

Lower Bounds on the Distortion of Embedding Finite Metric Spaces in Graphs

Y. Rabinovich¹ and R. Raz²

¹Computer Science Department, University of Toronto,
Toronto, Ontario, Canada M5S 1A1
yuri@cs.toronto.edu

²Department of Applied Mathematics, Weizmann Institute of Science,
Rehovot, Israel
ranraz@wisdom.weizmann.ac.il

Abstract. The main question discussed in this paper is how well a finite metric space of size n can be embedded into a graph with certain topological restrictions.

The existing constructions of graph spanners imply that any n -point metric space can be represented by a (weighted) graph with n vertices and $n^{1+O(1/r)}$ edges, with distances distorted by at most r . We show that this tradeoff between the number of edges and the distortion cannot be improved, and that it holds in a much more general setting. The main technical lemma claims that the metric space induced by an unweighted graph H of girth g cannot be embedded in a graph G (even if it is weighted) of smaller Euler characteristic, with distortion less than $g/4 - \frac{3}{2}$. In the special case when $|V(G)| = |V(H)|$ and G has strictly less edges than H , an improved bound of $g/3 - 1$ is shown. In addition, we discuss the case $\chi(G) < \chi(H) - 1$, as well as some interesting higher-dimensional analogues. The proofs employ basic techniques of algebraic topology.

1. Introduction

The geometry of finite metric spaces and, in particular, the questions related to their faithful representation by subsets of certain metric spaces of special interest, have been intensively studied and developed in recent years by researchers from different areas of Mathematics and Theoretical Computer Science. Let us mention but a few. Johnson and Lindenstrauss [7], and Bourgain [4] study embeddings in Hilbert spaces from the perspective of functional analysis. Motivated by questions in graph theory, Graham and Winkler [6] consider embeddings in \mathbb{Z}^d , cubes, etc. In Linial *et al.* [8] embeddings in low-dimensional real normed spaces are investigated and applied to the design of graph

algorithms. In the area of distributed computing, sparse subnetworks of a given network which faithfully represent the distances play a central role in various constructions, e.g., design of synchronizers. Such subnetworks, called *spanners*, were originated by Awerbuch [2], and Peleg and Ullman [12].

We continue this line of research and study embeddings in metric spaces induced by weighted graphs of bounded Euler characteristic.

Let (X, d) and (Y, δ) be finite metric spaces, $|X| = |Y|$. Let f be a one-to-one map from X onto Y . The *Lipschitz norm* of f , denoted $\|f\|_{\text{Lip}}$ or simply $\|f\|$, is

$$\|f\| = \max_{s \neq t \in X} \frac{\delta(f(s), f(t))}{d(s, t)}.$$

The *Lipschitz distance* (or *distortion*) between X and Y is

$$\text{dist}(X, Y) = \inf_{f: X \rightarrow Y} \|f\| \cdot \|f^{-1}\|,$$

where the infimum (actually minimum) is taken over all one-to-one maps. This distance is akin to the Banach–Mazur distance between linear normed spaces. The Lipschitz distance is a very natural measure for finite metric spaces, and is the one commonly used (see, e.g., Bourgain [4]). It has the following projective property:

$$\text{dist}(X, \alpha Y) = \text{dist}(X, Y),$$

where αY is the metric space on the same point set as Y , with all the distances multiplied by $\alpha > 0$. Indeed, if f is the least-distortion map from X onto Y , the same map f' from X onto αY has

$$\|f'\| = \alpha \|f\|; \quad \|f'^{-1}\| = \frac{1}{\alpha} \|f^{-1}\|.$$

In fact, the Lipschitz distance is closely related to the projective Hilbert metric (see, e.g., Bushell [5] for a beautiful application of this metric).

Let X be a finite metric space of size n . The central question discussed in this paper is how well X can be approximated by a subset S of cardinality n of vertices (with the induced metric) of an undirected graph G with positive weights on edges and bounded *Euler characteristic*. Recall that the Euler characteristic of G is $\chi(G) = |E(G)| - |V(G)| + 1$. In other words, we want to estimate the value of $\min_S \text{dist}(X, S)$, where the minimum is taken over all such $S \subseteq V(G)$. Naturally, the answer depends on X .

An important special case of the general question is when X is a metric space induced by a graph H with n vertices and e edges. Is it possible to approximate such X by a graph G with the same number of vertices, but strictly less edges? The later question, or more precisely the upper bound part of it, can be answered using *graph spanners*. A t -spanner ($t \geq 1$) of a graph H is a connected subgraph G of H , such that for every $i, j \in V$

$$\text{dist}_G(i, j) \leq t \cdot \text{dist}_H(i, j).$$

The key observation relating spanners to our framework is that $\text{dist}(H, G) \leq t$. Indeed, for the identity map on the vertices $I: V \mapsto V$, $\|I\| \leq t$ and $\|I^{-1}\| = 1$. Thus spanners with few edges provide some answers to our questions. The first explicit construction

of sparse spanners for arbitrary unweighted graphs was accomplished by Awerbuch [2], and Peleg and Schäffer [11]. A different construction appeared more recently in Althöfer *et al.* [1]; it improves the constants, and works for weighted graphs as well.

Theorem 1.1 [1]. *Let H be an arbitrary (weighted) graph with n vertices. Then, for all integer $t > 1$, H has a t -spanner with at most $n^{1+2/t}$ edges.*

Consequently, any metric space X of cardinality n can be t -approximated by such a graph. For the sake of completeness, here is a sketch of a short constructive proof of this theorem, a variant of the one appearing in [1]. It emphasizes the special role played by subgraphs of girth $> t + 1$. Recall that the girth of a graph is the smallest number of edges in a cycle.

Proof. Throw away (one by one) the inessential edges of H , i.e., the edges whose removal does not affect the distances in H . Consider the remaining edges one by one, in order of nondecreasing weight. At each stage add the current edge to the emerging structure if it does not close a cycle of (edge) size $\leq t + 1$ there; otherwise discard it. Call the obtained subgraph G . Obviously, G has girth $> t + 1$, and therefore, by a well-known theorem from extremal graph theory (see, e.g., Bollobás [3]), $|E(G)| \leq n^{1+2/t}$. On the other hand, the manner in which G was constructed implies that every nonchosen edge in H is approximated by some path in G with distortion $\leq t$. Consequently, the same holds for any path in H , and the theorem follows. \square

The above bound is tight (up to constants) for graph spanners. Indeed, if H is a simple unweighted (i.e., unit-weighted) graph of girth g , any proper subgraph of H distorts the distances by at least $g - 1$. Now, for all integer $g > 3$, there exist (explicitly constructible) graphs with n vertices and girth g , which have more than $\frac{1}{2}n^{1+4/(3g-6)}$ edges (see Margulis [9]). The lower bound follows.

