A Placement Scheme for Peer-to-Peer Networks based on Principles from Geometry

Michael Kleis Fraunhofer Institute FOKUS Competence Center Next Generation Network Integration (NGNI) Kaiserin-Augusta-Allee 31, D-10589 Berlin, Germany kleis@fokus.fraunhofer.de

Xiaoming Zhou Delft University of Technology Faculty of Electrical Engineering, Mathematics, and Computer Science P.O. Box 5031, 2600 GA Delft, The Netherlands X.Zhou@ewi.tudelft.nl

Abstract

Crucial for the performance of Peer-to-Peer networks based on geometric topologies is the measurement complexity and quality of the mapping function used to map a node in the network to a point in the geometric target space. In this paper we study how results from mathematics as well as data mining can be applied to this mapping problem. Using a metric space model for networks and results from mathematics a relation between the number of nodes to be mapped, the worst case error of the mapping and the dimension of the geometric target space is formulated. As a main result Geometric Cluster Placement (GCP) is presented, an improved and resilient placement algorithm based on GNP. An evaluation of GCP presented is based on measurement data from the RIPE NCC Test Traffic Measurement (TTM) Project.

1. Introduction

Similarity searching in databases for protein sequences or multimedia data tends to be a computationally expensive task. Due to the complexity of the data objects every comparison performed during the search process requires a not negligible amount of computing time[10]. *Multidimensional scaling* [7] and related approaches targeting in drastically lowering the time requirements of such a search process by defining the similarity or dissimilarity of two data objects in terms of a *distance function*. With the help of this distance function, a so called *embedding* of the data objects into a *k*-dimensional geometrical space is constructed with the property that the *distortion* of the inter object distances caused by this embedding is minimal. Using this property of "distance-preserving" with a bounded error, a similarity query can now be answered by performing a *range query* in the *k*-dimensional space that is comparatively cheap.

Peer-to-Peer (P2P) networks used as a platform for file sharing are very similar to distributed multimedia databases. Therefore it is not uncommon to consider methods related to multidimensional scaling and multidimensional indexing to be used in models for next generation Peer-to-Peer networks as well. Geometry offers a rich repository to choose an organisation principle for a P2P network (i.e. a geometric topology) from. In fact, several geometric principles to organise the nodes in Overlay Networks have already been studied ([13], [9], [8]). Viewing computer networks as geometric objects offers a new perspective on a number of problems. The pre-requisite of such an approach is the dispose of a function or an algorithm for mapping nodes in a network into a geometric space of a fixed dimension. As we will see later in an example, it is not possible to map every network into a geometric space of arbitrary dimension while expecting that the *distortion* of distances caused by the mapping can be kept small. Crucial for the performance of a P2P network based on a geometric topology (GP2P network) is to which extent the used mapping function preserves the network distances between any two hosts in the geometric representation of the network. In the case of using a geometric topology without taking the assignment of valid coordinates to the nodes into consideration, it is possible that two nodes are neighbours in the Geometric representation of the P2P network while their distance in the underlying network is high. Such a situation has to be avoided since leading to high communication costs.

In the area of computer networks the Global Network Positioning System (GNP) [11] is a popular approach towards a usable and general propose mapping function. Standard GNP has a low measurement complexity by requiring k + 1 measurements per node, where k denotes the dimension of the geometric target space. As a drawback, the error caused by GNP in the sense of distance preserving can be comparatively high as we will see later.

In the area of database applications (e.g. data mining) as well as mathematics several mapping functions and algorithms have been developed with the aim to have concrete bounds for the amount of distance *distortion* caused [4], [6], [10]. In this paper we study how parts of this results can be applied to the mapping of nodes in a network into a k-dimensional geometric target space. As the main contribution a Geometric Cluster Placement Scheme (GCP) is presented. Based on the GNP idea, GCP is using a two phase dynamic landmark selection principle in combination with a cluster based measurement approach, which is a novel principle to the best of the authors knowledge. A comparison of GNP and GCP is provided in the form of simulation results.

