The Complexity of Low-Distortion Embeddings Between Point Sets *

Christos Papadimitriou Shmuel Safra

Abstract

We prove that it is NP-hard to approximate by a ratio better than 3 the minimum distortion of a bijection between two given finite three-dimensional sets of points.

1 Introduction

Given a bijection f between two finite metrics of the same cardinality, the *expansion* of f is the maximum of $\frac{d(f(i),f(j))}{d(i,j)}$ over all pairs of points of the domain, and the *distortion* of f is the product of the expansion of f times the expansion of the inverse of f.

Embedding a metric space into another with low distortion has been over the past two decades a central and deep problem in Mathematics [13]. These ideas have been imported in Computer Science, where they have been used gloriously in approximation algorithms [3, 8], on-line algorithms [6, 10], learning [4], and many other applications (see these Computer Science-friendly surveys [12, 16, 18]). Metrics induced by low dimensional point sets have also been a staple in combinatorial optimization and computational geometry over decades, see, e.g., [17, 2]. And small-distortion embeddings between finite metrics have been recently proposed as an important algorithmic framework in Computer Vision [7].

Despite the widespread attention embeddings of bounded distortion have attracted in Computer Science, the literature concerns itself with worst-case bounds over all point sets; the issue is never raised about specific point sets (except in order to establish lower bounds). For example, it is quite astonishing that the computational problem of determining the minimum distortion of two finite metrics has been first formulated and attacked only very recently [14].

In their recent pioneering paper, [14] define this problem, and come up with a polynomial-time algorithmic solution, via dynamic programming, for the special case in which the point sets are one-dimensional, and the distortion is known to be at most $2+\sqrt{5}$; when the distortion is lower than this, the bijection f between two point sets on the real line must respect certain constraints on the relative ordering of the points, rendering the possible bijections computable inductively based on a finite number of indices, and resulting in a polynomial-time algorithm (with exponent 11). They go on to derive certain algorithms and hardness results for non-geometric metrics; they propose the complexity of minimum distortion in more general settings than those they solved as an important open question.

In this paper we prove the first complexity and inapproximability results for this problem. In particular we prove:

*Computer Science Division, UC Berkeley, CA 94720. {christos, safra} @cs.berkeley.edu. Research partially supported by NSF ITR Grant, and by a Miller Professorship.

Theorem 1 *For any fixed $\delta > 1$, it is NP-complete to decide whether two given 3-dimensional point sets have a bijection with expansion (or distortion) at most δ.*

In fact, the result holds even if δ is allowed to be a positive function of n, as long as it is less than an exponential.

Our reduction is from 3-COLORABILITY; since this problem is NP-complete even for planar graphs, it would appear that there is hope of extending our proof to two dimensions. Unfortunately, there are strong reasons why the rest of our construction is inherently 3-dimensional. The two-dimensional case remains a very challenging open problem.

As mentioned above, by $d(i, j)$ we mean throughout the Euclidean distance between two points; however, our construction can be modified to prove the same result for any L^p metric.

By an extension of our basic construction we can prove the following:

Theorem 2 *The minimum expansion and distortion of 3-dimensional point sets is NP-hard to approximate to within any factor smaller than* 3 *(and within any additive constant).*

Notice that our complexity result holds for any $\delta > 1$, however close to 1; the results in [14] suggest that the problem becomes harder when δ is large. Indeed, our proof of Theorem 2 leaves open the possibility that the optimum distortion may be approximated quite closely if it is small. Furthermore, it turns out that the special case $\delta = 1$ is a much easier problem, since it seeks a rigid isomorphism between the two pointsets:

Theorem 3 *The problem of deciding whether two d-dimensional point sets have a bijection with distortion 1 can be solved in polynomial time for any fixed dimension d, and is equivalent to* GRAPH ISOMORPHISM *when d is allowed to grow.*

