Inapproximability for planar embedding problems

Jeff Edmonds*   Anastasios Sidiropoulos†   Anastasios Zouzias†

October 11, 2009

Abstract

We consider the problem of computing a minimum-distortion bijection between two point-sets in $\mathbb{R}^2$. We prove the first non-trivial inapproximability result for this problem, for the case when the distortion is constant. More precisely, we show that there exist constants $0 < \alpha < \beta$, such that it is NP-hard to distinguish between spaces for which the distortion is either at most α, or at least β, under the Euclidean norm. This addresses a question of Kenyon, Rabani and Sinclair [KRS04], and extends a result due to Papadimitriou and Safra [PS05], who gave inapproximability for point-sets in $\mathbb{R}^3$.

We also apply similar ideas to the problem of computing a minimum-distortion embedding of a finite metric space into $\mathbb{R}^2$. We obtain an analogous inapproximability result under the ℓ_∞ norm for this problem. Inapproximability for the case of constant distortion was previously known only for dimension at least 3 [MS08].

1 Introduction

Over the recent years, metric embeddings have resulted in some of the most beautiful and powerful algorithmic techniques, with applications in many areas of computer science [LLR95, Ind01]. For two metric spaces (X, d), (Y, ρ), the key parameter that quantifies the usefulness of an embedding $f : X \to Y$ is the *distortion* which is defined to be

$$\mathsf{dist}(f) = \max_{x, y \in X} \frac{\rho(f(x), f(y))}{d(x, y)} \cdot \max_{x', y' \in X} \frac{d(x', y')}{\rho(f(x'), f(y'))}.$$

For many interesting classes of spaces the worst-case distortion of such a mapping can be very large. A way of bypassing this obstacle is the design of algorithms that approximate the minimum distortion. Several results of this kind have been recently obtained [KRS04, PS05, HP05, KS07, CMO+08, BDG+05, BCIS05, BCIS06, Edm07, BIS07, MS08]. They all fall in one of the following two categories:

- Given as an input two metric spaces of the same size, we are asked to approximate the minimum distortion *bijection* between them.

- Given one metric space, we are asked to approximate the minimum distortion *injection* into a fixed infinite metric space.

In almost all cases, either the ambient space of the point sets for the bijection problem, or the host space for the injection problem is $\mathbb{R}^d$ for a fixed d, equipped with some ℓ_p norm.

Although the bijection and injection problems have inherent differences, there are some striking similarities between the current approximability status of both. For the case $d = 1$ (i.e. the real line), both problems are known to be hard to approximate within polynomial factors [HP05, BCIS05], and there are exact algorithms when the distortion is sufficiently small [KRS04, BIR]. For $d \geq 3$ it is known that both problems are hard to approximate, even when the optimum distortion is constant [PS05, MS08].

The above discussion naturally leads to the following question:

> For $d = 2$, what is the approximability of the bijection and injection problems when the optimum distortion is constant?

We remark that for the case of bijections this question has been posed by Kenyon, Rabani and Sinclair [KRS04], and has also been emphasized by Papadimitriou and Safra [PS05], and by Khot and Saket [KS07]. We give the first non-trivial inapproximability bounds for this problem. More specifically, we prove the following two results.

THEOREM 1.1. (INAPPROXIMABILITY FOR BIJECTIONS) *It is NP-hard to decide whether the minimum distortion of a bijection between two finite point sets in $\mathbb{R}^2$ under the Euclidean norm is either at most $3.61 + \varepsilon$, or at least $4 - \varepsilon$, for any $\varepsilon > 0$.*

THEOREM 1.2. (INAPPROXIMABILITY FOR INJECTIONS) *It is NP-hard to decide whether the minimum distortion*

*York University, E-mail: jeff@cse.yorku.ca
†University of Toronto, E-mail: {tasoss,zouzias}@cs.toronto.edu

of an injection of a finite metric space into $\mathbb{R}^2$ under the ℓ_∞ norm is either at most 3, or at least $3.5 - \varepsilon$, for any $\varepsilon > 0$.

The approximability of both problems when the optimal distortion is arbitrarily close to 1 remains an important open problem.

1.1 Our techniques The common theme used in the proofs of both theorem 1.1 and 1.2 is as follows. We reduce from a Satisfiability problem. The main idea is to construct spaces for which there are essentially two possible low-distortion embeddings, locally. As is typical, this basic *binary* gadget is then used to encode necessary boolean constrains. We now elaborate on each particular reduction.

Bijections We build a binary gadget by constructing a pair of point sets such that locally there are (essentially) only two possible low-distortion bijections between them. This is very similar to the ideas used by Papadimitriou and Safra [PS05] to locally encode 3-Coloring. The extra difficulty in our case however is that we cannot avoid "crossings" between different gadgets. This problem does not appear in $\mathbb{R}^3$ since by appropriately "bending" the point sets one can easily avoid all such crossings. We bypass this obstacle by considering gadgets that have distances in two different scales. At places where two gadgets cross, one of them has distances in a small scale, while the other in a large scale. This allows us to show that the two crossing gadgets are sufficiently independent.

Injections Previous constructions of binary gadgets for the injection problem fall in one of two categories: In the first category, the lower bound on the distortion is obtained via a *volume argument* (e.g. [BDG$^+$05, BCIS06, BCIS05]). This means that the local volume growth of the input space is higher than growth of the host space, and therefore some distances have to expand. This technique however works only for super-constant distortion.

In the second category, the distortion bound is obtained by a *topological argument* (e.g. [Mat90, MS08]). The idea here is to obtain a metric space by discretizing a topological space. Then, topological non-embeddability can be used to prove that a pair of points has to contract, and therefore the distortion must be high. This method works even for constant distortion when $d \geq 3$, but due to the topological restrictions of the plane, it does not appear to be sufficient for our setting.

We overcome the limitations of the above techniques by designing gadgets that utilize both the volume and the topological argument. To the best of our knowledge, this is the first non-embeddability proof of this kind. In

the heart of our analysis is a characterization of low-distortion embeddings of d-grids into $\mathbb{R}^d$, what could be of independent interest.

1.2 Related work
Bijections The minimum-distortion bijection problem was introduced by Kenyon, Rabani and Sinclair [KRS04]. Among other results, they gave an exact polynomial-time algorithm for $d = 1$, provided that the distortion is less than $3 + 2\sqrt{2}$. This result was improved by Chandran et al. [CMO$^+$08] who gave an algorithm that works for distortion less than $5 + 2\sqrt{6}$. For $d = 3$, Papadimitriou and Safra [PS05] proved that the distortion is hard to approximate within a factor better than 3, even when the optimum is arbitrarily close to 1. Subsequently, it was shown by Hall and Papadimitriou [HP05] that even for $d = 1$, the problem is hard to approximate within polynomial factors, provided that the distortion is polynomially large. For bijections between general metric spaces, an $\Omega(\log^{1/4-\varepsilon} n)$-approximability result was obtained by Khot and Saket [KS07].

Injections For $d = 1$, Bădoiu et al. [BCIS05] showed that the problem is hard to approximate within polynomial factors, when the distortion is polynomial. They also gave an approximation algorithm when the input is a tree metric. A similar approximation result has also been obtained for unweighted graphs by Bădoiu et al. [BDG$^+$05].

For $d = 2$, it has been shown by Bădoiu et al. [BCIS06] that the problem is NP-hard. They also gave an approximation algorithm for ultrametrics.

For $d \geq 3$, Matoušek and Sidiropoulos [MS08] have shown that the problem is hard to approximate within polynomial factors, even when the distortion is constant. For $d = 2$ they obtained approximability within polynomial factors, when the distortion is polynomially large.

Edmonds [Edm07] proved that the problem under the ℓ_∞ norm is NP-hard for $d = 3$, and also gave a polynomial-time algorithm for $d = 2$.

For unbounded-dimensional spaces, it has been shown by Linial, London and Rabinovich [LLR95] that an optimal embedding into ℓ_2 can be computed via semidefinite programming. In contrast, deciding whether a metric can be isometrically embedded into ℓ_1, is NP-hard [DL97]. Lee, Naor and Peres [LNP06] have obtained a $O(1)$-approximation algorithm for embedding tree metrics into ℓ_p.

2 Inapproximability for bijections

In this section, we give the proof of theorem 1.1. We will reduce the problem one-in-three 3SAT to our

problem. One-in-three 3SAT is a variant of the 3-satisfiability problem. Like 3SAT, the input instance $\varphi = C_1 \wedge \ldots \wedge C_m$ is a collection of clauses, where each clause C_j is a collection of three literals, where each literals i either a variable χ_i or its negation. The *one-in-three 3SAT* problem is to determine whether there exists a truth assignment to the variables so that each clause has exactly one true literal. This problem is well known to be NP-complete [Sch78]. Given an instance $\varphi = C_1 \wedge \ldots \wedge C_m$ of one-in-three 3SAT, we construct an instance $\langle S, T \rangle$ of the bijection problem such that if φ is satisfiable, then there is a bijection between S and T with distortion at most α (Lemma 2.5) and if φ is not satisfiable (Lemma 2.4), then every bijection has distortion at least β, where $\beta > \alpha$. The constants $\alpha, \beta > 1$ will be specified later. Given an instance φ, we construct $\langle S, T \rangle$ as follows. As is typical, the instance $\langle S, T \rangle$ has a *variable gadget* for each variable χ_i in φ, a *clause gadget* for each clause C_j, and a *literal gadget* for each literal $\langle \chi_i, C_j \rangle$.