The remainder of the paper is organised as follows: In chapter two we provide a theoretical background for the mapping problem by collecting results from mathematics and data mining. It is further motivated under what assumptions it is possible to formulate a relation between the number of nodes in a GP2P network, the distortion of inter-node distances caused by the mapping function, as well as the dimension of the geometric target space. Chapter three is discussing basic aspects of landmark selection. In chapter four the principle of the Geometric Cluster Placement scheme is motivated and evaluated by experiments using measurement data from the RIPE NCC Test Traffic Measurement (TTM) Project. Standard GNP is used as the reference mapping function in the experiments. A conclusion is provided in chapter five followed by acknowledgements.

2. Theoretical background

During this paper the model for a computer network is a weighted undirected graph G = G(V, E) with positive edge weights, where V denotes the set of vertices in G and E the set of edges between the vertices. The distance between two vertices $a, b \in V$ is measured using the *shortest-path* metric $d_s(a, b) :=$ "length of the shortest path between a and b".

The two simplifications made by this model, namely that G is *undirected* and that d_s is defined as the *shortest-path*

metric allows us to identify the resulting pair (G, d_s) with a *Metric Space*[10] that is defined as:

Definition 1 (Metric Space) Let X be a nonempty set, the pair (X, d), where d is a mapping $d : X \times X \mapsto R_0^+$ is called a Metric Space if and only if:

$$d(x,y) = 0 \quad \Leftrightarrow \quad x = y \tag{1}$$

$$d(x,y) = d(y,x) \quad \forall \quad x,y \in X$$
(2)

$$d(x,y) \le d(x,z) + d(z,y) \quad \forall \quad x,y,z \in X.$$
(3)

To be able to use the *metric space* model for a computer network, one has to accept two requirements. First, asymmetric routing is not taken into account to accomplish requirement (2). Secondly, it has to be assumed that exclusively shortest-path routing is used in the network to exploit the fact that shortest-path routing in a network establishes the triangle inequality which is (3).

In the remainder of this document, unless otherwise stated, (G, d_s) will denote a computer network represented by a finite *metric space* equipped with a shortest-path metric d_s . The function $F : G \mapsto R^k$ represents the mapping function (e.g. GNP), that maps a node x in the network to an element $F(x) \in (R^k, d')$ in the geometric target space. The notation (R^k, d') is used to specify that d' is the geometric distance measure of the target space R^k . To achieve a maximum of freedom regarding this geometric distance function (i.e. *metric*), it is introduced in a general form as l^p , defined as:

$$l^{p}(x,y) := \left(\sum_{i=1}^{k} |x_{i} - y_{i}|^{p}\right)^{1/p}, 1 \le p < \infty$$
$$l^{\infty}(x,y) = \max_{i=1,\dots,k} |x_{i} - y_{i}|$$

In the euclidean case we have $d' = l^2$.

2.1. Mapping functions and the curse of dimensionality

A special group of mapping functions studied in the area of mathematics are the so called *isometries*. An *isometry* is characterised by the fact that it preserves distances. This means that if a mapping from (G, d_s) to (R^k, d') in the form of an *isometry* is found, one has

$$d_s(x,y) = d'(F(x), F(y)), \forall x, y \in (G, d_s).$$

Unfortunately it is usually not possible to find such a mapping if one poses constrains to the measurement complexity, the dimension k as well as the metric of the target space. For example there is no distance preserving mapping of the three dimensional simplex in figure 1 to (R^2, l^1) . More flexibility regarding the metric and dimension of the target space is gained by accepting a degree of distortion [10] of the inter-object distances caused by the mapping.





Definition 2 (Distortion) Let $c_1, c_2 \ge 0$ be two real numbers. The distortion of a mapping $F : (G, d_s) \mapsto (R^k, d')$ is defined as the product $c_1 \cdot c_2$ of the lowest values c_1, c_2 for which the inequality

$$\frac{1}{c_1} \cdot d_s(x,y) \le d'(F(x),F(y)) \le c_2 \cdot d_s(x,y)$$

holds for all vertices x, y in G.

The constants c_1 and c_2 denote the *contraction* respectively the *extension* caused by the mapping regarding the original distances between the nodes in the network and the distances after the mapping process in the geometric target space.