Sketch: In d dimensions, it suffices to check all bijections between subsets of size $d + 1$, yielding a polynomial exhaustive algorithm; see [5, 1] for faster algorithms. For unbounded dimensions, the problem can be reduced to a version of GRAPH ISOMORPHISM with labelled edges (where the labels correspond to the distances between points, square roots of rational numbers); that version is then reduced to graph isomorphism, simulating labels by distinguishing attachments to the edges. For a reduction in the other direction, any graph can be embedded in $n - 1$ dimensions so that adjacent vertices are at a distance 1, while non-adjacent vertices are at a distance $1 + \epsilon$, say. $\square$

2 The Construction

We start from a cubic graph $G = (V, E)$ that may or may not be 3-colorable. We shall produce two sets of points in three dimensions with rational coordinates, S and T (called the *solver* and the *target* point set), such that there is an embedding from S to T of expansion $\leq \delta$ iff G is three colorable. (The result for distortion follows from the observation that S and T will be such that in any minimum expansion map from S to T, the expansion of the inverse map achieves its minimum value $\frac{\sqrt{3+y^2}}{\sqrt{1+y^2}}$, where y is a parameter of the construction defined below.)

The two sets will differ only slightly, so in the next few pages we shall describe S. We start by embedding G in three dimensions (we omit the straightforward but tedious proof of the necessary lemma).

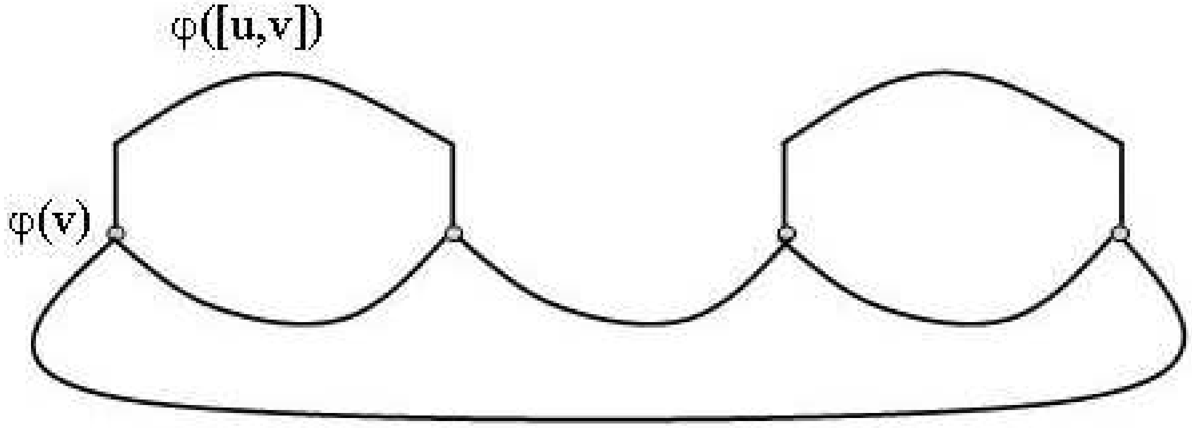

Figure 1: The embedding of Lemma 11

Lemma 1 *Given any $\delta > 1, \theta > 0$, and $\Delta = 4((\delta-1)^2 + 1/(\delta-1)^2)$, any cubic graph has an embedding ϕ in $\Re^3$ such that*

- *the vertices $\phi(v_1), \phi(v_2), \ldots$ are points on a straight line;*
- *each edge $[u,v]$ is mapped to a continuous piecewise linear curve $\phi([u,v])$ between $\phi(u)$ and $\phi(v)$, where the length of each segment is at least $1/theta$, and the angle between any two consecutive segments is at most θ;*
- *$\phi([u,v])$ has length at least 10Δ;*
- *$\phi([u,v])$ starts at u (and ends at v) as a straight line segment of length Δ; furthermore, the three segments emanating from v are co-planar and form $2\pi/3$ angles with each other;*
- *if a point of $\phi([u,v])$ is closer than Δ to a point of $\phi([w,z])$, then either $[u,v]$ and $[z,w]$ share an endpoint and the two points belong to the beginning segments, or $[u,v] = [w,z]$ and the points are within a distance of 2Δ of each other on the curve.*

At a very high level, S and T follow ϕ (see Figure 1). For each node $v \in V$ we have in the two point sets S and T the point $\phi(v)$. The intention of the reduction is that these points are mapped to one another in any low-distortion bijection. For small δ (say, less than 2), this is trivial, since the alignment of the $\phi(v)$'s is enough to guarantee this. For large values of δ, specialized tricks (attaching to these points uniquely identifiable configurations) can easily accomplish the same; we omit here the details.