We begin with an informal description.

Variable Gadgets: The gadget for the variable χ_i consists of two touching *gears*, one corresponding to χ_i and the other to $\bar{\chi}_i$. A gear spinning clockwise corresponds to χ_i (respectively $\bar{\chi}_i$) being set to true, while spinning counter clockwise corresponds to it being set to false. As is the case with two touching gears, they are designed so that they must spin in opposite directions. This ensures that χ_i and $\bar{\chi}_i$ have opposite values.

Gears: A gear consists of a circle of points alternating one from S and next from T (and a few extra anchoring points). Lemma 2.2 proves that any low distortion bijection needs to either map each point in S to the next clockwise point in T or map each to the next counter clockwise point. This can be visualized as the gear spinning either clockwise or counter clockwise. Lemma 2.3 proves that where the gears meet the points in S must be mapped in the same direction. This can be visualized as the teeth of the two touching gears moving in the same direction, causing the gears themselves to move in opposite directions.

Chains: A chain consists of a line of points alternating one from S and the next from T. It both begins and ends with a point from T. One is referred to as the *front end* and the other as the *back end*. If each point in S is mapped to the next point in T towards the front end, then we say that the front end is *pushing*, otherwise it is *pulling*. Similar, with the back end. Note that one end is pushing iff the other end is pulling. As with a gear, Lemma 2.2 proves that for any low distortion bijection, one of these is the case.

Literal gadgets: For each appearance of a variable χ_i in some clause C_j, the gadget for the literal $\langle \chi_i, C_j \rangle$ consist of a single chain. If χ appears positively in C_j, then the middle of this chain connects with the gear corresponding to χ_i otherwise it connects to that for $\bar{\chi}_i$. The chain is oriented on the gear so that the front end is pushing iff the corresponding literal χ_i or $\bar{\chi}_i$ is true, which is also iff the back end is pulling.

Clause gadgets: The gadget for the clause $C_j = \chi_{\langle j,1 \rangle} \wedge \chi_{\langle j,2 \rangle} \wedge \chi_{\langle j,3 \rangle}$ has a spot where the front ends of its three literal gadgets meet and a spot where their three back ends meet. Their front ends meet by having their front end point, which is from T, be in fact the same point and having the three chains pointing towards it each separated by an angle of 120 degrees. We refer to this common T point as the *front end* of the clause gadget. Note that exactly one of the three chains has its front-most S point be mapped (pushed) to this middle T point. Hence, exactly one of the literals in the clause is true. In a complementary way, the three back ends of the chains also meet at 120 degrees. However, in this case, there is one extra point from S in the middle, which we refer to as the *back end* of the clause gadget. This center point can be mapped (pulled out) to the T point ending of exactly one of the three chains. Again, this corresponds to there being exactly one of the literals being true. Since the chain for each literal pushes at one of its ends iff it is pulling at the other, which of the three literals is true is consistent at the two ends of the clause gadget.

Crossings: The chains must lie in the plane and must travel from a variable gadget to the two spots for each clause gadget. If one did not want these chains to cross each other, one could imagine doing the reduction to planar one-in-three 3SAT, which is in fact NP-complete [MR06]. However, by looking at Figure 5, one can see that simply for one clause, there needs to be a crossing of the chains to get the three front ends and the three back ends to the correct spots in the clause gadget. Luckily, chains can cross without effecting each other. The one slowly has its points further and further apart and the other has its points closer and closer together until they can cross without conflict.

2.1 Gears Let $k = \Omega(1/\varepsilon^5)$ be an integer. A *gear gadget of size k* is a gadget g with points $\{p_0, \ldots, p_{2k-1}\}$ arranged in a circle so that consecutive points are distance one apart. The even indexed points are in S and the odd indexed points are in T. See figure 1. The gadget also has ten extra points x_i, y_i, z_i, v_i, and w_i as depicted in the figure. We denote by $S(g)$ and $T(g)$ the points of g in S and T, respectively.
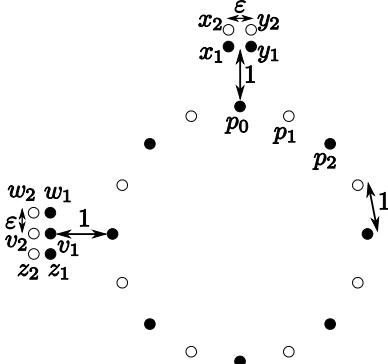
Figure 1: The gear gadget. Black dots denote points in S and white dots denote points in T.

LEMMA 2.1. *Let g be a gear gadget. Let $f : S(g) \to T(g)$ be a bijection with distortion less than $4 - o(1)$. Then, there exists a* spin $s \in \{-1, 1\}$, *such that for any $p_i \in S$, we have $f(p_i) = p_{i+s \mod 2k}$. We denote this s by* $\mathsf{spin}_f(g)$.

Proof. Initially ignoring the extra x, y, and z points, the circle of points from S could rotate any amount and flipped in order to be mapped to those in T. This would result to an isometry because the distance between any points p_i and p_j in S would be the same as the distance between the corresponding points $f(p_i)$ and $f(p_j)$ in T. We will argue now that this is the only way to map the points without creating $4 - o(1)$ distortion. By way of contradiction, consider a different mapping. Such a mapping would require two consecutive points p_l and p_{l+2} in S that are mapped to points $f(p_l)$ and $f(p_{l+1})$ in T that are distance more than two apart. In fact, the next option is that they are distance $4 - o(1)$. The $o(1)$ arises from the curvature of the circle, which is made small by making k large. This gives that $\mathsf{exp}(f) \geq \|f(p_{l+1}) - f(p_l)\|_2 / \|p_{l+1} - p_l\|_2 \geq \frac{4-o(1)}{2-o(1)} = 2 - o(1)$. Similarly, $\mathsf{exp}(f^{-1}) \geq 2 - o(1)$ and $\mathsf{dist}(f) = \mathsf{exp}(f) \times \mathsf{exp}(f^{-1}) \geq 4 - o(1)$. This contradicts the distortion being only 4 minus a smaller amount. Hence, the S circle can only be rotated by some amount and possibly flipped.

To ensure that it does not rotate by more than 1 or -1, we add the z, v and w points. To ensure that it does not flip, we add the x, and y points. Because z_1, v_1 and w_1 are so close together, they must be mapped close together or else the distortion would be way too big, i.e. $f(\{z_1, v_1, w_1\}) = \{z_2, v_2, w_2\}$. Similarly, $f(\{x_1, y_1\}) = \{x_2, y_2\}$.

We now argue that $f(p_0)$ is either p_1 or p_{2k-1}. This determines the spin $s \in \{1, -1\}$. Assume for the

sake of contradiction that $f(p_0) = p_t$, where $|t| > 1$. Then, $\mathsf{exp}(f) \geq \frac{\|f(p_0)-f(x_1)\|_2}{\|p_0-x_1\|_2} > \frac{\sqrt{3^2+1^2}}{1} = \sqrt{10}$. Also, $f^{-1}(p_1) \neq p_0$ means that $f^{-1}(p_1)$ at best is p_2. This gives that $\mathsf{exp}(f^{-1}) \geq \frac{\|f^{-1}(p_1)-f^{-1}(x_2)\|_2}{\|p_1-x_2\|_2} \geq \frac{\sqrt{2^2+1^2}}{\sqrt{1^2+1^2}} = \sqrt{5/2}$. Therefore the distortion of f is at least 5, a contradiction. By similar considerations we conclude that $f(p_{3k/2})$ is either $p_{3k/2+1}$ or $p_{3k/2-1}$. This ensures that the S circle does not flip before mapping to the T circle.

2.2 Chains A *chain* is similar to a gear except open. It is a path of points $\langle p_0, p_1, p_2, \ldots, p_{2k} \rangle$ with the odd indexed points in S and the even indexed points in T, beginning and ending with a point in T. We call the set of these points the *body* of the chain. Also there are *anchor* points x_1, x_2, y_1, y_2 placed nearest to the point $p_{k/3}$. We assume w.l.o.g. that $p_{k/3} \in S$. As with gears, for a chain g we denote by $S(g)$ and $T(g)$ the points of g in S and T, respectively. See figure 2.
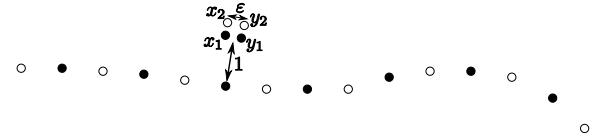

Figure 2: A chain.

