

THE COMPLEXITY OF GENERALIZED DOMINO TILINGS

IGOR PAK* AND JED YANG*

ABSTRACT. Tiling planar regions with dominoes is a classical problem in which the decision and counting problems are polynomial. We prove a variety of hardness results (both NP- and #P-completeness) for different generalizations of dominoes in three and higher dimensions.

1. INTRODUCTION

The study of *domino tilings* is incredibly rich in its history, applications, and connections to other fields. A geometric version of the *perfect matching* problem in grid graphs, the problem is a special case of a general study of tilings of finite regions, with its own set of problems and generalization directions, very different from graph theory. In this paper we study the decision and the counting problems for the tilings of a given region. Both problems have a large literature and long history in combinatorics and computational complexity, as well as in the recreational literature, but much of the work concentrated on tilings in the plane. Significantly less is known in three and higher dimensions, where much of the intuition and many tools specific to two dimensions break down.

To summarize the results of this paper in one sentence, we show that the classical domino tiling problems become computationally hard already in three dimension, even when the topology of regions is restricted. This further underscores the fundamental difficulty of obtaining even the most basic (positive) tiling results in higher dimensions. We should mention that dimension three is critically important for applications ranging from enumerative combinatorics to probability, to statistical physics, and to solid-state chemistry (see e.g. [KRS, Ken, LP, Til]).

The (single tile) *tileability problem* can be stated as follows. Given a tile T , decide whether a region $\Gamma \subset \mathbb{R}^d$ is tileable with copies of T (rotations and reflections are allowed). When T is a domino $[2 \times 1]$, the tileability problem can be solved in polynomial time, via the reduction to perfect matching [LP], and the result generalizes to any $d \geq 2$. On the other hand, when T is a bar $[3 \times 1]$, the tileability problem is NP-complete [BNRR] in the plane.

Our first result is on the tileability with a brick $T = [2 \times 2 \times 1]$, which we call a *slab*. Viewed as half-cubes, slabs can be thought of as a natural generalization of 2-dimensional dominoes.

Theorem 1.1 *Tileability of 3-dimensional regions with slabs is NP-complete.*

Now, when the region $\Gamma \subset \mathbb{R}^2$ is simply connected (s.c.), the tileability problem is simpler in many cases. For example, when T is a domino, the s.c. tileability problem can be solved in linear time in the area $|\Gamma|$, while quadratic for general regions [Thu] (see also [Cha]). Moreover, when T is any rectangle, the s.c. tileability problem is polynomial [Rém] (see also [KK]), as opposed to NP-complete for general regions. Interestingly, this phenomenon does not extend to higher dimensions.

Theorem 1.2 *Tileability of contractible 3-dimensional regions with slabs is NP-complete.*

For the definition of contractible regions in higher dimension, generalizing s.c. regions in the plane, see Section 2. Note that in the plane, finding a tile such that the tileability problem of simply connected regions is NP-complete remains an open problem (see Section 8).

*Department of Mathematics, UCLA, Los Angeles, CA 90095, USA; {pak, jedyang}@math.ucla.edu.

The next set of results concerns *counting problems*: given Γ , compute the number of tilings of Γ with copies of tile T (again, rotations and reflections are allowed). It is well known that in the plane, the number of domino tilings of Γ can be computed in polynomial time (see e.g. [Ken, LP]). This is perhaps surprising, since computing the number of perfect matchings is classically $\#P$ -complete [Val1]. The following result is one of the historically first applications of $\#P$ -completeness.

Theorem 1.3 (Valiant [Val2]) *Number of domino tilings of 3-dimensional regions is $\#P$ -complete.*

Let us note that Valiant’s result is strongly related, but slightly different (see Subsection 8.4). We re-prove Valiant’s theorem both for completeness and as a stepping stone towards stronger results. The proof uses a reduction from the perfect matching problem in cubic bipartite graphs. This construction allows us to prove the following far-reaching extension of Theorem 1.3.

Theorem 1.4 *Number of domino tilings of contractible 3-dimensional regions is $\#P$ -complete.*

We conclude with the following counting version of theorems 1.1 and 1.2. This is obtained essentially for free by using the same construction as in the proof of the theorems.

Theorem 1.5 *Number of slab tilings of 3-dimensional regions is $\#P$ -complete. Moreover, the result holds when considering only contractible regions.*

Much of the rest of the paper consists of proofs of these results. In Section 7, we present generalizations to dimensions $d \geq 4$. We conclude with final remarks and open problems in Section 8, where we also continue a historical discussion of these and related tiling results.

2. DEFINITIONS AND BASIC RESULTS

Consider \mathbb{Z}^3 as the union of closed unit cubes in \mathbb{R}^3 centered at the integer lattice points. We will suggestively refer to the elements of \mathbb{Z}^3 as *cubes*. A *region* Γ is a finite subset of \mathbb{Z}^3 , and is naturally identified with $\bar{\Gamma} \subset \mathbb{R}^3$, the union of the corresponding closed unit cubes. We say that a region $\Gamma \subset \mathbb{Z}^3$ is *contractible* if $\bar{\Gamma} \subset \mathbb{R}^3$ is contractible and has contractible interior. Neither of




Figure 1: Examples of non-contractible regions of \mathbb{Z}^3 .

the requirements is redundant. For example, the set on the left in Figure 1 consisting of six closed cubes is contractible but its interior is not, since the center becomes a hole. The region on the right with seven closed cubes is not contractible but its interior is. We wish to be strict and consider these as non-contractible.

Let T be a *tile*, which is simply a contractible region. A T -tiling of a region Γ is a partition of Γ into pairwise disjoint (as subsets of \mathbb{Z}^3) isomorphic copies of T . Here we allow T to be rotated and reflected—though all the tiles we will consider are bricks, which have mirror symmetries. When the tile T is understood, we simply call it a tiling. This leads us to the following two decision problems:

T -TILEABILITY

Instance: A region Γ .

Decide: Does the region Γ admit a T -tiling?

CONTRACTIBLE T -TILEABILITY**Instance:** A contractible region Γ .**Decide:** Does the region Γ admit a T -tiling?

To analyze the complexities of these problems, we will follow the usual line of attack by embedding known problems as tiling problems. As such, let us define two more decision problems we will use. A graph is *cubic* if each vertex has degree 3. A *perfect matching* is a collection of pairwise-disjoint edges that cover all the vertices.

