_{ij} - 1}} E_{ij}^{(3)} \\ &= \frac{3}{2} \sqrt{d_{ij} - 1} \begin{cases} \frac{1}{\Delta_{ij}}, & \text{if } \|v_{ji}\| > \Delta_{ij} \\ \frac{\Delta_{ij}}{\|v_{ji}\|^2}, & \text{if } \|v_{ji}\| < \Delta_{ij} \end{cases} \end{aligned} \quad (34)$$

$$F_{ij}^{(4)} = F_{ij}^{(2)} E_{ij}^{(4)}. \quad (35)$$

E. Simulation Friction

The attenuating force that affects on particle i is given by

$$\vec{F}_T i = \begin{cases} \mu_s N_i \hat{\vec{x}}_i, & \text{if } \dot{v}_i = 0 \\ \mu_k N_i \dot{v}_i, & \text{otherwise} \end{cases} \quad (36)$$

where \hat{x} denotes the particle velocity. The system static and dynamic friction coefficients are constant for all particles. When the particle velocity approaches zero, the friction force used in movement equation is only the dynamic one. In order to prevent particle to switch its direction during an iteration due to friction, the friction is affecting only at the approximately part of iteration time interval where the velocity is opposite to it.

The normal force size N_i represents the relative weight of the particle. Particles with larger weight are less affected by small changes in the positions of the rest of the particles. In the first phase, it is given by

$$N_i = \frac{n}{\text{adjust factor}} \text{avg}_{\substack{j=1 \\ j \neq i}}^n \Delta_{ij} \quad (37)$$

whereas at phases $p = 2, 3, \dots$ all particles have equal relative weight

$$N_i = \frac{n}{\text{normalized adjust factor}} \forall i = 1, 2, \dots, n \quad (38)$$

since the pair error functions are normalized by the pair network distance (10) in those phases. The adjust factors are empirical constants set to

$$\begin{cases} \text{adjust factor} = 7500 \\ \text{normalized adjust factor} = 100. \end{cases} \quad (39)$$

ACKNOWLEDGMENT

The authors would like to thank U. Ashkenazi for implementing the SDP embedding algorithm and for his insight of SDP. They also thank C. Jin for providing the IDMaps simulation code and for helping to integrate new features into it.

REFERENCES

- [1] "Private Network—Network Interface Specification Version 1.0 (PNNI)," ATM Forum Technical Committee, Tech. Recommendation AF-PNNI-0055.000, Mar. 1996.
- [2] B. Awerbuch and Y. Shavitt, "Topology aggregation for directed graphs," *IEEE/ACM Trans. Networking*, vol. 9, pp. 82–90, Feb. 2001.
- [3] W. Theilmann and K. Rothermel, "Dynamic distance maps of the Internet," in *Proc. IEEE INFOCOM*, Tel Aviv, Israel, Mar. 2000, pp. 275–284.