Lemma 2 *In any bijection of expansion δ or less, and for any $v \in V$, the point $\phi(v)$ of S is mapped to the point $\phi(v)$ of T.*

Each curve $\phi([u,v])$ is fleshed out as a "prismoid," such as is shown in Figure 2.

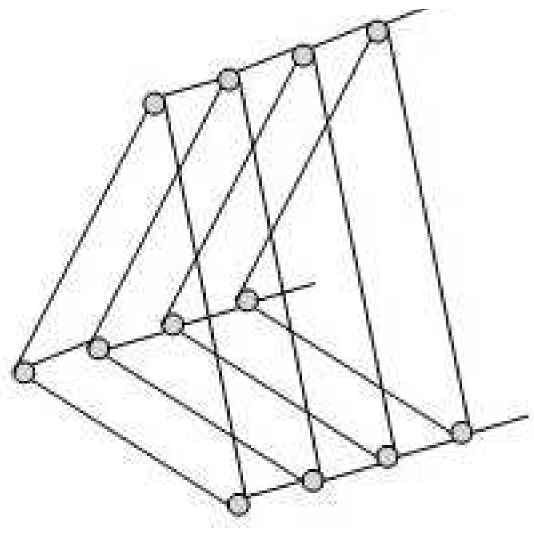

Figure 2: Prismoid

A prismoid is a sequence of sets of three points, each set forming an equilateral triangle.[1] The edge lengths of the triangles are one, and the plane of each equilateral triangle is normal to the curve $\phi([u,v])$ at points spaced z away from each other, starting at the end of the initial and final segments of $\phi([u,v])$, Δ away from $\phi(u)$ and $\phi(v)$. Near the derivative discontinuities of $\phi([u,v])$, the triangles are arranged so that the angles formed by the three rays with the plane of the two segments are preserved. By choosing θ small enough we ensure that for any two consecutive triangles the distances between corresponding points are within a δ factor of one another. The distance z between two consecutive triangles is chosen so that $\sqrt{1+z^{-2}} > 2\delta - 1$. This will imply that, in any low-distortion bijection, the points on the same column of a prismoid of S are all mapped to the corresponding points of some other column of the homologous prismoid of T.

The three prismoids corresponding to node $v \in V$ start therefore at a distance of Δ away from $\phi(v)$. The prismoid corresponding to edge $[u,v]$ is divided into two half-prismoids, one for each endpoint. The two half-prismoids of the same edge meet at an area in the middle, called the *verification region* (this is where the colors of the two vertices are verified to be different, and where S and T differ), to be described in detail soon.

Consider a half-prismoid, corresponding to an edge $[u,v]$ out of v. It has three columns of points. We call these three columns a, b, and c (intuitively, the colors will be encoded as permulations of these labels). At the v end of the prismoid there are certain geometric features (sets of points) that ensure consistency between the labels across the three prismoids emanating from v (so that v has one color), as well as the correct orientation of the labels (this issue will be explained later). Even though the prismoids end at a distance Δ away from $\phi(v)$, the columns do not. They continue individually, as sequences of points z apart, eventually diverging from each other (see Figure 3) and ending at three symmetrically placed points v_a, v_b, v_c, respectively. Column a first makes a turn that brings it close (closer than $1/\delta$) to column b and then continues to v_a; column b makes

[1]Our description will require the use of fractional and irrational numbers; in the final construction these will be approximated by rationals, and then scaled to integers. If δ is large, then any reasonable approximation would do; otherwise, an approximation within $(\delta-1)^2$ is needed.