CUBIC BIPARTITE PERFECT MATCHING

Instance: A cubic bipartite graph G .**Decide:** Does the graph G admit a perfect matching?

Suppose we are given a set of boolean variables. A *literal* is a variable or the negation thereof. A *3-SAT expression* \mathcal{C} is a conjunction of clauses, each a disjunction of three literals. We will write $\mathcal{C} = \{C_1, \dots, C_t\}$ as a set of triples C_i of literals, where the negation of a variable v is denoted \bar{v} . An assignment of boolean values to the variables is *1-in-3 satisfying* if each clause C_i has precisely one true literal. This leads to the following natural decision problem:

1-IN-3 SAT

Instance: A 3-SAT expression \mathcal{C} .**Decide:** Does the expression \mathcal{C} admit a 1-in-3 satisfying assignment?

For each decision problem, we also have an associated counting problem, where we count the number of witnesses. In particular, for the problems above, we need to count the number of tilings, perfect matchings, and satisfying assignments, respectively. By abuse of language, we say a decision problem is $\#P$ -complete if its associated counting problem is $\#P$ -complete. We reduce the tiling problems from the following two well-known problems.

Theorem 2.1 ([DL]) CUBIC BIPARTITE PERFECT MATCHING is $\#P$ -complete.

Theorem 2.2 ([Sch, CH]) 1-IN-3 SAT is NP-complete and $\#P$ -complete.

For convenience and without loss of generality, we will assume that each variable v_k appears unnegated somewhere in the boolean expression. Indeed, if a variable occurs only negated, we may simply replace it by its negation. If v_k does not occur, add the clause (v_k, \bar{v}_k, f) , where f is a new variable. The number of satisfying assignments is obviously preserved.

3. PROOF OF THEOREM 1.3

3.1. Reduction construction. We will reduce the counting of perfect matchings in cubic bipartite graphs to the counting of tilings with dominoes. Given a cubic bipartite graph, we want to create a region, such that the number of tilings of the region is the number of perfect matchings of the given graph.

Given two cubes $u, v \in \mathbb{Z}^3$, a *wire* is a sequence of distinct cubes $u = c_1, c_2, \dots, c_t = v$ in \mathbb{Z}^3 , such that c_i and c_j are adjacent if and only if $|i - j| = 1$. We call c_1 and c_t the *endpoints* of the wire, $\{c_2, \dots, c_{t-1}\}$ the *interior*, and t the *length*. A collection of wires is *proper* if no wires intersect the interior of other wires, and any two cubes in the union of the wires are adjacent if and only if they are consecutive elements in (precisely) one such wire.

Let G be a connected graph. We will create a region $\Gamma \subset \mathbb{Z}^3$ representing this graph G as follows. Let $f : V(G) \rightarrow \mathbb{Z}^3$ be an injective map, and identify each vertex with its image. For each edge $uv \in E(G)$, we connect $f(u)$ and $f(v)$ by a wire, which is identified with the edge. Let Γ be the union of these wires (and thus contains the vertices). If the collection of the wires is proper, we call Γ a *lattice drawing* of G .

Color the space $\chi : \mathbb{Z}^3 \rightarrow \mathbb{Z}_2$ in a checkerboard fashion, and call $\chi(x, y, z) = x + y + z \pmod{2}$ the *parity* of $(x, y, z) \in \mathbb{Z}^3$.

Lemma 3.1 *Every connected simple graph G with maximum degree at most 6 has a lattice drawing. Moreover, if G is bipartite, it can be drawn with each edge having endpoints of opposite parity.*

Proof. It is obvious that if we pick f so that the images are sufficiently spaced out, we can connect the vertices with proper wires. The maximum degree condition is due to the limitation that a cube has only 6 neighbors. The second part of the lemma is immediate. \square

3.2. Proof of correctness. Fix a lattice drawing Γ constructed above, and consider a tiling by dominoes. The neighborhood of a vertex is shown in Figure 2. The solid blue square represents the vertex, with three emanating wires representing its incident edges.




Figure 2: A vertex with three incident wires.

Consider a wire corresponding to some edge. Since the collection of wires is proper, there are no other cubes adjacent to the interior of this wire. Thus in any domino tiling, the dominoes must tile the wire along its length. Since the endpoints have opposite parity, this wire has even number of cubes. Hence either it is partitioned into dominoes under the tiling, or the dominoes it contains are all in its interior, in which case its endpoints are covered by dominoes that are contained in other wires.

Lemma 3.2 *If M is the collection of edges in Γ that are partitioned into dominoes, then M is a perfect matching of G .*

Proof. Pick a vertex and consider the domino covering it: that domino is contained in precisely one of the incident edges. Thus for every vertex precisely one incident edge is in M , implying that M is a perfect matching. \square

We just extracted a perfect matching of G from a tiling of Γ . Notice that this map is bijective, that is, every perfect matching induces a unique tiling. Indeed, for each edge in the matching, tile the corresponding wires by dominoes covering the endpoints, then finish the tiling in the obvious manner. This means that the number of tilings is exactly the number of perfect matchings, concluding the proof of Theorem 1.3.

4. PROOF OF THEOREM 1.4

4.1. Reduction construction. The main idea is to create a large contractible “plate” that admits a unique domino tiling, and then place the construction from the proof of Theorem 1.3 adjacent to this plate, while being careful as to not introduce new tilings for the plate. Given a region Γ , a subregion $S \subset \Gamma$ is called *frozen* if it is tiled the same way in all possible tilings of Γ .

We will call the $z = k$ plane *Layer k* . The entire construction of Γ will lie in four layers, from Layer -1 to Layer 2 . When considering cubes on a layer, we will employ the usual convention of

referring to the $+x$, $-x$, $+y$, $-y$ directions as right, left, above, and below. When referring to an adjacent cube in a different layer, we will specify the layer specifically, as to avoid confusion.

Let us start with a *plate* that lies in Layer 0, which is a region whose columns have even lengths and are offset with each other, causing *jagged* borders. It is obvious that it admits a unique tiling,




Figure 3: A plate with a unique tiling.

which is shown in Figure 3. We will modify this plate carefully as to introduce only local changes in the tilings while preserving the unique global tiling.

