

Continuous Combinatorics of Abelian Group Actions

Su Gao

Steve Jackson

Edward Krohne

Brandon Seward

Author address:

SCHOOL OF MATHEMATICAL SCIENCES AND LPMC, NANKAI UNIVERSITY,
TIANJIN 300071, P.R. CHINA

Email address: `sgao@nankai.edu.cn`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF NORTH TEXAS, 1155 UNION
CIRCLE #311430, DENTON, TX 76203, U.S.A.

Email address: `jackson@unt.edu`

Email address: `krohneew@gmail.com`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA SAN DIEGO,
SAN DIEGO, CA 92093, U.S.A.

Email address: `b.m.seward@gmail.com`

Contents

Introduction	7
Chapter 1. Basic Definitions and Results	13
1.1. Bernoulli shift actions	13
1.2. Hyper-aperiodic elements	15
1.3. Toast, an application of hyper-aperiodicity	17
1.4. A hyper-aperiodic element with partial periodicity	19
1.5. The continuous chromatic number of $F(2^{\mathbb{Z}^2})$	20
1.6. CBS line sections and single line sections	23
1.7. Continuous cocycles	26
1.8. Marker distortion	29
Chapter 2. The Twelve Tiles Theorem	35
2.1. From proper colorings to subshifts of finite type	36
2.2. The two tiles theorem in one dimension	39
2.3. Statement of the twelve tiles theorem	42
2.4. Proof of the negative direction	49
2.5. Proof of the positive direction	54
2.6. Continuous structurings	61
2.7. Higher dimensions	66
Chapter 3. Applications of the Twelve Tiles Theorem	69
3.1. Basic applications of the Twelve Tiles Theorem	69
3.2. Continuous subactions and regular tilings	75
3.3. Undirected graph homomorphisms and the reduced homotopy group	81
3.4. Negative conditions for graph homomorphisms	85
3.5. A positive condition for graph homomorphisms	90
Chapter 4. Applications to Decidability Problems	99
4.1. Three general problems	99
4.2. The subshift problem for $F(2^{\mathbb{Z}})$	100
4.3. The subshift problem for $F(2^{\mathbb{Z}^2})$	102
4.4. The graph homomorphism problem	113
4.5. The tiling problem	121
Bibliography	125

Abstract

This paper develops techniques which are used to answer a number of questions in the theory of equivalence relations generated by continuous actions of abelian groups. The methods center around the construction of certain specialized hyper-aperiodic elements, which produce compact, free subflows with useful properties. For example, we show that there is no continuous proper 3-coloring of the Schreier graph on $F(2^{\mathbb{Z}^2})$, the free part of the shift action of \mathbb{Z}^2 on $2^{\mathbb{Z}^2}$. With earlier work of [11] this computes the continuous chromatic number of $F(2^{\mathbb{Z}^2})$ to be exactly 4. Combined with marker arguments for the positive directions, our methods allow us to analyze continuous homomorphisms into graphs, and more generally equivariant maps into subshifts of finite type. We present a general construction of a finite set of tiles for $2^{\mathbb{Z}^n}$ (there are 12 for $n = 2$) such that questions about the existence of continuous homomorphisms into various structures reduce to finitary combinatorial questions about the tiles. This tile analysis is used to deduce a number of results about $F(2^{\mathbb{Z}^n})$.

2010 *Mathematics Subject Classification.* Primary 54H05, 05C15; Secondary 03E15, 05C99.

Key words and phrases. equivalence relation, hyperfinite, Cayley graph, Schreier graph, chromatic number, graph homomorphism, tiling, subshift of finite type, hyper-aperiodic.

Research of the first author was partially supported by the U.S. NSF grants DMS-1201290 and DMS-1800323, and NSFC grants 12250710128 and 12271263. Research of the second author was partially supported by the NSF grants DMS-1201290 and DMS-1800323. Some results in this paper were included in the third author's Ph.D. dissertation at the University of North Texas, with the first and second authors serving as co-advisors, and the research was partially supported by the NSF grant DMS-0943870. The authors thank the anonymous referee for two rounds of careful reading of earlier versions of the manuscript and for useful comments which resulted in significant improvements.

Introduction

The last several decades have seen much development of *invariant descriptive set theory*, or the theory of definable equivalence relations E on standard Borel spaces X . Most often the term “definable” is taken to mean “Borel,” that is, E is a Borel subset of $X \times X$. This is a broad field of investigation, overlapping with much of topological dynamics, functional analysis, ergodic theory, and model theory, and interacting with computability theory, geometric group theory, and in general algebra and combinatorics (see [10] for an introduction). Understanding of these definable equivalence relations is often accomplished by imposing and investigating additional structures on the equivalence classes. In recent years the study of combinatorial problems relating to such structures on definable equivalence relations has emerged as a field of its own interest, which is known as the field of *Borel combinatorics*. The recent survey papers of Kechris [20] and Kechris, Marks [21] provide an excellent account of this subject. The results and methods of the present paper can be viewed as a contribution to this emerging field.

This paper is chiefly concerned with *countable* Borel equivalence relations, that is, Borel equivalence relations each of whose equivalence classes is countable. By a fundamental theorem of Feldman–Moore [9], every countable Borel equivalence relation E is the *orbit equivalence relation* of a Borel action by a countable group G , that is, there is an action $\cdot : G \curvearrowright X$ such that xEy if and only if $\exists g \in G (g \cdot x = y)$. In view of this theorem, it is natural to study countable Borel equivalence relations “group by group,” that is, to begin with the simplest groups and work our way up to progressively more complicated groups. The simplest infinite group is $G = \mathbb{Z}$, and its actions are the subject of study in the classical theory of dynamical systems. This paper deals with actions of the next group $G = \mathbb{Z}^2$ and, to some extent, the other finitely generated abelian groups $G = \mathbb{Z}^n$ (the general finitely generated case $G = \mathbb{Z}^n \times A$, where A is finite abelian, is not essentially different). Along these lines, the next steps would concern general abelian groups, nilpotent groups, solvable groups, amenable groups, etc. In [11] it was shown that Borel actions of countable abelian groups are hyperfinite, in [25] this was extended to countable locally nilpotent groups, and in [5] this was extended further to all countable polycyclic groups (which includes some groups of exponential growth). In a related direction, several authors have considered the Borel combinatorics of actions of groups of subexponential growth [6, 7, 27].

This approach should not strike anyone as very ambitious, but we will demonstrate that there are reasons to slow down and examine thoroughly the properties of actions of seemingly simpler groups before we consider the more complicated ones. Even for actions of simpler groups such as \mathbb{Z}^n , many fundamental questions remain open, not to mention the significant new difficulties when we attempt to

generalize known results for simpler groups to more complicated ones. Two important examples are the *hyperfiniteness problem* and the *union problem*. A countable Borel equivalence relation E is *hyperfinite* if it can be written as the increasing union $E = \bigcup E_n$ of finite subequivalence relations E_n (by “finite” we mean each equivalence class of E_n is finite). Kechris and Weiss asked for which groups G must all their Borel actions be hyperfinite. Specifically, they asked whether this is true for all amenable groups G . The union problem asks whether the increasing union of hyperfinite equivalence relations is hyperfinite. In view of recent progress, both of these questions are now understood to be closely related to the question of what types of *marker structures* can be put on equivalence relations generated by actions of G . By a “marker structure” we mean a Borel set $M \subseteq X$ which is a *complete section* (i.e., M meets every E -class) and often also *co-complete* (i.e., $X - M$ is also a complete section). Usually we are interested in M having some geometrical or combinatorial structure which can be used to get useful results about E . We thus see that several of the fundamental problems in the theory of countable Borel equivalence relations are closely connected with the Borel combinatorics of these equivalence relations.

The problems considered in the present paper are motivated by the observation that the most straightforward constructions of marker structures on the equivalence relations often yield *relatively clopen* structures instead of more complex Borel structures. However, as we impose more requirements on the marker structures, whether this phenomenon will continue becomes a question. In this paper we will concentrate on questions concerning *continuous combinatorics*, which is about relatively clopen structures and the existence of continuous mappings, as opposed to Borel structures and the existence of Borel mappings considered in Borel combinatorics. The decision to focus on this set of questions is largely driven by the richness of methods and results that are already present in this sub-area. Our results about the general Borel combinatorics of abelian group actions will be presented elsewhere (the forthcoming [13] and [12]).

We explain in some more detail the connections and differences of considerations in Borel combinatorics and continuous combinatorics of countable abelian group actions. A good example of this concerns the currently known results on the hyperfiniteness problem. Slaman and Steel [26] and independently Weiss [29] proved the basic result that E is hyperfinite if E can be generated by a Borel action of \mathbb{Z} . To do this they defined a sequence $M_1 \supseteq M_2 \supseteq \dots$ of Borel marker sets which are decreasing and $\bigcap_n M_n = \emptyset$. We call such a sequence a *Slaman–Steel marker sequence*. Some years later, Weiss showed that every Borel action of \mathbb{Z}^n is hyperfinite by showing that every free Borel action of \mathbb{Z}^n admits a marker set with a bounded geometry. More precisely, we have the following basic marker lemma (a proof can be found in [11]).