Let us return to our original question: What is the smallest possible distortion in representing a given finite metric space X by a graph G of the same size as X , but with a bounded number of edges? More generally, what is the smallest possible distortion in representing X by a subset of vertices (with the induced metric) of a weighted graph G of bounded Euler characteristic $\chi(G)$? Theorem 1.1 provides an upper bound for this question. However, the simple argument used for obtaining a lower bound for spanners does not apply in the general case, as the class of the approximating graphs is not restricted to subgraphs of H . Nevertheless, somewhat surprisingly, we shall see that (almost) the same lower bound holds. Weaker existential related results were obtained in Althöfer *et al.* [1] (see the section on Steiner spanners).

The central result of this paper, which we shall refer to as the Main Lemma, claims that any embedding of an unweighted graph H with girth g in a graph G with $\chi(G) < \chi(H)$, has a distortion of at least $g/4 - \frac{3}{2}$. In the special case when $|V(H)| = |V(G)|$, this lower bound is improved to $g/3 - 1$.

The Main Lemma implies:

Theorem 1.2. *Let X be a metric space associated with an unweighted graph H of size n and girth $10 < g \leq n$, with $|E(H)|$ close to the maximal possible under these*

restrictions (i.e., Margulis graphs or better). Then X is at distance at least $g/4 - \frac{3}{2}$ from any subset of size n of any graph G with $\chi(G) \leq \frac{1}{2}n^{1+4/(3g-6)} - n + 1$.

In addition to the main question, we briefly discuss the case when $\chi(G) < \chi(H) - 1$. We also discuss an interesting extension of a special case of the Main Lemma to higher dimensions.

2. The Main Lemma: Special Case

In this section we prove the Main Lemma for the special case when the size (the number of vertices) of the approximating graph is equal to that of the approximated. There are two reasons for treating this important special case separately. The first is that it permits a clear exposition of the main ideas of the proof. Second, the result in this case is stronger.

All the graphs under discussion are undirected; the edges may have (strictly) positive weights. In an unweighted graph all edges have a unit weight. Sometimes the graphs will be identified with the metric spaces they induce on the vertices. The *girth* of a graph is the smallest number of edges in a cycle.

Lemma 2.1. *Let H be a simple unweighted connected graph, and let G be an arbitrary (weighted) graph with the same number of vertices, but strictly less edges. Then*

$$\text{dist}(H, G) \geq g/3 - 1,$$

where g is the girth of H .

Proof. For an arbitrary one-to-one mapping f from $V(H)$ to $V(G)$, it will be shown that necessarily $\|f\| \cdot \|f^{-1}\| \geq g/3 - 1$. It will be convenient to use h for f^{-1} . Without loss of generality, G is assumed to be simple, i.e., with no double edges and self-loops, and connected. Let n be the size of H and G , i.e., $|V(H)| = |V(G)| = n$.

Instead of discrete graphs and functions, we shall work with their continuous counterparts. First, let us associate our graphs with certain one-dimensional simplicial complexes endowed with a metric. Think of each edge $e = (v, u)$ of weight w of the graph as an interval of length w equipped with the usual line metric, with one endpoint identified with the vertex v and the other with u . Edge metrics naturally induce a metric on the entire structure. The structures obtained in this fashion from H and G will be denoted \tilde{H} and \tilde{G} , respectively, and called sometimes the *continuous* H, G . Note that the distances between the vertices are preserved, i.e., $\text{dist}_H(v, u) = \text{dist}_{\tilde{H}}(v, u)$. Every path in \tilde{H} (and \tilde{G}) has a naturally defined length. The distance between any pair of points $x, y \in \tilde{H}$ (or \tilde{G}) is equal to the length of a shortest path connecting them (there might be several). Such paths will be called *geodesic paths*.

Next, extend f and h to continuous maps $\tilde{f}: \tilde{H} \mapsto \tilde{G}$ and $\tilde{h}: \tilde{G} \mapsto \tilde{H}$, respectively, in the following piecewise-linear manner: For each edge $e = (v, u)$ of H , mark some geodesic path $P_e = P_{(v,u)}$ from $f(v)$ to $f(u)$ in \tilde{G} . Let $x \in \tilde{H}$ be a point in \tilde{H} , and

assume that it belongs to the edge (a, b) . Let

$$\alpha = \frac{\text{dist}_{\tilde{H}}(a, x)}{\text{dist}_{\tilde{H}}(a, b)}.$$

Now $\tilde{f}(x)$ is defined as the (unique) point $y \in P_{(a,b)}$ such that

$$\frac{\text{dist}_{\tilde{G}}(f(a), y)}{\text{dist}_{\tilde{G}}(f(a), f(b))} = \alpha.$$

Define \tilde{h} analogously. Notice, that unlike f and h , their continuous counterparts \tilde{f}, \tilde{h} , do not have to be one-to-one. Also, there might be more than one way to obtain the extensions, as they depend on the choice of the geodesic paths. In general, $\tilde{h} \circ \tilde{f}$ need not be the identity function on \tilde{H} anymore.

We note that \tilde{f} is indeed an extension of f , i.e., for any vertex $v \in H$, $\tilde{f}(v) = f(v)$. Another observation is that the Lipschitz norm of \tilde{f} is equal to that of f ,

$$\|\tilde{f}\| = \max_{s \neq t \in \tilde{H}} \frac{\text{dist}_{\tilde{G}}(\tilde{f}(s), \tilde{f}(t))}{\text{dist}_{\tilde{H}}(s, t)} = \|f\|.$$

This follows directly from the structure of \tilde{H} and \tilde{G} , and the obvious fact that for $s, t \in \tilde{H}$ belonging to the same edge $e = (a, b)$,

$$\frac{\text{dist}_{\tilde{G}}(\tilde{f}(s), \tilde{f}(t))}{\text{dist}_{\tilde{H}}(s, t)} = \frac{\text{dist}_G(f(a), f(b))}{\text{dist}_H(a, b)}.$$

Respectively, $\|\tilde{h}\| = \|h\|$.