Suppose we are given a cubic bipartite graph G with bipartition A and B . We will place the vertices in A on the left side of a plate, and vertices in B on the right side. Now we must connect the corresponding vertices with wires. We can achieve the desired connections by interchanging adjacent wires in steps. Indeed, think of the correspondence as a permutation and write it as a product of transpositions. The resulting schematic is known as a *wiring diagram*. These wires will, for the most part, stay in Layer 1. In between each pair of wires, we will add a *tension line*. We then modify the region locally with a *crossover X-gadget* at each location indicated in the wiring diagram.

As an example, we see the general picture in Figure 4. There is a big plate, with offset columns




Figure 4: Overview of the layout.

causing jagged borders on the top and bottom of the region, indicated by the short dashes. There are four vertices, each one with three wires coming out of it, indicated by the solid lines. The tension lines between the wires are indicated by the long dashed lines. The red boxes are the X-gadgets.

The wire and the tension line are shown in Figure 5. The schematic is drawn in Layer 0. The unique tiling of the frozen region in Layer 0 is shown: dominoes that are entirely in this layer are drawn as 2×1 rectangles in the plane. When a $+$ (resp. $-$) sign is present, it means the region includes the adjacent cube in Layer $+1$ (resp. -1).

Here we show half of a vertex V-gadget in Figure 6a. A single square in the schematic denotes that the cube in Layer 0 is always tiled with a domino sticking out to Layer 1 or -1 . The symbol \pm signifies that the adjacent cubes in both Layers ± 1 are included. The blue circle is the center of


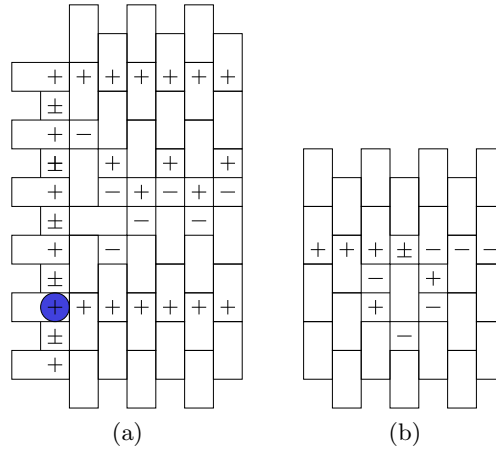


Figure 5: A wire (left) and a tension line (right).

Figure 6: Half of a vertex V -gadget (left) and a hole H -gadget (right).

the V -gadget, and needs to satisfy positional parity as defined in the proof of the previous theorem. Besides some isolated cubes, Layer 1 contains the wires. Notice that there are three wires coming out of the center of the V -gadget. The third wire that is not drawn completely in the figure should mirror the other wire in the obvious way, including the initialization of the tension line. The wires (sequences of $+$) extend towards the right. Between each pair, we have a tension line (the sequence of jagged $+$ and $-$). Notice that we can easily make the V -gadgets taller as to have more separation between the wires and tension lines exiting to the right. This will be used to align all these things together correctly. We use the mirror image when putting the V -gadget on the right side of the plate. Between pairs of V -gadgets, we also add tension lines running horizontally across the entire plate.

To make the crossover X -gadget, we combine the hole H -gadget (Figure 6b) and the splitter Y -gadget (Figure 7). We use the same schematic convention as above in these figures. Here the symbol \oplus denotes that there is a cube in Layer 1, but not in Layer 0. Similarly, symbol \boxplus signifies the inclusion of the cubes on Layers 1 and 2 (in addition to Layer 0). First use the H -gadget to guide the lower wire to Layer -1 , and use the Y -gadget to merge and align them together. Then we reverse the process using a rotated copy of the splitter, as to make the wires crossover each other. Finally, bring the wire back to Layer 1 with another H -gadget. This completes the construction of the crossover X -gadget.




Figure 7: A splitter Y-gadget.

It is clear that the X-gadget takes two wires on the left, separated by a tension line, and outputs two wires on the right, again separated by a tension line. Notice that in the construction of the X-gadget, we could make the output wires further away from each other, which is also the case for the V-gadget, as mentioned above. As such, we could align all these wires and tension lines so they feed into each other perfectly across the entire construction. Thus using X-gadgets, we are able to correctly connect the wires coming from V-gadgets. This complete the construction of the region, which we shall call Γ .

4.2. Proof of correctness. It remains to check that Γ is contractible, and its domino tilings are in bijection with perfect matchings in graph G .

Lemma 4.1 *The region Γ is contractible.*

Proof. Consider Layer 0. It has a hole for each Y-gadget. However, the adjacent cube in Layer 1 is present, and is contained in a 3×3 region. Furthermore, the adjacent cube in Layer -1 is not present. As such, filling in the hole in Layer 0 does not alter contractibility. Now that there are no holes in Layer 0, we may deformation retract everything in other layers to Layer 0, and then contract it to a point. \square

It remains to find a frozen subregion $\Gamma' \subset \Gamma$ such that $\Gamma_0 = \Gamma \setminus \Gamma'$ is a lattice drawing of the given graph. Let us start with $\Gamma' = \emptyset$ (thus $\Gamma_0 = \Gamma$) and alter Γ' inductively. A cube in a region is *isolated* if it has precisely one neighbor (out of the six adjacent possibilities). Notice that an isolated cube must be tiled by a domino covering its unique neighbor. We will *remove* the isolated cubes of Γ_0 by moving them from Γ_0 to Γ' . We repeat this process inductively until there are no more isolated cubes.

Lemma 4.2 *At any step, Γ' is a frozen subregion of Γ . When the process finishes, Γ_0 is a lattice drawing of the given graph.*

Proof. It follows immediately by construction that Γ' is frozen. The rest of this proof is devoted to the second part of the lemma.

Let us perform another induction on the number of X-gadgets. For the base case, suppose we did not have any X-gadgets. That is, we start with the boring cubic graph where disjoint pairs of vertices are joined by three edges each.

Consider each V-gadget. After the first step, everything in Layers 0 and -1 on the two left-most columns in Figure 6a are removed, leaving the wire in Layer 1. The tension lines are all removed as well. It is clear that in the next step, the remaining horizontal domino will also be removed, along with the domino above it.

Thus after two steps, what we are left with in Layer 0 is a collection of disconnected plates, each of which has a wire across it horizontally in Layer 1. These wires protrude to the left and right.

