Space lower bounds for low-stretch greedy embeddings*

Ioannis Caragiannis and Christos Kalaitzis

Computer Technology Institute and Press "Diophantus" &
Department of Computer Engineering and Informatics,
University of Patras, 26504 Rio, Greece.

Abstract. Greedy (geometric) routing is an important paradigm for routing in communication networks. It uses an embedding of the nodes of the network into points of a metric space (e.g., $\mathbb{R}^d$) equipped with a distance function (e.g., the Euclidean distance ℓ_2) and uses as address of each node the coordinates of the corresponding point. The embedding has particular properties so that the routing decision for a packet is taken greedily based only on the addresses of the current node, its neighbors, and the destination node and the distances of the corresponding points. In this way, there is no need to keep routing tables at the nodes. Embeddings that allow for this functionality are called greedy embeddings. Even though greedy embeddings do exist for several metric spaces and distance functions, they usually result in paths of high stretch, i.e., significantly longer than the corresponding shortest paths.
In this paper, we show that greedy embeddings in low-dimensional Euclidean spaces necessarily have high stretch. In particular, greedy embeddings of n-node graphs with optimal stretch requires at least $\Omega(n)$ dimensions for distance ℓ_2. This result disproves a conjecture by Maymounkov (2006) stating that greedy embeddings of optimal stretch are possible in Euclidean spaces with polylogarithmic number of dimensions. Furthermore, we present trade-offs between the stretch and the number of dimensions of the host Euclidean space. Our results imply that every greedy embedding into a Euclidean space with polylogarithmic number of dimensions (and Euclidean distance) has stretch $\Omega\left(\frac{\log n}{\log\log n}\right)$. We extend this result for a distance function used by an $O(\log n)$-stretch greedy embedding presented by Flury, Pemmaraju, and Wattenhofer (2009). Our lower bound implies that their embedding has almost best possible stretch.

1 Introduction

Greedy routing utilizes a particular assignment of node addresses so that routing of packets can be performed using only the address of the current node of a traveling packet, the addresses of its neighbors, and the address of the destination node. The node addresses are usually defined using a *greedy embedding*. Formally,

* This work was partially supported by the EC-funded STREP project EULER.

a greedy embedding of a graph $G = (V, E)$ into a host metric space (X, dist) is a function $f : V \to X$ so that the following property holds: for any two nodes u, t of G, there exists a node v in the neighborhood $\Gamma(u)$ of u in G so that $\mathrm{dist}(f(v), f(t)) < \mathrm{dist}(f(u), f(t))$. A typical example of a host metric space is the d-dimensional Euclidean space $\mathbb{R}^d$ equipped with the Euclidean distance ℓ_2. Given a greedy embedding f, the coordinates of point $f(u)$ can be used as the address of node u. Then, when node u has to take a decision about the next hop for a packet with destination address $f(t)$, it has to select among its neighbors a node with address $f(v)$ such that $\mathrm{dist}(f(v), f(t)) < \mathrm{dist}(f(u), f(t))$. It is clear that, in this way, the packet is guaranteed to reach its destination within a finite number of steps.

Greedy embeddings were first defined by Papadimitriou and Ratajczak [15]. They proved that any 3-connected planar graph can be greedily embedded into the 3-dimensional Euclidean space using a non-Euclidean distance function. They also conjectured that every such graph can be greedily embedded in the Euclidean space with Euclidean distance. The conjecture was proved by Moitra and Leighton [14]. Kleinberg [10] showed that any tree can be greedily embedded in the 2-dimensional hyperbolic space. This immediately yields a greedy embedding for any graph (by just embedding a spanning tree). Epstein and Woodrich [6] observed that the coordinates of the nodes in Kleinberg's embedding require too much space and modified it so that each coordinate is represented with $O(\log n)$ bits (where n is the size of the graph). Greedy embeddings into $O(\log n)$-dimensional Euclidean spaces (with ℓ_∞ distance) are also known [10] and exploit an isometric embedding of trees into Euclidean spaces due to Linial et al. [11].

Note that the approach of computing the greedy embedding of a spanning tree ignores several links of the network. Hence, it may be the case that even though a packet could potentially reach its destination with a few hops using a shortest path, it is greedily routed through a path that has to travel across a constant fraction of the nodes of the whole network. The measure that can quantify this inefficiency is the stretch of a greedy routing algorithm, i.e., the ratio between the length of the path used by the algorithm over the length of the corresponding shortest path.