In the lemma below, for a free action $\mathbb{Z}^n \curvearrowright X$ we define an extended metric $\rho: X \times X \rightarrow [0, +\infty]$ by the rule: $\rho(x, y) = \infty$ if $[x] \neq [y]$, and otherwise if $y = g \cdot x$, where $g = (a_1, \dots, a_n)$, then $\rho(x, y) = \|g\| = \max\{|a_1|, \dots, |a_n|\}$. We will also occasionally refer to the extended metric ρ_1 on $X \times X$ where we use instead $\|g\|_1 = \sum_{i=1}^n |a_i|$. The following lemma holds (by the same proof) for both metrics ρ and ρ_1 .

LEMMA 0.0.1. *For every integer $d > 1$ there is a relatively clopen set $M_d \subseteq F(2^{\mathbb{Z}^n})$ (where $F(2^{\mathbb{Z}^n})$ is the free part of the shift action, see §1 for the precise definition) satisfying:*

- (1) *For all $x \neq y$ in M_d , $\rho(x, y) > d$.*
- (2) *For all $x \in X$, there is a $y \in M_d$ with $\rho(x, y) \leq d$.*

Extending this hyperfiniteness result to the general (infinitely generated) abelian group case runs into extra difficulties as this type of marker structure does not seem as useful for studying these group actions, and this case of the hyperfiniteness problem remained open for some time. In [11] it was shown that the Borel action of any countable abelian group is hyperfinite. The key ingredient in the argument is a stronger type of marker structure on the equivalence relation, which we refer to as an *orthogonal marker structure*. First, the “marker regions” (finite subequivalence classes associated to each marker point) are essentially rectangles (the actual construction first puts down rectangles and then modifies them in finitely many steps to fractal-like regions). Second, the rectangular regions for different levels M_{d_1} and M_{d_2} must intersect in “general position”, that is, no close parallel faces (this is where the terminology “orthogonal” comes from). The fact that the marker regions have a more regular geometric shape is important in the arguments. Thus, the question of what types of regular marker structures can be put on equivalence relations generated by actions of G has deep connections with the theory of these equivalence relations.

Many interesting combinatorial questions about the equivalence relation E also end up being connected with delicate questions about marker structures; indeed this is one of the main themes of this paper. The orthogonal marker method mentioned above is one recent technique which can be used to produce intricate Borel structures on E . There is one important consideration, however, which applies to constructions that rely upon a sequence of marker sets (such as the orthogonal marker method). Although the basic marker construction for producing M_d satisfying (1) and (2) above produces *clopen* sets M_d , in passing to the “limit” the final construction generally only produces Borel sets. Thus, this method often cannot be used to produce clopen or continuous structures on E . In fact, it was realized some time ago that there are limitations on what continuous structures can be attained. A simple example of this occurs in the case of Slaman–Steel markers on actions of \mathbb{Z} . We noted above that a Borel sequence $M_1 \supseteq M_2 \supseteq \dots$ in $F(2^{\mathbb{Z}})$ with $\bigcap_n M_n = \emptyset$ can always be found, but in [11] it was shown that there cannot be a sequence of (relatively) clopen marker sets $M_d \subseteq F(2^{\mathbb{Z}})$ with this property. The basic method is that of a *hyper-aperiodic element* (also called *2-coloring* in earlier papers), which will be defined in §1. In [14] and [15] it was shown that hyper-aperiodic elements exist for any countable group G , and a number of consequences of this are deduced. However, for many of the more delicate questions, the mere existence of hyper-aperiodic elements is not sufficient. One of the main themes of this paper is that by constructing hyper-aperiodic elements on G with certain specialized properties, one can extend the method greatly to answer many of these more delicate questions. We emphasize that these hyper-aperiodic element methods are used to rule out continuous structures on E , and thus are complementary to the marker construction methods which are used to produce clopen or Borel structures on E . Alternatively, we may view the hyper-aperiodic element methods as showing that

particular orbits with specified properties must exist given any continuous assignment of structure (e.g., a proper 3-coloring of the Schreier graph) to the E classes. In [12] it was shown that the hyper-aperiodic element method can be extended in some cases to the Borel context via a forcing construction. However, not all of the “negative” (that is, non-structure) results at the continuous level obtained by the hyper-aperiodic element arguments can be extended to the Borel level, as the orthogonal marker methods can be used to obtain “positive” Borel constructions. An example of this is the chromatic number problem for $F(2^{\mathbb{Z}^n})$ where we seek to determine both the continuous chromatic number $\chi_c(F(2^{\mathbb{Z}^n}))$ and the Borel chromatic number $\chi_b(F(2^{\mathbb{Z}^n}))$ for $F(2^{\mathbb{Z}^n})$ (these are the minimum number of colors needed to properly color the Schreier graphs of $F(2^{\mathbb{Z}^n})$ with continuous, respectively Borel, functions). For $n > 1$ we show in this paper that there is no continuous proper 3-coloring of $F(2^{\mathbb{Z}^n})$, while recently the authors have shown using the orthogonal marker methods that there is a Borel proper 3-coloring of $F(2^{\mathbb{Z}^n})$ (the Borel result will appear in [13]). Thus, for $n > 1$ we have:

$$3 = \chi_b(F(2^{\mathbb{Z}^n})) < \chi_c(F(2^{\mathbb{Z}^n})) = 4.$$

This result shows a remarkable difference between the continuous and Borel combinatorics of actions of \mathbb{Z}^n .

For a number of results, it is still not known whether the negative results given by hyper-aperiodic element methods can be extended to the Borel context. There is a clear tension between attempting to do so (say by forcing arguments) and the orthogonal marker method. Delineating the exact boundaries of these methods remains an important problem.

As surveyed in [20], there is a rich theory of *descriptive graph combinatorics* developed by many authors which deals with the Borel combinatorics of graph structures on equivalence relations. The first systematic study of this subject was done by Kechris, Solecki, and Todorcevic [22], in which they studied problems such as Borel and measurable chromatic numbers, edge chromatic numbers, perfect matchings, etc. For example, it is implied by results of [22] that the Borel chromatic number of $F(2^{\mathbb{Z}^n})$ is no more than $2n + 1$ and the Borel edge chromatic number of $F(2^{\mathbb{Z}^n})$ is no more than $4n - 1$. As mentioned above, we will show in this paper that the continuous chromatic number of $F(2^{\mathbb{Z}^n})$ is 4. We will also show that the continuous edge chromatic number of $F(2^{\mathbb{Z}^2})$ is 5 and that there is no continuous perfect matching of $F(2^{\mathbb{Z}^n})$ for any $n \geq 1$. With these results, the present paper can be viewed as a contribution to the theory of descriptive graph combinatorics.

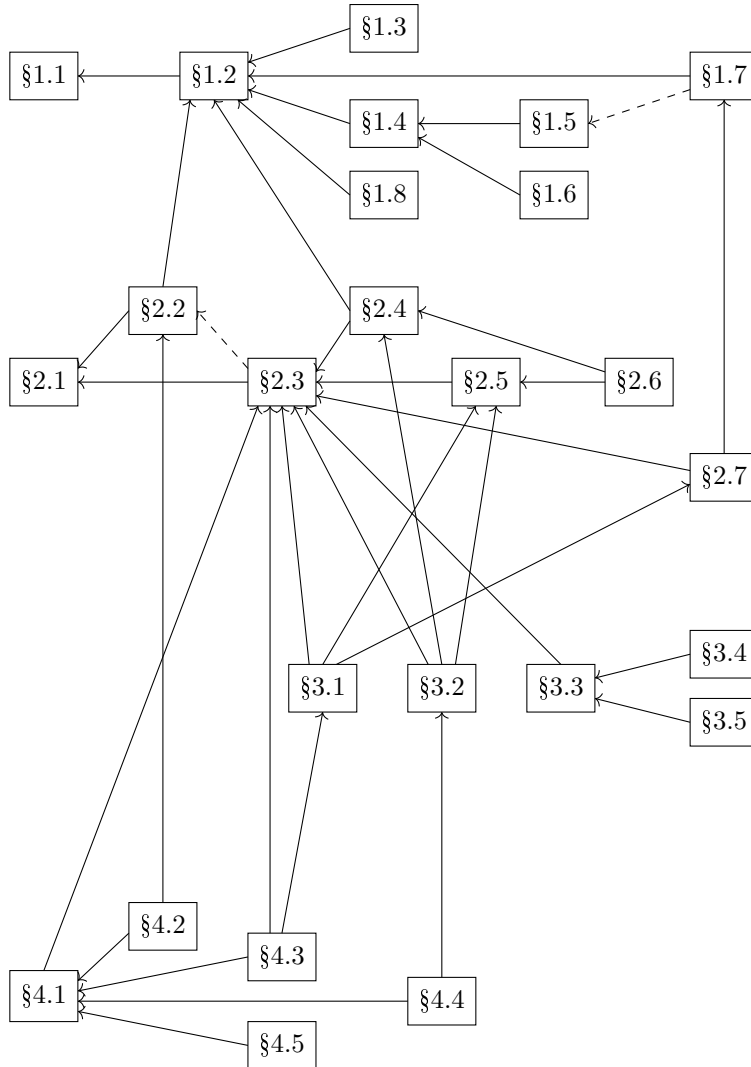
The rest of the paper is organized as follows. We give some background and motivations for the problems studied in this paper in the Introduction. Then in Chapter 1 we present the basic concepts, methods, and results in the continuous combinatorics of \mathbb{Z}^n actions. Specifically, in §1.1 and §1.2 we recall basic notation and terminology including the notion of hyper-aperiodic element on a general countable group G , and for the sake of completeness give simple proofs that these exist for $G = \mathbb{Z}^n$, the main case of interest for this paper. In §1.3 we give a quick application of the hyper-aperiodic elements in a problem about a marker structure called *toast*. In §1.4 and §1.5 we present the first main application of our method to the continuous chromatic number problem for $F(2^{\mathbb{Z}^n})$. Here we give a self-contained proof that the continuous chromatic number of $F(2^{\mathbb{Z}^n})$ is at least 4. Later in the paper this result will be proved twice more, and each proof is shorter