There are two fundamental results from mathematics that provide theoretical bounds for the distortion of a mapping. The Theorem of *Bourgain* [2] provides a relation between the number of elements of a metric space, the distortion of the mapping and the dimension of the target space. We cite the version of this Theorem provided in [10].

Theorem 1 (Bourgain) Every *n*-point metric space (X, d) can be embedded in an $O(\log n)$ -dimensional Euclidean space with an $O(\log n)$ distortion.

The statement of Theorem 1 provides an answer to a general problem in the area of landmark based distance estimation schemes, that is for example formulated in [12] as: "Could we assume that a vector space with 5 to 7 dimensions can model any network performance metric ?". Accepting the simplifications necessary for the identification of a computer network with a metric space, the Theorem states that there is a functional relation between the dimension of the geometric target space, the worst case error of the mapping function and the number of nodes to be placed. While being a quantitative result, in the case of a network with n nodes to be mapped, the resulting dimension of the target space has to be selected in an order of $\log(n)$ expecting a worst case distortion of the inter-node distances caused by the mapping also in the order $\log(n)$. For example a Graph with 100.000 vertices can be embedded into $R^{c_1 \cdot 11}$ with a distortion of $c_2 \cdot 11$ where $c_1, c_2 > 0$ are constants.

Using a geometric topology for a Peer-2-Peer network, a high dimensional geometric space can have an impact to the number of neighbours each member of the GP2P network has to serve. Therefore low dimensional geometric target spaces can be considered as a benefit. In this context the question to which extent a further dimension reduction of the target space has an impact to the distortion of the mapping function is of interest. For the purpose of such a dimension reduction (i.e. a mapping from a high dimensional geometric space to a lower dimensional geometric space) the following Theorem, from *Johnson-Lindenstrauss*¹, provides an answer. We cite the version of the Theorem provided in [3].

Theorem 2 (Johnson-Lindenstrauss) For any $(0 < \epsilon < 1)$ and any integer n, let k be a positive integer that

$$k \ge 4(\epsilon^2/2 - \epsilon^3/3)^{-1}\log n.$$

Then for any set V of points in \mathbb{R}^d , there is a mapping $F: \mathbb{R}^d \mapsto \mathbb{R}^k$ such that $\forall u, v \in V$,

$$(1-\epsilon)||u-v||^2 \le ||F(u)-F(v)||^2 \le (1+\epsilon)||u-v||^2.$$

Further this map can be found in randomized polynomial time.

In the case $\epsilon \to 1$ we get from Theorem 2 as a lower bound for the dimension of the target space, $k \ge 24 \cdot \log n$. A further result presented in [3] states that any weighted graph can be embedded into a $k \le C \log n$ dimensional geometric space with a distortion of $O(n^{2/k}(\log n)^{3/2}/\sqrt{k})$.

2.2. Evaluation Framework

Based on the theoretical results collected in the previous section one can expect to realise a GP2P network for *n* nodes in a space of dimension $k \leq O(\log n)$ with an distortion of the inter-node distances in the order of $O(n^{2/k}(\log n)^{3/2}/\sqrt{k})$ [3]. Targeting in GP2P networks with a high number of nodes (e.g. planet scale P2P networks), this offers two main approaches:

1. Using a high dimensional target space for the GP2P network: For a geometric topology, a high dimensional target space has in general an impact to the connectivity of the GP2P network (i.e. the size of the neighbour table each node has to maintain). A high

¹ The Theorem is also known as the Johnson-Lindenstrauss Lemma

number of node neighbours can result in a high complexity in the form of management as well as control traffic required to maintain the GP2P structure in a distributed manner. Examples for such management and control requirements are alive-messages to neighbours and communication as well as computations required to set up the GP2P structure (e.g. in the case of triangulation based schemes [9]). For example in the case of a Content Addressable Network (CAN) [13] containing *n* nodes owning different zones of the same size, a *k*-dimensional space is leading to an average number of 2*k* individual node neighbours and an average routing path length of $k/4 \cdot (n^{\frac{1}{k}})$.