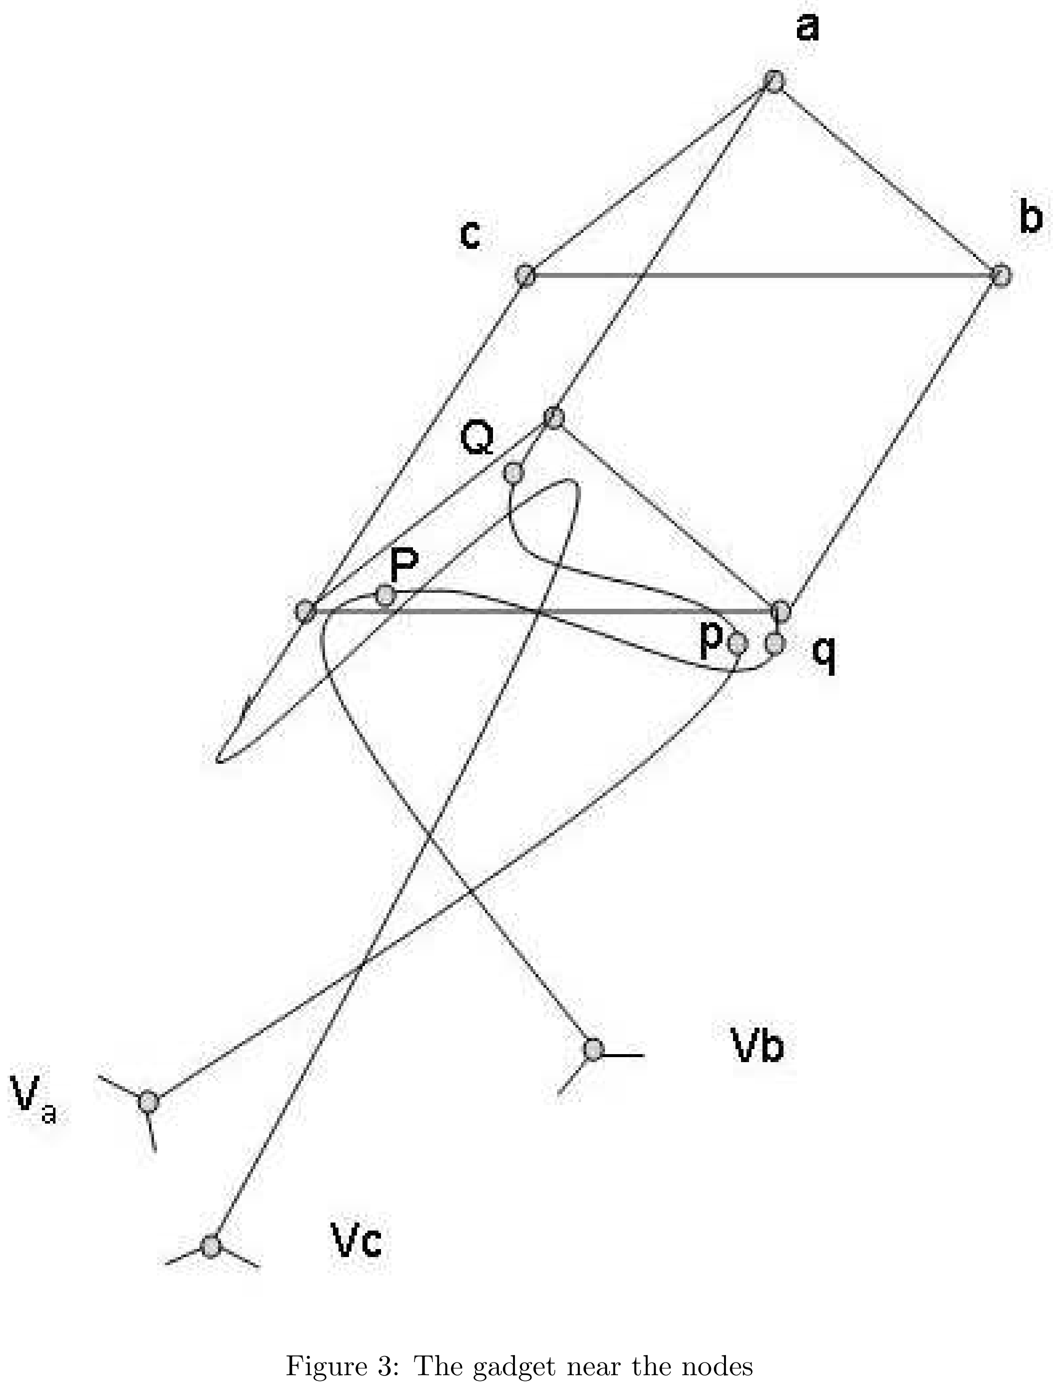

Figure 3: The gadget near the nodes

a turn that brings it close to column c before heading to v_b, and c gets close to a (see the figure). This is done by a smooth curve (we omit here details similar to those in Lemma 1) and with full rotational symmetry of $2\pi/3$.

The purpose of this is that, intuitively, the three labels are now cyclically ordered. Formally, consider the three cyclic permutations of $\{a, b, c\}$, $\chi_1 = (abc), \chi_2 = (bca), \chi_3 = (cab)$.

Lemma 3 *In any bijection between S and T of expansion $\leq \delta$, and for any vertex $v \in V$ and edge $[u, v] \in E$, there is an $i \leq 3$ such that the a column (resp., b, c column) of the half-prismoid corresponding to v and $[u, v]$ of S is mapped to the $\chi_i(a)$ (resp., $\chi_i(b)$, $\chi_i(c)$) column of the corresponding half-prismoid of T.*

Proof: By Lemma 2 the points $\phi(v)$ and $\phi(u)$ will be correctly mapped. This implies that the prismoid $\phi([u, v])$ will be mapped onto itself, triangle by triangle. Therefore, each of the three columns a, b, c in S will be mapped to a different one of the columns a, b, c in T.

It remains to show that the resulting permutation will be cyclic. Suppose it is not, that is, a is mapped to b and b to a. But this distorts the distance between the two points named p, q in Figure 3 beyond $1 + \delta$, mapping them to P, Q. □

There is one small complication we must deal with before describing the verification regions of the prismoids. The construction so far assigns to each prismoid an orientation (clockwise or couterclockwise) via the names of the columns of the prismoid. In order for these labels to be consistent at the two half-prismoids (and thus handled appropriately in the verification region) the orderings must be