since it is based on more powerful general theorems. In §1.6 and §1.7 we give some more applications of the hyper-aperiodic elements in problems involving *clopen line sections* and *continuous cocycles*. In §1.8 we prove some results about *marker distortion* for marker sets $M \subseteq F(2^{\mathbb{Z}})$. Aside from being of independent interest, these results serve as a warm-up and motivation for the results in the rest of the paper. The notion of marker distortion measures the deviation from the marker set being regularly spaced. We show that the best possible marker distortion differs for clopen versus Borel marker sets (the Borel proof will be given in [13]). The clopen marker distortion result uses a specialized hyper-aperiodic element which incorporates the notion of a “tile.” This notion plays an important role in Chapter 2, which analyzes the existence of continuous equivariant maps from $F(2^{\mathbb{Z}^n})$ into subshifts of finite type in terms of the existence of such maps on a certain set of finite tiles. The main result of the paper, known as the Twelve Tiles Theorem, is presented in this chapter. In §2.1 we fix some notation on subshifts of finite type. In §2.2 we present a much simpler case of the tile analysis for dimension $n = 1$. In §2.3 we state the Twelve Tiles Theorem, which is the tile analysis for dimension $n = 2$. The theorem is proved in §2.4 and §2.5. In §2.6 we give a restatement of the Twelve Tiles Theorem in terms of *continuous structurings*. To some, this version of the theorem might be easier to apply as it does not involve the notion of subshifts of finite type. Our tile analysis techniques can be extended to higher dimensions. However, a general theorem for arbitrary dimensions is hard to state precisely, and we do not have any specific application for such a result at this time. So in §2.7 we state and prove a weak version of the tile theorem for higher dimensions, which we will use in Chapter 3.

Chapters 3 and 4 are devoted to various applications of the Twelve Tiles Theorem. In §3.1 we deduce several basic corollaries of the Twelve Tiles Theorem, such as the non-existence of a continuous perfect matching on $F(2^{\mathbb{Z}^n})$ and the computation of the edge chromatic number of $F(2^{\mathbb{Z}^2})$. In §3.2 we prove the fact that there is no continuous proper free subaction of \mathbb{Z}^2 on $F(2^{\mathbb{Z}^2})$. This can be used to answer in the negative a question raised in [11]. In §3.3, §3.4 and §3.5 we delve more specifically in the question of characterizing graphs Γ for which there exists a continuous homomorphism from $F(2^{\mathbb{Z}^2})$ to Γ . The result that $\chi_c(F(2^{\mathbb{Z}^2})) > 3$ is the special case of this analysis for which $\Gamma = K_3$, the complete graph on 3 vertices. The Twelve Tiles Theorem gives, in principle, an answer to the continuous graph homomorphism problem, but not necessarily in a way which is easy to apply. So, we prove “positive” and “negative” conditions for the existence or non-existence of such continuous homomorphisms in §3.5 and §3.4, respectively. We illustrate these conditions with a number of specific examples. In Chapter 4 we consider the decidability of three basic problems concerning the continuous combinatorics of $F(2^{\mathbb{Z}^n})$: the subshift, graph homomorphism, and tiling problems. It follows from the Twelve Tiles Theorem that all of these problems are semi-decidable, that is, they correspond to Σ_1^0 sets (coding the various objects in question, e.g., finite graphs, by integers). The graph homomorphism and tiling problems are special instances of the more general subshift problem. These problems are stated in §4.1. We show in §4.2 that the general subshift problem in one dimension, that is $n = 1$, is computable (i.e., corresponds to a Δ_1^0 set of integers). Thus, there is an algorithm which when given an integer coding a subshift of finite type $Y \subseteq 2^{\mathbb{Z}}$ will decide if there is a continuous, equivariant map from $F(2^{\mathbb{Z}})$ to Y . It follows immediately

that the graph homomorphism and tiling problems are also decidable for $n = 1$. In contrast, we show in §4.3 that the subshift problem for $n \geq 2$ is Σ_1^0 -complete, and in particular not decidable. In §4.4 we show that for $n \geq 2$ even the continuous graph homomorphism problem is Σ_1^0 -complete. Finally, in §4.5 we discuss the continuous tiling problem for $n = 2$. Here we do not yet know whether the problem is decidable. In fact, we do not know the answer to the tiling question even for some very specific simple sets of tiles. We mention some computer experiments which may shed some light on this question.

For the convenience of the reader we include below a diagram of logical dependence and references between sections. The arrows illustrate logical dependence between sections, and the dash arrows illustrate references and analogs.



Basic Definitions and Results

1.1. Bernoulli shift actions

If G is a group with identity 1_G and X is a set, a (left) *action* of G on X is a map $(g, x) \mapsto g \cdot x$ from $G \times X$ to X satisfying:

- (1) $g \cdot (h \cdot x) = (gh) \cdot x$ for all $g, h \in G, x \in X$.
- (2) $1_G \cdot x = x$ for all $x \in X$.

We also write $G \curvearrowright X$ to denote that G acts on X . When there is more than one action of G on X being considered, we employ notation such as $g \cdot x, g \bullet x, g * x$ to distinguish them.

When G acts on X and $x \in X$, we write $[x]$ for the orbit equivalence class of x , that is, $[x] = \{y \in X : \exists g \in G (g \cdot x = y)\}$. Of course, there is a danger of confusion with this notation as it does not specify the action. When there are different actions considered concurrently, we will introduce appropriate ad hoc notations to distinguish them.

If G acts on both X and Y , we say a map $\pi: X \rightarrow Y$ is *equivariant* if $\pi(g \cdot x) = g \cdot \pi(x)$ for all $g \in G$ and $x \in X$. We say π is an *equivariant embedding* if π is also injective.

Throughout the rest of the paper, X will henceforth denote a *Polish space*, i.e., a separable, completely metrizable topological space. When only the Borel structure of X is relevant, we refer to X as a *standard Borel space*. When G is also a topological group, we say the action is *continuous* or *Borel* if the map $(g, x) \mapsto g \cdot x$ from $G \times X$ to X is continuous or Borel. When G is a countable discrete group (which is the case for this paper), the action is continuous or Borel if and only if for each $g \in G$ the map $x \mapsto g \cdot x$ from X to X is continuous or Borel.

For G a countable group, the set 2^G is a compact zero-dimensional Polish space with the usual product of the discrete topology on $2 = \{0, 1\}$. There are two natural (left) actions of G on 2^G , the *left-shift*, and the *right-shift* (also called the *Bernoulli shift* actions). The left-shift action is given by

$$g \cdot x(h) = x(g^{-1}h),$$

for $g, h \in G$ and $x \in 2^G$. The right-shift action is given by

$$g \cdot x(h) = x(hg).$$

It is more common in the literature to use the left-shift action, and either one could be used throughout this paper. In fact, the left-shift and right-shift actions on 2^G are isomorphic G -actions under the map $\pi: 2^G \rightarrow 2^G$ given by $\pi(x)(h) = x(h^{-1})$. However, we prefer to use the right-shift action throughout.

An action of G on X is *free* if $g \cdot x \neq x$ for all $g \neq 1_G$ in G and all $x \in X$. If G acts on X , then the *free part* $F(X)$ is $\{x \in X : \forall g \neq 1_G (g \cdot x \neq x)\}$. $F(X)$ is the

largest invariant subset of X on which G acts freely. When G is a countable group acting continuously, $F(X)$ is also a G_δ subset of X , and thus a Polish space with the subspace topology. Except in trivial cases, however, $F(X)$ will not be compact.