Let us proceed with a formal definition of the stretch of a greedy embedding.

Definition 1. *Let f be a greedy embedding of a graph G into a metric space $(X, dist)$. A path $\langle u_0, u_1, ..., u_t \rangle$ is called a greedy path if $u_{i+1} \in argmin_{v \in \Gamma(u_i)}\{dist(f(v), f(u_t))\}$, where $\Gamma(u_i)$ denotes the neighborhood of u_i in G. We say that f has stretch ρ for graph G if, for every pair of nodes u, v of G, the length of every greedy path from u to v is at most ρ times the length of the shortest path from u to v in G. The stretch of f is simply the maximum stretch over all graphs.*

We use the terms no-stretch and optimal stretch to refer to embeddings with stretch equal to 1.

Maymounkov [13] considers the question of whether no-stretch greedy embeddings into low-dimensional spaces exist. Among other results, he presents a lower bound of $\Omega(\log n)$ on the dimension of the host hyperbolic space for greedy embeddings with optimal stretch. Furthermore, he conjectures that any graph can be embedded into Euclidean or hyperbolic spaces with a polylogarithmic number of dimensions with no stretch. We remark that a proof of this conjecture would probably justify greedy routing as a compelling alternative to compact routing [16].

Flury et al. [8] present a greedy embedding of any n-node graph into an $O(\log^2 n)$-dimensional Euclidean space that has stretch $O(\log n)$. Each coordinate in their embedding uses $O(\log n)$ bits. They used the min-max$_c$ distance function which views the d-dimensional Euclidean space as composed by d/c c-dimensional spaces and, for a pair of points x, y, takes the ℓ_∞ norm of the projections of x and y into those spaces, and finally takes the minimum of those ℓ_∞ distances as the min-max$_c$ distance between them. Their embedding uses an algorithm of Awerbuch and Peleg [4] to compute a tree cover of the graph and the algorithm of Linial et al. [11] to embed each tree in the cover isometrically in a low-dimensional Euclidean space. In practice, greedy embeddings have been proved useful for sensor [8] and internet-like networks [5].

In this paper, we present lower bounds on the number of dimensions required for low-stretch greedy embeddings into Euclidean spaces. We first disprove Maymounkov's conjecture by showing that greedy embeddings into $(\mathbb{R}^d, \ell_2)$ have optimal stretch only if the number of dimensions d is linear in n. The proof uses an extension of the hard crossroads construction in [13] and exploits properties of random sign pattern matrices. Namely, we make use of a linear lower bound due to Alon et al. [3] on the minimum rank of random $N \times N$ sign pattern matrices. We also obtain an $\Omega(\sqrt{n})$ lower bound through an explicit construction that uses Hadamard matrices. These results are stated in Theorem 2.

Furthermore, we present trade-offs between the stretch of greedy embeddings into $\mathbb{R}^d$ and the number of dimensions d for different distance functions. Namely, for every integer parameter $k \geq 3$, we show that greedy embeddings into $\mathbb{R}^d$ with ℓ_p distance have stretch smaller than $\frac{k+1}{3}$ only if $d \in O\left(\frac{n^{1/k}}{\log p}\right)$ (Theorem 5). This implies that the best stretch we can expect with a polylogarithmic number of dimensions is $\Omega\left(\frac{\log n}{\log\log n}\right)$. Our arguments use a result of Erdös and Sachs [7] on the density of graphs of high girth and a result of Warren [17] that upper-bounds the number of different sign patterns of a set of polynomials. In particular, starting from a dense graph with high girth, we construct a family of graphs and show that, if d is not sufficiently large, any embedding f fails to achieve low stretch for some graph in this family.

We extend our lower bound arguments to greedy embeddings into $\mathbb{R}^d$ that use the min-max$_c$ distance function that has been used in [8]. Note that the lower bound does not depend on any other characteristic of the embedding of [8] and applies to every embedding in $(\mathbb{R}^d, \text{min-max}_c)$. We show that the best stretch we can hope for with $d \in \text{polylog}(n)$ is $\Omega\left(\frac{\log n}{\log\log n}\right)$ (Theorem 7). This lower bound indicates that the embedding of Flury et al. [8] is almost optimal among all greedy embeddings in $(\mathbb{R}^d, \text{min-max}_c)$. Furthermore, our proof applies to a larger class of distance functions including ℓ_∞.