2. Optimise the mapping for a low dimensional target space: As stated before, keeping the dimension of the target space low is leading in general to a computationally less expensive construction of the geometric GP2P topology. As a drawback we can expect to find no high quality mapping function in the case of a high number of nodes in the GP2P network, as motivated before.

For the purpose of reducing measurement traffic as well as overhead for setup and maintainment of a GP2P network, the optimisation of the mapping function for a geometric space of dimension two is central for us in this paper.

As a prerequisite for being able to optimise a mapping function regarding their quality, a suitable quality measure is needed. Based on general considerations and [5] we are using the following quality measures:

1. Scalability

Measurement complexity below $O(n^2)$ (where *n* is the number of nodes in the GP2P network) as well as a target space R^k with a low *k*.

2. Stress [5]

Definition 3 (Stress) *The* **stress** *caused by a mapping F is defined as*

$$\frac{\sum_{x,y} \left(d^{'}(F(x),F(y))-d(x,y)\right)^{2}}{\sum_{x,y} d(x,y)^{2}}$$

Stress is a measure for the overall deviation in distances caused by a mapping.

3. Average error

If G(V, E) is a graph, we define for two nodes $x, y \in V$ the relative error caused by the mapping F as $\frac{|(d_s(x,y)-d^{'}(F(x),F(y))|}{d_s(x,y)}$. The average error is defined as the average of the relative errors for all nodes $x, y \in V$ being members of the GP2P network.

Stress is used since it is a common quality measure e.g. in the context of *Multidimensional Scaling*. Common variants of the *Multidimensional Scaling* approach are seeking to minimise stress as defined in definition 3. A further frequently used quality measure is average error [12]. The average error values are also provided to illustrate the values of stress.

3. Landmark selection

For the realisation of a mapping function, GNP uses a landmark based measurement principle. If the dimension of the geometric target space is k, GNP requires a minimum of k+1 landmark computers (i.e. measurement points). First the coordinates of the landmarks $L = \{l_1, l_2, ..., l_{(k+1)}\}$ are calculated by measuring the inter landmark distances. If pathological cases are avoided (e.g. the coordinates of all landmarks lie on a line) the mapping of the landmarks into the geometric target space is possible without an error. If now the coordinates of a node x in the network have to be calculated, the network distance of this node to the k + 1 landmarks is measured by using for example round trip time measurement. This measurement results are now used to calculate the node coordinates by minimising an error function utilising for example the *simplex-downhill* algorithm. In our experiments we used the error function defined in [11]:

$$f_e(x, F(x), L, k) := \sum_{i=1}^{k+1} (d'(F(x), F(l_i)) - d_s(x, l_i))^2$$

While in [11] the landmarks are chosen static, we performed experiments using a dynamic landmark principle, with the aim to realise a more flexible as well as resilient placement scheme.

In general the way to select the landmarks used for the measurements has an impact to the quality of the resulting mapping. One of the most critical situations in this context is described in figure 2. If for two nodes X, Y the distances to the landmarks out of the set $L = \{l_1, l_2, l_3\}$ are equal (i.e. $d_s(X, l_i) = d_s(Y, l_i) = w_i$, i = 1, 2, 3), both points will get the same coordinate in the target space, while having a network distance unequal to zero.

One possibility to avoid such a situation is illustrated in figure 3 and one of the main arguments in the proof of Theorem 1. If the distance $D_s(x, L)$ of a point x to the landmark set L is defined as $\min_{l \in L} d_s(x, l)$, the situation of figure 3 is the following: Since $D_s(x, L) < r_1 < r_2$ and $D_s(y, L) > r_2$, it is to be expected that the GNP mapping function will calculate coordinates F(x) and F(y) with

$$l^{2}(F(x), F(y)) > r_{2}^{'} - r_{1}^{'} > 0.$$



We are using this argument heuristically, by assuming $r_2^{'} > r_1^{'} > 0$, depending on the *distortion* introduced by GNP.