The shift action of a countable group G is of particular significance as it is essentially universal for all actions of G according to the following theorem of Becker–Kechris [2] and Dougherty–Jackson–Kechris [8].

THEOREM 1.1.1. *The shift action of G on $(2^\omega)^G \cong 2^{G \times \omega}$ given by $(g \cdot x)(h, n) = x(hg, n)$ is a universal Borel G -action. That is, if $G \curvearrowright X$ is a Borel action, then there is a Borel equivariant embedding $\pi: X \rightarrow (2^\omega)^G$. Furthermore, if the action of G on X is continuous and X is zero-dimensional, then π can be chosen to be continuous.*

If S is a generating set of G , the *Cayley graph* $\Gamma = \Gamma(G, S)$ is a labeled directed graph with the vertex set $V(\Gamma) = G$ and the edge set $E(\Gamma)$ consisting of pairs of the form (g, sg) for $g \in G$ and $s \in S$, where the edge (g, sg) is assigned the label s . In addition, if G acts on a space X , the *Schreier graph* $\Gamma = \Gamma(X, G, S)$ is a labeled directed graph with the vertex set $V(\Gamma) = X$ and the edge set $E(\Gamma)$ consisting of pairs of the form $(x, s \cdot x)$ for $x \in X$ and $s \in S$, where the edge $(x, s \cdot x)$ is assigned the label s . In this sense the Cayley graph on G is a special case of a Schreier graph, where the action of G on itself is by left multiplication. When connectedness is considered, we consider the unlabeled and undirected Schreier graphs. Thus, since S is always a generating set of G , the Cayley graph of G is always connected, and each connected component of the Schreier graph is exactly an orbit equivalence class of the action. In the case that the action of G on X is free, there is a natural graph isomorphism between the Cayley graph on G and each connected component of the Schreier graph. Specifically, if the action is free and $x \in X$, then the map $g \mapsto g \cdot x$ induces a graph isomorphism from G to $[x]$.

In this paper we only deal with finitely generated groups, and the generating set in our study will always be finite and minimal (in the sense that no proper subset is a generating set). In particular, a generating set does not contain the identity. In addition, for all the Schreier graphs Γ we consider in this paper, the vertex set $V(\Gamma) = X$ will be a topological space, and the edge set $E(\Gamma) = \{(x, s \cdot x) \in X \times X : s \in S\}$ will be a closed subset of $X \times X$, and thus a topological space in its own right. In this sense, we refer to these Schreier graphs as *topological graphs*.

For a graph Γ (which could be labeled, or directed, or both) and a set K of colors, a *proper (K -)coloring* is a map $\kappa: V(\Gamma) \rightarrow K$ such that if $(x, y) \in E(\Gamma)$, then $\kappa(x) \neq \kappa(y)$. The *chromatic number* of Γ , denoted by $\chi(\Gamma)$, is the least cardinality of a set K such that there exists a proper K -coloring for Γ . When the graph Γ is a topological graph, we may consider *continuous proper colorings* and *Borel proper colorings*, and correspondingly its *continuous chromatic number*, denoted as $\chi_c(\Gamma)$, and, respectively, *Borel chromatic number*, denoted as $\chi_B(\Gamma)$.

The main case of interest for this paper is $G = \mathbb{Z}^n$ and $X = F(2^{\mathbb{Z}^n})$, the free part of the shift action of \mathbb{Z}^n on $2^{\mathbb{Z}^n}$. We work with the standard generating set $\{e_1, \dots, e_n\}$, where $e_i = (0, \dots, 0, 1, 0, \dots, 0) = (a_1, \dots, a_n)$ with $a_i = 1$ and $a_j = 0$ for all $j \neq i$. The Cayley graph on \mathbb{Z}^n is therefore an n -dimensional grid, which is isomorphic to the Schreier graph on each orbit equivalence class.

Recall from the Introduction that for $x, y \in F(2^{\mathbb{Z}^n})$, the “distance” $\rho(x, y)$ between them is the ℓ_∞ norm of the unique $g \in \mathbb{Z}^n$ such that $g \cdot x = y$. Formally, $\rho(x, y) = \infty$ if $[x] \neq [y]$, and if $y = g \cdot x$ with $g = (a_1, \dots, a_n) \in \mathbb{Z}^n$, then $\rho(x, y) =$

$\|g\| = \max\{|a_1|, \dots, |a_n|\}$. Occasionally we will use an alternative distance function ρ_1 on $X \times X$ where instead of the ℓ_∞ norm $\|g\|$ we use the ℓ_1 norm $\|g\|_1 = \sum_{i=1}^n |a_i|$.

1.2. Hyper-aperiodic elements

We next recall the definition of a hyper-aperiodic element from [15].

DEFINITION 1.2.1. Let G be a countable group. A point $x \in 2^G$ is *hyper-aperiodic* if the closure of its orbit is contained in the free part, i.e. if $\overline{\{x\}} \subseteq F(2^G)$.

The significance of this notion is that the closure of the orbit of a hyper-aperiodic element provides an invariant compact subset of $F(2^G)$ which allows us to use topological arguments more fully when analyzing continuous marker structures on $F(2^G)$. Hyper-aperiodicity can be tested via the following combinatorial characterization (see [15]).

LEMMA 1.2.2. *Let G act on 2^G by right-shifts. Then $x \in 2^G$ is hyper-aperiodic if and only if for all $s \neq 1_G$ in G there is a finite set $T \subseteq G$ such that*

$$\forall g \in G \exists t \in T x(tg) \neq x(tsg).$$

The finite set T in the above lemma is said to *witness* the hyper-aperiodicity of x for s .

REMARK 1.2.3. When using the left-shift action, the combinatorial condition of the previous lemma becomes $\forall g \in G \exists t \in T x(gt) \neq x(gst)$.

A related notion is that of $x, y \in 2^G$ being orthogonal.

DEFINITION 1.2.4. Let $x, y \in 2^G$. We say that x and y are *orthogonal*, denoted $x \perp y$, if there exists a finite $T \subseteq G$ such that for any $g, h \in G$ there is $t \in T$ with $x(tg) \neq y(th)$.

The topological significance of orthogonality is the following lemma from [15].

LEMMA 1.2.5. *Let $x, y \in 2^G$. Then $x \perp y$ if and only if $\overline{\{x\}} \cap \overline{\{y\}} = \emptyset$.*

In [15] the existence of hyper-aperiodic elements, in fact that of large families of pairwise orthogonal hyper-aperiodic elements, is shown for any countable group G .

THEOREM 1.2.6 ([15]). *Let G be a countable group. Then there is a hyper-aperiodic element $x \in 2^G$. In fact there is a perfect set $P \subseteq 2^G$ of pairwise orthogonal hyper-aperiodic elements.*

The proof of Theorem 1.2.6 is nontrivial, as it gives hyper-aperiodic elements $x \in 2^G$ for an arbitrary countable group G . For specific groups like $G = \mathbb{Z}^n$ these elements are actually easy to construct directly. Since we will be constructing specialized hyper-aperiodic elements for $G = \mathbb{Z}^n$ throughout the paper, for the sake of completeness we give a new simple argument that they exist in this case. This argument will be similar to others used later in the paper.

LEMMA 1.2.7. *Let G be a countable group and $x \in 2^G$ be hyper-aperiodic. Then for any $s \neq 1_G$ in G there is a finite $T \subseteq G$ such that*

$$\forall g \in G [(\exists t_1 \in T x(t_1g) \neq x(t_1sg)) \wedge (\exists t_2 \in T x(t_2g) = x(t_2sg))].$$

PROOF. Let $x \in 2^G$ be a hyper-aperiodic element. Let $s \neq 1_G$. Let $T \subseteq G$ witness the hyper-aperiodicity of x for s as in Lemma 1.2.2, and we may assume $1_G \in T$. Let T' witness the hyper-aperiodicity for s^2 . Let $T'' = T \cup T' \cup T's$. Then we claim that T'' witnesses the desired property for s . For suppose $g \in G$. If $x(g) = x(sg)$, then since there is a $t \in T$ with $x(tg) \neq x(tsg)$, we can take $t_1 = t$, $t_2 = 1_G$. Suppose $x(g) \neq x(sg)$. Let $t \in T'$ be such that $x(tg) \neq x(ts^2g)$. Consider $x(tg)$, $x(tsg)$, and $x(ts^2g)$, which are elements of $\{0, 1\}$. Either $x(tg) = x(tsg)$ or $x(tsg) = x(ts^2g)$. So we can take $t_1 = 1_G$ and either $t_2 = t$ or $t_2 = ts$. \square

If G, H are countable groups and $x \in 2^G, y \in 2^H$, let $x \oplus y \in 2^{G \times H}$ be given by $(x \oplus y)(g, h) = x(g) + y(h) \bmod 2$.