Each of these plates have jagged columns thanks to the initial boundary and the tension lines. As such, in each step, the top- and bottom-most cubes are isolated, hence removed with their neighbors, creating another jagged boundary. This process continues until all cubes in Layer 0 are removed, including the ones adjacent to the wire in Layer 1. Note that half of the cubes adjacent to said wire are removed from above, and the other from below. Thus it is important that both sides were jagged, which is the motivation for introducing the tension lines. We are left with the lattice drawing of the aforementioned contrived cubic graph.

Now for the inductive step, suppose we add an extra X-gadget. We want to show that the global situation is not altered, and that the local situation is what we wanted. Indeed, the X-gadget consists of two H-gadgets and two Y-gadgets. Having analyzed the V-gadget, one should be able to inspect the H-gadget construction in Figure 6b, and easily see that after finitely many steps, we will be left with the wires in Layers ± 1 , plus the connecting cube between them in Layer 0.

The Y-gadget is a bit more complicated, but can be carefully analyzed in the same manner. Indeed, after the first step, all the cubes in Layer -1 are removed except for the wire. Similarly, the Layers 1 and 2 cubes for \boxplus are removed, so are all the cubes in Layer 1, save the wire. As for Layer 0, we get a small island of six vertical dominoes in the middle, and a giant plate with enough tension lines such that a similar inductive removal process as above will remove it completely. Hence we are left with precisely the wires, thus completing the proof of the lemma. \square

5. PROOF OF THEOREM 1.1 AND THE FIRST PART OF THEOREM 1.5

5.1. Reduction construction. Here we reduce 1-IN-3 SAT to a tiling problem. The reduction will be parsimonious, thus proving both NP-completeness and #P-completeness.

Consider a sequence of cubes in the $z = 0$ plane, where two cubes are adjacent if and only if they are consecutive elements in the sequence (this is called a wire in the proof of Theorem 1.3). Consider the sequence as a region and duplicate it also in the $z = 1$ plane. We now call a *wire* the union of these two copies. In a wire, a pair of adjacent cubes, one with $z = 0$ and one with $z = 1$, is called a *cube pair*. In any tiling by slabs, each cube pair must be covered by the same slab. If we translate the wire to the $z = k$ and the $z = k + 1$ planes, we say the wire *lives* in the $z = k$ *biplane*. Color the biplane $\chi(x, y, z) = x + y \pmod{2}$ in a checkerboard fashion, ignoring the z -coordinate. We call the color of a cube (pair) its *parity*, which can be *even* or *odd*. Similarly, we may rotate and consider wires living in $x = k$ or $y = k$ biplanes. In those cases, the biplane coloring would ignore the x - or the y -coordinate, respectively. Notice that a cube $(x, y, z) \in \mathbb{Z}^3$ might have different parities depending on the biplane being considered. However, each cube pair lives in a unique biplane: thus the parity of a cube pair is unambiguous. When we draw diagrams consisting of wires in each biplane, we simply draw the two-dimensional regions used to form the wires.

The variable V-gadget, shown in Figure 8, is made by removing two diagonal cubes of a $2 \times 2 \times 2$ region, and then attaching six wires. Notice that the wires come in pairs, and each pair lives in




Figure 8: A variable V-gadget, with six wires coming out of it.

a biplane. We label the wires with their initial coordinate directions. Each wire has an endpoint in the V-gadget, called the *variable endpoint*; the other one is called the *free endpoint*. We now describe where to place the V-gadgets and how to join these free endpoints together.

Suppose we are given a 3-SAT expression $\mathcal{C} = \{C_1, \dots, C_t\}$ on the variable set $X = \{v_1, \dots, v_n\}$, where each clause $C_i = (c_{i1}, c_{i2}, c_{i3})$ consists of three literals, each either a variable v_k or a negation \bar{v}_k , for some $k \in [n] = \{1, \dots, n\}$. We now construct a region Γ .

If c_{ij} is v_k or its negation \bar{v}_k , place a V-gadget at $(10k, 10(3i + j), 0)$. That is, the wires live in the $x = 10k$, $y = 10(3i + j)$, and $z = 0$ biplanes, respectively. Moreover, the variable endpoint of each wire has odd parity with respect to the coloring of the biplane in which the wire lives.

Consider the $x = 10k$ plane for each $k \in [n]$, which is drawn schematically in Figure 9. Notice




Figure 9: Synchronization of variables: This is a diagram in some $x = 10k$ (bi)plane. The squares represent the V-gadgets and the lines represent the wires.

that all the $\pm z$ wires corresponding to the variable v_k are on this biplane. For each V-gadget, we will bend the $+z$ wire to the left and down as shown. We link up the first wire with the last, and also the remaining adjacent pairs.

Figure 10 shows a clause C-gadget, which is simply a joining of three wires. The parity of the gadget is the parity of the blue cube pair at the center where the three wires meet.

Now let us work in the $z = 0$ biplane, schematically shown in Figure 11. For each clause, there are three corresponding V-gadgets. From each V-gadget, take the $+y$ (resp. $-y$) wire if the corresponding variable appears in the clause unnegated (resp. negated). Extend these three wires to $x = 10(n + 1)$ and join them together at a C-gadget with even parity. Similarly, extend the three




Figure 10: A clause C-gadget, where three wires meet.




Figure 11: Diagram in the $z = 0$ (bi)plane. The circles, where the lines meet, indicate the C-gadgets, which are at two different parities.

remaining $\pm y$ wires to $x = -10$ and join together at a C-gadget with odd parity. We call these the even and odd C-gadgets corresponding to each clause, respectively.

Finally, for each V-gadget, let us join the remaining $+x$ wire and the $-x$ wire together to form a loop.

Lemma 5.1 *The wires can be joined together such that no cubes are adjacent unless they are in fact adjacent cubes from a single wire.*

Proof. As the V-gadgets are sufficiently spaced out, it is very obvious that we can make the wires avoid each other, besides when they meet at the V- and C-gadgets. \square

This finishes the construction.

5.2. Proof of correctness. Fix a region Γ constructed above from a 3-SAT expression \mathcal{C} . It remains to show that the tilings of Γ and the 1-in-3 satisfying assignments of \mathcal{C} are in bijective correspondence.

First, consider a tiling of Γ . We say that the *phase* of an endpoint of a wire is *positive* if the wire contains the tile that covers that endpoint, and *negative* otherwise.