opposite at the two sides. This is accomplished by assigning a *direction* to each edge of G; that is, we think of G as a directed graph. Now, for each node v and each edge $[u, v]$ adjacent to it, if the edge leaves v then we orient the three labels clockwise, as in Figure 3, otherwise counterclockwise. (Notice that the above does not affect the construction, just our accounting of the labels.)

We are now ready to complete the construction by describing certain important details in the prismoids of S and T. For each edge of G, the verification region of the corresponding prismoid is simple: the two half-prismoids are simply separated by a gap of y that is longer than the gap z between two consecutive triangles (see the sketch pertaining to S in Figure 4). In particular, y must satisfy $\sqrt{\gamma} \leq y < \sqrt{2\gamma}$, and also $\sqrt{2\gamma - 1} \leq y$ when $\gamma \leq 1/2$, where $\gamma = \frac{1}{3((\delta-1)^2+2(\delta-1))}$.

In the verification region of T, however, the triangles of the two half-prismoids (still at a distance y apart) are "up-side down" (see the last two parts of Figure 4). This is accomplished by "slow rotation" of the triangles in one of the half-prismoids, by an angle that is small enough (at most $z \cdot (\delta - 1)$) so that it does not distort the distance between the points on each column by too much. The intuitive justification of the inverted triangles (a key element of our proof) is the following: Two inverted (*à la* star of David) triangles are "in between" two of the three possible alignments of the triangles. In this sense, they can be used to prohibit the third. Since colors are triangle rotations, an alignment of triangles is a difference between colors, and disallowing one value for this difference can be valuable in simulating 3-coloring.