We remark that our proofs are not information-theoretic in the sense that we do not prove lower bounds on the number of bits required in order to achieve low stretch embeddings. Instead, we prove that the particular characteristics of the host metric space (e.g., small number of dimensions) and the distance function do not allow for low stretch greedy embeddings even if we allow for arbitrarily many bits to store the coordinates.

We also remark that there is an extensive literature on low-distortion embeddings where the objective is to embed a metric space (e.g., a graph with the shortest path distance function) into another metric space so that the distances between pairs of points are distorted as less as possible (see [12] for an introduction to the subject and a coverage of early related results). Greedy embeddings are inherently different than low-distortion ones. Ordinal embeddings (see [2]) where one aims to maintain the relative order of intra-point distances are conceptually closer to our work. Actually, the use of Warren's theorem has been inspired by [2].

The rest of the paper is structured as follows. We present the extension of the hard crossroads construction in Section 2. In Section 3, we present the trade-off for greedy embeddings into Euclidean spaces using the distance function ℓ_p. The optimality of the embedding of Flury et al. [8] is proved in Section 4. We conclude in Section 5.

2 Hard crossroads

In this section we extend the hard crossroads construction of Maymounkov [13] by exploiting sign pattern matrices. We begin with some necessary definitions. Throughout the text, we use $[N]$ to denote the set $\{1, 2, ..., N\}$. We also make use of the signum function $\text{sgn} : \mathbb{R} \setminus \{0\} \to \{-1, 1\}$.

Definition 2. *A square $N \times N$ matrix S is called a sign-pattern matrix if $S_{i,j} \in \{-1, 1\}$ for every $i, j \in [N]$. The minimum rank of a sign-pattern matrix S, denoted by $mr(S)$, is the minimum rank among all $N \times N$ matrices M with non-zero entries for which it holds that $sgn(M_{i,j}) = S_{i,j}$ for every $i, j \in [N]$.*

In our construction, we use sign-pattern matrices of high minimum rank. Such matrices do exist as the following result of Alon et al. [3] indicates.

Theorem 1 (Alon, Frankl, and Rödl [3]). *For every integer $N \geq 1$, there exist $N \times N$ sign-pattern matrices of minimum rank at least $N/32$.*

The proof of this statement uses the probabilistic method. As such, the result is existential and does not provide an explicit construction (for many different values of N). A slightly weaker lower bound on the minimum rank is obtained by the so-called Hadamard matrices.

Definition 3. *A Hadamard matrix H_N with N nodes (where N is a power of 2) is defined as $H_1 = [1]$ and*

$$H_N = \begin{bmatrix} H_{N/2} & H_{N/2} \\ H_{N/2} & -H_{N/2} \end{bmatrix}.$$

It is known (see [9]) that the sign-pattern matrix H_N has $\mathrm{mr}(H_N) \geq \sqrt{N}$.

We are ready to present our hard crossroads construction. We will use an $N \times N$ sign pattern matrix S. Given S, we construct the hard crossroads graph $G(S)$ as follows. G has three levels of nodes:

– There are N nodes at level 0: u_1, u_2, ..., u_N.
– There are $2N$ nodes at level 1: nodes u_i^+ and u_i^- for each node u_i of level 0.
– There are also N nodes v_1, ..., v_N at level 2.

The set of edges of $G(S)$ is defined as follows. Each node u_i is connected to both nodes u_i^+ and u_i^- at level 1. For each pair of integers $1 \leq i, j \leq N$, there is an edge between u_i^+ and v_j if $S_{ij} = 1$ and an edge between u_i^- and v_j otherwise. See Figure 1 for an example that uses the sign pattern matrix

$$S = \begin{bmatrix} 1 & -1 & -1 \\ -1 & 1 & 1 \\ 1 & -1 & 1 \end{bmatrix}.$$

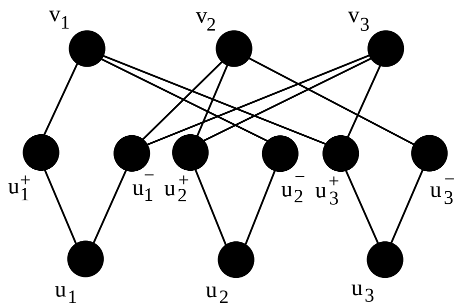

Fig. 1. An example of a hard crossroad.