Lemma 5.2 *A wire of even length has the same phase at both endpoints, while a wire of odd length has opposite phases at its two endpoints.*

Proof. As in the proof of Theorem 1.3, the slabs must line up along the length of the wire, since the various wires do not intersect or come near each other (except at the endpoints). \square

The following is immediate by observing the way slabs can come together to tile the V-gadget:

Lemma 5.3 *There are two ways to tile each V-gadget locally. In either tiling, the three + wires have the same phase at their variable endpoints, and the three – wires have the opposite phase at their variable endpoints as the + wires.*

Let the *phase* of the V-gadget be the common phase of the + wires at their variable endpoints.

Lemma 5.4 *The V-gadgets corresponding to v_k (or its negation \bar{v}_k) all have the same phase.*

Proof. Consider two adjacent V-gadgets in the construction shown in Figure 9. Suppose the gadget on the left has positive phase (the opposite situation is similar). Then by Lemma 5.3, its $-z$ wire has negative phase at the variable end. This wire is linked with the $+z$ wire of the gadget on the right. Since both endpoints have odd parity, the wire has odd length, thus by Lemma 5.2, the variable endpoint of the $+z$ wire of the gadget on the right (hence the gadget) has positive phase. \square

This allows us to assign a truth value to the variable v_k :

Lemma 5.5 *Let the boolean assignment of v_k be the common phase of the V-gadgets corresponding to v_k . This defines a 1-in-3 satisfying assignment.*

Proof. Fix a clause (c_1, c_2, c_3) and consider the corresponding even C-gadget. There are three wires connecting the corresponding V-gadgets to the center of the C-gadget. The variable endpoints of the wires have odd parity, and the free endpoints are all at the center of the C-gadget, which has even parity. As such, the wires have even lengths; thus by Lemma 5.2, each wire has the same phase on both its endpoints.

The center of the C-gadget is tiled by some slab, which is contained in precisely one of its three wires. Without loss of generality, suppose only the wire corresponding to c_1 is of positive phase (at both of its endpoints). Recall that the phase of a V-gadget is the phase of its + wires at their variable endpoints. If $c_i = v_k$ is an unnegated variable, then the $+y$ wire was used, and thus v_k (and hence c_i) is positive if and only if $i = 1$. On the other hand, if $c_i = \bar{v}_k$ is a negated variable, then the $-y$ wire was used. Thus v_k is *negative* if and only if $i = 1$, but then c_i is again positive if and only if $i = 1$. \square

This shows that the assignment is indeed 1-in-3 satisfying. We thus have a map from the tilings of Γ to the satisfying assignments of \mathcal{C} .

For the inverse, suppose we have a 1-in-3 satisfying assignment. Simply tile each V-gadget according to the assignment in the obvious way. It only remain to tile the wires. However, by construction, once we choose the phase of a V-gadget, the tiling of all six wires emanating from that region is forced.

The $\pm z$ wires fit together because the phases of the V-gadgets are consistently assigned, in accordance to the truth value of the variable in the given satisfying assignment.

Recall that for each V-gadget, its $+x$ and $-x$ wires form a loop. Since the two endpoints both have odd parity, the length of the wire is odd. By Lemma 5.2, the two endpoints have opposite phases in a tiling, which is precisely what the V-gadget would produce as in Lemma 5.3.

It is obvious that the $\pm y$ wires fill up the even C-gadgets perfectly because the boolean assignment was 1-in-3 satisfying. For the odd C-gadgets, notice that since we take the leftover wires, precisely one of the three is *negative* at its variable endpoint. However, since we join them at a cube pair with odd parity, the wires are of odd lengths, thus precisely one of the three is positive on the free endpoint at the C-gadget, as needed.

This establishes the desired bijective correspondence, and implies that the construction is indeed a parsimonious reduction from 1-IN-3 SAT to the slab tiling problem, concluding the proof of

Theorem 1.1. Note that this also implies that the associated counting problem is $\#P$ -complete, yielding the first part of Theorem 1.5; in the next section we prove a stronger result.

6. PROOF OF THEOREM 1.2 AND THE SECOND PART OF THEOREM 1.5

6.1. **Reduction construction.** Given a region Γ_0 as constructed from the previous section, we will enlarge it to a contractible region $\Gamma \supset \Gamma_0$, such that $\Gamma' = \Gamma \setminus \Gamma_0$ is a frozen subregion of Γ .




Figure 12: A variable V-gadget, shown in two perspectives.

First we modify the variable V-gadget. The following description is relative to a V-gadget placed at $(0, 0, 0)$. We must make these adjustments to each V-gadget used. Recall that in the proof of Theorem 1.1, each $+z$ wire was bent as in Figure 9. Let us make this precise by making it bend towards the $-y$ direction when it hits $z = 3$, and bend towards $-z$ when it hits $y = -5$ (see Figure 12). We define a region X , shown in Figure 13a, that will fill the hole this loop makes. Indeed, let $X = \{-1\} \times [-5, 0] \times [0, 3] \cup \{(0, -1, 0), (0, -1, 1)\} \setminus \{(-1, 0, 0), (-1, 0, 1)\}$. The unique




Figure 13: Regions used to modify the V-gadget.

tiling is indicated in the figure; the labels will be used in the proof of Lemma 6.2. Similarly, the $\pm x$ wires are linked together to form a loop. Let us bend them towards the $+z$ direction when they hit 4 and -3 , respectively, and bend towards each other when they hit $z = 5$. Define a region that will fill this hole, shown in Figure 13b: Take

$$[-3, 4] \times \{2\} \times [0, 5],$$

add the following six cubes

$$(-4, 2, 0), (-4, 2, 1), (1, 3, 0), (2, 3, 0), (1, 1, 1), (2, 1, 1),$$

and remove these two cubes

$$(0, 2, 0), (0, 2, 1).$$

Call this region Y . Note that the slab labeled 4 in the figure extends behind and is hidden from view. We will add regions X and Y to the V-gadget. This makes the V-gadget contractible.