This concludes the construction. The proof of Theorem 1 follows from this result:

Lemma 4 *There is a bijection between S and T of expansion δ or less if and only if G is three-colorable.*

Proof: Suppose that a 3-coloring of G exists; we describe a bijection f of disortion $\leq \delta$ from S to T as follows: We map all points on the prismoids of S to the homologous points of T; the only thing that remains to be specified is the mapping between individual columns of these prismoids.

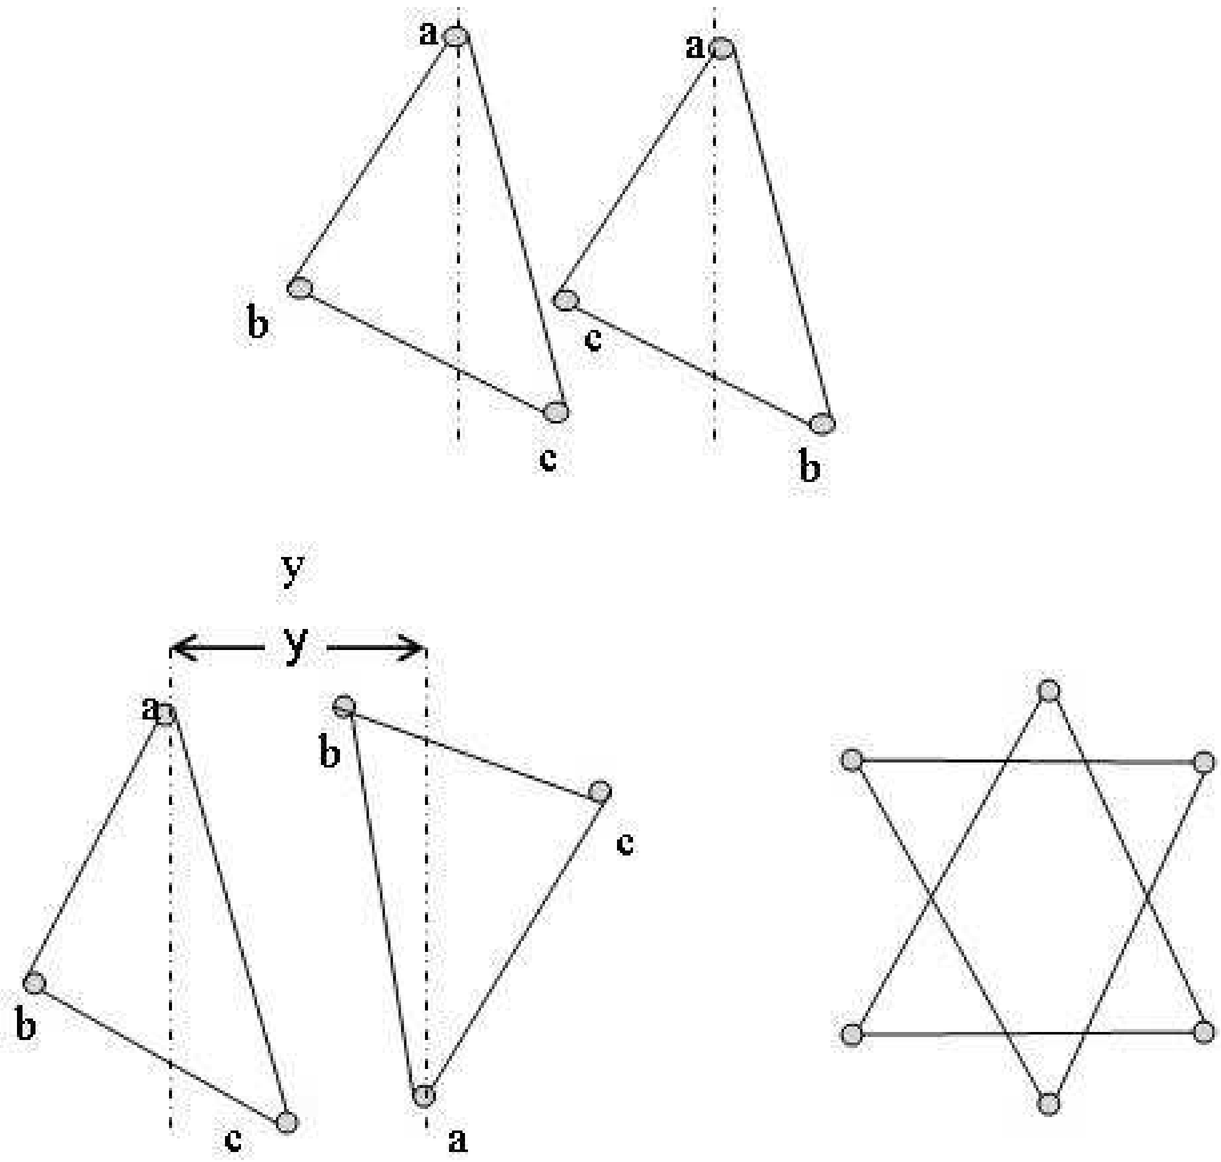

Figure 4: David

We define this map as follows: Take a node v of G, and suppose that, in the assumed coloring, its color is i, where $1 \leq i \leq 3$. Then in any half-prismoid out of $\phi(v)$, we map the a column to the $\chi_i(a)$ column, the b column to $\chi_i(b)$, and the c column to $\chi_i(c)$. This completes the definition of f. We must show that for any points $p, q \in S$, $\frac{d(p,q)}{\delta} \leq d(f(p), f(q)) \leq d(p,q)\delta$. Because the two sets are identical except for the verification regions (and the slight rotations of the triangles, which however by their definition do not distort the metric by more than δ), we need to only check the verification regions of all edges. Consider then the verification region of the set S corresponding to an edge $[u, v]$, where u has color i and v has color j, and consider two vertices p in the side of u and q on the side of v with labels $\ell(p), \ell(q) \in \{a, b, c\}$, respectively. The distance $d(p,q)$ is equal to y if $\ell(p) = \ell(q)$, and $\sqrt{1+y^2}$ otherwise. Now going to T, the distance $d(f(p), f(q))$ is either $\sqrt{y^2+4/3}$ if $\chi_i(\ell(p)) = \chi_j(\ell(q))$, or $\sqrt{y^2+1/3}$ if $\chi_i(\ell(p)) \neq \chi_j(\ell(q))$. However, since we assume a legal 3-coloring of G, $\ell(p) = \ell(q)$ implies $\chi_i(\ell(p)) \neq \chi_j(\ell(q))$, and hence $\frac{d(f(p),f(q))}{d(p,q)}$ is either $\frac{\sqrt{y^2+1/3}}{y}$ (if $\ell(p) = \ell(q)$) or $\frac{\sqrt{y^2+4/3}}{\sqrt{y^2+1}}$. It is easy to check that both ratios are between $\frac{1}{1+\delta}$ and $1+\delta$, and hence f is a low-expansion bijection.

Conversely, suppose that there is no legal coloring of G, and still there is a low-expansion bijection f. In order to have low expansion, f must map vertices to vertices and half-prismoids to half-prismoids, and it therefore implicitly chooses for each vertex a cyclic permutation of a, b, c. Since there is no 3-coloring of G, there must be an edge such that the two permutations are identical. Therefore, there are two points p and q at the verification region of that edge on the S side with the same label and with $d(p,q) = y$, mapped by f to two points with the same label, with $d(f(p), f(q)) = \sqrt{y^2+4/3}$. But for the given value of y, this has expansion greater than δ, proving the lemma and the theorem. □

3 Inapproximability

Our construction, with a little more care in the selection of z and y, immediately gives an inapproximability result: It is NP-hard to approximate distortion (and expansion) within a factor of $2-\epsilon$, for any $\epsilon > 0$ (and within any additive constant). To see why, let us just note that the "2" in this bound is the ratio between the longest to the shortest distance between two different points in a regular hexagon — equivalently, the "star of David" verification region with y taken in the limit to be zero.

It would appear that 2 is the limit of this construction: The star of David is the placement of two triples of points, call them a, b, c and A, B, C that maximizes the minimum of all distances of the form aA/aB —this is the key quantity for proving inapproximability.