We will show that any greedy embedding f of $G(S)$ into $\mathbb{R}^d$ with distance function ℓ_2 has stretch smaller than 2 only if d is large. Observe that, for every pair of integers i, j, the unique shortest path (of length 2) from node u_i to node v_j goes either through node u_i^+ (if $S_{ij} = 1$) or node u_i^- (if $S_{ij} = -1$); any other path connecting these two nodes has length at least 4. Since f has stretch smaller than 2, this means that the only greedy path from node u_i to node v_j defined by f is their shortest path. Hence, f should satisfy that $\mathrm{dist}(f(u_i^+), f(v_j)) <$

$\mathrm{dist}(f(u_i^-), f(v_j))$ if $S_{i,j} = 1$ and $\mathrm{dist}(f(u_i^+), f(v_j)) > \mathrm{dist}(f(u_i^-), f(v_j))$ otherwise. In other words, depending on whether $S_{i,j} = 1$ or $S_{i,j} = -1$, the embedding should guarantee that the point $f(v_j)$ lies at the "left" or the "right" side of the hyperplane that bisects points $f(u_i^+)$ and $f(u_i^-)$.

For $i \in [N]$, denote by $U_i^+, U_i^- \in \mathbb{R}^d$ the vectors of coordinates of points $f(u_i^+)$ and $f(u_i^-)$. Also, for $j \in [N]$, denote by $V_j \in \mathbb{R}^d$ the vector of coordinates of point $f(v_j)$. Finally, denote by $(a_i, b_i) = \{x \in \mathbb{R}^d : a_i^T x = b_i\}$ the hyperplane that bisects points $f(u_i^+)$ and $f(u_i^-)$. Without loss of generality, assume that $a_i^T U_i^+ - b_i > 0$ (and $a_i^T U_i^- - b_i < 0$).

By the property of the greedy embedding f, the hyperplane (a_i, b_i) partitions the set of points corresponding to nodes of level 2 into two sets depending on whether node v_j is connected to u_i^+ (i.e., when $S_{i,j} = 1$) or u_i^- (when $S_{i,j} = -1$). We have $a_i^T V_j - b_i > 0$ for each $i, j \in [N]$ such that $S_{i,j} = 1$ and $a_i^T V_j - b_i < 0$, otherwise.

Now, denote by V the $d \times N$ matrix with columns $V_1, V_2, ..., V_N$, by A the $N \times d$ matrix with rows $a_1^T, ..., a_N^T$, and by b the vector with entries $b_1, ..., b_N$. Let $M = A \cdot V - b \cdot 1^T$. Observe that $\mathrm{sgn}(M_{i,j}) = S_{i,j}$, for every $i, j \in [N]$ and, hence, $r(M) \geq \mathrm{mr}(S)$. Using the fact that d is not smaller than the rank of matrix V and well-known facts about the rank of matrices, we obtain

$$d \geq \mathrm{rank}(V) \geq \mathrm{rank}(A \cdot V) = \mathrm{rank}(M + b \cdot 1^T) \geq \mathrm{rank}(M) - 1 \geq \mathrm{mr}(S) - 1.$$

The next statement follows using either a sign pattern matrix with minimum rank at least $N/32$ (from Theorem 1, such graph do exist) or the Hadamard matrix with N nodes (and by observing that the number of nodes in $G(S)$ is $n = 4N$).

Theorem 2. *Let f be a greedy embedding of n-node graphs into the metric space $\mathbb{R}^d$ with distance function ℓ_2. If f has stretch smaller than 2, then $d \geq n/128 - 1$. Furthermore, for every value of n that is a power of 2, there exists an explicitly constructed n-node graph which cannot be embedded into $(\mathbb{R}^d, \ell_2)$ with stretch smaller than 2 if $d < \sqrt{n}/2 - 1$.*

3 A lower bound based on high-girth graphs

Our second lower bound argument exploits graphs with high girth. The main observation behind the construction described below is that, in a graph of girth g, if the greedy path between two very close nodes (say, of original distance 2) defined by an embedding is not the corresponding shortest path, then this embedding has stretch at least $\Omega(g)$.

Given a girth-g graph $G = (V, E)$ with N nodes and m edges, we will construct a family $\mathcal{H}$ of 4^m graphs so that every greedy embedding in a low-dimensional Euclidean space with distance ℓ_p has high stretch for some member of this family.