An additional complication is that in order to allow different chains to cross each other, the points in a chain starting at distance 1 apart, need to be able to get further and further apart exponentially until there is a gap between them of some large size M and then get closer and closer together again until they are again distance 1 apart. More formally, we require that for some fixed and sufficiently large parameters $N \gg M > 0$, we have for any $i, j \in \{0, \ldots, 2k\}$:

- $\|p_i - p_{i+1}\|_2 \in [1, M]$, ensuring that the distances are within a reasonable range,

- With $|i - j| \leq N$, $\|p_i - p_j\|_2 \geq \frac{1}{1+\varepsilon} \sum_{t=1}^{j-1} \|p_t - p_{t+1}\|_2$, ensuring that the chain does not curve too quickly so that locally it looks like a straight path,

- With $|i - j| > N$, $\|p_i - p_j\|_2 > N$, ensuring that it does not circle back too close to itself (or to another gadget except where specified).

Finally, we say that some odd index $\gamma \in \{N, \ldots, 2k-N\}$ is a *gap* if $\|p_\gamma - p_{\gamma+1}\|_2 = M$.

- If $\gamma_1 \neq \gamma_2$ are gaps, then $|\gamma_2 - \gamma_1| > 2N$, ensuring that gaps are sufficiently far apart.

- For each $t \in \{1, \ldots, N\}$, we have

$$\|p_{\gamma+t} - p_{\gamma+t+1}\|_2 = \|p_{\gamma-t} - p_{\gamma-t+1}\|_2$$
$$= \max\left\{1, \frac{M}{2.2(1+\varepsilon)^{t-1}}\right\},$$

ensuring exponential decay on either side of the gap. The factor of 2.2 will become clear later in the analysis.

- For i such that $|\gamma - i| > N$ for all gaps γ, we have $\|p_i - p_{i+1}\|_2 = 1$, ensuring that distances return to one between gaps.

- Finally, for any gap γ we have $|\gamma - k/3| > N$, ensuring that all gaps are far enough from the anchor point.

The proof of the following lemma is very similar to the one for lemma 2.1, so we only sketch its proof. An additional complication is that we can only enforce a consistent motion on the chain for points that are sufficiently far apart from the gaps.

LEMMA 2.2. *Let g be a chain gadget, and let $f : S(g) \to T(g)$ be a mapping with distortion less than $4 - o(1)$. Then, there exists a spin $s \in \{-1, 1\}$ satisfying the following property. For any $p_i \in S$ such that $\gamma - 10 \le i < \gamma + 10$ for all gaps γ, we have $f(p_i) = p_{i+s}$. We denote this s by $\mathsf{spin}_f(g)$. (Chains actually push or pull, but we keep the notation of spin to be consistent with gears.)*

Proof. [Proof sketch] Let $\gamma_1, \ldots, \gamma_t$ be the gaps of the chain. Let $c_0, c_1, \ldots, c_t$ be the sequences of points obtained after removing points p_j with $\gamma_l - 10 \le j < \gamma_l + 10$ for some gap γ_l.

Since f has distortion less than $4 - o(1)$, we can argue as in the case for gears (lemma 2.1) that for every c_i, consecutive points in S should be mapped to consecutive points in T. Also observe that since each c_i is sufficiently long, the mapping f cannot flip c_i. Therefore, for every c_i there exists $s_i \in \mathbb{Z}$ such that for every $p_j \in c_i$ we have $f(p_j) = p_{j+s}$.

We next argue that all the s_i must be equal. For each $i \in \{0, \ldots, t-1\}$ let c_i' be the set of points between c_i and c_{i+1}. These are exactly the points that are near the gap γ_i. Since the distortion is small, the only points that can be mapped in $T(c_i')$ are those from $S(c_i) \cup S(c_i') \cup S(c_{i+1})$. By counting we therefore have that the number of points from $S(c_i)$ that are mapped to $T(c_i')$ must be equal to the number of points from $S(c_i')$ that are mapped to $T(c_{i+1})$. Thus, $s_i = s_{i+1}$ for any i, implying that all s_i are equal.

Since the anchor point is close to some c_{i^*}, we deduce that $s_{i^*} \in \{-1, 1\}$, finishing the proof.

2.3 Crossing chains We will now describe how two chains c_1 and c_2 can cross while keeping their spins $s \in \{-1, 1\}$ independent. Let $p_0, \ldots, p_{2k}$ and $p_0', \ldots, p_{2k'}'$ be the points in the body of c_1 and c_2, respectively. Let Γ_1, Γ_2 be the polygonal curves obtained by connecting consecutive points in the body of each one of c_1 and c_2. Then, we say that the chains c_1 and c_2 *cross* if Γ_1 and Γ_2 intersect at a point p^* such that (figure 3)

- p^* lies at the middle of a gap segment $\langle p_i, p_{i+1} \rangle$, with $\|p_i - p_{i+1}\|_2 = M$,

- p^* lies at the middle of a non-gap segment $\langle p_j', p_{j+1}' \rangle$, with $\|p_j' - p_{j+1}'\|_2 = 1$, and

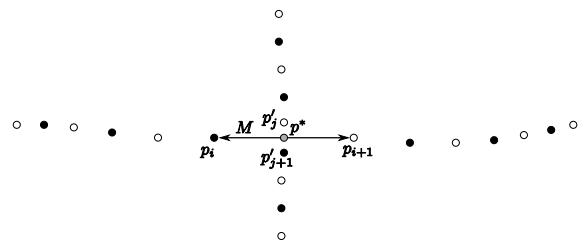- the segments $\langle p_i, p_{i+1} \rangle$ and $\langle p_j', p_{j+1}' \rangle$ are orthogonal.

Figure 3: Crossing chains.

2.4 Connecting gears and chains In our construction, a variable gadget consists of two gear gadgets that are connected so that they have opposite spins. Also the literal gadget for $\langle \chi_i, C_j \rangle$ consists of a chain which must connect to the gear gadget for the literal χ_i or $\bar{\chi}_i$ so that it pushes or pulls depending on the spin of the gear.

We now describe a way to arrange a gear gadget g_1 and either a chain or a gear gadget g_2 so that in any embedding f with distortion less than $4 - o(1)$, they must have opposite spins $\mathsf{spin}_f(g_1) = -\mathsf{spin}_f(g_2)$. For the case of two gears we simply place the centers of g_1 and g_2 at distance $r_1 + r_2 + \varepsilon$, where $r_1, r_2 = \Omega(1/\varepsilon^5)$ are their respective radii. We also rotate them so that they meet at a point from S. See figure 4. For the case of a gear and a chain we apply the same construction by bending the chain so that locally it looks like a circular arc of appropriate radius. We also make sure that all pairs of consecutive points of the chain on this circular arc are at distance 1. In both cases we say that g_1 and g_2 are *connected*.

LEMMA 2.3. *Let g_1, g_2 be two gear/chain gadgets that are connected as above and let $f : S(g_1) \cup S(g_2) \to$*
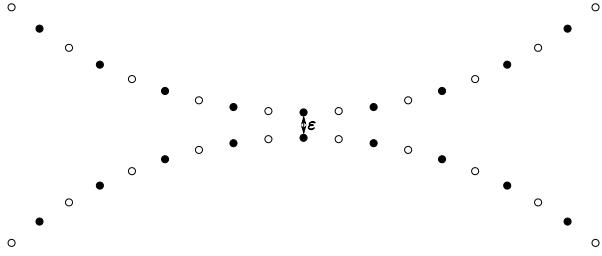
Figure 4: A connection.

$T(g_1) \cup T(g_2)$ be a mapping with distortion at most $4 - o(1)$. Then, $\mathsf{spin}_f(g_1) = -\mathsf{spin}_f(g_2)$.

Proof. Suppose $\mathsf{spin}_f(g_1) \neq -\mathsf{spin}_f(g_2)$. Let $p \in S(g_1)$, $p' \in S(g_2)$ be the closest pair of points between g_1 and g_2. We have $\mathsf{exp}(f) \geq \|f(p) - f(p')\|_2 / \|p - p'\|_2 = \Omega(1/\varepsilon)$. It is also clear that $\mathsf{exp}(f^{-1}) = \Omega(1)$, a contradiction.

2.5 Putting everything together We are now ready to describe how all the different pieces of the reduction can be combined to construct the instance to the embedding problem. For every variable we introduce a copy of the variable gadget. Similarly, for every clause we introduce its front end in T and its back end in S. We make sure that all these gadgets are sufficiently far from each other.

Finally, for every literal $\langle \chi_i, C_j \rangle$ we have a chain $c_{\langle i,j \rangle}$ that starts from the front end of C_j and ends at the back end of C_j. More precisely, the front end of the chain is the same as the one of C_j, and its back end points towards the back end of C_j (figure 5). The three front ends of the chains of a clause meet at 120 degree angles, and similarly for the back ends. The exact configuration is given in figures **??** and **??**. The chain also bends appropriately so that it connects to the appropriate gear of the variable gadget for χ_i. Whenever two chains need to cross, we simply introduce a crossing as explained in section 2.3.

Figure 5 gives a complete example for the clause $(\chi_1 \vee \bar{\chi_2} \vee \chi_3)$. Note that in the figure, the chain corresponding to the literal χ_3 is actually connected to the gear gadget for the literal $\bar{\chi_3}$. This is done only to reduce the number of crossings, and make the figure easier to read. The resulting construction is still correct because the chain is connected to the gear with an opposite (i.e. clockwise instead of anti-clockwise) orientation.