Claim 2.2. *If there exist $x, y \in \tilde{H}$ such that*

$$\text{dist}_{\tilde{H}}(x, y) \geq \frac{g}{3} \quad \text{and} \quad \tilde{f}(x) = \tilde{f}(y),$$

then $\|f\| \cdot \|h\| \geq g/3 - 1$.

Proof. Let v and u be the vertices closest to x and y , respectively. Since the set of vertices constitutes a $\frac{1}{2}$ -net for \tilde{H} , $\text{dist}_H(v, u) \geq g/3 - 1$. On the other hand, since $\tilde{f}(x) = \tilde{f}(y)$,

$$\text{dist}_G(f(v), f(u)) \leq \text{dist}_{\tilde{G}}(\tilde{f}(x), \tilde{f}(v)) + \text{dist}_{\tilde{G}}(\tilde{f}(y), \tilde{f}(u)) \leq \|\tilde{f}\| \cdot (\frac{1}{2} + \frac{1}{2}) = \|f\|.$$

Combining the two inequalities we get $\|h\| \geq (g/3 - 1)/\|f\|$. \square

Thus, under the conditions of Claim 2.2 the lemma is true. It remains to take care of the case when no two points at distance $\geq g/3$ in \tilde{H} are mapped by \tilde{f} to the same point.

Assume by contradiction that indeed $\|f\| \cdot \|h\| < g/3 - 1$. Since $\|\tilde{f}\| = \|f\|$ and $\|\tilde{h}\| = \|h\|$, the same also applies to $\|\tilde{f}\| \cdot \|\tilde{h}\|$. Let T be a mapping from \tilde{H} to itself, defined by $T(x) = \tilde{h} \circ \tilde{f}(x)$. Clearly, T is continuous. It will be shown that T is homotopic to identity, which, together with some basic facts from algebraic topology, contradicts the condition $|E(H)| > |E(G)|$.

Claim 2.3. For any $x \in \tilde{H}$, the distance between x and $T(x)$ is $< g/2$.

Proof. Suppose that x belongs to the edge e of \tilde{H} , and $\tilde{f}(x)$ belongs to the edge $c = (u, v)$ of \tilde{G} . By definition of \tilde{f} , there exist points $p, q \in e$ such that x lies between them, and $\tilde{f}(p) = u$, $\tilde{f}(q) = v$. (If $\tilde{f}(x)$ is a vertex of G , the proof is simplified). Let $P_c \subseteq \tilde{H}$ be the geodetic path to which the edge $c \in \tilde{G}$ is mapped by \tilde{h} . The endpoints of P_c are the vertices $h(u)$ and $h(v)$. By definition, $T(x) = \tilde{h} \circ \tilde{f}(x)$ belongs to this path.

Since p and $h(u)$ are mapped by \tilde{f} to the same point u , the distance between them is, by our first assumption, smaller than $g/3$. The same applies to q and $h(v)$.

The distance between $h(u)$ and $h(v)$ is bounded by $g/3 - 1$. Indeed,

$$\frac{\text{dist}_H(h(u), h(v))}{\text{dist}_{\tilde{H}}(p, q)} = \frac{\text{dist}_H(h(u), h(v))}{\text{dist}_G(u, v)} \cdot \frac{\text{dist}_{\tilde{G}}(u, v)}{\text{dist}_{\tilde{H}}(p, q)} \leq \|h\| \cdot \|\tilde{f}\| = \|h\| \cdot \|f\|,$$

and the last term is less than $g/3 - 1$ (by the second assumption).

Consider the following piecewise geodetic circular path in \tilde{H} :

$$p \rightarrow h(u) \xrightarrow{P_c} h(v) \rightarrow q \rightarrow p.$$

By the previous assertions, its length is bounded by $g/3 + (g/3 - 1) + g/3 + \text{dist}_{\tilde{H}}(p, q) \leq g$. On the other hand, both x and $T(x)$ belong to it. Subsequently, $\text{dist}_{\tilde{H}}(x, T(x)) < g/2$. \square

The next step in the proof of the lemma is to establish a homotopy between T and $\text{Id}_{\tilde{H}}$. Observe that there is a *unique* geodetic path $P(x) \subseteq \tilde{H}$ between x and $T(x)$, since the distance between them is less than half the girth of H . Define

$$M[t, x] = (1 - t)x + tT(x), \quad t \in [0, 1],$$

where by the right-hand side we mean the *unique* point $y \in P(x)$ such that $\text{dist}_{\tilde{H}}(x, y) / \text{dist}_{\tilde{H}}(x, T(x)) = t$. Note that if two points $x, y \in \tilde{H}$ are close enough, so are $T(x)$ and $T(y)$, and moreover, the two paths $P(x)$ and $P(y)$ must be nearly identical. Therefore $M[t, x]$ is continuous, and thus it is the required homotopy.

To complete the proof we need some basic definitions and facts from algebraic topology (see, e.g., [10]):

First, recall the notion of the *first homology group* $H_1(K)$ of a (continuous) graph K . It can be defined as the abelian group of all *flows* in K . Given an arbitrary orientation of the edges, a flow in (oriented) K is an integer-valued function on $E(K)$ such that for all vertices v of K , the sum of values of the incoming edges is equal to that of the outgoing. It is readily checked that $H_1(K)$ is isomorphic to the additive group $\mathbb{Z}^{\chi(K)}$, where $\chi(K) = |E(K)| - |V(K)| + 1$ is the Euler characteristic of K . Note that the Euler characteristic of a discrete graph is equal to that of its continuous counterpart.

Next, we need the following fundamental fact (which, for future use, is stated here for arbitrary simplicial complexes and the i th homology groups):

Fact. Let K and L be simplicial complexes, $f: K \mapsto L$, $h: L \mapsto K$, continuous maps, and assume that $h \circ f$ is homotopic to Id_K . Then these mappings induce group

homomorphisms $f^*: H_i(K) \mapsto H_i(L)$ and $h^*: H_i(L) \mapsto H_i(K)$ such that $h^* \circ f^*$ is an automorphism of $H_i(K)$. Consequently, $H_i(K)$ is isomorphically embeddable in $H_i(L)$.