The family $\mathcal{H}$ is constructed as follows. Each graph in this family has a supernode u of three nodes u_0, u_1, and u_2 for each node u of G. Denote by

$n = 3N$ the number of nodes of the graphs in $\mathcal{H}$. Node u_0 is connected with nodes u_1 and u_2. For each edge (u, v) of G, there are four different ways to connect the supernodes u and v: one of the nodes u_1 and u_2 is connected to one of the nodes v_1 and v_2. See Figure 2 for an example. Note that the 4^m different graphs of the family $\mathcal{H}$ are constructed by considering the different combinations on the way we connect supernodes corresponding to the endpoints of the edges in G.

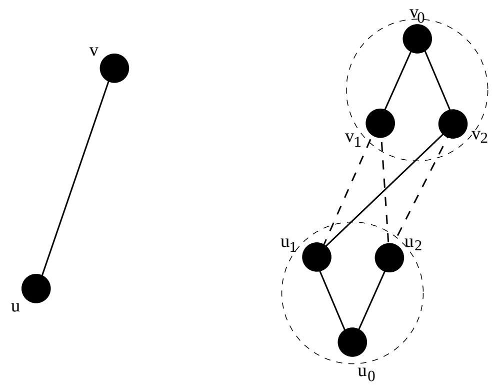

Fig. 2. The construction of graphs in $\mathcal{H}$. The figure shows two supernodes (the dotted circles) corresponding to two adjacent nodes u, v of G. There is an edge between nodes u_1 and v_2; the three dashed edges denote the alternative ways to connect the supernodes.

Consider the nodes u_0 and v_0 of the graphs in $\mathcal{H}$ that correspond to two adjacent nodes u and v of G. By our construction, in every graph of $\mathcal{H}$, the distance between u_0 and v_0 is 3. Any other path connecting these two nodes has length at least $g + 1$; to see this, observe that, since G has girth g, any path that connects some of the nodes u_1 and u_2 with some of the nodes v_1 and v_2 (in any graph of $\mathcal{H}$) has length at least $g - 1$. So, in any graph of $\mathcal{H}$, the greedy path from node u_0 to node v_0 defined by any embedding with stretch smaller than $\frac{g+1}{3}$ should be identical to the corresponding (unique) shortest path. This is the only requirement we will utilize in our proof; clearly, there are additional necessary requirements which we will simply ignore.

For an ordered pair of adjacent nodes u and v in G, consider the three nodes u_1, u_2, and v_0 of the graphs in $\mathcal{H}$. Let x, y, z be points of $\mathbb{R}^d$ corresponding to these nodes. We will view the coordinates of x, y, and z as real variables. What an embedding f has to do is simply to set the values of these variables, possibly selecting different values for different graphs of family $\mathcal{H}$. Define the quantity $\mathrm{dist}(x, z)^p - \mathrm{dist}(y, z)^p$ and observe that, due to the definition of the distance ℓ_p, this is a polynomial of degree p over the coordinates of x, y, and z. Let us call this polynomial $\mathcal{P}^{u_0 \to v_0}$. Clearly, the sign of $\mathcal{P}^{u_0 \to v_0}$ indicates the result of the comparison of $\mathrm{dist}(x,z)$ with $\mathrm{dist}(y,z)$. Let f be an embedding of stretch smaller than $\frac{g+1}{3}$ and observe that the greedy path from u_0 to v_0 defined by f is identical to the corresponding shortest path only if

– $\mathcal{P}^{u_0\to v_0}$ has sign -1 for all graphs in $\mathcal{H}$ that connect supernodes u and v through an edge with endpoint at u_1 and
– $\mathcal{P}^{u_0\to v_0}$ has sign 1 for all graphs in $\mathcal{H}$ that connect supernodes u and v through an edge with endpoint at u_2.

Denote by $\mathcal{P}$ the set of all polynomials $\mathcal{P}^{u_0\to v_0}$ for every ordered pair of adjacent nodes u and v in G and recall that the graphs of $\mathcal{H}$ are produced by selecting all possible combinations of ways to connect supernodes corresponding to adjacent nodes of G. By repeating the above argument, we obtain that f has stretch smaller than $\frac{g+1}{3}$ for every graph in $\mathcal{H}$ only if the set of polynomials $\mathcal{P}$ realizes all the 4^m different sign patterns.

Here is where we will use a result of Warren [17] which upper-bounds the number of different sign patterns of a set of polynomials.