4. Experiments

For the experiments performed in this paper, we have used delay data provided by RIPE NCC TTM [1] (Test TrafficMeasurement) project. The TTM project has a number of goals including Internet end-to-end one-way delay analysis. The TTM infrastructure consists of approximately 60 measurement boxes scattered over Europe (and a few in the US and Asia). Between each pair of measurement boxes, IP packets of a fixed length (100 bytes), called probe-packets, are continuously transmitted with interarrival times of about 40 seconds, resulting in a total of about 2160 probe-packets per day. Due to the synchronizations with GPS in all testboxes, RIPE TTM achieves a delay accuracy within 10 μ s. The end-to-end delay (D) of the IP packet is determined as the difference between the timestamps of depar-

ture at the source and arrival at the destination. We have analyzed the data collected by TTM on May 15, 2003, at which time there were 58 active boxes, where 48 hosts are located in Europe, 7 in the US, and 1 in Japan, Australia and New Zealand. In order to know the congestion-free delay, for each sender-destination pair we computed the minimum end-to-end delay over 24 hours (that is approximately 2160 probe-packets). In this paper we ignore asymmetry, by always considering the symmetrised network distance, defined as the sum of the network distances in both directions. The result of the delay can be considered as a roundtrip time. We omitted pairs for which the delay in one of the directions is missing, leaving in total 1503 pairs and 1024 within Europe. These measurements provide us with the network distance matrix used as the basis for the experiments.

4.1. Impact of the selection of landmarks to Stress and Average Error

Using a simulation environment, we compared different approaches for selecting the GNP landmarks out of a set of nodes with the standard GNP approach. In the simulation the first three nodes joining a GP2P network are used for defining the two dimensional target space R^2 in the following way:

The first node l_1 is mapped to the origin (0,0) of \mathbb{R}^k . The second node l_2 is mapped to the point $(d_s(l_1, l_2), 0)$. To calculate the coordinate $\vec{l_3}$ of the third node l_3 , we measure $d_s(l_1, l_3)$, $d_s(l_2, l_3)$ and define $\vec{l_3}$ as the result of minimising the error function $f_e(l_3, (0,0), \{l_1, l_2\}, 2)$ (c.f. Section 3) using the simplex-downhill algorithm. After the target space is established, the following join procedure for a new node was simulated:

A node willing to join, has to contact an arbitrary member of the GP2P network (i.e. a bootstrap node). This bootstrap node initiates a search query for the landmark nodes about to be used for the measurement process. For this search function the following schemes are evaluated:

- 1. **Furthest Node Placement in Network (FNP)** The GP2P members with the highest network distance to the node that is about to be placed are used as landmarks.
- 2. Closest Nodes Placement in Network (CNP) The GP2P members with the lowest network distance to the node to be placed are used as landmarks.
- 3. Random Node Placement (RNP)

Random GP2P members are used for the measurements.

4. Random Cluster Placement (RCP)

A static set L of landmarks as in standard GNP is used. If the landmarks receive a measurement request, every landmark l_i starts to collect randomly a set A_i of assist nodes from the current GP2P network members. The set A_i containing also l_i , is referred to as the measurement cluster of landmark l_i or just the measurement cluster A_i . The node x willing to join is now measured by calculating

$$d(x, A_i) := \min_{a \in A_i} d(x, a)$$

for every measurement cluster A_i . During the simulation, we used a static cluster size of five nodes per cluster.

For every landmark selection scheme, 50 measurements are performed selecting randomly 50 nodes out of the test network for being mapped to R^2 . The test network was modelled by a fully meshed graph G = G(V, E), generated using the distance matrix based on the RIPE NCC TTM data. Since G is fully meshed, the routing used in the simulation is reflecting the internet routing by selecting always the path of topological length one (and not the shortest path), from any source to any destination. The vertices $v \in V$ of the graph are the RIPE NCC TTM measurement boxes.

4.2. Evaluation

As the result of the simulation we compare the stress and average error of the used schemes with standard GNP in two dimensions. In table 1 the average values of *average error* in percent and *stress* are shown. Table 2 shows the variance of the *average error* and *stress* for the 50 simulation runs. To enhance the legibility of the results *bezier curves* are used for a visualisation of the characteristics of the compared landmark selection schemes in figure 4.