LEMMA 1.2.8. *Let $x \in 2^G$ and $y \in 2^H$ be hyper-aperiodic elements. Then $x \oplus y \in 2^{G \times H}$ is a hyper-aperiodic element.*

PROOF. Let $s = (s_1, s_2) \neq (1_G, 1_H)$. Suppose first that $s_1 \neq 1_G$ and $s_2 = 1_H$. Let T witness the hyper-aperiodicity of x for s_1 . Let $(g, h) \in G \times H$. Then there is a $t \in T$ with $x(tg) \neq x(ts_1g)$, and thus

$$\begin{aligned} (x \oplus y)[(t, 1_H)(s_1, s_2)(g, h)] &= x(ts_1g) + y(h) \bmod 2 \\ &\neq x(tg) + y(h) \bmod 2 = (x \oplus y)[(t, 1_H)(g, h)]. \end{aligned}$$

So, $T \times \{1_H\}$ witnesses the hyper-aperiodicity of $x \oplus y$ for s in this case. The case $s_1 = 1_G, s_2 \neq 1_H$ is similar.

Suppose that $s_1 \neq 1_G$ and $s_2 \neq 1_H$. Let T_1, T_2 witness the enhanced hyper-aperiodicity in Lemma 1.2.7 respectively for x for s_1 and for y for s_2 . Let $t_1 \in T_1$ be such that $x(t_1g) \neq x(t_1s_1g)$, and let $t_2 \in T_2$ be such that $y(t_2h) \neq y(t_2s_2h)$. Then

$$\begin{aligned} (x \oplus y)[(t_1, t_2)(g, h)] &= x(t_1g) + y(t_2h) \bmod 2 \\ &\neq x(t_1s_1g) + y(t_2s_2h) \bmod 2 = (x \oplus y)[(t_1, t_2)(s_1, s_2)(g, h)]. \end{aligned}$$

Thus $T_1 \times T_2$ witnesses the hyper-aperiodicity of $x \oplus y$ for s . \square

LEMMA 1.2.9. *For each $n \geq 1$, there is a hyper-aperiodic element on \mathbb{Z}^n .*

PROOF. From Lemma 1.2.8 it is enough to show that there is a hyper-aperiodic element on \mathbb{Z} . This can be done in many ways. For example, let $x(a) = \beta(a) \bmod 2$, where $\beta(a)$ is the number of 1's in the base 2 expansion of $|a|$. Then $x \in 2^{\mathbb{Z}}$ is hyper-aperiodic. For suppose $s \neq 0$ and $|s|$ has highest non-zero base 2 digit in the ℓ th position counting from the right (i.e., $\ell = \lfloor \log_2(|s|) \rfloor$), then we may define $T = \{n \in \mathbb{Z} : |n| \leq 4 \cdot 2^\ell\}$. Now for any $a \in \mathbb{Z}$ we can find $n \in T$ such that the binary expansion of $a + n$ has 0s in positions lower than the ℓ th digit, and that the ℓ th and $\ell + 1$ st digits of $a + n$ can be taken to be either $(0, 0)$ or $(1, 0)$. In the first case, adding s changes the parity of the number of 1s by the parity of $\beta(s)$, and in the second case by $1 + \beta(s)$. Thus there is an $n \in T$ with $x(a+n) \neq x(a+n+s)$. \square

We also note that once we have a hyper-aperiodic element on \mathbb{Z}^n , then we may easily construct orthogonal such elements.

LEMMA 1.2.10. *There are $y_0, y_1 \in 2^{\mathbb{Z}^n}$ which are hyper-aperiodic and with $y_0 \perp y_1$.*

PROOF. Let $y \in 2^{\mathbb{Z}^n}$ be hyper-aperiodic. Let p_1, p_2, q_1, q_2 be distinct primes and let $m \geq 2 \max\{p_1, p_2, q_1, q_2\}$. Let $M = 2m + 1$, and let y_0 be obtained from y by replacing each point $(a_1, \dots, a_n) \in \mathbb{Z}^n$ with a copy of $[0, 2M]^n$, and if

$y(a_1, \dots, a_n) = 0$ we make y_0 have period $(p_1, 0, \dots, 0)$ on the corresponding copy of $[0, 2M]^n$, and if $y_0(a_1, \dots, a_n) = 1$ we make y_0 have period $(p_2, 0, \dots, 0)$ on this copy. For example, we could use

$$y_0(a_1M + r_1, \dots, a_nM + r_n) = 1 \text{ iff } p_{y(\bar{a})} | r_1,$$

where $0 \leq r_j < M$. Likewise, y_1 is defined from y using q_1, q_2 . It is straightforward to check that y_0, y_1 are hyper-aperiodic and that $y_0 \perp y_1$. \square

1.3. Toast, an application of hyper-aperiodicity

We give here a quick application of hyper-aperiodic elements, which only requires Lemma 1.2.9. We show that there does not exist a continuous *toast structure* on $F(2^{\mathbb{Z}^n})$. The concept of toast was introduced by B. Miller, and has proved to be a useful concept (e.g., for the construction of Borel proper 3-colorings of $F(2^{\mathbb{Z}^n})$, a result which the authors will present in [13]).

This easy argument shows how hyper-aperiodic elements are used to show certain continuous structurings of shift actions cannot exist. By a *structuring* we mean a way to assign a certain kind of structure on $F(2^{\mathbb{Z}^n})$. A precise general definition will be given later in Definition 2.6.4 (the formal definition is not necessary for this section, though toast is an example of a structuring with an infinitary language). For example, a *linear order structuring* would be a way to assign a linear order to each class of $F(2^{\mathbb{Z}^n})$.

Here we define directly what a toast structuring is, and what it means for it to be continuous. In §2.6 we give the general definition of a structuring and what it means for the structuring to be continuous. In the terminology below, we will sometimes for simplicity drop the word “structuring,” for example, we will refer to an unlayered toast structuring as just an unlayered toast.

REMARK. The proof Theorem 1.3.2 below uses only the existence of a hyper-aperiodic element in $F(2^{\mathbb{Z}^d})$. For most of the results of the current paper, however, the mere existence of hyper-aperiodic elements is not enough. For these results, including our main “12 tiles” theorem (Theorems 2.3.5 and 2.6.8) certain specialized hyper-aperiodic elements must be constructed. Indeed, the construction of specialized hyper-aperiodic elements, and their use in studying continuous structurings on shift actions, is one of the main themes of this paper. Nevertheless, the toast result will serve as motivation for the later results.

Toast comes in two flavors, layered and unlayered (general) toast. The following definition of toast and accompanying figure is taken from [12].

Below, for a set $C \subseteq F(2^{\mathbb{Z}^d})$ we write ∂C for the set of boundary points of C , meaning those points in C that are adjacent to a point not in C :

$$\partial C = \{x \in C : \exists y \in [x] \setminus C \text{ } \rho_1(x, y) = 1\}.$$

DEFINITION 1.3.1. Let $\{T_n\}$ be a sequence of subequivalence relations of the shift equivalence relation on $F(2^{\mathbb{Z}^d})$ on some subsets $\text{dom}(T_n) \subseteq F(2^{\mathbb{Z}^d})$ with each T_n -equivalence class finite. Assume $\bigcup_n \text{dom}(T_n) = F(2^{\mathbb{Z}^d})$. We say $\{T_n\}$ is an (*unlayered*) *toast* if:

- (1) For each T_n -equivalence class C , and each T_m -equivalence class C' where $m > n$, if $C \cap C' \neq \emptyset$ then $C \subseteq C'$.

- (2) For each T_n -equivalence class C there is $m > n$ and a T_m -equivalence class C' such that $C \subseteq C' \setminus \partial C'$.

We say $\{T_n\}$ is a *layered toast* if, instead of (2) above, we have

- (2') For each T_n -equivalence class C there is a T_{n+1} -equivalence class C' such that $C \subseteq C' \setminus \partial C'$.

We call the toast *continuous* if for every n and for every $g \in \mathbb{Z}^d$ the set $\{x \in F(2^{\mathbb{Z}^d}) : (x, g \cdot x) \in T_n\}$ is relatively clopen in $F(2^{\mathbb{Z}^d})$. Similarly, we call the toast *open* if for every $g \in \mathbb{Z}^d$ this set is relatively open in $F(2^{\mathbb{Z}^d})$.

Figure 1 illustrates the definitions of layered and unlayered toast. The term ‘toast’ is simply motivated by the visual shapes of these regions in Figure 1 together with how the regions are stacked (specifically condition (1) above).




FIGURE 1. (a) layered toast

(b) general toast

In [12] a forcing argument was used to show that a Borel layered toast does not exist for $F(2^{\mathbb{Z}^d})$. The following result shows that a continuous (unlayered) toast does not exist for $F(2^{\mathbb{Z}^d})$. Thus, the only remaining possibility for $F(2^{\mathbb{Z}^d})$ is a Borel unlayered toast, which in fact does exist for $F(2^{\mathbb{Z}^d})$ (a result of the authors to be presented in [13]).