Let $Z = [-10, 10n + 11] \times [0, 30(t + 1) + 1] \times \{-1\} \subset \mathbb{Z}^3$, where n and t are the number of variables and clauses, respectively. We then modify Z with a *hole H-gadget* around each V-gadget to allow its $\pm z$ wires to pass through. Indeed, Figure 14 shows the construction in the $z = -1$




Figure 14: A hole H-gadget.

plane viewed from below. The surrounding slabs on the boundary match that of the unique tiling of Z . The two cube pairs missing allow the $\pm z$ wires to pass through (spaced out according to the precise construction of the V-gadget in Figure 12). The cube pairs in $z = -2$ plane are necessarily tiled with its unique neighbor in the $z = -1$ plane. Let $X_k = \{10k - 1\} \times [0, 30(t + 1) + 1] \times [-9, -2]$ for each $k \in [n]$. Here we assume that the linking of the $\pm z$ wires (see Figure 9) are done at z -coordinates, say, -5 and -7 , which are in $[-9, -2]$. Let Γ' be the disjoint union of the X and Y for each V-gadget, the region Z , and the X_k for each k . This finishes the construction.

6.2. Proof of correctness. It remains to check that that the construction works. The following two lemmas easily imply the result.

Lemma 6.1 *The region Γ constructed above is contractible.*

Proof. Notice that there are no holes in the $z = -1$ plane. In $z < -1$, we may deformation retract the wires onto the X_k plates, and then deformation retract these plates onto the $z = -1$ plane.

We now check that the modified V-gadget is contractible. The center of an original V-gadget was an example of a non-contractible region shown in Figure 1. Indeed, the interior will leave a point hole in the middle; however, it is filled by the slab labeled 4 in Figure 13b. The loop formed by the $+z$ wire bending down is filled by X , while the loop formed by joining the $\pm x$ wires is filled by Y . Moreover, the $\pm y$ wires are all lying along Z so we may deformation retract everything to the $z = -1$ plane. We omit the (easy) details. \square

Lemma 6.2 *The subregion $\Gamma' \subset \Gamma$ is frozen.*

Proof. Consider a cube $c \in \Gamma$. If it forms a $3 \times 1 \times 1$ region with a cube pair, then no slab containing this cube may contain its neighbor that is part of the cube pair. Not counting such cube pairs, if c only has two neighbors left, then c must be tiled by a slab covering both these neighbors. If

this happens, we call c *isolated* and say that it *forces* the unique slab containing c . Just like in the proof of Theorem 1.4, we may inductively remove these forced slabs, which are obviously frozen.

Consider region X_k . Notice that the cube at the corner $(10k-1, 1, -9)$ is isolated. After removing the slab forced by it, we may consider the new corner next to it at $(10k-1, 3, -9)$. Inductively, we may remove the cubes with $z = -9$ or -8 . Now when looking at $z = -7$, we may run into cubes that neighbor $\pm z$ wires. However, these wires are made of cube pairs that form $3 \times 1 \times 1$ regions with the cubes in question. As such, we may continue the removal process unhindered. It is clear that in this manner, we may remove all cubes from X_k .

Similarly, we remove most slabs from Z , starting from the boundary and working our way in. The ones sticking down out of the $z = -1$ plane can also be removed, leaving us with precisely two slabs around the hole for the $-z$ wire from the H-gadget (labeled with 1 in Figure 14). As for each V-gadget, we remove all the cubes from X and Y . In Figure 13, the forced slabs are labeled with a sample removal order, where each number is positioned on the isolated cube used at each step. Finally, we may remove the leftover slabs from Z . \square

Now, given a 1-IN-3 SAT expression \mathcal{C} , we may take a region Γ_0 constructed in the proof of Theorem 1.1 and enlarge it to a contractible region Γ as described above. Since $\Gamma' = \Gamma \setminus \Gamma_0$ is frozen, Γ is tileable if and only if Γ_0 is tileable. These are tileable if and only if \mathcal{C} is 1-in-3 satisfiable, thus Theorem 1.2 is proved. Moreover, since both reductions are parsimonious, we conclude the second counting result in Theorem 1.5 as well.

7. GENERALIZED DOMINOES IN HIGHER DIMENSIONS

Let $\mathcal{D}_r^d \subset \mathbb{Z}^d$ be a $2 \times \dots \times 2 \times 1 \times \dots \times 1$ region with r number of 2's and $d - r$ number of 1's. We call this a *generalized domino* in d dimensions of rank r . Call the r coordinate directions that are 2 cubes wide *fat*. Previously we were concerned with tiling by dominoes \mathcal{D}_1^3 and slabs \mathcal{D}_2^3 in three dimensions.

Theorem 7.1 *For $2 \leq r < d$, tiling (contractible) d -dimensional regions with \mathcal{D}_r^d is NP-complete. Similarly, for $1 \leq r < d$ and $d \geq 3$, tiling (contractible) d -dimensional regions with \mathcal{D}_r^d is #P-complete.*

Note that in other cases of r and d , these problems are in P (see next section).

Proof. We reduce specific tiling problems in \mathbb{Z}^3 to higher dimensions. Let us first prove the statements without the contractibility constraint.

Suppose $2 \leq r < d$. Let $\Gamma \subset \mathbb{Z}^3$ be a region as afforded by the proof of Theorem 1.1. Let $\Gamma' = \Gamma \times \{0, 1\}^{r-2} \times \{0\}^{d-r-1}$, that is, $\{(x_1, \dots, x_d) \in \mathbb{Z}^d : (x_1, x_2, x_3) \in \Gamma, x_4, \dots, x_{r+1} \in \{0, 1\}, x_{r+2} = \dots = x_d = 0\}$. Consider a tiling of Γ' by \mathcal{D}_r^d . Notice that there are no $2 \times 2 \times 2$ subregion in Γ , thus each tile must be oriented with two of its fat directions in the first three coordinates, the remaining $r - 2$ fat directions in the next $r - 2$ coordinates. This means that any \mathcal{D}_r^d -tiling of Γ' induces a \mathcal{D}_2^3 -tiling of Γ . It is easy to see that this correspondence is actually a bijection, finishing the proof of NP-completeness, and also #P-completeness when $r \geq 2$.

The case of $r = 1$ is straightforward. Take Γ from the proof of Theorem 1.3. Define $\Gamma' = \Gamma \times \{0\}^{d-3}$, and follow the argument above.