2.6 Sufficiently gadget-respecting maps The previous lemmas all assume that the maps $f : S \to T$

are *gadget-respecting*, i.e. the points in S from one gadget are only mapped to points in T from the same gadget. We will not be able to prove that this is completely true. However, we will prove that low distortion bijections are *sufficiently* gadget-respecting for these previous lemmas to still hold.

We don't have a problem for pairs of gadgets that are sufficiently far apart, because we cannot have the S points of one gadget mapped to the T points of more than one far apart gadgets. Also the gadgets will be of different sizes k and hence the S points of one must be mapped to the correct one in T.

The difficulty arises when two gadgets g_1 and g_2 connect. To ensure that our bijections are gadget-respecting with respect to such pairs, we will use the same trick as we did with the anchors x_1 and x_2. We assign these gadgets two different *duplication integers* $d(g_1)$ and $d(g_2)$. Then each point in the gadget g_i is replaced by a *cloud* of $d(g_i)$ points that have ε^2 distance between them. Any low distortion bijection cannot map two S points in the same cloud to T points in different clouds or to points in a cloud with a different duplication number $d(g_i)$.

To ensure that all low distortion bijections are (sufficiently) gadget-respecting, it is sufficient to set the duplication numbers as follows. For each variable χ, its gadget consists of two connecting gears. Give one $d(g_i) = 1$ and the other $d(g_i) = 2$. These same duplication numbers can be used for every variable because the variable gadgets are far apart. For all chain gadgets set $d(g_i) = 3$. This ensures gadget respectiveness between the gears and the chains. Chains for literal gadgets from different clauses are far enough apart so that they do not interact with each other.

The only places remaining where gadgets with the same duplication number are close enough to be non-gadget-respecting is at chain crossings and at the ends of the three chains for literals within the same clause. Observe that near the point where two chains cross, consecutive points are within distance 1 in one of the chains, and within distance $M \gg 1$ in the other. Similarly with the case of anchors, this ensures that the map is sufficiently gadget-respecting around the crossing point.

It remains to deal with the front and back ends of the three chains for literals within the same clause that meet at two spots in the clause gadget. These three chains cannot have different duplication numbers because the middle point where the three meet does in fact need to be able to map to any of the three chains. We argue that any non-gadget-respectiveness at the ends of the chains is sufficiently localized not to effect the proof of Lemma 2.2 stating that the spin of
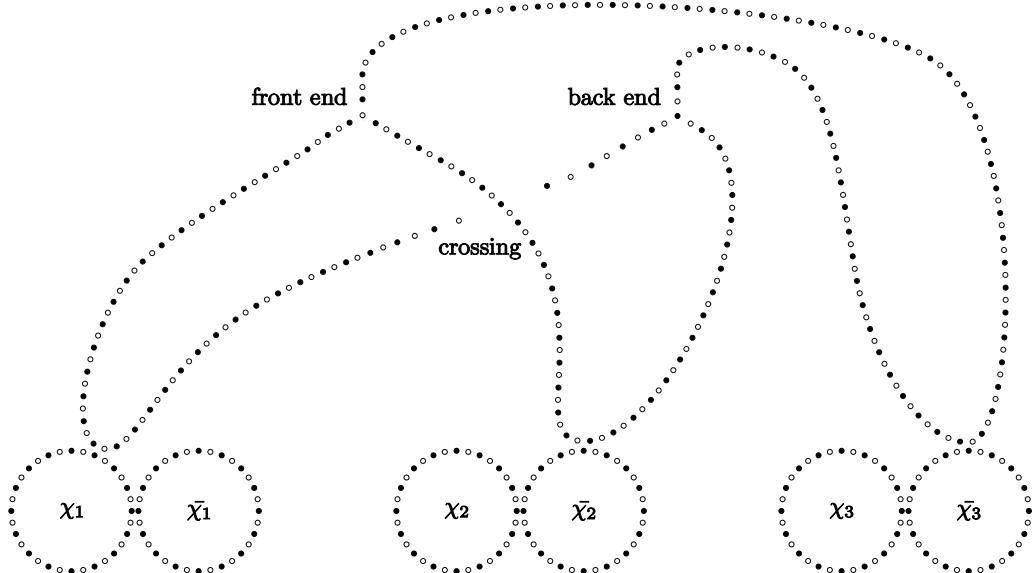
Figure 5: A clause gadget connected to the variable gadgets. The anchor points are omitted.

the chain (at least in the middle of the chain away from these connecting spots) is either 1 or -1.

2.7 No instances

LEMMA 2.4. *Let $\langle S, T \rangle$ be as described for the instance φ. Let $f : S \to T$ be a bijection with distortion less than $4 - o(1)$. Then from f, we can construct a satisfying assignment for φ.*

Proof. Given such a bijection f, Lemma 2.1 proves that each gear has a spin in $\{-1, 1\}$. Hence, it is well defined to construct an assignment by setting χ_i to be true iff the first gear in the corresponding variable gadget has clockwise spin 1. Lemma 2.2 proves each chain has a spin in $\{-1, 1\}$ and Lemmas 2.3 proves connecting gears/chains have opposite spins. The discussion during the first part of the reduction in Section 2 then argues that the front ends of literal gadgets push iff the corresponding literals are satisfied. It goes on to argue that for each clause, at exactly one of its three literal gadgets pushes. This would complete the proof that the constructed assignment satisfies φ.

The only difficulty is that as the three literal gadgets converge at this spot in the clause gadget, the bijection may cease to be gadget-respecting. We solve this problem by drawing a circle around each of the meeting spots sufficiently big that the literal gadgets have become far enough apart by the time they leave the circle that we can be assured that the bijection is gadget-respecting outside these circles and from one circle to the next. We do not need to worry at all about

how S points in one circle get mapped to T points in the same circle. Recall that the front ends of the literal gadgets for the three literals in a clause meet by having their front most point, which is from T, be in fact the same point. We can also set the circle to intersect each literal gadget at a segment $\langle p_i, p_{i+1} \rangle$ with $p_i \in T$ and $p_{i+1} \in S$. Together these mean that the number of T points in the circle is two fewer than the number of S points. A simple counting argument then gives that two of the three literal gadgets must have a point from S mapped from inside the circle to a point outside the circle. This must be a mapping from $p_{i+1} \in S$ to $p_i \in T$. By definition, this is pulling, not pushing. This proves that exactly one of the three literal gadgets is pushing. Hence, exactly one of the literals in the clause is true.

2.8 Yes instances

LEMMA 2.5. *Let $\langle S, T \rangle$ be as described for the instance φ. Given a satisfying assignment for φ, we can construct a bijection $f : S \to T$ with distortion at most $3.61 + o(1)$.*

The proof is deferred to the full version of this paper.

References

[BCIS05] M. Bǎdoiu, J. Chuzhoy, P. Indyk, and A. Sidiropoulos. Low-distortion embeddings of general metrics into the line. In *Proc. 37th ACM Symposium*

on Theory of Computing, 2005. Full version available at `http://www.mit.edu/~tasos/papers.html`. (Cited on pages 1 and 2)

[BCIS06] M. Bădoiu, J. Chuzhoy, P. Indyk, and A. Sidiropoulos. Embedding ultrametrics into low-dimensional spaces. In *22nd Annual ACM Symposium on Computational Geometry*, 2006. (Cited on pages 1 and 2)

[BDG$^+$05] M. Bădoiu, K. Dhamdhere, A. Gupta, Y. Rabinovich, H. Raecke, R. Ravi, and A. Sidiropoulos. Approximation algorithms for low-distortion embeddings into low-dimensional spaces. In *Proc. 16th ACM-SIAM Symposium on Discrete Algorithms*, 2005. (Cited on pages 1 and 2)

[BIR] M. Bădoiu, , P. Indyk, and Y. Rabinovich. Embedding into the line. Manuscript. (Cited on page 1)

[BIS07] M. Bădoiu, , P. Indyk, and A. Sidiropoulos. Approximation algorithms for embedding general metrics into trees. In *Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms*, 2007. (Cited on page 1)

[CMO$^+$08] N. Chandran, R. Moriarty, R. Ostrovsky, O. Pandey, M. Safari, and A. Sahai. Improved algorithms for optimal embeddings. *ACM Trans. Algorithms*, 4(4):1–14, 2008. (Cited on pages 1 and 2)

[DL97] M. M. Deza and M. Laurent. *Geometry of Cuts and Metrics*. Algorithms and Combinatorics 15. Springer-Verlag, Berlin etc., 1997. (Cited on page 2)

[Edm07] J. Edmonds. Embedding into ℓ_∞^2 is easy embedding into ℓ_∞^3 is NP-complete. In *Proc. 18th ACM-SIAM Symposium on Discrete Algorithms*, 2007. (Cited on pages 1 and 2)

[HP05] A. Hall and C. H. Papadimitriou. Approximating the distortion. In *APPROX-RANDOM*, 2005. (Cited on pages 1 and 2)

[Ind01] P. Indyk. Algorithmic applications of low-distortion embeddings. In *Proc. 42nd IEEE Symposium on Foundations of Computer Science*, 2001. (Cited on page 1)