But it turns out that the natural limit of our construction is 3, not 2. Here we outline the required extension.

First, we start with $2k+1$-colorability (our construction so far is the $k = 1$ special case). The embedding proceeds as before, except that the prismoids have now a cross-section that is a regular $2k+1$-gon. The consistency and orientation gadget is now a simple generalization of the $k = 1$ case.

Finally, for the verification regions of T: One of the two half-prismoids of each edge accomplishes, as before, the slow π rotation that creates an "up-side-down $2k+1$-gon, while the other is now divided into k segments. The interface between any two consecutive segments is very much like the verification region: the two regular $2k+1$-gons are placed as in Figure 4, in a half-shifted position relative to one another. This allows the current rotation of the $2k+1$ labels to shift by one on either side. By shifting appropriately at these $k-1$ interfaces, any alignment of the $2k+1$-gons from S can be matched — with the unique exception of the identical alignment (a against a etc.) that corresponds to the two nodes having the same color.

Now, in order for this alignment to be achieved, at one of the interfaces the rotation must be larger, at least as large as the $1-3$ distance shown in Figure 4. The ratio of this distance to $1-2$ is easily seen to be $\frac{\sin(3\pi/4k+2)}{\sin(\pi/4k+2)}$, which goes to 3 as k increases.

References

[1] H. Alt and K. Mehlhorn and H. Wagener and Emo Welzl. Congruence, similarity and symmetries of geometric objects, Discrete Comput. Geom., 3, pp. 237–256, 1988.

[2] Sanjeev Arora "Polynomial Time Approximation Schemes for Euclidean Traveling Salesman and other Geometric Problems," J. ACM 45(5), pp. 753-782, 1998.

[3] Sanjeev Arora, Satish Rao, Umesh Vazirani, "Expander flows and a sqrtlog n-approximation to sparsest cut," STOC 2004.

[4] R. I. Arriaga, Santosh Vempala. An Algorithmic Theory of Learning: Robust Concepts and Random Projection, FOCS 1999: 616-623

[5] M. D. Atkinson. An optimal algorithm for geometrical congruence. J. Algorithms, 8, pp. 159–172, 1987.

[6] Y. Bartal: On Approximating Arbitrary Metrices by Tree Metrics. STOC 1998: 161-168

[7] Serge Belongie, Jitendra Malik, Jan Puzicha "Shape Matching and Object Recognition Using Shape Contexts," IEEE Transactions on Pattern Analysis and Machine Intelligence 24(4), pp. 509-522, 2002.

[8] U. Feige: Approximating the Bandwidth via Volume Respecting Embeddings. J. Comput. Syst. Sci. 60(3): 510-539 (2000)

[9] James Haralambides, Fillia Makedon "Approximation Algorithms for the Bandwidth Minimization Problem for a Large Class of Trees" Theory Comput. Syst. 30(1), pp. 67-90, 1997.

[10] M. Herlihy,S. Tirthapura and R. Wattenhofer, Competitive concurrent distributed queuing, PODC 2001.

[11] P. Indyk and R. Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality. STOC 1998, pp. 604-613.

[12] P. Indyk. Algorithmic Applications of Low-Distortion Geometric Embeddings. Tutorial at FOCS 2001, pp. 10-33.

[13] W. B. Johnson and J. Lindenstrauss. Extensions of Lipschitz mappings into a hilbert space. In Conference in modern analysis and probability, pages 189–206, 1984.

[14] Claire Kenyon, Yuval Rabani, Alistair Sinclair "Low Distortion Maps Between Point Sets," STOC 2004, to appear.

[15] J. M. Kleinberg. Two algorithms for nearest-neighbor search in high dimensions. STOC 1997, pp. 599-608.

[16] J. Matousek *Lectures on Discrete Geometry,* Springer, 2002.

[17] Christos H. Papadimitriou: The Euclidean Traveling Salesman Problem is NP-Complete. Theor. Comput. Sci. 4(3), pp. 237-244, 1977.

[18] S. Vempala *The Random Projection Methods*, SIAM, 2004.