To prove the result for contractible regions, proceed in the same manner, and take the regions constructed in the proof of the corresponding theorems. For $2 \leq r < d$, take the region Γ from the proof of Theorem 1.2, and define Γ' in the same way as above. Notice, however, that the above strategy does not yield a reduction from \mathcal{D}_2^3 -tilings to \mathcal{D}_r^d -tilings in general. Indeed, we used the fact that the specific $\Gamma \subset \mathbb{Z}^3$ did not contain $2 \times 2 \times 2$ subregions, which is no longer the case. The outlined argument actually still works, but more care must be taken. A tile might now be oriented

with three of its fat directions in the first three coordinates. This means a \mathcal{D}_r^d -tiling of Γ' will, *a priori*, induce a tiling of Γ with slabs and the $[2 \times 2 \times 2]$ cube. However, in the proof of Theorem 1.2, we see that no tilings of Γ by slabs contain a $2 \times 2 \times 2$ subregion tiled by two slabs. As such, all \mathcal{D}_r^d tiles in Γ' are still constrained to be oriented as before, with precisely two fat directions in the first three coordinates. This implies the result in this case.

The same approach works for the $r = 1$ case as well. Take Γ as in the proof of Theorem 1.4, and define Γ' as above. Now notice that in any tiling of Γ , there is no $2 \times 2 \times 1$ subregion of Γ that is tiled by two dominoes, which proves the result. \square

8. FINAL REMARKS AND OPEN PROBLEMS

8.1. Historically, the tiling problems played a crucial role in the developments of modern theoretical computer science. The tileability of the whole plane with a set of tiles was shown to be undecidable [Ber, Rob]. A version of the finite tileability problem was stated to be NP-complete in Levin's original paper [Lev], where he (independently) defined NP-completeness. Finally, Theorem 1.3 is one of the first few applications of #P-completeness, developed by Valiant [Val1].

8.2. In the context of statistical physics, the domino tiling problem is called the *dimer problem*, and has a long history. The classical results of Fisher [Fis] and Kasteleyn [Kas] express the number of domino tilings of finite regions in the plane as a certain Pfaffian, equal to a square root of a determinant. A closely related *monomer-dimer model* is #P-complete already in the plane [Jer] (see also [Vad]). Let us mention that both models can be polynomially approximated (see [JSV, KRS]). We refer to [Tho] for graph-theoretic reasons precluding the Pfaffian method in higher dimensions, and to [DG, HK] (see also [HN]) for the general hardness results on *graph homomorphisms*, a concept generalizing perfect matchings.

8.3. In the *general tileability problem*, a set \mathbf{T} of tiles is fixed, and one considers tilings with parallel translations of copies of tiles $T \in \mathbf{T}$. For a single tile T , the tileability problem is a special case, where \mathbf{T} consists of all reflections and rotations of T . Let us briefly elaborate on the state of art of these tiling problems as they pertain to our results.

In the plane, when $|\mathbf{T}| = 1$, *i.e.*, when translates of a single tile are used, the tileability problem is linear in the area. However, already for $\mathbf{T} = \{2 \times 1, 1 \times 3\}$, the tileability is NP-complete [BNRR]. It is not known whether the corresponding counting problem is #P-complete. The smallest set \mathbf{T} for which #P-completeness is known is the set of four rotations of the L -tromino and a $[2 \times 2]$ square [MR].

For simply connected regions in the plane, the authors recently found a set of 23 Wang tiles, which they showed to be NP-complete and #P-complete [PY]. This construction can be further decreased to 15 tiles [Yang], and the authors conjecture that only three rectangles suffices (see [PY]). In contrast to the decision problem, it is still unclear how much the simple connectivity of regions helps counting the number of tilings, other than making the reductions harder and more technical.

8.4. In [Val2], Valiant's goal is to show the hardness of the number of perfect matchings of grid graphs, defined as subgraphs (of the unit distance graph) of \mathbb{Z}^3 . This is slightly weaker than the hardness of domino tilings of 3-dim regions, since the latter problem corresponds to *induced* subgraphs of \mathbb{Z}^3 . However, Valiant's proof can be slightly modified, from subgraphs of $[2 \times n \times n]$ to induced subgraphs of $[3 \times n \times n]$, to achieve the same goal. The proof we present in Section 3 is in fact a variation of the argument in [Val2].

8.5. For tiling in three dimensions, we consider the domino and the slab. Here each tile has three distinct orientations. If we only allow two out of the three orientations in either case, each problem becomes equivalent to the ordinary domino tiling problem in two dimensions, and hence is in P. Of course, two orientations suffices for the tromino tile $I_3 = [3 \times 1 \times 1]$ (or the $[3 \times 3 \times 1]$ tile) to guarantee NP-completeness, since we have NP-completeness for the $[3 \times 1]$ tile.

Note that the *augmentation approach*, mentioned in the introduction, allows us to show the NP-completeness of the tileability problem with I_3 of contractible regions in \mathbb{R}^3 . To see this, start with a flat region and color the holes in a checkerboard fashion. Now fill each unit cube with an up or down vertical tromino, depending on the color (see Figure 15). The contractible region thus obtained is tileable with I_3 if and only if the starting region is tileable by $[3 \times 1]$.




Figure 15: A construction of contractible 3-dim regions for tromino tilings.

8.6. The conditions in Theorem 7.1 are best possible. Indeed, if $r = 0$ or $r = d$, we are tiling with cubes of side length 1 or 2, respectively, both of which are linear in the volume. The only remaining cases are the decision problem with \mathcal{D}_1^d , and the counting problem of ordinary dominoes in dimension $d = 2$, both of which are in P (see e.g. [LP]).

8.7. The idea of graph embedding into a grid is quite standard. In this context of domino tilings it was used in [DKLM, Val2], and for other small sets of tiles in [BNRR, MR, PY]. Of course, the technical details are somewhat different in each case.

8.8. It would be interesting to see if Theorem 1.4 holds for contractible regions inside a $[2 \times n \times n]$ brick, as they have a rather simple geometric structure. Our proof gives only a width 4 bound, but we believe that a width 3 modification can be made without difficulty. We should mention that for every fixed c , the counting problem is polynomial for regions inside a $[c \times c \times n]$ brick.