[KRS04] C. Kenyon, Y. Rabani, and A. Sinclair. Low distortion maps between point sets. *Proceedings of the Symposium on Theory of Computing*, 2004. (Cited on pages 1 and 2)

[KS07] S. Khot and R. Saket. Hardness of embedding metric spaces of equal size. In *APPROX '07/RANDOM '07: Proceedings of the 10th International Workshop on Approximation and the 11th International Workshop on Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 218–227, 2007. (Cited on pages 1 and 2)

[LLR95] N. Linial, E. London, and Yu. Rabinovich. The geometry of graphs and some its algorithmic applications. *Combinatorica*, 15:215–245, 1995. (Cited on pages 1 and 2)

[LNP06] J. R. Lee, A. Naor, and Y. Peres. Trees and Markov convexity. In *Proc. 17th ACM-SIAM Symposium on Discrete Algorithms*, 2006. (Cited on page 2)

[Mat90] J. Matoušek. Bi-Lipschitz embeddings into low-dimensional Euclidean spaces. *Comment. Math. Univ. Carolinae*, 31:589–600, 1990. (Cited on page 2)

[MR06] W. Mulzer and G. Rote. Minimum weight triangulation is NP-hard. In *SCG '06: Proceedings of the twenty-second annual symposium on Computational geometry*, pages 1–10, New York, NY, USA, 2006. ACM. (Cited on page 3)

[MS08] J. Matoušek and A. Sidiropoulos. Inapproximability for metric embeddings into $\mathbb{R}^d$. In *Proceedings of 49th Annual IEEE Symposium on Foundations of Computer Science*, 2008. (Cited on pages 1, 2 and 11)

[PS05] C. Papadimitriou and S. Safra. The complexity of low-distortion embeddings between point sets. *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pages 112–118, 2005. (Cited on pages 1 and 2)

[Sch78] T. Schaefer. The complexity of satisfiability problems. In *Proc. 10th Annual ACM Symposium on Theory of Computing*, pages 216–226, 1978. (Cited on page 3)

A Inapproximability for injections

In this section we give the proof of theorem 1.2. We will reduce the problem Planar-3SAT to our problem. Recall that an instance of Planar-3SAT is a CNF formula $\varphi = C_1 \wedge \ldots \wedge C_m$, on n variables $\chi_1, \ldots, \chi_n$, such that the incidence bipartite graph for φ is planar. Let G_φ be the incidence graph of φ, with $V(G_\varphi) = \{\chi_1, \ldots, \chi_n, C_1, \ldots, C_m\}$, and such that $\{\chi_i, C_j\} \in E(G_\varphi)$ iff the variable χ_i appears in the clause C_j. Given such a formula φ, we will construct a finite metric space M_φ. If φ is satisfiable, we will produce an embedding of M_φ into ℓ_∞^2 with distortion 3. If φ is not satisfiable, we show that every such embedding has distortion of at least $7/2 - \varepsilon$, for $\varepsilon > 0$ that can be made arbitrarily small.

The metric space M_φ will be produced in stages by switching between an *embedding* representation, a *graph* representation, and the standard *set of distances* representation. We start by defining within $\mathbb{R}^2$ an $L \times L$ checker board, where each $\langle u, v \rangle \in [L]^2$ indexes a radius one (i.e. 2×2) square of the board. (It is the domino pieces not the board that is coloured black and white). Over this board, a *fine grid* of points $X \subseteq \mathbb{R}^2$ is laid. Note that X is an *embedding* of these points in that it specifies a location for each of them. Next we define a matrix $H_\varphi \in \{0,1\}^{L \times L}$. When $H_\varphi[u, v] = 1$, then the 1×1 square of grid is cut out of the center of the 2×2 square indexed by $\langle u, v \rangle$. Denote the new set of embedded points by $X - H_\varphi$. From this embedding, we can define the complete graph $G(X - H_\varphi)$ which has a vertex for each point in $X - H_\varphi$ such that the *length* of each edge is given by the ℓ_∞^2 distance between them in this embedding. Then we will change this graph into

$G(X - H_\varphi) + V_\varphi + E_\varphi$ by adding some new vertices V_φ and some new edges E_φ. Some of these new edges will be between already existing vertices from $X - H_\varphi$. In this case, the new edges are shorter than the original ones. Finally, $M_\varphi = M(G(X - H_\varphi) + V_\varphi + E_\varphi)$ will be the metric space whose points consist of the vertices of the graph $G(X - H_\varphi) + V_\varphi + E_\varphi$ and whose distances are given by the length of the shortest-path between these vertices in this graph. We now repeat these steps a little more formally.

A.1 The grid The initial fine grid is

$$X = \{0, \varepsilon, 2\varepsilon, ..., 2L - \varepsilon, 2L\}^2.$$

where ε is the distance between each point and its four nearest neighbors. For pair of points $p, q \in \mathbb{R}^2$, we denote by $X[p, q]$ the rectangular part of X between p and q, by $X^\circ[p, q]$ its *interior*, and by $X'[p, q]$ its *boundary*, namely

$$X[p, q] = \{x \in X : p_1 \le x_1 \le q_1 \text{ and } p_2 \le x_2 \le q_2\}$$

$$X^\circ[p, q] = X[\langle p_1 + \varepsilon, p_2 + \varepsilon \rangle, \langle q_1 - \varepsilon, q_2 - \varepsilon \rangle]$$

$$X'[p, q] = X[p, q] \setminus X^\circ[p, q]$$

For example, for each $\langle u, v \rangle \in [L]^2$, the checker board with center $\langle u, v \rangle$ is covered by the grid $X[\langle u - 1, v - 1 \rangle, \langle u + 1, v + 1 \rangle]$. When $H_\varphi[u, v] = 1$, the interior $X^\circ[\langle u - \frac{1}{2}, v - \frac{1}{2} \rangle, \langle u + \frac{1}{2}, v + \frac{1}{2} \rangle]$ of the 1×1 square of grid in its center is cut out. Cutting out all of these center square gives

$$X - H_\varphi = X - \bigcup_{\langle u, v \rangle \in H_\varphi^{-1}(1)} X^\circ[\langle u - \frac{1}{2}, v - \frac{1}{2} \rangle, \langle u + \frac{1}{2}, v + \frac{1}{2} \rangle]$$

Let

$$X'(u, v) = X'[\langle u - \frac{1}{2}, v - \frac{1}{2} \rangle, \langle u + \frac{1}{2}, v + \frac{1}{2} \rangle]$$

denote the rim of the hole within the $\langle u, v \rangle^{th}$ square of the checker board.

A.2 The domino gadget A *domino gadget* $K_{\langle u, v \rangle}$ is defined to consist of an adjacent pair of squares $\langle u, v \rangle$ and $\langle u + 2, v \rangle$ of the checker board, both of whose center square has been cut out. See Figure 6(a). To this, we add a new vertex $x_{\langle u, v \rangle} \in V_\varphi$ and an edge in E_φ of length $1\frac{1}{2}$ between this node and every point in the two hole rims $X'(u, v) \cup X'(u + 2, v)$. A useful image is that embedding the metric space M_φ requires placing a *domino piece* on top of this domino gadget. A domino piece consists of one black square and one white square. Placing it with the black square on the left corresponds

to embedding the extra vertex $x_{\langle u, v \rangle}$ in the center of the left hole, while placing it with the black square on the right corresponds to embedding $x_{\langle u, v \rangle}$ in the right hole. We will prove that every lower distortion embedding does one of these two acts. Being forced to choose between these binary options will correspond to setting the variable x_i to zero or one.

A.3 The clause gadget For every clause C_j, we define a *clause gadget* $\mathcal{C}_j$ which is similar to the domino gadget, except that it uses three squares of the checker board arranged like an 'L' as in figure 6(b). Let $\mathcal{C}_j^1, \mathcal{C}_j^2, \mathcal{C}_j^3$ denote its three hole rims and denote the single additional vertex by y_j.

A.4 The variable gadget For every variable χ_i in φ, we will have a *variable gadget* $\mathcal{V}_i$. Assume that χ_i appears (either as a positive or a negative literal) in n_i clauses in φ. The gadget $\mathcal{V}_i$ consists of $2n_i$ copies of the domino gadget arranged in a rectangular line. Corners are turned by having the adjacent pair of consecutive domino gadgets share a vertex. See Figure 6(d).

We connect every pair of consecutive domino gadgets, for example $K_{\langle u, v \rangle}$ and $K_{\langle u+4, v \rangle}$, by adding an edge in E_φ of length 2 between every vertex in the right hole rim $X'(u + 2, v)$ of the left $K_{\langle u, v \rangle}$ domino gadget and every vertex in the left hole rim $X'(u + 4, v)$ of the right $K_{\langle u+4, v \rangle}$ domino gadget. The effect of these edges is to force that the dominoes placed on these two adjacent gadgets cannot have their black squares touching, i.e. both black on the left, both on the right, or pointing outward.