- [4] P. Francis, S. Jamin, C. Jin, Y. Jin, D. Raz, Y. Shavitt, and L. Zhang, "IDMaps: a global Internet host distance estimation service," *IEEE/ACM Trans. Networking*, vol. 9, pp. 525–540, Oct. 2001.
- [5] T. Ng and H. Zhang, "Predicting Internet network distance with coordinates based approaches," in *Proc. IEEE INFOCOM*, June 2002, pp. 170–179.
- [6] E. Cronin, S. Jamin, C. Jin, A. Kurc, D. Raz, and Y. Shavitt, "Constrained mirror placement on the Internet," *IEEE J. Select. Areas Commun.*, vol. 20, pp. 1369–1382, Sept. 2002.
- [7] W. Togerson, "Multidimensional scaling of similarity," *Psychometrika*, vol. 30, pp. 379–393, 1965.
- [8] G. Zigelman, R. Kimmel, and N. Kiryati, "Texture mapping using surface flattening via multidimensional scaling," *IEEE Trans. Visual Comput. Graphics*, pp. 4–6, 2002.
- [9] G. Yona, N. Linial, and M. Linial, "ProtoMap: automatic classification of protein sequences and hierarchy of protein families," *Nucleic Acids Res.*, vol. 28, pp. 49–55, 2000.
- [10] P. Eades, "A heuristic for graph drawing," *Congressus Numerantium*, vol. 42, pp. 149–160, 1984.
- [11] T. Kamada and S. Kawai, "An algorithm for drawing general undirected graphs," *Inform. Process. Lett.*, vol. 31, pp. 7–15, 1989.
- [12] T. M. J. Fruchterman and E. M. Reingold, "Graph drawing by force-directed placement," *Software—Practice and Experience*, vol. 21, pp. 1129–1164, 1991.
- [13] R. Davidson and D. Harel, "Drawing graphs nicely using simulated annealing," *ACM Trans. Graphics*, vol. 15, pp. 301–331, 1996.
- [14] I. Herman, G. Melancon, and M. Marshall, "Graph visualization and navigation in information visualization: a survey," *IEEE Trans. Visual Comput. Graphics*, vol. 6, pp. 24–43, 2000.
- [15] N. Linial, E. London, and Y. Rabinovich, "The geometry of graphs and some of its algorithmic applications," *Combinatorica*, vol. 15, pp. 215–245, 1995.
- [16] L. Subramanian, S. Agarwal, J. Rexford, and R. Katz, "Characterizing the Internet hierarchy from multiple vantage points," in *Proc. IEEE INFOCOM*, June 2002, pp. 618–627.
- [17] C. Labovitz, G. R. Malan, and F. Jahanian, "Internet routing instability," in *Proc. ACM SIGCOMM*, Aug. 1997.
- [18] V. Paxson, "End-to-end routing behavior in the Internet," *IEEE/ACM Trans. Networking*, vol. 5, pp. 601–615, Oct. 1997.
- [19] B. Waxman, "Routing of multipoint connections," *IEEE J. Select. Areas Commun.*, vol. 6, pp. 1617–1622, 1988.
- [20] R. Albert and A.-L. Barabási, "Topology of evolving networks: local events and universality," *Phys. Rev. Lett.*, pp. 5234–5237, Dec. 11, 2000.
- [21] B. Awerbuch, Y. Du, B. Khan, and Y. Shavitt, "Routing through networks with hierarchical topology aggregation," *J. High-Speed Networks*, vol. 7, no. 1, pp. 57–73, 1998.
- [22] W. B. Johnson and J. Lindenstrauss, "Extensions of Lipschitz mappings into a Hilbert space," *Contemporary Math.*, vol. 26, pp. 189–206, 1984.
- [23] E. Ashkenazi, Experiments with low-distortion low-dimensional embeddings of graphs into normed spaces, Electr. Eng.-Systems Dept. Tel-Aviv Univ., Apr. 2002.
- [24] B. Borchers. (2000, Dec.) CDSP 3.2, A Library for Semidefinite Programming. [Online]. Available: <http://www.nmt.edu/~borchers/csdp.html>
- [25] J. Nelder and R. Mead, "A simplex method for function minimization," *Computer J.*, vol. 7, pp. 308–313, 1965.



Yuval Shavitt (S'88–M'97–SM'00) received the B.Sc. degree in computer engineering (*cum laude*), the M.Sc. degree in electrical engineering, and the D.Sc. degree from the Technion—Israel Institute of Technology, Haifa, in 1986, 1992, and 1996, respectively.

From 1986 to 1991, he served in the Israel Defense Forces first as a System Engineer and the last two years as a Software Engineering Team Leader. After graduation, he spent a year as a Postdoctoral Fellow with the Department of Computer Science, Johns Hopkins University, Baltimore, MD. Between 1997 and 2001, he was a Member of Technical Staff with the Networking Research Laboratory, Bell Labs, Lucent Technologies, Holmdel, NJ. In October 2000, he became a Faculty Member with the Department of Electrical Engineering, Tel-Aviv University, Tel Aviv, Israel. His recent research focuses on Internet measurement, mapping, and characterization.

Dr. Shavitt served as TPC member for INFOCOM 2000–2003, IWQoS 2001 and 2002, ICNP 2001, and MMNS 2001, and on the executive committee of INFOCOM 2000, 2002, and 2003. He is an Editor of *Computer Networks* and served as a Guest Editor for the IEEE JOURNAL OF SELECTED AREAS IN COMMUNICATIONS and the *World Wide Web Journal*.



Tomer Tankel (S'03) received the B.Sc. degree in scientific computation and the M.Sc. degree in electrical engineering from Tel-Aviv University, Tel Aviv, Israel, in 1986 and 1992, respectively, where he is currently working toward the Ph.D. degree.

Between 1992 and 2002, he was a Software Architect with subsidiaries of Comverse Technology. Since 2002, he has been with the Computer Network Laboratory, Tel-Aviv University. His research focuses on Internet mapping and peer-to-peer networks.