We return now to our proof. Since $T = \tilde{h} \circ \tilde{f}$ is homotopic to identity, the above fact implies that the first homology group $H_1(\tilde{H})$ is embeddable in $H_1(\tilde{G})$. On the other hand, $\chi(\tilde{H}) = |E(H)| - n + 1 > |E(G)| - n + 1 = \chi(\tilde{G})$, and $\mathbb{Z}^{\chi(\tilde{H})}$ cannot be isomorphically embedded in $\mathbb{Z}^{\chi(\tilde{G})}$. Contradiction. \square

3. The Main Lemma: General Case

In this section we address the general case, i.e., when the sole restriction on the approximating graph G is $\chi(G) < \chi(H)$. In what follows we will often refer to the proof of Lemma 2.1 presented in the previous section.

The proof of the special case of the Main Lemma proceeded differently depending on whether the conditions of Claim 2.2 were satisfied. However, is it at all possible for these conditions to fail? In this section we show that a slightly weaker statement always holds, and continues to hold under more general circumstances.

Theorem 3.1. *Let H be an unweighted graph of girth g , and let G be a (finite, weighted) graph satisfying $\chi(G) < \chi(H)$. Then, for any continuous $\tilde{f}: \tilde{H} \rightarrow \tilde{G}$, there exist $x, y \in \tilde{H}$ such that $\tilde{f}(x) = \tilde{f}(y)$ and $\text{dist}_{\tilde{H}}(x, y) \geq g/4 - \frac{1}{2}$.*

Remark. Notice that here, unlike before, \tilde{f} is an arbitrary continuous function, and not an extension of some discrete f .

Proof. The general scheme of the proof resembles that of Lemma 2.1; the main difference lies in finding a suitable \tilde{h} .

Without loss of generality, we may assume that \tilde{f} is onto; if it is not, we shall consider its image $\tilde{f}(\tilde{H}) \subseteq \tilde{G}$ instead of \tilde{G} . Obviously, $\chi(\tilde{f}(\tilde{H})) \leq \chi(\tilde{G}) = \chi(G) < \chi(H)$. The vertices of $\tilde{f}(\tilde{H})$ are, naturally, those points of $\tilde{f}(\tilde{H})$ whose neighborhoods are *not* locally homeomorphic to \mathbb{R} , plus the images of vertices of \tilde{H} .

Here comes the only fine point in the proof which did not occur previously. We shall formally (i.e., without effecting the actual structure) extend the set of vertices $V(\tilde{G})$ of \tilde{G} by adding to it a finite set of points $Ext \subset \tilde{G}$. The extended vertex set defines a new edge set for \tilde{G} . The goal of the extension is to achieve the following property: For any edge $e \subseteq \tilde{H}$ and a point $x \in e$, either the image $\tilde{f}(x)$ is a vertex of \tilde{G} , or $\tilde{f}(x)$ belongs to an edge $c = (u, v) \subseteq \tilde{G}$ and there exist points $p, q \in e$ such that $x \in [p, q]$, and $\tilde{f}(p) = u$, $\tilde{f}(q) = v$.

The desired extension can be achieved by taking $Ext = \{w \mid w \text{ is an extreme point for some } \tilde{f}(e), \text{ where } e \subseteq \tilde{H} \text{ is an edge of } \tilde{H}\}$. An extreme point w is defined as an interior point of some edge c in \tilde{G} , such that any neighborhood of w in c contains points *not* in $\tilde{f}(e)$. Clearly, no edge $c \subseteq \tilde{G}$ may contain more than two extreme points associated with an edge $e \subseteq \tilde{H}$.

The rest of the proof is very close to the one we had before. Define a function

$h: V(\tilde{G}) \rightarrow \tilde{H}$ as follows: For $u \in V(\tilde{G})$, let $h(u)$ be *some* point $p \in \tilde{H}$ such that $\tilde{f}(p) = u$. The function $\tilde{h}: \tilde{G} \rightarrow \tilde{H}$ is defined as before, by extending h to the whole \tilde{G} in a piecewise linear manner using geodetic paths. Define also $T: \tilde{H} \rightarrow \tilde{H}$, by $T(x) = \tilde{h} \circ \tilde{f}(x)$.

Now, assume by contradiction that the theorem is wrong, i.e., the distance between any $x, y \in \tilde{H}$ with $\tilde{f}(x) = \tilde{f}(y)$, is (strictly) less than $g/4 - \frac{1}{2}$.

Claim 3.2. *Under the above assumption, for any $x \in \tilde{H}$ the distance between x and $T(x)$ is (strictly) less than $g/2$.*

Proof. Suppose that x belongs to edge e of \tilde{H} , and $\tilde{f}(x)$ belongs to edge $c = (u, v)$ of \tilde{G} . We assume without loss of generality, that $\tilde{f}(x)$ is *not* a vertex of \tilde{G} (otherwise, the proof becomes simpler). By the manner in which the old set of vertices $V(\tilde{G})$ was extended, there exist points $p, q \in e$ such that $x \in [p, q]$, and $\tilde{f}(p) = u$, $\tilde{f}(q) = v$. Let $P_c \subseteq \tilde{H}$ be the geodetic path to which the edge $c \in \tilde{G}$ was mapped by \tilde{h} . The endpoints of P_c are $h(u)$ and $h(v)$. By the definition of \tilde{h} , $T(x) = \tilde{h} \circ \tilde{f}(x)$ belongs to this path.

By the definition of h , the points p and $h(u)$ are mapped by \tilde{f} to the same point u , and therefore the distance between them is (by our assumption) less than $g/4 - \frac{1}{2}$; the same applies to q and $h(v)$. Also, $\text{dist}(p, q) \leq 1$. Clearly, the length of the geodetic path P_c is at most $\text{dist}_{\tilde{H}}(h(u), p) + \text{dist}_{\tilde{H}}(p, q) + \text{dist}_{\tilde{H}}(q, h(v)) < g/2$.

Consider the following piecewise geodetic circular path in \tilde{H} :

$$p \rightarrow h(u) \xrightarrow{P_c} h(v) \rightarrow q \rightarrow p.$$

It contains both x and $T(x)$; by the previous assertions, its length is (strictly) less than $(g/4 - \frac{1}{2}) + g/2 + (g/4 - \frac{1}{2}) + 1 = g$. Thus we conclude $\text{dist}_{\tilde{H}}(x, T(x)) < g/2$. \square

The rest of the proof proceeds exactly as before. The contradiction is reached by directly comparing the Euler characteristics of \tilde{G} and \tilde{H} , rather than the number of edges in these graphs. \square

Arguing as in the proof of Claim 2.2, we obtain the general case of the Main Lemma as a corollary to Theorem 3.1.