A.5 The literal gadget Recall that G_φ is the incidence graph of φ. Fix a planar drawing of G_φ into $\mathbb{R}^2$. Consider a variable χ_i. We order the clauses in which χ_i appears by $C_{\langle i, 0 \rangle}, \dots, C_{\langle i, n_i - 1 \rangle}$, so that for every $j : 0 \le j < n_i$, the edge $\{\chi_i, C_{\langle i, j \rangle}\}$ appears immediately before $\{\chi_i, C_{\langle i, j+1 \mod n_i \rangle}\}$ in a clockwise ordering of the edges incident to χ_i in the drawing of G_φ.

For every variable-clause $\langle \chi_i, C_j \rangle$ appearing together in φ, we introduce a *literal gadget* $\mathcal{L}_{\langle i, j \rangle}$ consisting of a path of dominoes gadgets from the variable gadget $\mathcal{V}_i$ to the clause gadget $\mathcal{C}_j$. See Figure 6(c). This path $\mathcal{L}_{\langle i, j \rangle}$ connects perpendicularly to the variable gadget $\mathcal{V}_i$ at its $2j - 1^{st}$ domino gadget. If χ_i appears in C_j as a positive literal, then it is connected to clockwise most square in this domino gadget otherwise the counter clockwise most one. Figure 6(d) depicts an example of such a construction. The three path $\mathcal{L}_{\langle i, j \rangle}$ from the three literals in a clause C_j connect with the clause gadgets $\mathcal{C}_j$ as depicted in Figure 6(e). Just as done in the variable gadgets, we add an edge in E_φ of length 2
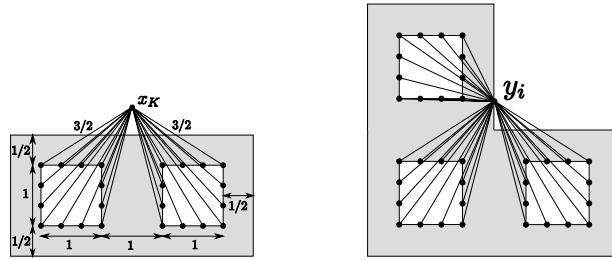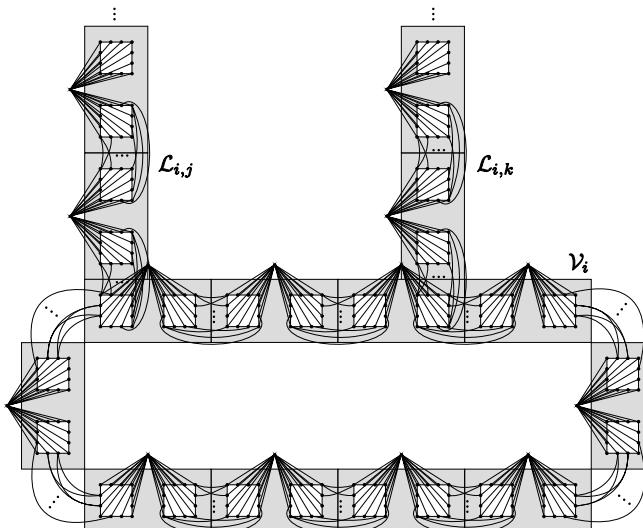
(a) The domino gadget. The shaded area corresponds to points from $X - H_\varphi$.
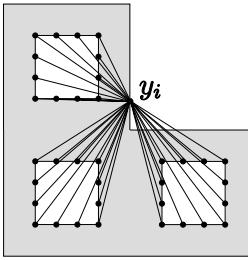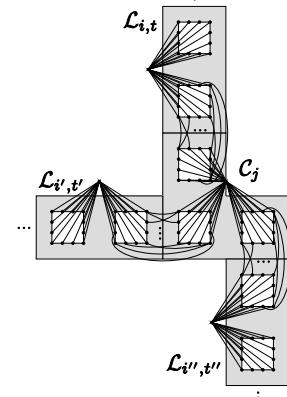
(b) The clause gadget.
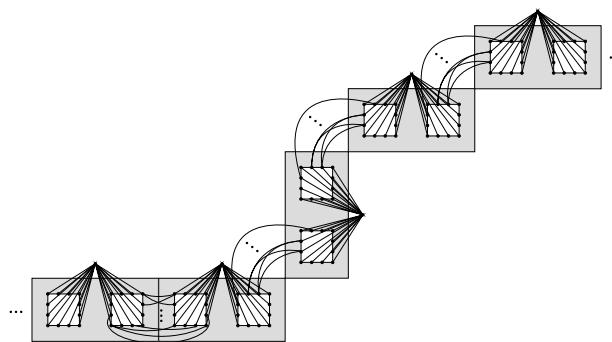
(c) Segment of a literal gadget. For clarity, only a few of the edges between different domino gadgets are depicted.

(d) Connecting literal gadgets to a variable gadget.

(e) Connecting literal gadgets to a clause gadget.

Figure 6: The gadgets.

between every vertex in one hole rim and every vertex in the adjacent hole rim of every pair of touching domino gadgets.

A.6 Satisfiable instances

We now show that if the formula φ is satisfiable, then M embeds into ℓ_∞^2 with small distortion.

LEMMA A.1. *If φ is satisfiable, then M embeds into ℓ_∞^2 with distortion at most 3.*

Proof. Let χ be a satisfying assignment. We start forming the embedding of the metric space by placing domino pieces on each domino gadget so that no black squares are adjacent. To begin, if $\chi_i = 1$, then each domino gadget in the variable gadget $\mathcal{V}_i$ is covered with a domino piece with the black square of the piece is pointing counter clockwise, otherwise clockwise. This means that if χ_i satisfies the clause C_j, then the literal gadget $\mathcal{L}_{\langle i,j \rangle}$ feeds into the variable gadget at a square that has been coloured white. Hence, the domino pieces covering the literal gadget can all have their black squares pointing towards the variable gadget and away from the clause gadget. In contrast, if χ_i does not satisfy the clause C_j, then the literal gadget $\mathcal{L}_{\langle i,j \rangle}$ feeds into the variable gadget at a square that has been coloured black. Hence, the domino pieces covering the literal gadget must have their black squares pointing towards the clause gadget. Finally, the clause gadget $\mathcal{C}_j$, having only one extra vertex y_j, need only have one of its three squares coloured black. We have already seen that the literal gadget corresponding to its satisfying literal has a white end point adjacent to the clause gadget. Hence, the clause square it is adjacent to can be the one coloured black.

Having covered each domino gadget $K_{\langle u,v \rangle}$ with a domino piece so that no black squares are touching, we independently embed the points for each square of the checker board into its own square. Instead of giving an embedding that never shrinks the required distances and expands them by at most a factor of 3, we will give one that never shrinks them by more than a factor of 2 and never expands them by more than a factor of $1\frac{1}{2}$. To begin, each square that does not have a hole is embedded unchanged as described in $X \subseteq \mathbb{R}^2$.

If the square is coloured black, we begin by embedding the extra vertex $x_{\langle u,v \rangle} \in V_\varphi$ at the center $\langle u,v \rangle$ of its square. The hole rim $X'(u,v)$ for this square is expanded from being of radius $\frac{1}{2}$ to being of radius $\frac{3}{4}$ located in $\mathbb{R}^2$ at $X'[\langle u-\frac{3}{4}, v-\frac{3}{4} \rangle, \langle u+\frac{3}{4}, v+\frac{3}{4} \rangle]$. This expand distances across the hole by a factor of $\frac{3}{4}/\frac{1}{2} = 1\frac{1}{2}$. It also squishes the grid points towards the outside of its square in the natural way. This space between the outside of the square and the hole rim of hole had been

of width $\frac{1}{2}$ but now it has narrowed to $\frac{1}{4}$. This involves shrinking the distances between grid points by a factor of 2.

On the other hand, if the squares coloured white, the hole rim $X'(u, v+2)$ is shrunk from being of radius $\frac{1}{2}$ to being of radius $\frac{1}{4}$. This shrinks distances across the hole by a factor of $\frac{1}{2}/\frac{1}{4} = 2$. It also expands the grid points away from the outside of its square in the natural way. This space between the outside of the square and the hole rim of hole had been of width $\frac{1}{2}$ but now it has expanded to $\frac{3}{4}$. This involves stretching the distances between grid points by a factor of $\frac{3}{4}/\frac{1}{2} = 1\frac{1}{2}$.

The extra edge in E_φ between the vertex $x_{\langle u,v \rangle}$ and this expanded hole rim has ℓ_∞^2 distance $\frac{3}{4}$, but it is supposed to have length $1\frac{1}{2}$. This gives a shrinking distortion by a factor of $(1\frac{1}{2})/(\frac{3}{4}) = 2$. The extra edge between this same vertex $x_{\langle u,v \rangle}$ and the opposite side of the shrunken hole rim of the adjacent square has ℓ_∞^2 distance $2 + \frac{1}{4}$, but it is supposed to have length $1\frac{1}{2}$. This gives a stretching distortion by a factor of $(2\frac{1}{4})/(1\frac{1}{2}) = 1\frac{1}{2}$.