THEOREM 1.3.2. *There does not exist an open toast structure on $F(2^{\mathbb{Z}^d})$ for any $d \geq 1$.*

PROOF. Suppose $\{T_n\}$ were an open toast structure on $F(2^{\mathbb{Z}^d})$. Let $x \in F(2^{\mathbb{Z}^d})$ be a hyper-aperiodic element. Let $K = \overline{[x]}$, so $K \subseteq F(2^{\mathbb{Z}^d})$ is compact. For each $m \in \mathbb{N}$, consider the open set U_m consisting of those points $y \in K$ having the property that there is $n \leq m$ and a T_n -class C with $y \in C \setminus \partial C$. Then the U_m 's are increasing, open, and cover K . So there is a least m with $U_m = K$. Since m is least, $T_m \neq \emptyset$. Consider a T_m -class C' . By definition of toast, C' is finite. So we can find a point $y \in \partial C'$. Since $y \in U_m$, there is $n \leq m$ and a T_n -class C with $y \in C \setminus \partial C$. Moreover, since $C \neq C'$ and T_m is an equivalence relation, we must have $n < m$. However, this implies that $C' \cap C \neq \emptyset$ but $C \not\subseteq C'$, contradicting Definition 1.3.1.(1). \square

1.4. A hyper-aperiodic element with partial periodicity

We will construct a hyper-aperiodic element on \mathbb{Z}^2 that possesses some partial vertical periodicity. We will use heavily the definitions of hyper-aperiodic and orthogonal given in Definitions 1.2.1 and 1.2.4.

The following lemma is a warm-up.

LEMMA 1.4.1. *Let x, y_0, y_1 be hyper-aperiodic elements of $2^{\mathbb{Z}}$ with $y_0 \perp y_1$. Then the element $z \in 2^{\mathbb{Z}^2}$ defined by*

$$z(u, v) = y_{x(u)}(v)$$

is hyper-aperiodic.

PROOF. Let $s = (a, b)$ be a non-identity element in \mathbb{Z}^2 . First assume $a = 0$. Then $b \neq 0$. Let T_0 and T_1 be finite sets witnessing the hyper-aperiodicity of y_0 and y_1 , respectively, for b . Let $T = \{0\} \times (T_0 \cup T_1)$. We claim that T witnesses the hyper-aperiodicity of z for s . To see this, let $g = (u, v) \in \mathbb{Z}^2$ be arbitrary. If $x(u) = 0$, then there is $t \in T_0$ such that

$$z(u, v+t) = y_{x(u)}(v+t) = y_0(v+t) \neq y_0(v+b+t) = y_{x(u)}(v+b+t) = z(u, v+b+t),$$

or if we let $\tau = (0, t) \in T$,

$$z(g + \tau) \neq z(g + s + \tau).$$

If $x(u) = 1$, then similarly there is $t \in T_1$ such that $z(g + \tau) \neq z(g + s + \tau)$, where $\tau = (0, t) \in T$. This proves the lemma in the case $a = 0$.

Now assume $a \neq 0$. Let S be the finite set witnessing the hyper-aperiodicity of x for a . Let R be a symmetric (i.e. $r \in R$ implies $-r \in R$) finite set witnessing the orthogonality of y_0 with y_1 . Note that R also witnesses the orthogonality of y_1 with y_0 . Let $T = S \times R$. We claim that T witnesses the hyper-aperiodicity of z for s . To see this, let $g = (u, v) \in \mathbb{Z}^2$ be arbitrary. Let $t \in S$ be such that $x(u+t) \neq x(u+a+t)$. Let $r \in R$ be such that $y_0(v+r) \neq y_1(v+b+r)$ (if $x(u+t) = 0$) or $y_1(v+r) \neq y_0(v+b+r)$ (if $x(u+t) = 1$). Letting $\tau = (t, r)$, we have

$$z(u+t, v+r) = y_{x(u+t)}(v+r) \neq y_{x(u+a+t)}(v+b+r) = z(u+a+t, v+b+r)$$

or $z(g + \tau) \neq z(g + s + \tau)$. This proves the lemma in the case $a \neq 0$. \square

Next we modify the construction of the hyper-aperiodic element z by allowing it to possess some partial vertical periodicity. The partial vertical periodicity is governed by a function $f: \mathbb{Z} \rightarrow \mathbb{Z}^+$ in the sense that for each $u \in \mathbb{Z}$, $f(u)$ is a vertical period for $z(u, \cdot)$.

Let Λ be an infinite set of prime numbers. Let $f: \mathbb{Z} \rightarrow \mathbb{Z}^+$ be a function such that

- (i) For all $u \in \mathbb{Z}$ there are $p \in \Lambda$ and $n \in \mathbb{Z}^+$ such that $f(u) = p^n$;
- (ii) For all $p \in \Lambda$ and $m \in \mathbb{N}$, there are $a \in \mathbb{Z}$ and $k \in \mathbb{Z}^+$ such that $f(i) = p^k$ for all $i \in [a, a+m]$;
- (iii) $f(u)$ is monotone increasing for $u > 0$, monotone decreasing for $u < 0$, and $f(u) \rightarrow \infty$ as $|u| \rightarrow \infty$.

Such functions f can be constructed easily by diagonalization.

LEMMA 1.4.2. *Let x, y_0, y_1 be hyper-aperiodic elements of $2^{\mathbb{Z}}$ with $y_0 \perp y_1$. Let Λ be an infinite set of prime numbers and $f: \mathbb{Z} \rightarrow \mathbb{Z}^+$ be a function as above. Then the element $z \in 2^{\mathbb{Z}^2}$ defined by*

$$z(u, v) = y_{x(u)}(v \bmod f(u))$$

is hyper-aperiodic.

PROOF. Let $s = (a, b)$ be a non-identity element in \mathbb{Z}^2 . First assume $a = 0$. Then $b \neq 0$. Let T_0 and T_1 be finite sets witnessing the hyper-aperiodicity of y_0 and y_1 , respectively, for b . Let $B \in \mathbb{Z}^+$ be greater than $|b|$ or $|t|$ for all $t \in T_0 \cup T_1$. Let $A \in \mathbb{Z}^+$ be such that for all $|u| \geq A$, $f(u) > 3B$. Let $T = [-A, A] \times [-3B, 3B]$. We claim that T witnesses the hyper-aperiodicity of z for s . To see this, let $g = (u, v) \in \mathbb{Z}^2$ be arbitrary. First, there is $c \in [-A, A]$ such that $f(u+c) > 3B$. Without loss of generality assume $x(u+c) = 0$. Next, there is $d \in [-2B, 2B]$ such that

- the difference between $v+d$ and any multiple of $f(u+c)$ is greater than B , and
- the difference between $v+b+d$ and any multiple of $f(u+c)$ is greater than B .

Finally, there is $t \in T_0$ such that

$$z(u+c, v+d+t) = y_{x(u+c)}(v+d+t) = y_0(v+d+t)$$

$$\neq y_0(v+d+b+t) = y_{x(u+c)}(v+d+b+t) = z(u+c, v+d+b+t).$$

If we let $\tau = (c, d+t) \in T$, we have

$$z(g+\tau) \neq z(g+s+\tau).$$

This proves the lemma in the case $a = 0$. The case $a \neq 0$ is similarly handled by a combination of the above shifting technique and the proof of the corresponding case of Lemma 1.4.1. \square

Note that in the proof of Lemma 1.4.2 only condition (iii) of the function f is used to guarantee that z is a hyper-aperiodic element. The other two conditions will be used in the proof of the following theorems.

1.5. The continuous chromatic number of $F(2^{\mathbb{Z}^2})$

In this section we give a self-contained proof for the fact that there are no continuous proper 3-colorings on $F(2^{\mathbb{Z}^2})$. It has been proved in [11] that there is a continuous proper 4-coloring on $F(2^{\mathbb{Z}^2})$. Combining these results, we conclude that the continuous chromatic number for $F(2^{\mathbb{Z}^2})$ is exactly 4.

THEOREM 1.5.1. *There is no continuous proper 3-coloring of $F(2^{\mathbb{Z}^2})$. Consequently, the continuous chromatic number of $F(2^{\mathbb{Z}^2})$ is 4.*

As an immediate corollary, we conclude that for any $n \geq 2$, the continuous chromatic number for $F(2^{\mathbb{Z}^n})$ is exactly 4.