Lemma 3.3. *Let H be a simple unweighted connected graph of size n and girth g , and let G be a (finite, weighted) graph of size $\geq n$, such that $\chi(G) < \chi(H)$. Then, for any subset S of cardinality n of vertices of G equipped with the induced metric, $\text{dist}(H, S) \geq g/4 - \frac{3}{2}$.*

An interesting question is whether the (multiplicative) constant $\frac{1}{4}$ appearing in Lemma 3.3 is the best possible. We conjecture that it can be improved to $\frac{1}{3}$. When the approximated graph is the n -cycle (C_n), the constant is indeed $\frac{1}{3}$, as will be shown in Section 5. Here we show an upper bound arbitrarily close to $\frac{1}{3}$ for this special case.

Consider the embedding of C_n in the tree T_n as illustrated by Fig. 1. Assign to all the outer edges of the tree (associated with the leaves) weight 1, and to all the inner edges




Fig. 1. An embedding of C_n in a tree.

weight 2δ , where δ is an arbitrary small positive constant. It is readily checked that in this case $\|f\| = 2 + 2\delta$, while for sufficiently large n , $\|f^{-1}\| = n/6$. Therefore, the distortion is $(1 + \delta)n/3$.

4. When $\chi(G) < \chi(H) - 1$

Given a graph H define its *approximation pattern* as a function $A_H(i), i \in \mathbb{N} \cup \{0\}$, where $A_H(i)$ is the minimum possible distortion in an embedding of H in a graph G with $\chi(G) \leq i$. Theorem 1.1 provides an upper bound on the *envelope* of the approximation patterns of all graphs of size n . That is, any graph of size n can do at least as well. We also know that for any particular i there exists a graph H of size n which cannot do much better. The problem that we consider in this section—and, in fact, in the entire paper—has to do with providing bounds on the approximation pattern of a *fixed* graph H .

Let H be a simple unweighted graph (not a tree). Clearly, by omitting one edge in a shortest cycle, one arrives at a $(\text{girth}(H) - 1)$ -spanner of H with $|E(H)| - 1$ edges. Combining this simple fact with the Main Lemma yields $A_H(\chi(H) - 1) = \Theta(\text{girth}(H))$. But what about $A_H(\chi(H) - k)$, where $k > 1$? The above bound obviously becomes loose. In this section, we propose some approaches toward bounding these numbers.

The first approach seems to be appropriate for small k ; it extends the one used for $k = 1$. We demonstrate it for $k = 2$, where the bound is sharp, and the method is not obscured by technicalities.

Let H be a simple unweighted graph of size n , and let C be a cycle of minimal length in it (in general, not unique). Define g_2 as the length of the second (after C) shortest simple cycle in H . Notice that the definition does not depend on the choice of C . Obviously, by omitting two properly chosen edges from H , one can obtain a $(g_2 - 1)$ -spanner of H .

Theorem 4.1. *Let H be as above, and let \tilde{G} be an arbitrary (without loss of generality continuous) graph with finitely many vertices and $\chi(\tilde{G}) \leq \chi(H) - 2$. Then, for any subset $S \subseteq \tilde{G}, |S| = n$, with the induced metric, $\text{dist}(H, S) \geq g_2/4 - \frac{3}{2}$.*

Proof. If $g_2 = \text{girth}(H)$, the theorem holds in a trivial manner; in what follows we assume that $g_2 > \text{girth}(H)$, or, equivalently, there exists a *unique* cycle of minimal length in H , namely C .

Put for convenience, $r = \chi(H)$, $s = \chi(\tilde{G})$. Proceeding as in the proof of Lemma 3.3 (including the assumption that without loss of generality \tilde{f} is onto), we conclude that our theorem can potentially fail only when for all $x \in \tilde{H}$,

$$\text{dist}_{\tilde{H}}(x, T(x)) < g_2/2,$$

where, as always, $T(x) \equiv \tilde{h} \circ \tilde{f}(x)$. Assume this is the case. Contracting C (i.e., gluing together all the vertices in C and deleting the $|C|$ self-loops created by this), we obtain a new (simple) graph H' . Let I be the projection map of \tilde{H} on \tilde{H}' defined by the contraction. We claim that under the assumptions, the maps I and $I \circ T: \tilde{H} \rightarrow \tilde{H}'$ are homotopic.

Let $x, y \in \tilde{H}$ be any two points at distance smaller than $g_2/2$, and let $P_1(x, y)$, $P_2(x, y)$ be any two geodesic paths between them. We claim that under I both P_1 and P_2 are mapped to the *same* path $P(x', y') \subseteq \tilde{H}'$, where $x' = I(x)$, $y' = I(y)$. Indeed, by the uniqueness of the minimal length cycle in H , the circular path $x \xrightarrow{P_1} y \xrightarrow{P_2} x$ can contain no simple cycle but C , and the latter is contracted by I to a point. After this observation the task of establishing a homotopy between I and $I \circ T$ becomes simple: for every $x \in \tilde{H}$, $I(x)$ can be transferred to $I \circ T(x)$ in \tilde{H}' in a continuous and uniform manner along the path P'_x , which is the image (under I) of all geodesic paths $P(x, T(x)) \subseteq \tilde{H}$. Clearly, the dependence of P'_x on x is a continuous one.

However, I and $I \circ T$ cannot be homotopic. If they were, the induced maps I^* , $(I \circ T)^*: H_1(\tilde{H}) \rightarrow H_1(\tilde{H}')$ would be equivalent. It is easy to verify that I^* is onto; it is a projection map with kernel generated by the element of $H_1(\tilde{H})$ associated with C . But $(I \circ T)^*$ cannot be onto! Indeed, the homotopy groups are, respectively, $H_1(\tilde{H}) = \mathbb{Z}^r$, $H_1(\tilde{G}) = \mathbb{Z}^s$, and, by a direct calculation, $H_1(\tilde{H}') = \mathbb{Z}^{r-1}$. Thus

$$(I \circ T)^*: \mathbb{Z}^r \xrightarrow{\tilde{f}^*} \mathbb{Z}^s \xrightarrow{\tilde{h}^*} \mathbb{Z}^r \xrightarrow{I^*} \mathbb{Z}^{r-1}.$$

But $s < r - 1$. Contradiction. □

Remark. In the case when G is an n -vertex graph, and $S = V(G)$, the above bound can be improved, as in Lemma 2.1, to $\text{dist}(H, S) \geq g_2/3 - 1$.