Theorem 3 (Warren [17]). *Let $\mathcal{P}$ be a set of K polynomials of degree δ on ℓ real variables. Then, the number of different sign patterns $\mathcal{P}$ may have is at most $\left(\frac{4e\delta K}{\ell}\right)^{\ell}$.*

Observe that $\mathcal{P}$ consists of $2m$ polynomials of degree p over $3dN$ real variables. Using Theorem 3, we have that the number of different sign patterns the polynomials of $\mathcal{P}$ can realize is $\#\mathrm{sp} \le \left(\frac{8emp}{3dN}\right)^{3dN}$.

We will now assume that $d \le \frac{m}{N\log 8ep}$ in order to show that $\#\mathrm{sp} < 4^m$ and obtain a contradiction. Observe that the bound we obtained using Warren's theorem can be expressed as a function $f(d) = \left(\frac{A}{dB}\right)^{dB}$ where $A = 8emp$ and $B = 3N$. Also, observe that f is increasing in $[0, \frac{A}{eB}]$, i.e., for $d \le \frac{8mp}{3N}$. Clearly, our assumption satisfies this inequality. Hence, we have

$$\#\mathrm{sp} = \left(\frac{8emp}{3dN}\right)^{3dN} \le (8ep\log 8ep)^{\frac{m}{\log 8ep}} < (8ep)^{\frac{2m}{\log 8ep}} = 4^m.$$

We have obtained a contradiction. Hence, $d > \frac{m}{N\log 8ep}$. In order to obtain a lower bound on d only in terms of the number of nodes and p, it suffices to use a girth-g graph G that is as dense as possible. The following well-known result indicates that dense girth-g graphs do exist.

Theorem 4 (Erdös and Sachs [7]). *For every $g \ge 3$ and for infinitely many values of N, there are N-node graphs of girth at least g with at least $\frac{1}{4}N^{1+1/g}$ edges.*

We remark that the parameters of this theorem can be slightly improved. For example, the complete N-node graph has girth 3.

Hence, by selecting G to have $m \ge \frac{1}{4}N^{1+1/g}$ edges and recalling that $n = 3N$, we obtain that $d > \frac{N^{1/g}}{4\log 8ep} \ge \frac{n^{1/g}}{12\log 8ep}$. The next statement summarizes the discussion in this section.

Theorem 5. *Let $k \geq 3$ be an integer and let f be a greedy embedding of n-node graphs to the metric space $\mathbb{R}^d$ with distance function ℓ_p. If f has stretch smaller than $\frac{k+1}{3}$, then $d > \frac{n^{1/k}}{12 \log 8ep}$.*

An immediate corollary of this statement (for $k \in O\left(\frac{\log n}{\log \log n}\right)$) is that, in general, the best stretch we can achieve using Euclidean spaces with polylogarithmic number of dimensions and the distance ℓ_2 is $\Omega\left(\frac{\log n}{\log \log n}\right)$. Also, note that using the N-node complete graph (of girth 3) as G, we obtain a lower bound of $\Omega(n/\log p)$ on the number of dimensions for no-stretch greedy embeddings that use the distance ℓ_p. For $p = 2$, we essentially obtain the same bound we have obtained in Section 2.

4 A lower bound for the embedding of Flury et al.

We will now adapt the lower bound of the previous section for the greedy embeddings presented by Flury et al. [8]. The embedding of [8] uses the Euclidean space $\mathbb{R}^d$ and the distance function min-max$_c$ defined as follows. Let c be a factor of d and let $(s_1, s_2, ..., s_d)$ and $(t_1, t_2, ..., t_d)$ be points of $\mathbb{R}^d$. For each $j \in [d/c]$, let

$$D_j = \max_{(c-1)j+1 \leq i \leq cj} |s_i - t_i|.$$

Then, the function min-max$_c$ is defined as

$$\text{min-max}_c(s, t) = \min_{j \in [d/c]} D_j.$$

We remark that our arguments below can be extended (without any modification) to any distance function such that $\text{dist}(s, t)$ is computed by applying operators min and max on the quantities $|s_i - t_i|$ for $i \in [d]$. For example, the ℓ_∞ distance belongs in this category.

In the previous section, the sign of a single polynomial indicated the result of the comparison between two distances. This was possible because of the definition of the ℓ_p distance function. With distance function min-max$_c$, this is clearly impossible. However, the argument of the previous section can be adapted using the sign pattern of a set of polynomials as an indication of the result of the comparison between two distances.

Again, we use the same construction we used in Section 3. Starting from a girth-g N-node graph G with m edges, we construct the family of graphs $\mathcal{H}$. Our argument will exploit the same minimal requirement we exploited in the previous section. Namely, for every ordered pair of adjacent nodes u and v in G, in any graph of $\mathcal{H}$ the greedy path from node u_0 to node v_0 defined by any embedding with stretch smaller than $\frac{g+1}{3}$ should be identical to the corresponding shortest path.