The extra edge in E_φ between the shortest distance between points in the stretched hole rim of one square and a point in the shrunken hole rim of the adjacent square is ℓ_∞^2 distance $(1 - \frac{3}{4}) + (1 - \frac{1}{2}) = 1$, but it is supposed to have length 2. This gives a shrinking distortion by a factor of 2. The largest such distance is $(1 + \frac{3}{4}) + (1 + \frac{1}{2}) = 3$, but it is also supposed to have length 2. This gives a stretching distortion by a factor of $3/2 = 1\frac{1}{2}$. If two white squares are adjacent then these distances are between two shrunken hole rims and the distortions are better.

A.7 Embeddings of d-grids into $\mathbb{R}^d$ without big holes

In order to analyze the reduction for the case of unsatisfiable instances, we need to gain some understanding of the structure of low-distortion embeddings of d-grids into $\mathbb{R}^d$. To that extend, we show that in any embedding of sufficiently small distortion of a d-grid into $\mathbb{R}^d$, the image of the grid induces a net on a large ball around the image of the center of the grid. This basic property will be later used in our analysis to show that in any low-distortion embedding, the image of the wall induces a net in $\mathbb{R}^d$.

We begin by stating the main topological lemma that will be used in the proof. We only need the 2-dimensional version of the lemma, but it is in fact convenient to prove it for any fixed dimension. Its proof follows easily by the *nesting lemma* from [MS08] (corollary 2.2).

LEMMA A.2. (MAIN TOPOLOGICAL LEMMA) *Let $d \geq 2$, and let B^d be the unit ball in $\mathbb{R}^d$. Let $f : B^d \to \mathbb{R}^d$*

be a continuous map, such that $f(0)$ is in the closure of the unbounded connected component of $\mathbb{R}^d \setminus f(\partial B^d)$. Then, there exist $x, x' \in B^d$, with $\|x - x'\|_2 \geq \xi_d$, and $f(x) = f(x')$.

Since the metrics we are dealing with are not isometric to d-grids, we need to prove the property for a slightly more general class of graphs, that we call *central contractions of grids*. Intuitively, a central contraction is obtained from a grid, by adding an arbitrary set of edges between vertices that are close to the center of the grid.

DEFINITION A.1. (CENTRAL CONTRACTION OF A GRID) *Let J be a d-dimensional grid, with $V(J) = \{v_i : \forall j \in [d], 0 \leq i_j \leq k_j\}$, and $E(J) = \{\{v_i, v_{i'}\} : \|i - i'\|_1 = 1\}$, with each edge having unit length. Let J' be a graph obtained from J by adding a finite set of edges $\{v_i, v_j\}$, each having an arbitrary positive length, and such that for each $t \in [d]$, $i_t, j_t \in \{k_t(\frac{1}{2} - \frac{\xi_d}{8d^2}), \ldots, k_t(\frac{1}{2} + \frac{\xi_d}{8d^2})\}$. Then, J' is called a* central contraction *of J.*

We prove the above statement using lemma A.2. Informally, the argument is as follows. Consider a low distortion non-contracting embedding f of a central contraction of a grid, in which there is a large empty ball close to the image of the centroid of the grid. Let B_1 be the largest such ball. We can find a vertex u^* which is close to the centroid, and its image lies on the boundary of B_1. Since u^* is close to the centroid, there is a sufficiently large sub-grid Q centered at u^*. We extend f linearly to a continuous mapping $\tilde{g}$ of an appropriate d-dimensional simplicial complex $\tilde{Q}$ for which Q is a net. The complex $\tilde{Q}$ is chosen to be homeomorphic to a solid d-cube, which is in turn homeomorphic to the unit ball. Since the expansion is small, the image of each simplex of $\tilde{Q}$ is small, relative to the radius of B_1. Therefore, we can slightly modify the mapping $\tilde{g}$, so that the image of the complex avoids the interior of B_1. By applying a suitable homeomorphism on a subset of $\mathbb{R}^d$, we obtain a continuous map of $\tilde{Q}$ into $\mathbb{R}^d$, such that $\tilde{g}(u^*)$, lies on the boundary of the unbounded connected component of $\mathbb{R}^d \setminus \tilde{g}(\partial \tilde{Q})$. Since $\tilde{Q}$ is homeomorphic to the unit ball, we can apply lemma A.2, to obtain two points in $\tilde{Q}$ that are far from each other, and have the same image under $\tilde{g}$. Since Q is a net on $\tilde{Q}$, we can find vertices that are far from each other in Q, and their images are very close under f, contradicting the non-contraction hypothesis.

LEMMA A.3. (FROM GRIDS TO NETS) *Let $d \geq 2$, and let $J = (V, E)$ be a d-dimensional grid with $V(J) = \{v_i : \forall j \in [d], 0 \leq i_j \leq k_j\}$, and $E(J) = \{\{v_i, v_{i'}\} : \|i - i'\|_1 = 1\}$, such that for each $j \in [d]$, $k_j \geq R$, for some $R \geq c \cdot \frac{128 \cdot d}{\xi_d}$. Assume that each edge of J has*

unit length. Let J' be a central contraction of J. Let $f : V(J') \to \mathbb{R}^d$ be a non-contracting embedding of J' into $\mathbb{R}^d$ with expansion c. Then, for any $p \in \mathbb{R}^d$, with $\|p - f(v_{k_1/2,\ldots,k_d/2})\|_2 \leq R/16$, there exists $u \in V(J')$, with $\|p - f(u)\|_2 \leq 2 \cdot c$.

Proof. Let f be a non-contracting embedding of J' into $\mathbb{R}^d$ with expansion c. Assume for the sake of contradiction that there exists a point $p \in \mathbb{R}^d$, with $\|p - f(v_{k_1/2,\ldots,k_d/2})\|_2 \leq R/16$, such that for any $u \in V(J)$, $\|p - f(u)\|_2 > 2 \cdot c$.

Let u^* be the vertex of J which is nearest to p under f. That is, $u^* = \operatorname{argmin}_{u \in V(J)} \|p - f(u^*)\|_2$. Since f is non-contracting, we have

$$
\begin{aligned}
D_J(u^*, v_{k_1/2,\ldots,k_d/2}) &\leq \|f(u^*) - f(v_{k_1/2,\ldots,k_d/2})\|_2 \\
&\leq \|f(u^*) - f(p)\|_2 \\
&+ \|f(p) - f(v_{k_1/2,\ldots,k_d/2})\|_2 \\
&\leq 2 \cdot \|f(p) - f(v_{k_1/2,\ldots,k_d/2})\|_2 \\
&\leq R/8
\end{aligned}
$$

Therefore, there exist $i^* \in \mathbb{Z}^d$, such that for each $j \in [d]$, $i_j^* \in \{k_j/8, \ldots, 7k_j/8\}$, with $u^* = v_{i^*}$.

Let $A = \{v_{i_1,\ldots,i_d} \in V(J) : \|i - i^*\|_\infty \leq 3R/8\}$, and define the vertex-induced subgraph $Q = J[A]$. We construct a d-dimensional simplicial complex $\tilde{Q}$ corresponding to the graph Q as follows. The set of 0-simplices of $\tilde{Q}$ is A. For each $i \in \mathbb{Z}^d$, such that for each $j \in [d]$, $i_j \in \{i_j^* - 3R/8, i_j^* + 3R/8 - 1\}$, let T_i be the triangulation of the hypercube $A_i = Q[\{v_j : j \in \{i_1, i_1 + 1\} \times \ldots \times \{i_d, i_d + 1\}\}]$. We add to $\tilde{Q}$ all the simplices in T_i.

Let g be the restriction of f on $V(Q)$. Recall that for a simplicial complex K, $|K|$ denotes the union of all its simplices. Let $\tilde{g}$ be the linear extension of g on $|\tilde{Q}|$.

Let $B_1 = B(p, \|p - f(u^*)\|_2)$. Note that $f(V(J)) \cap \operatorname{Int}(B_1) = \emptyset$. We proceed to define a map $\tilde{h} : |\tilde{Q}| \to \mathbb{R}^d \setminus \operatorname{Int}(B_1)$. For each point $x \in |\tilde{Q}|$ with $\tilde{g}(x) \notin \operatorname{Int}(B_1)$, we set $\tilde{h}(x) = \tilde{g}(x)$. For each point $x \in |\tilde{Q}|$ with $\tilde{g}(x) \in \operatorname{Int}(B_1)$, let r_x be the ray starting at p and passing through $\tilde{g}(x)$. We set $\tilde{h}(x)$ to be the point where r_x intersects ∂B_1.

Define $\varphi : |\tilde{Q}| \to \mathbb{R}^d$ where for each 0-simplex $v_{i_1,\ldots,i_d} \in \tilde{Q}$, $\varphi(v_i) = (i_1 - i_1^*, \ldots, i_d - i_d^*)$, and for all other points $x \in \tilde{Q}$, $\varphi(x)$ is defined via a linear extension.

Let C denote the unit ball in $\mathbb{R}^d$ under the ℓ_2 norm. We define a map $\mu : C \to |\tilde{Q}|$ as follows. Let $\mu(0) = u^*$, and for each $x \in \mathbb{R}^d \setminus \{0\}$, let $\mu(x) = \varphi^{-1}(\frac{3R}{8} \cdot x \cdot \frac{\|x\|_2}{\|x\|_\infty})$.