The second approach we discuss is much simpler, and seems to be more suitable for large values of k . The idea is to find a large-girth subgraph of H , whose (inner) distances are close to those in H . Currently, we do not clearly understand the limits of this approach. We demonstrate how it works by considering a concrete example.

Let $M_{n \times n}$ denote a discrete $(n + 1) \times (n + 1)$ grid. How well can it be approximated by a subset of a tree? The boundary cycle Q of $M_{n \times n}$ with its inner metric is at a distance 2 from C_{4n} . The latter is, by Corollary 5.3 (to be proven in the next section) at a distance $\geq \frac{4}{3}n - 1$ from any subset S' , $|S'| = 4n$, of a tree. Therefore for any subset S of size $(n + 1)^2$ of a tree, $\text{dist}(H, S) \geq \frac{2}{3}n - 1/2$. The bound is quite tight: it is easy to find a subtree in $M_{n \times n}$ which distorts the distances by at most $n + 3$.

5. Approximating Cycles; Multidimensional Generalizations

In this section we give a different proof for Theorem 3.1 for the special case when H is a cycle. The statement and the proof are then generalized to higher dimensions.

Let $\text{dist}_{\mathbb{S}}(\cdot, \cdot)$ denote the geodesic metric on the unit sphere \mathbb{S} of a prescribed dimension.

Proposition 5.1. *Let \tilde{T} be a topological tree, i.e., a one-dimensional simplicial complex associated with a discrete tree T . Then, for any continuous map \tilde{f} from \mathbb{S}^1 to \tilde{T} there exist points $x, y \in \mathbb{S}^1$ such that $\text{dist}_{\mathbb{S}}(x, y) \geq 2\pi/3$, and $\tilde{f}(x) = \tilde{f}(y)$.*

This proposition is a consequence of the following:

Proposition 5.2. *Let $\tilde{f}: \mathbb{S}^1 \mapsto \tilde{T}$ be a continuous map, and let $\{I_1, I_2, I_3\}$ be an arbitrary partition of \mathbb{S}^1 to three intervals with mutually disjoint interiors. Then there exists $c \in \tilde{T}$ such that $\tilde{f}^{-1}(c)$ has a representative in each of these intervals.*

Proof. Recall the well-known simple fact that for any three points in \tilde{T} there exists a (unique) *central* point c , such that for any pair of these points, c lies on the path connecting them. Let s_1, s_2, s_3 be the points in \mathbb{S}^1 such that $I_1 = [s_1, s_2]$, $I_2 = [s_2, s_3]$, and $I_3 = [s_3, s_1]$, and let c be the central point of $\tilde{f}(s_1), \tilde{f}(s_2), \tilde{f}(s_3)$. By definition, $\tilde{f}^{-1}(c)$ has a nonempty intersection with each of the three intervals. \square

If the intervals I_1, I_2 , and I_3 are chosen to be of the (same) length $2\pi/3$, then among any three points $x_1 \in I_1, x_2 \in I_2$, and $x_3 \in I_3$ there exists a pair x_i, x_j with $\text{dist}_{\mathbb{S}}(x_i, x_j) \geq 2\pi/3$. Combining this observation with Proposition 5.2, we derive Proposition 5.1. \square

Arguing as in the proofs of Theorem 3.1 and Claim 2.2, the above propositions yield:

Corollary 5.3. *Let S be an n -point finite metric space defined by a subset of vertices of some tree T . Then*

$$\text{dist}(C_n, S) \geq \frac{n}{3} - 1.$$

Next, we discuss the analogues of Propositions 5.2 and 5.1 in higher dimensions. We start with the former. Let Δ_n be the boundary of a simplex σ_{n+1} of dimension $n + 1$. By definition, Δ_n is a union of the $(n + 2)$ -faces of σ_{n+1} . Denote this family of faces by $\mathcal{F} = \{F_i\}_{i=1}^{n+2}$. Each F_i is an n -dimensional simplex; different faces have mutually disjoint interiors. In what follows, it will be convenient to view Δ_n as the following equivalent set:

$$\Delta_n = \left\{ (x_1, x_2, \dots, x_{n+2}) \in \mathbb{R}^{n+2} \mid x_i \geq 0 \text{ for } i = 1, \dots, n + 2; \right. \\ \left. \sum_{i=1}^{n+2} x_i = 1; \text{ at least one } x_i \text{ is } 0 \right\}.$$

In this representation the face F_i corresponds to the subset of Δ_n of all points with $x_i = 0$.

For a simplicial complex K , let $|K|$ denote its geometrical realization. For the sake of convenience we shall identify $|K|$ with K .

Theorem 5.4. *Let K be a simplicial complex of dimension $\leq n$ such that the n th homology group $H_n(K)$ is trivial. Then, for any continuous map $f: \Delta_n \mapsto |K|$, there exists a point $c \in |K|$ such that $f^{-1}(c)$ has a representative in every $F_i \in \mathcal{F}$.*

Proof. Let $U_i \subseteq |K|$ be the image of F_i under f , for $i = 1, \dots, n+2$. The union of the U_i -s, i.e., $f(\Delta_n)$, will be called U . Notice that all the U_i -s are compact, and so is U . For $u \in U$, let $d(u)$ be the vector of distances

$$d(u) = (\text{dist}_{|K|}(u, U_1), \text{dist}_{|K|}(u, U_2), \dots, \text{dist}_{|K|}(u, U_{n+2})) \in \mathbb{R}^{n+2}.$$

Assume by contradiction that the theorem is false. Then it is possible to define a map $h: U \mapsto \Delta_n$,

$$h(u) = \frac{d(u)}{\|d(u)\|_1}.$$

Let us check that h is indeed well defined. Since by our assumption the intersection $\bigcap_{i=1}^{n+2} U_i$ is empty, the denominator is never zero. The coordinates of $d(u)$, being distances, are nonnegative. Also, since every $u \in U$ belongs to some U_i , at least one of these coordinates is 0.