COROLLARY 1.5.2. *For any $n \geq 2$ the continuous chromatic number of $F(2^{\mathbb{Z}^n})$ is 4.*

PROOF. It has been proved in [11] that there is a continuous proper 4-coloring on $F(2^{\mathbb{Z}^n})$ for all $n \geq 2$. To see that the continuous chromatic number of $F(2^{\mathbb{Z}^n})$ is not 3, it suffices to notice that $F(2^{\mathbb{Z}^2})$ can be naturally identified with a \mathbb{Z}^2 -invariant subset of $F(2^{\mathbb{Z}^n})$, meaning there is a continuous map $\varphi: F(2^{\mathbb{Z}^2}) \rightarrow F(2^{\mathbb{Z}^n})$ such that whenever $(a_1, a_2) \cdot x = y$ we have $(a_1, a_2, 0, \dots, 0) \cdot \varphi(x) = \varphi(y)$. Granting such a φ , if $c: F(2^{\mathbb{Z}^n}) \rightarrow 3$ were a proper 3-coloring, then $\varphi \circ c: F(2^{\mathbb{Z}^2}) \rightarrow 3$ would also be a proper 3-coloring. The identification φ can be defined by sending $x \in F(2^{\mathbb{Z}^2})$ to the function $\varphi(x)$ satisfying

$$\varphi(x)(a_1, a_2, \dots, a_n) = \begin{cases} x(a_1, a_2), & \text{if } a_3 = \dots = a_n = 0, \\ 0, & \text{otherwise.} \end{cases}$$

Then if $x \in F(2^{\mathbb{Z}^2})$ one can easily check $\varphi(x) \in F(2^{\mathbb{Z}^n})$. Obviously φ also satisfies the equivariant requirement. \square

The rest of this subsection is devoted to a proof of Theorem 1.5.1. Toward a contradiction, assume $c: F(2^{\mathbb{Z}^2}) \rightarrow \{0, 1, 2\}$ is a continuous proper 3-coloring of $F(2^{\mathbb{Z}^2})$.

In the following we denote $e_1 = (1, 0)$ and $e_2 = (0, 1)$. We say $x, y \in F(2^{\mathbb{Z}^2})$ are *adjacent* if there is $i = 1, 2$ with $x = e_i \cdot y$ or $y = e_i \cdot x$ (i.e., x and y are adjacent in the Schreier graph $\Gamma(F(2^{\mathbb{Z}^2}))$). For adjacent x and y , let

$$\theta(x, y) = \begin{cases} +1, & \text{if } c(y) = c(x) + 1 \pmod{3}, \\ -1, & \text{if } c(y) = c(x) - 1 \pmod{3}. \end{cases}$$

We define a potential function $P(x, y)$ for $x, y \in F(2^{\mathbb{Z}^2})$ with $[x] = [y]$. Let $x_0 = x, x_1, \dots, x_{k+1} = y$ be a *walk* in the Schreier graph $\Gamma(F(2^{\mathbb{Z}^2}))$ from x to y , i.e., x_i and x_{i+1} are adjacent for all $i = 0, \dots, k$ (we do not require the x_i 's to be distinct). Then let (where the following sum is taken in \mathbb{Z} , not $\mathbb{Z}/3\mathbb{Z}$)

$$P(x, y) = \sum_{i=0}^k \theta(x_i, x_{i+1}).$$

We first verify that $P(x, y)$ is well-defined.

LEMMA 1.5.3. *$P(x, y)$ does not depend on the path chosen from x to y .*

PROOF. For a walk $\pi = (x_0, \dots, x_{k+1})$ in the Schreier graph $\Gamma(F(2^{\mathbb{Z}^2}))$, set $P(\pi) = \sum_{i=0}^k \theta(x_i, x_{i+1})$. We write $\pi^{-1} = (x_{k+1}, \dots, x_0)$ for the reversal of π . If π ends at y and σ begins at y , then we write $\pi\sigma$ for the concatenated path. Note that $P(\pi^{-1}) = -P(\pi)$ and that $P(\pi\sigma) = P(\pi) + P(\sigma)$. We say that π is *backtracking* if there is i with $x_i = x_{i+2}$. Note that two paths π and π' are homotopy equivalent if and only if they become equal after removing all instances of back-tracking (this can be taken as a definition if desired), and that $P(\cdot)$ is invariant under homotopy equivalence. Let G be the homotopy group of the Schreier graph at the point x .

Write K_3 for the triangle graph having vertex set $\{0, 1, 2\}$. Note that the coloring c provides a graph homomorphism from $F(2^{\mathbb{Z}^2})$ to K_3 . Write σ_z for the path $(z, (1, 0) \cdot z, (1, 1) \cdot z, (0, 1) \cdot z, z)$. Since K_3 has no 4-cycles, the image of σ_z under c must be trivial up to homotopy equivalence. It follows $P(\sigma_z) = 0$. Similarly, if $\ell_{w,z}$ is a walk from w to z then $P(\ell_{w,z} \sigma_z \ell_{w,z}^{-1}) = P(\ell_{w,z}) + 0 - P(\ell_{w,z}) = 0$.

Now suppose that π_1 and π_2 are paths from x to y . We claim that $\ell = \pi_1 \pi_2^{-1}$ is homotopy equivalent to a concatenation of paths of the form $\ell_{x,z} \sigma_z^{\pm 1} \ell_{x,z}^{-1}$. Let

N be the normal subgroup of G consisting of products of elements of the form $\ell_{x,z}\sigma_z^{\pm 1}\ell_{x,z}^{-1}$. We show that every path ℓ in G is homotopy equivalent to an element of N .

View x as the origin in \mathbb{Z}^2 , and identify ℓ as a sequence of directed edges in the Cayley graph of \mathbb{Z}^2 . Let $z = (a, b)$ be a vertex of the path ℓ which as a point in \mathbb{Z}^2 is of maximal distance from the origin. If $z = (0, 0)$, then ℓ is the trivial path and we are done. Otherwise, let u be the vertex of ℓ immediately preceding z on the path, and let v be the vertex immediately following z . One of the coordinates of u is strictly less in absolute value than the corresponding coordinate of z , and the other coordinate of u is equal to the corresponding coordinate of z , and likewise for v . If $u = v$, then we may delete the part of the path uzv from ℓ to a shorter path ℓ' which is homotopy equivalent to ℓ . By induction on the length of ℓ , the result holds for ℓ' and we are done. If $u \neq v$, then say (the other cases being similar) $u = (a, b - 1)$, $v = (a - 1, b)$. Then we may replace uzv with $uz'v$ where $z' = (a - 1, b - 1)$. The new path ℓ' is equal to ℓ modulo an element of N , has the same length as ℓ , and either has a strictly smaller value for the maximum distance of a vertex from $(0, 0)$ or else has the same maximum and a strictly smaller number of vertices attaining this maximum. By induction we may assume the result holds for ℓ' and we are done. \square

Let z be the hyper-aperiodic element on \mathbb{Z}^2 as constructed in Lemma 1.4.2. Then $\overline{[z]} \subseteq F(2\mathbb{Z}^2)$ and $c \upharpoonright \overline{[z]} : \overline{[z]} \rightarrow \{0, 1, 2\}$ is a continuous function on a compact space. Thus there is $N \in \mathbb{Z}^+$ such that for all $x \in \overline{[z]}$, $x \upharpoonright [-N, N]^2$ completely determines $c(x)$.

Fix two distinct prime numbers $p, q \in \Lambda$ where $p, q \neq 2, 3$. By condition (ii) of the function f , there are $a, b \in \mathbb{Z}$ and $n, k \in \mathbb{Z}^+$ such that $f(i) = p^n$ for all $i \in [a, a + 2N + 1]$ and $f(i) = q^k$ for all $i \in [b, b + 2N + 1]$. Let $x = (a + N + 1, 0) \cdot z$ and $y = (b + N + 1, 0) \cdot z$. By the construction of z , we have that for all $i \in [-N, N]$, $i \cdot x(0, \cdot)$ is vertically periodic with a period p^n and $i \cdot y(0, \cdot)$ is vertically periodic with a period q^k . By the choice of N , we conclude that $c((0, t) \cdot x)$ is periodic with respect to t with a period p^n , i.e.,

$$c((0, t) \cdot x) = c((0, t + p^n) \cdot x)$$

for all $t \in \mathbb{Z}$; similarly, $c((0, t) \cdot y)$ is periodic with respect to t with a period q^k .

Consider the horizontal path from x to y . Its length is $|a - b|$. Thus $|P(x, y)| \leq |a - b|$. Set $D = p^n q^k$ and let $m \in \mathbb{Z}^+$ be arbitrary. Let $x' = (0, mD) \cdot x$ and $y' = (0, mD) \cdot y$. Then the horizontal path from x' to y' is still of length $|a - b|$, and we also have $|P(x', y')| \leq |a - b|$.