For an ordered pair of adjacent nodes u and v in G, consider the three nodes u_1, u_2, and v_0 of the graphs in $\mathcal{H}$. Let x, y, z be points of $\mathbb{R}^d$ corresponding to these nodes. Again, we view the coordinates of x, y, and z as variables; recall that what an embedding has to do is simply to set the values of these variables. Let us now define the following set of polynomials $\mathcal{Q}^{u_0 \to v_0}$ over the coordinates of x, y, and z:

- The $d(d-1)/2$ polynomials $(x_i - z_i)^2 - (x_j - z_j)^2$ for $1 \leq i < j \leq d$.
- The $d(d-1)/2$ polynomials $(y_i - z_i)^2 - (y_j - z_j)^2$ for $1 \leq i < j \leq d$.
- The d^2 polynomials $(x_i - z_i)^2 - (y_j - z_j)^2$ for $i, j \in [d]$.

Notice that there is no way to restrict all the above polynomials to have strictly positive or strictly negative values. So, the term sign in the following corresponds to the outcome of the signum function sgn0 : $\mathbb{R} \to \{-1, 0, +1\}$. Note that the sign of each polynomial above does not change if we replace the squares by absolute values. Hence, for all assignments for which the sign pattern of $\mathcal{Q}^{u_0 \to v_0}$ is the same, the relative order of the absolute values of the difference at the same coordinate between the pairs of points x, z and y, z is identical. Now assume that two embeddings f and f' are such that $\text{min-max}_c(f(u_2), f(v_0)) > \text{min-max}_c(f(u_1), f(v_0))$ and $\text{min-max}_c(f'(u_2), f'(v_0)) < \text{min-max}_c(f'(u_1), f'(v_0))$. This means that the embeddings f and f' correspond to two assignments of values to the coordinates of x, y, and z so that the set of polynomials $\mathcal{Q}^{u_0 \to v_0}$ has two different sign patterns.

Let $\mathcal{Q}$ be the union of the $2m$ sets of polynomials $\mathcal{Q}^{u_0 \to v_0}$ and $\mathcal{Q}^{v_0 \to u_0}$ for every adjacent pair of nodes u, v in G. The above fact implies that in order to find greedy embeddings of stretch smaller than $\frac{g+1}{3}$ for every graph in $\mathcal{H}$, the set of polynomials $\mathcal{Q}$ should have at least 4^m different sign patterns. Again, we will show that this is not possible unless d is large. Since signs are defined to take values in $\{-1, 0, 1\}$ now, we cannot use Warren's theorem in order to upper-bound the number of sign patterns of $\mathcal{Q}$. We will use a slightly different version that applies to our case that is due to Alon [1].

Theorem 6 (Alon [1]). *Let $\mathcal{Q}$ be a set of K polynomials of degree δ on ℓ real variables. Then, the number of different sign patterns $\mathcal{Q}$ may have (including zeros) is at most $\left(\frac{8e\delta K}{\ell}\right)^{\ell}$.*

Observe that $\mathcal{Q}$ contains at most $2d^2m$ degree-2 polynomials in total and the number of variables is $3dN$. By applying Theorem 6, the different sign patterns the polynomials of $\mathcal{Q}$ may have are $\#\text{sp} \leq \left(\frac{64emd}{3N}\right)^{3dN}$.

We will now assume that $d \leq \frac{m}{3N \log N}$ in order to obtain a contradiction. Observe that the bound we obtained using Theorem 6 is increasing in d. Hence, we have

$$\#\text{sp} \leq \left(\frac{64emd}{3N}\right)^{3dN} \leq \left(\frac{64em^2}{9N^2 \log N}\right)^{\frac{m}{3 \log N}} < \left(\frac{16eN^2}{9 \log N}\right)^{\frac{m}{\log N}} \leq n^{\frac{2m}{\log N}} = 4^m,$$

where the second inequality follows since $m < N^2/2$ and the third one follows since N is sufficiently large (i.e., $N > 2^{16e/9}$). Again, we have obtained a contradiction. Hence, $d > \frac{m}{3N \log N}$. By setting G to have at least $\frac{1}{4}N^{1+1/g}$ edges (using Theorem 3) and since $n = 3N$, we have $d > \frac{n^{1/g}}{36 \log n}$.