Consider the map $\psi : C \to \mathbb{R}^d \setminus \operatorname{Int}(B_1)$ defined by $\psi(x) = \tilde{h}(\mu(x))$. The map ψ is clearly continuous. Furthermore, $\psi(C)$ is homeomorphic to a subset of the unit ball in $\mathbb{R}^d$, under a homeomorphism that sends 0 to

the boundary of the unit ball. We can thus apply lemma A.2 and obtain points $y, y' \in C$, such that $\|y - y'\|_2 \geq \xi_d$, and $\psi(y) = \psi(y')$. Let $\mu(y) = x$, and $\mu(y') = x'$, for some $x, x' \in |\tilde{Q}|$. Let σ, σ' be simplices of $\tilde{Q}$ such that $x \in \sigma$, $x' \in \sigma'$. Pick vertices $w, w' \in \tilde{Q}$, with $w \in \sigma$, $w' \in \sigma'$.

Observe that for each $p \in |\tilde{Q}|$, $\|\tilde{h}(p) - \tilde{g}(p)\|_2 < c$. Thus,

$$
\begin{aligned}
\|f(w) - f(w')\|_2 &= \|\tilde{h}(w) - \tilde{h}(w')\|_2 \\
&\leq \|\tilde{h}(w) - \tilde{h}(x)\|_2 \\
&+ \|\tilde{h}(x') - \tilde{h}(w')\|_2 \\
&< \|\tilde{g}(w) - \tilde{g}(x)\|_2 \\
&+ \|\tilde{g}(x') - \tilde{g}(w')\|_2 + 4 \cdot c \\
\text{(A.1)} \quad &\leq 6 \cdot c
\end{aligned}
$$

Since J' is a central-contraction of J, it follows that for each $u, v \in V(J')$, $D_{J'}(u, v) \geq D_J(u, v) - R \cdot \frac{\xi_d}{4d}$. Thus,

$$
\begin{aligned}
D_{J'}(w, w') &\geq \|\varphi(w) - \varphi(w')\|_1 - R \cdot \frac{\xi_d}{4d} \\
&\geq \|x - x'\|_1 - \|\varphi(w) - x\|_1 \\
&- \|\varphi(w') - x'\|_2 - R \cdot \frac{\xi_d}{4d} \\
&\geq \frac{1}{d} \cdot \frac{3R}{8} \cdot \|y - y'\|_2 - 2 \cdot d - R \cdot \frac{\xi_d}{4d} \\
&\geq \frac{1}{d} \cdot \frac{3R}{8} \cdot \xi_d - 2 \cdot d - R\frac{\xi_d}{4d} \\
\text{(A.2)} \quad &> \frac{\xi_d}{8d} \cdot R - 2 \cdot c \geq 14 \cdot c
\end{aligned}
$$

Combining (A.1) and (A.2) we obtain a contradiction of the fact that f is non-contracting.

A.8 Unsatisfiable instances We will now show that if there exists an embedding of M into ℓ_∞^2 with small enough distortion, then φ is satisfiable. Fix an embedding f of M into ℓ_∞^2 with distortion less than $7/2 - 100\varepsilon$. Let $r = 10 \cdot \varepsilon$, and let $c > 0$ be a sufficiently large constant, to be determined later. Let A be a rim of the input graph. Let $\gamma(A)$ denote the union of all balls of radius r in $\mathbb{R}^2$ centered in the images of points that are at distance at most $c \cdot \varepsilon$ from A. I.e.

$$
\gamma(A) = \bigcup_{x : D(x, A) \leq c \cdot \varepsilon} B(x, r)
$$

DEFINITION A.2. (VERTEX INSIDE A RIM) *We say that a vertex v is* inside *a rim A if $f(v)$ is contained in a bounded path-connected component of $\mathbb{R}^2 \setminus \gamma(A)$.*

LEMMA A.4. *For every $v \in V_\varphi$, there exists a rim A such that v is inside A.*

Proof. Consider the set $T \subset \mathbb{R}^2$ defined by $T = \bigcup_{x \in X - H_\varphi} B(f(x), r)$. By the non-contraction of f we have $f(v) \notin T$.

Let J be the union of all rims together with the boundary of the grid X, and let $J' = \{x \in X - H_\varphi : D(x, J) \leq c \cdot \varepsilon\}$. Let $p \notin J'$. By lemma A.3 we have that $f(X - H_\varphi)$ induces a r-net in a ball of radius $\Omega(c\varepsilon)$ around $f(p)$. Therefore, $B(p, r) \cap \partial T = \emptyset$. Thus, for all points $q \in X - H_\varphi$ such that $B(q, r) \cap \partial T \neq \emptyset$, we have $q \in J'$.

Let Y be the closure of the component of $\mathbb{R}^2 \setminus T$ containing $f(v)$. By the non-contraction of f it follows that all the points $p \in J'$ such that $B(f(p), r) \cap \partial Y \neq \emptyset$, must be at distance at most $c \cdot \varepsilon$ from either the boundary of X, or exactly one of the rims. It clearly cannot be the case that these points in J' are close to the boundary of X, since then the distance between $f(v)$ and all the rims would be $\Omega(L)$, which contradicts the fact that f has expansion $O(1)$. Therefore, these points are at distance at most $c \cdot \varepsilon$ from exactly one rim A.

It remains to show that $\gamma(A)$ does not intersect the boundary of an unbounded component of $\mathbb{R}^2 \setminus T$. Indeed, observe that if this happens then $f(A)$ must have diameter $\Omega(L)$, a contradiction. We have thus obtained that $f(v)$ is in a bounded component of $\mathbb{R}^2 \setminus \gamma(A)$, concluding the proof.

LEMMA A.5. *Let K be a copy of the domino gadget. Then, x_K is inside one of its two rims, i.e. a domino piece must either be placed colouring at least one of the two squares of the domino gadget black.*

Proof. Assume for the sake of contradiction that x_K is not inside a rim of K. By lemma A.4 it follows that x_K is inside some rim A not in K. By inspecting the construction it follows that $D(x_K, A) \geq 5/2$. Moreover, there exists a rim A' in K such that for any $a \in A$, $a' \in A'$, we have $D(a, a') \geq 4$. Let $J_A = \{x \in X - H_\varphi : D(x, A) \leq c \cdot \varepsilon\}$. We have that $f(x_K)$ is contained in a bounded component of $\mathbb{R}^2 \setminus \bigcup_{x \in J_A} B(x, r)$. It follows that the component that contains $f(x_K)$, contains also a ball of radius $5/2 - O(\varepsilon)$ centered at $f(x_K)$, that does not intersect T. It follows that $\|f(x_K) - f(A')\|_\infty > 4 + 5/2 - O(\varepsilon) = 11/2 - O(\varepsilon)$. However, $D(x_K, A') = 3/2$, implying that the distortion of f is at least $11/3 - O(\varepsilon) > 7/2$, a contradiction.

The proofs of lemmas A.6, A.7, and A.8 bellow are similar to the proof of lemma A.5, and are omitted from this extended abstract.

LEMMA A.6. *Let $\mathcal{C}_i$ be a clause gadget. Then y_i is inside one of its three rims, i.e. at least one of the three squares of each clause gadget must be black.*

LEMMA A.7. *Let K_1, K_2 be distinct copies of the domino gadget. Let A_1 be a rim of K_1 that is attached to a rim A_2 of K_2. If x_{K_1} is inside A_1, then x_{K_2} is not inside A_2, i.e. the domino pieces can not be placed so that two adjacent squares are both black.*

LEMMA A.8. *Let K is a copy of the domino gadget, and let $\mathcal{C}_i$ be a clause gadget. Let A_1 be a rim of K that is attached to a rim A_2 of $\mathcal{C}_i$. If y_i is inside A_2, then x_K is not inside A_1.*

LEMMA A.9. *If M embeds into ℓ_∞^2 with distortion less than $7/2 - O(\varepsilon)$, then φ is satisfiable.*

Proof. Considering any way of covering the domino gadgets with domino pieces and colouring at least one square of each clause gadget black. From this we will produce a satisfying assignment χ. Consider the variable χ_i. Because no two adjacent squares are both black, the domino pieces are placed on the rectangular path of the variable gadget $\mathcal{V}_i$ either with each piece's black square pointing counter clockwise or each pointing clockwise. Set $\chi_i = 1$ iff it is counter clockwise. Now consider a clause C_j. We must prove that it is satisfied by this assignment. At least one of the three squares in the clause gadget $\mathcal{C}_j$ must be black. Let χ_i be the variable corresponding to the literal gadget $\mathcal{L}_{\langle i,j \rangle}$ adjacent to this black square. Because no two adjacent squares are both black, the black square of the domino pieces placed on this literal gadget must face towards the variable gadget $\mathcal{V}_i$ ending in a black square. Hence, the square in the variable gadget $\mathcal{V}_i$ adjacent to this literal gadget must be white. If χ_i appears in C_j as a positive literal, then this white square is the clockwise most square in this domino gadget. Hence, we set $\chi_i = 1$ satisfying the clause C_j. Similarly, if χ_i appears in C_j as a negative literal.