The rest is simple. Consider the map $T: \Delta_n \mapsto \Delta_n$ defined as $T \equiv h \circ f$. By the definition of h , T maps each F_i into itself. Proceeding as in the proof of Claim 2.3, we construct a homotopy between T and the identity map Id_{Δ_n} :

$$M[t, x] = (1-t)x + tT(x), \quad t \in [0, 1],$$

where the convex combination is well defined, as both x and $T(x)$ lie in the same face. Such homotopy is, however, impossible. As we have already mentioned before, its existence would imply that the n th homology group $H_n(\Delta_n)$ can be embedded in $H_n(U)$. But $H_n(U)$ is trivial, as U is a subset of an n -dimensional $|K|$ with $H_n(k) = 0$, whereas $H_n(\Delta_n)$ is not. \square

Remark 5.5. The same result can be shown to hold in a slightly more general setting, when K is an arbitrary n -dimensional simplicial complex and f is homotopic to a map that sends the whole Δ_n to a single point in $|K|$.

Before discussing the high-dimensional analogue of Proposition 5.1, let us view the problems we have been dealing with from a more general perspective. Suppose that H and G are metric spaces such that for some topological reasons H cannot be continuously embedded in G . Then, for any continuous map $f: H \mapsto G$ there exist some $x, y \in H$ such that $f(x) = f(y)$. Define $\alpha(H, G) = \inf_f \sup_{x,y} \text{dist}_H(x, y)$ over all such x, y . What can be said about $\alpha(H, G)$? The same question can be asked for f -s of a given homotopy type. The problem seems to be appealing and mathematically interesting. Although some classic results in this direction exist (e.g., the Borsuk–Ulam theorem can

be viewed as related), we are unaware of any general method for obtaining sharp bounds on $\alpha(H, G)$ in the general case.

Let α_n be the infimum of $\alpha(\mathbb{S}^n, |K|)$, taken over all simplicial complexes K of dimension $\leq n$ such that the n th homology group $H_n(K)$ of K is trivial. Then:

Theorem 5.6. α_n is equal to the (geodetic) diameter of the regular simplex inscribed in the unit sphere \mathbb{S}^n , i.e.,

$$\alpha_n = \arccos\left(-\frac{1}{n+1}\right).$$

Proof. Let us identify \mathbb{S}^n with the unit sphere in \mathbb{R}^{n+1} centered at the origin; Δ_n can then be identified (up to scaling) with the boundary of the regular simplex σ_{n+1} inscribed in \mathbb{S}^n .

The upper bound is easy, and is given by the following construction:

Let $\{v_i\}_{i=1}^{n+2} = \Delta_n \cap \mathbb{S}^n$ be the vertices of Δ_n , so that $v_i \notin F_i$. Let \mathcal{D} be the cone of all facets of dimension $n - 1$, i.e.,

$$\mathcal{D} = \{x \in \mathbb{R}^{n+1} \mid x = tp, \text{ where } t \geq 0, \text{ and } p \text{ belongs to the boundary of some } F_i\}.$$

Clearly, \mathcal{D} is contractible and of dimension n .

Let $F'_i \subset \mathbb{S}^n$, $i = 1, \dots, n + 2$, be the central projections of the F_i 's on the sphere \mathbb{S}^n . Define the map $T: \mathbb{S}^n \mapsto \mathcal{D}$ as follows. For $x \in F'_i$, let $T(x) = y$, where $y \in \mathcal{D}$ is such that $y - x$ is a multiple of v_i , and is as short as possible. In other words, T projects F'_i on \mathcal{D} parallel to v_i . It is easily verified that T is continuous, and that, for any vertex v_i of Δ_n , $-v_i$ is mapped to the origin. Furthermore, for any $x_1, x_2 \in \mathbb{S}^n$ with $T(x_1) = T(x_2)$, the Euclidean distance between them is equal to $\|t_i v_i - t_j v_j\|_2$ for some pair of vertices v_i, v_j , and scalars $0 \leq t_i, t_j \leq 1$. Since $\langle v_i, v_j \rangle < 0$, this distance is bounded by $\|v_i - v_j\|_2$. Consequently, the geodetic distance between x_1 and x_2 is at most $\text{dist}_{\mathbb{S}}(v_i, v_j) = \arccos(-1/(n + 1))$.

A careful analysis shows that the best lower bound on α_n one can get from Theorem 5.4 is only $\arccos((n - 2)/2n)$, which is too weak for $n > 1$. The proof of the stronger lower bound is based on the following different statement:

Lemma 5.7. *Let K be a simplicial complex of dimension $\leq n$ such that $H_n(K) = 0$. Then, for any continuous map $f: \mathbb{S}^n \mapsto K$, there exists a point $c \in K$ such that the convex hull of $f^{-1}(c)$ contains the center of \mathbb{S}^n .*

Proof. Let K be endowed with a metric induced by the Euclidean distance. Since the n th homology group of the image of f is also trivial (it is a subcomplex of K), we may assume without loss of generality that f is onto, and K is compact.

Assume by contradiction that the statement is false. Define a map $h: K \mapsto \mathbb{S}^n$ as follows. For a point $a \in K$, let $p(a)$ be the (uniquely defined) closest-to-the-center point of the convex hull of $f^{-1}(a)$. Let

$$h(a) = \frac{p(a)}{\|p(a)\|_2}.$$

Due to our assumptions, h is well defined. Call two unit vectors $x, y \in \mathbb{S}^n$ *compatible* if $\langle x, y \rangle > 0$. The key property of h is that, for every $x \in \mathbb{S}^n$, the vectors x and $h \circ f(x)$ are compatible. The reason for this is that $p \circ f(x)$ is the closest-to-the-origin point in the convex hull of $f^{-1} \circ f(x)$, and as such has a positive inner product with any point in $\text{conv}\{f^{-1} \circ f(x)\}$, including x . We leave the verification of this simple geometrical statement to the reader. This property alone would have yielded a contradiction if h were continuous. But it is not. To overcome this difficulty we construct a new function $h_1: K \mapsto \mathbb{S}^n$, such that h_1 is continuous, and, for every $a \in K$, $h(a)$ and $h_1(a)$ are compatible.