	Average error	Stress	
CNP	307.40503	0.54648435	
RCP	358.80453	0.6869309	
GNP	642.54254	1.1906508	
RNP	895.65173	1.3859375	
FNP	2251.3467	5.356254	

Table 1. Average values

Confirming the observations of section 3, the CNP approach provides the best results regarding average error and stress, followed by the cluster based placement scheme RCP. In general GCP can be considered as not practical for large scale GP2P networks, because of the required initial n measurements for a node joining a GP2P network containing n nodes. In contrast, the RCP scheme requires c*(k+1)

	Average error	Stress	
CNP	28382.758	0.037409663	
RCP	26124.48	0.45122776	
GNP	122362.1	0.11339183	
RNP	195577.45	0.05818931	
FNP	2011521.7	20.806479	

Table 2. Variance

measurements, only depending on the dimension of the target space and the cluster size. The computational complexity of RCP regarding the minimisation of the error function f_e is equal to that of GNP, because just the k + 1 minimal measurement results are used in the error minimisation process.

4.3. The Geometric Cluster Placement (GCP) framework

Utilising the results above, the aim of GCP is to find landmarks with a low network distance to a joining node x while having a constant measurement complexity. To accomplish this task GCP performs a two step mapping, targeting in the usage of geometric distance information already available in the GP2P network.

In the first step, the coordinates of the node x about to join the GP2P network are calculated using RNP (i.e. k + 1random nodes from the GP2P network as landmarks). With the help of this first coordinate, the node joins the GP2P network. After the join procedure, the node calculates the set of its k + 1 geometric closest neighbours, that are usually direct neighbours or reachable via these direct neighbours. Each l_i of this new landmark nodes calculates now a cluster G_i out of its c (i.e. c denotes the cluster size) geometric closest neighbours with the property that

$$\bigcap_{i=1}^{k+1} G_i = \emptyset. \tag{4}$$

The node is now measured again and replaced with the help of the k + 1 measurement results

$$d_s(x,G_i) := \min_{g \in G_i} d_s(x,g).$$

The cluster based measurement principle in the second mapping step is used, since a placement selecting random landmarks produces a relative high error compared to standard GNP (c.f. figure 4). By using the minimum result of the cluster measurements an enhancement of the accuracy of the second mapping step is expected.

For a comparison of GNP and GCP the same simulation environment was used. For the realisation of requirement 4 a simple first come first served principle was simulated (i.e. the first landmark that requests a node to be in G_i



Figure 4. Average error and stress comparison

gets that node). An analysis of different message and time efficient clustering algorithms for GCP is beyond the scope of this paper. In a real world implementation for example an expanding ring search can be used to build the measurement clusters. The maximum Time To Live (TTL) value for the ring search has to be selected in a way that the time complexity of GCP meets the application requirements. As the result of the simulation, we compare the stress and average error of GCP using a cluster size of five with standard GNP in two dimensions. In table 3 the average values of average error in percent and stress are shown. Table 4 shows the variance of the average error and stress for the 50 simulation runs. The values of the x-axis in figure 5 represent the number of the simulation run, while the y-axis is used for displaying the stress respectively average error measured in the corresponding simulation run. Table 4 and figure 5 show that the variance in the average error and the stress using GCP is definitely lower then using standard GNP. Regarding the average values for stress and average error, table 3 shows that the values of GCP in two dimensions are sig-



Figure 5. GNP, GCP average error and stress comparison

	Average error	Stress
GCP	356.48734	0.7253786
GNP	642.54254	1.1906508

Table 3. Average Values

nificant better then the values of standard GNP in two dimensions. A further benefit of GCP is that if it is used in the context of GP2P networks, GCP does not rely on static landmarks. As a possible drawback, there is a cooperation of the nodes in the GP2P network required for performing the cluster measurements.