8.9. In the plane, there are several other generalizations of domino tilings. Notably, the first author introduced and studied *ribbon tilings* (see [Pak]). Another interesting set of generalized dominoes $\mathbf{T}_n = \{2^k \times 2^{n-k}, 0 \leq k \leq n\}$ was studied in [Korn]. Both sets satisfy the *local move* property: every two tilings of a s.c. region Γ can be obtained by a sequence of *flips*, each involving a fixed number of tiles (2 in both cases, see [Pak]). For the (usual) domino tilings this is a classical property going back to Kasteleyn and famously proved by Thurston via the *height functions* [Thu].

Of course, for the domino tilings of general regions there is no local move property, as large cycles can involve an unboundedly many tiles. However, for contractible regions, one would naturally assume that domino tilings and slab tilings are connected by flips corresponding to different tilings of $[2 \times 2 \times 2]$. This is, in fact, false. As an early prequel to the constructions in the proofs of theorems 1.2 and 1.4, we state the following easy result.

Proposition 8.1 *Domino tilings of contractible regions in \mathbb{R}^3 does not have the local move property. The same result holds for slab tilings.*

The proof is apparent from Figure 16 (cf. Figure 15). In the first case, the middle dominoes alternate between up and down. In the second case, slabs are all vertical, with the middle slabs all one layer below.



Figure 16: Cross sections of large regions in \mathbb{R}^3 with exactly two tilings, by dominoes and by slabs, respectively.

Acknowledgements. We are very grateful to Cris Moore for helpful conversations; in particular, Theorem 1.1 arose as a question in our discussions. The first author is partially supported by the NSF and BSF grants. The second author is supported by the NSF under Grant No. DGE-0707424.

REFERENCES

- [BNRR] D. Beauquier, M. Nivat, É. Rémila and M. Robson, Tiling figures of the plane with two bars, *Comput. Geom.* **5** (1995), 1–25.
- [Ber] R. Berger, The undecidability of the domino problem, *Mem. AMS* **66** (1966), 72 pp.
- [Cha] T. Chaboud, Domino tiling in planar graphs with regular and bipartite dual, *Theor. Comp. Sci.* **159** (1996), 137–142.
- [CH] N. Creignou and M. Hermann, On #P-completeness of some counting problems, Research Report 2144, INRIA, 1993.
- [DL] P. Dagum and M. Luby, Approximating the permanent of graphs with large factors, *Theoret. Comput. Sci.* **102** (1992), 283–305.
- [DKLM] S. Datta, R. Kulkarni, N. Limaye and M. Mahajan, Planarity, determinants, permanents, and (unique) matchings, *J. ACM Trans. Comput. Theory* **1** (2010), No. 3.
- [DG] M. Dyer and C. Greenhill, The complexity of counting graph homomorphisms, *Random Structures Algorithms* **17** (2000), 260–289.
- [Fis] M. E. Fisher, Statistical mechanics of dimers on a plane lattice, *Phys. Rev. (2)* **124** (1961), 1664–1672.
- [HK] P. Hell and D. G. Kirkpatrick, On generalized matching problems, *Inform. Process. Lett.* **12** (1981), 33–35.
- [HN] P. Hell and J. Nešetřil, Counting list homomorphisms for graphs with bounded degrees, in *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.* **63**, AMS, Providence, RI, 2004, 105–112.
- [Jer] M. Jerrum, Two-dimensional monomer-dimer systems are computationally intractable, *J. Statist. Phys.* **48** (1987), 121–134; Erratum in **59** (1990), 1087–1088.
- [JSV] M. Jerrum, A. Sinclair and E. Vigoda, A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries, *J. ACM* **51** (2004), 671–697.
- [Kas] P. W. Kasteleyn, The statistics of dimers on a lattice I, *Physica* **27** (1961), 1209–1225.
- [KK] C. Kenyon and R. Kenyon, Tiling a polygon with rectangles, in *Proc. 33rd FOCS* (1992), 610–619.
- [KRS] C. Kenyon, D. Randall and A. Sinclair, Approximating the number of monomer-dimer coverings of a lattice, *J. Statist. Phys.* **83** (1996), 637–659.
- [Ken] R. Kenyon, An introduction to the dimer model, in *ICTP Lect. Notes XVII*, Trieste, 2004.
- [Korn] M. Korn, *Geometric and algebraic properties of polyomino tilings*, MIT Ph.D. thesis, 2004; available at <http://dspace.mit.edu/handle/1721.1/16628>
- [Lev] L. Levin, Universal sorting problems, *Problems Inf. Transm.* **9** (1973), 265–266.
- [LP] L. Lovász and M. D. Plummer, *Matching theory*, AMS, Providence, RI, 2009.
- [MR] C. Moore and J. M. Robson, Hard tiling problems with simple tiles, *Discrete Comput. Geom.* **26** (2001), 573–590.
- [Pak] I. Pak, Tile invariants: New horizons, *Theor. Comp. Sci.* **303** (2003), 303–331.
- [PY] I. Pak and J. Yang, Tiling simply connected regions with rectangles, *J. Combin. Theory, Ser. A*, to appear.
- [Rém] É. Rémila, Tiling a polygon with two kinds of rectangles, *Discrete Comput. Geom.* **34** (2005), 313–330.
- [Rob] R. M. Robinson, Undecidability and nonperiodicity for tilings of the plane, *Invent. Math.* **12** (1971), 177–209.
- [Sch] T. Schaefer, The complexity of satisfiability problems, in *Proc. 10th STOC* (1978), 216–226.
- [Tho] R. Thomas, A survey of Pfaffian orientations of graphs, in *Proc. ICM Madrid*, Vol. III, 963–984, Zürich, 2006.
- [Thu] W. P. Thurston, Conway’s tiling groups, *Amer. Math. Monthly* **97** (1990), 757–773.
- [Til] B. de Tilière, *The dimer model in statistical mechanics*, Lecture notes at Swiss Doctoral Program in Mathematics and the EPFL doctoral school, University of Neuchâtel, September 2008; available at http://proba.jussieu.fr/~detiliere/Cours/polycop_Dimeres.pdf
- [Vad] S. P. Vadhan, The complexity of counting in sparse, regular, and planar graphs, *SIAM J. Comput.* **31** (2001), 398–427.
- [Val1] L. G. Valiant, The complexity of enumeration and reliability problems, *SIAM J. Comput.* **8** (1979), 410–421.
- [Val2] L. G. Valiant, Completeness classes in algebra, in *Proc. 11th STOC* (1979), 249–261.
- [Yang] J. Yang, Ph.D. thesis, in preparation.