The next statement summarizes the discussion in this section.

Theorem 7. *Let $k \geq 3$ be an integer and let f be a greedy embedding of n-node graphs to the metric space $\mathbb{R}^d$ with distance function* min-max$_c$. *If f has stretch smaller than $\frac{k+1}{3}$, then $d > \frac{n^{1/k}}{36 \log n}$.*

Again, this statement implies that the best stretch that can be achieved using Euclidean spaces with a polylogarithmic number of dimensions is $\Omega\left(\frac{\log n}{\log \log n}\right)$. Hence, the $O(\log n)$-stretch greedy embedding of [8] is almost best possible.

5 Concluding remarks

Note that an alternative proof of Theorem 2 could consider the 2^{N^2} different hard crossroads graphs (for all different $N \times N$ sign pattern matrices) and use Warren's theorem and similar reasoning with that we used in Section 3 in order to prove that greedy embeddings in $\mathbb{R}^d$ using distance ℓ_2 that have stretch smaller than 2 require $d \in \Omega(n)$. However, we find it interesting that the same result (and with slightly better constants) follow by adapting the original argument of Maymounkov [13]. Finally, we remark that even though we have focused on Euclidean spaces, our work has implications to embeddings in multi-dimensional hyperbolic spaces as well. We plan to discuss this issue in the final version of the paper.

Acknowledgement. We thank Tasos Sidiropoulos for helpful discussions.

References

1. N. Alon. Tools from higher algebra. In *Handbook of Combinatorics*, R. L. Graham, M. Grötschel, and L. Lovász, Eds. Vol. 2. MIT Press, Cambridge, pages 287–324, 1996.
2. N. Alon, M. Badoiu, E. D. Demaine, M. Farach-Colton, M. Hajiaghayi, and A. Sidiropoulos. Ordinal embeddings of minimum relaxation: general properties, trees, and ultrametrics. *ACM Transactions on Algorithms*, 4(4), art. 46, 2008.
3. N. Alon, P. Frankl, and V. Rödl. Geometrical realization of set systems and probabilistic communication complexity. In *Proceedings of the 26th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 277–280, 1985.
4. B. Awerbuch and D. Peleg. Sparse partitions. In *Proceedings of the 30th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 503–513, 1990.
5. M. Boguna, F. Papadopoulos, and D. Krioukov. Sustaining the internet with hyperbolic mapping. *Nature Communications*, 1, art. 62, 2010.
6. D. Eppstein and M. T. Goodrich. Succinct greedy geometric routing using hyperbolic geometry. *IEEE Transactions on Computers*, 60(11): 1571–1580, 2011.
7. P. Erdös and H. Sachs. Reguläre Graphen gegebener Taillenweite mit minimaler Knotenzahl. Wiss. Z. Martin-Luther-Univ. Halle-Wittenberg Math.-Natur. Reihe 12, 251-257, 1963.

8. R. Flury, S. V. Pemmaraju, and R. Wattenhofer. Greedy routing with bounded stretch. In *Proceedings of the 28th IEEE International Conference on Computer Communications (INFOCOM)*, pages 1737–1745, 2009.
9. J. Forster. A linear lower bound on the unbounded error probabilistic communication complexity. *Journal of Computer and System Sciences*, 65(4): 612–625, 2002.
10. R. Kleinberg. Geographic routing using hyperbolic space. In *Proceedings of the 26th IEEE International Conference on Computer Communications (INFOCOM)*, pages 1902–1909, 2007.
11. N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2): 215–245, 1995.
12. J. Matoušek. Lectures on Discrete Geometry. *Graduate Texts in Mathematics*, Springer, 2002.
13. P. Maymounkov. Greedy embeddings, trees, and Euclidean vs. Lobachevsky geometry. Manuscript, 2006.
14. A. Moitra and T. Leighton. Some results on greedy embeddings in metric spaces. *Discrete & Computational Geometry*, 44(3): 686–705, 2010.
15. C. H. Papadimitriou and D. Ratajczak. On a conjecture related to geometric routing. *Theoretical Computer Science*, 341(1): 3–14, 2005.
16. D. Peleg and E. Upfal. A trade-off between space and efficiency for routing tables. *Journal of the ACM*, 36(3): 510-530, 1989.
17. H. E. Warren. Lower bounds for approximation by nonlinear manifolds. *Transactions of the American Mathematical Society*, 138(1): 167–178, 1968.