GCP requires (k+1)*(c+1) measurements using a cluster size of c after the first c*(k+1) nodes have joined the GP2P network. While the selection of target dimension two for the evaluation makes management and evaluation intuitive, the stress and error values of GCP are comparatively

	Average error	Stress
GCP	39079.03	0.044214338
GNP	122362.1	0.11339183

Fable	4.	Var	rian	ice

high. In the case of a file sharing application, without strict realtime requirements a GP2P network in a two dimensional geometric space build up using GCP can be considered as an option. If the target application of the GP2P network has more strict realtime requirements (e.g. in the case of multimedia streaming) a higher dimensional target space is recommended. For example in the case of a three dimensional target space using the same simulation parameters, the average error of GCP reduces to 144,5 % with a stress value of 0.15160233. As motivated in section 2.2, in general a low dimensional target space results in lower measurement and management complexity with the drawback of higher placement errors, while a high dimensional target space results in lower placement errors but demands a higher measurement and management complexity.

5. Conclusions

In this paper we studied how results from mathematics as well as data mining can be applied to the problem of mapping nodes in a computer network into a low dimensional geometric space. A theoretical background for the mapping problem as well as the key points of an evaluation framework for the quality of a mapping function are provided. Accepting the heuristics required to identify a computer network with a metric space it is possible to provide a relation between the dimension of the geometric target space, the number of nodes to be mapped and the worst case error to be expected. This relation, provided by the Theorem of Bourgain[2], answers the question of "the curse of dimensionality" in the context of mapping functions for example posed in [12]. Different landmark selection schemes for being used with GNP are evaluated. The best results regarding stress and average error are obtained by using the Closest Nodes Placement (CNP) principle. As a practical mapping function a novel Geometric Cluster Placement (GCP) scheme is presented and compared to GNP. As a result it is stated, that the variation in the stress as well as the average error using GCP is lower then using the standard GNP approach. In GCP the requirement of a static landmark system is relaxed by using a more resilient landmark selection algorithm that relies on distance information already available in a Peer-to-Peer network based on a geometric topology.

6. Acknowledgments

The authors would like to thank the anonymous reviewers for their insightful comments which helped to improve the paper. Further thanks to Mikhail Smirnov and Christoph Reichert for many suggestions, comments and discussions.

References

- [1] Test traffic measurement project, http://www.ripe.net/ttm.
- [2] J. Bourgain. On lipschitz embedding of finite metric spaces in hilbert space. *Israel J. Math.*, 52:46–52, 1985.
- [3] S. Dasgupta and A. Gupta. An elementary proof of the johnson-lindenstrauss lemma. Technical Report TR-99-006, Berkeley, CA, 1999.
- [4] C. Faloutsos and K.-I. Lin. Fastmap: a fast algorithm for indexing, data-mining and visualization of traditional and multimedia datasets. In *Proceedings of ACM SIGMOD*, 1995.
- [5] G. Hjaltason and H. Samet. Contractive embedding methods for similarity searching in metric spaces. Technical Report TR-4102, Computer Science Department, University of Maryland, College Park, MD, 2 2000.
- [6] G. Hristescu and M. Farach-Colton. Cluster-preserving embedding of proteins. Technical report, Rutgers Univ., Piscataway, New Jersey, 1999.
- [7] J. B. Kruskal and M. Wish. Multidimensional scaling. Technical report, Sage University Series, Beverly Hills, 1978.
- [8] J. Liebeherr and T. K. Beam. Hypercast: A protocol for maintaining multicast group members in a logical hypercube topology. In *Networked Group Communication*, pages 72– 89, 1999.
- [9] J. Liebeherr, M. Nahas, and W. Si. Application-layer multicasting with delaunay triangulation overlays. Technical Report CS-2001-26, University of Virginia, 5 2001.
- [10] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15:215–245, 1995.
- [11] E. Ng and H. Zhang. Predicting internet network distance with coordiantes-based approaches. In *Proceedings of IEEE INFOCOM*, 6 2002.
- [12] M. Pias, J. Crowcroft, S. Wilbur, T. Harris, and S. Bhatti. Lighthouses for scalable distributed location. In *Proceedings of IPTPS*, 2003.
- [13] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A scalable content addressable network. In *Proceedings of ACM SIGCOMM*, 2001.