Claim 5.8. *Let $B_\varepsilon(a) \subseteq K$ denote the ball of radius ε around $a \in K$. There exists $\varepsilon > 0$, such that for every point $a \in K$ there exists a direction $d(a) \in \mathbb{S}^n$, such that, for every $x \in B_\varepsilon(a) \subseteq K$, $h(x)$ and $d(a)$ are compatible.*

Proof. Assume the statement is wrong. Then there exists a sequence of balls B_{ε_i} , $i = 1, 2, \dots$, with $\varepsilon_i \rightarrow 0$, such that none of their images under h has a common compatible direction. Equivalently, the convex hull of every $h(B_{\varepsilon_i})$, and therefore also of every $p(B_{\varepsilon_i})$, contains the center of the sphere. Since $p(B_{\varepsilon_i}) \subseteq f^{-1}(B_{\varepsilon_i})$, the convex hull of the latter set also contains the center. Finally, by the Carathéodory Theorem, there exist some $n + 2$ vectors in $f^{-1}(B_{\varepsilon_i})$ whose convex hull already contains the center. Consider a subsequence of the balls-sequence which converges to some point c of K . Using the standard compactness argument on the corresponding $(n + 2)$ -tuples in \mathbb{S}^n and keeping in mind the continuity of f , one arrives at a conclusion that the convex hull of $f^{-1}(c)$ contains the center of the sphere, contrary to our assumptions. \square

Let $\varepsilon > 0$ be as provided by Claim 5.8. Let TR be a triangulation of K such that the diameter of any simplex in TR is at most ε . For a vertex a of TR , define $h_1(a)$ to be a direction $d(a) \in \mathbb{S}^n$ as defined in Claim 5.8. Extend the function to the entire K in a piecewise linear manner using the triangulation TR . Clearly, the obtained h_1 is continuous. We claim that it is pointwise compatible with h . Indeed, it is true by definition for $a \in K$ which is a vertex of TR . If a is not a vertex, it is a convex combination of vertices $\{a_1, a_2, \dots, a_m\}$ of some simplex in TR . By the choice of TR , $\text{dist}_K(a, a_i) \leq \varepsilon$ for every a_i . By the manner in which h_1 was defined on the vertices of TR , $h(a)$ is compatible with all $h_1(a_i)$'s, and therefore it must be compatible with their convex combination $h_1(a)$ as well.

We return to the proof of the lemma. Define, as usual, $T: \mathbb{S}^n \mapsto \mathbb{S}^n$ by $T \equiv h_1 \circ f$. The map is continuous. Moreover, no $x \in \mathbb{S}^n$ is mapped to its antipode, since $T(x) = h_1 \circ f(x)$ is compatible with $h \circ f(x)$ which is in turn compatible with x . Therefore, for every $x \in \mathbb{S}^n$ there exists a *unique* geodetic path between x and $T(x)$. Arguing along the same lines as in the proof of Lemma 2.1, these geodetic paths can be used to construct a homotopy between T and the identity map, and conclude that $H_n(\mathbb{S}^n)$ is embeddable in $H_n(K) = 0$. The obtained contradiction establishes Lemma 5.7. \square

All that remains to complete the proof of Theorem 5.6 is to show the following simple geometrical fact:

Claim 5.9. *Let $S \subset \mathbb{S}^n$ be a set of points whose convex hull contains the center of the sphere. Then the (geodesic) diameter of S is at least that of the regular simplex inscribed in \mathbb{S}^n .*

Proof. Clearly, it suffices to show that the Euclidean diameter of S is at least that of a regular simplex inscribed in \mathbb{S}^n . By the Carathéodory Theorem, there exists a subset Y of S of size $m \leq n + 2$ such that $\text{conv}\{Y\}$ also contains the center of the sphere. Let M be an $(n + 1) \times m$ matrix whose i -column, viewed as a vector, is the i th member of Y . Consider $N = M^T M$. N is a symmetric quadratic matrix of size m , with 1's on the diagonal. Reformulating the claim on the diameter of Y in terms of N , we arrive at the equivalent statement that one of the off-diagonal entries of N is less than or equal to $-1/(n + 1)$. To prove the latter statement, notice that, by the choice of Y , there exists a nonnegative nonzero vector w such that $Mw = 0$, and consequently $Nw = 0$. Let w_i be the largest entry of w , and let c be the i th row of N . Since $\langle c, w \rangle = 0$, and $c_i = 1$, it holds that

$$\sum_{j \neq i} (w_j/w_i) c_j = -1.$$

Observing that there are $m - 1$ off-diagonal elements in the row c , and that, for every j , $0 \leq w_j/w_i \leq 1$, one concludes that there must exist an off-diagonal entry of c which is at most $-1/(m - 1)$. However, $m \leq n + 2$, and the statement follows. \square

This completes the proof of Theorem 5.6. \square

Acknowledgments

We would like to thank Leizhen Cai for helpful discussions and numerous references. Yuri Rabinovich would also like to thank Nathan Linial for discussion of the initial questions that led to this paper.

References

1. I. Althöfer, G. Das, D. Dobkin, D. Joseph, and J. Soares, On generating sparse spanners for weighted graphs, *Discrete Comput. Geom.* **9** (1993), 81–100.
2. B. Awerbuch, Complexity of network synchronization, *J. Assoc. Comput. Mach.* **32** (1985), 804–823.
3. B. Bollobás, *Extremal Graph Theory*, Academic Press, London, 1978.
4. J. Bourgain, On Lipschitz embedding of finite metric spaces in Hilbert space, *Israel J. Math.* **52** (1985), 46–52.
5. P. J. Bushell, Hilbert's metric and positive contraction mappings in a Banach space, *Arch. Rational Mech. Anal.* **52** (1973), 330–338.
6. R. L. Graham and P. M. Winkler, On isometric embeddings of graphs, *Trans. Amer. Math. Soc.* **288** (1985), 527–536.
7. W. B. Johnson and J. Lindenstrauss, Extensions of Lipschitz mappings into a Hilbert space, *Contemp. Math.* **26** (1984), 189–206.
8. N. Linial, E. London, and Y. Rabinovich, The geometry of graphs and some of its algorithmic applications, *Combinatorica* **15** (1995), 215–245.

9. G. A. Margulis, Explicit group-theoretical constructions of combinatorial schemes and their application to the design of expanders and superconcentrators, *Problems of Information Transmission* **24** (1988), 39–46. (Translated from Russian, *Problemy Peredachi Informatsii* **24** (1988), 51–66.)
10. J. R. Munkres, *Elements of Algebraic Topology*, Addison-Wesley, Reading, MA, 1984.
11. D. Peleg, A. Schäffer, Graph spanners, *J. Graph Theory* **13** (1989), 99–116.
12. D. Peleg and J. Ullman, An optimal synchronizer for the hypercube, *SIAM J. Comput.* **18** (1989), 740–747.
13. I. M. Yaglom and V. G. Boltyanskii, *Convex Figures*, Holt, Rinehart, and Winston, New York, 1961. (Translated from Russian.)

Received September 26, 1995, and in revised form March 18, 1996.