Define

$$r = \frac{P(x, (0, p^n) \cdot x)}{p^n} \text{ and } s = \frac{P(y, (0, q^k) \cdot y)}{q^k}.$$

Intuitively, r is slope of the change of the potential from x to $(0, p^n) \cdot x$, and s is the slope of the change of the potential from y to $(0, q^k) \cdot y$. We claim that $|r| < 1$ and $r \neq 0$. Similarly for s . To see the claim, note first that if $r = 1$ then $P(x, (0, p^n) \cdot x) = p^n$. Consider the vertical path x_0, \dots, x_{p^n} from x to $(0, p^n) \cdot x$. Its length is p^n , which implies that for all $i = 0, \dots, p^n$, $\rho(x_i, x_{i+1}) = +1$. It follows that $c(x_i) = c(x_j)$ if and only if $i \equiv j \pmod{3}$. Since $p \neq 3$ is prime, this implies $c((0, p^n) \cdot x) = c(x_{p^n}) \neq c(x_0) = c(x)$, contradicting the fact that by

design $t \mapsto c((0, t) \cdot x)$ is p^n -periodic. A similar argument rules out the possibility that $r = -1$. To see that $r \neq 0$, just note that in order for $r = 0$ we must have $P(x, (0, p^n) \cdot x) = 0$, and this cannot happen unless p^n is even. Since p is an odd prime, we are done.

Finally, since $p \neq q$, we conclude in addition that $r \neq s$.
By periodicity, we have

$$P(x, x') = mDr \text{ and } P(y, y') = mDs.$$

Since

$$P(x, x') + P(x', y') = P(x, y) + P(y, y'),$$

we have

$$P(x', y') - P(x, y) = P(y, y') - P(x, x') = mD(s - r).$$

Since $P(x, y), P(x', y') \in [-|a - b|, |a - b|]$ and $r \neq s$, this is a contradiction when m is large enough.

We remark that in the above proof we did not use condition (i) of the function f . This condition will be used in the proofs of the theorems in the next two sections.

1.6. CBS line sections and single line sections

A *line section* is a symmetric Borel relation S , which we think of as a graph, which is a subset of the Schreier graph $\Gamma(F(2^{\mathbb{Z}^2}))$ which is acyclic and which satisfies $\deg_S(x) \in \{0, 2\}$ for all $x \in F(2^{\mathbb{Z}^2})$. We write $V(S) = \{x \in F(2^{\mathbb{Z}^2}) : \deg_S(x) = 2\}$. Intuitively, the intersection of a line section S with an orbit is a collection of lines in the graph of $F(2^{\mathbb{Z}^2})$. We call these lines *S-lines*. We say that S is *clopen* if for $i = 1, 2$ the set $\{x \in F(2^{\mathbb{Z}^2}) : (x, e_i \cdot x) \in S\}$ is clopen, where $e_1 = (1, 0)$ and $e_2 = (0, 1)$. We say that S is a *complete section* if $[x] \cap V(S) \neq \emptyset$ for all x , and we call S *co-complete* if $[x] \cap (F(2^{\mathbb{Z}^2}) \setminus V(S)) \neq \emptyset$ for all x . If x and y are on the same S -line, then the *S-path-distance* between x and y , denoted as $\rho_S(x, y)$, is the length of the shortest S -path between x and y . Thus $\rho_S(x, y) = 1$ if and only if x and y are adjacent elements of an S -line.

For example, if S is defined as

$$(x, y) \in S \iff (1, 0) \cdot x = y \text{ or } (-1, 0) \cdot x = y,$$

then S is a clopen complete line section, but not co-complete. In other words, the entire space can be viewed as a collection of horizontal S -lines. In general, we call a line section S a *horizontal line section* if every edge in S is a horizontal (e_1) edge. We remark that a simple category argument gives the following result.

THEOREM 1.6.1. *There does not exist any Borel horizontal line section of $F(2^{\mathbb{Z}^2})$ which is both complete and co-complete.*

PROOF. Assume S is a Borel complete and co-complete horizontal line section of $F(2^{\mathbb{Z}^2})$. Note that $F(2^{\mathbb{Z}^2}) \setminus S$ is also such a set. By considering vertical translations of S , we note that countably many homeomorphic images of S cover the entire $F(2^{\mathbb{Z}^2})$. It follows that S is non-meager, and likewise for $F(2^{\mathbb{Z}^2}) \setminus S$. On the other hand, by considering only the horizontal translations, which is a continuous action of \mathbb{Z} on $F(2^{\mathbb{Z}^2})$, we have that S is a Borel invariant subset of $F(2^{\mathbb{Z}^2})$. It follows from the First Topological 0-1 Law ([19] Theorem 8.46) that S must be either meager or comeager, a contradiction. \square

In other words, the only complete Borel horizontal line section is the above trivial one that we demonstrated. In the following we consider line sections with bounded slopes from a large-scale perspective, which we call *line sections with coarsely bounded slopes* or *CBS line sections*.

DEFINITION 1.6.2. A line section S of $F(2^{\mathbb{Z}^2})$ has *coarsely bounded slopes* or is a *CBS line section* if for any $x \in F(2^{\mathbb{Z}^2})$ there exist $D, M \in \mathbb{Z}^+$ such that for any $x_1, x_2 \in [x] \cap V(S)$ lying on the same S -line, if $\rho_S(x_1, x_2) > D$ and $(u, v) \cdot x_1 = x_2$, then $|v| \leq M|u|$.

Thus, on each equivalence class there is a bound M on the slopes of the lines in S , at least for pairs of points which are at least D apart with respect to ρ_S . We have the following negative result about CBS line sections, whose proof is very similar to that of Theorem 1.5.1.

THEOREM 1.6.3. *There does not exist any clopen CBS line section of $F(2^{\mathbb{Z}^2})$ which is both complete and co-complete.*

PROOF. Toward a contradiction, assume S is a clopen CBS line section of $F(2^{\mathbb{Z}^2})$ which is both complete and co-complete. Let z be the hyper-aperiodic element on \mathbb{Z}^2 constructed in Lemma 1.4.2. Let D, M be given by the definition of CBS line sections with $x = z$. By assumption the sets $\{x \in \overline{[z]} : (x, e_i \cdot x) \in S\}$, $i = 1, 2$, are clopen. It follows from compactness of $\overline{[z]}$ that there is $N \in \mathbb{Z}^+$ such that for every $x \in \overline{[z]}$, $x \upharpoonright [-N, N]^2$ determines whether $(x, e_1 \cdot x)$ and $(x, e_2 \cdot x)$ are in S .

Note that by the boundedness property of S , every S -line meets every vertical line in $[z]$, i.e., for every $a \in \mathbb{Z}$, the vertex set of each S -line has a non-empty intersection with the set $\{(a, v) \cdot z \mid v \in \mathbb{Z}\}$. Indeed, if $(a_n, b_n) \cdot z$ enumerates a bi-infinite path in S then $|a_{n+1} - a_n| \leq 1$ for every n and the quantities $|a_{n+D+1} - a_n|$ and $|a_{n+D+2} - a_n|$ are never 0, so the signs of $a_{n+D+1} - a_n$ and $a_{n+D+2} - a_n$ never change and thus $\{a_n : n \in \mathbb{Z}\} = \mathbb{Z}$.

Fix two distinct prime numbers $p, q \in \Lambda$. By condition (ii) of the function f from §1.4, there are $a, b \in \mathbb{Z}$ with $|b - a| \geq D$ and $n, k \in \mathbb{Z}^+$ such that $f(i) = p^n$ for all $i \in [a, a + 2N + 2D + 1]$ and $f(i) = q^k$ for all $i \in [b, b + 2N + 2D + 1]$. We assume $a < b$, as the other case is essentially identical. Let $x = (a + N + D + 1, 0) \cdot z$ and $y = (b + N + D + 1, 0) \cdot z$. By the construction of z , we have that for all $i \in [-N - D, N + D]$, $(i, 0) \cdot x(0, \cdot)$ is vertically periodic with a period p^n and $(i, 0) \cdot y(0, \cdot)$ is vertically periodic with a period q^k . By the choice of N , we conclude that $S \upharpoonright \{(i, t) \cdot x : |i| \leq D, t \in \mathbb{Z}\}$ is vertically periodic with period p^n ; similarly $S \upharpoonright \{(i, t) \cdot y : |i| \leq D, t \in \mathbb{Z}\}$ is vertically periodic with period q^k . Note that for all $v, v' \in \mathbb{Z}$, if $(0, v) \cdot x$ and $(0, v') \cdot x$ lie on the same S -line then their S -path-distance must be at most D , and hence the S -path connecting them is contained in the vertical strip $\{(i, t) \cdot x : |i| \leq D, t \in \mathbb{Z}\}$. Writing $[S]$ for the equivalence relation given by the connected components of S , it follows that the restriction of $[S]$ to the vertical line extending from x is p^n -periodic and the restriction of $[S]$ to the vertical line extending from y is q^k -periodic.

Define ℓ to be the number of equivalence classes of the restriction of $[S]$ to the vertical line extending from x over some/any interval of size p^n (that is the number of distinct S -classes in a set of the form $(0, s) \cdot x, \dots, (0, (s + p^n - 1)) \cdot x$ divided by p^n . Intuitively, ℓ is the density of S -lines along the vertical line above

x . Similarly define r to be the number of equivalence classes of the restriction of $[S]$ to the vertical line extending from y in some/any interval of size q^k divided by q^k . So, r is the density of the S -lines above y . We note that $\ell, r \neq 0$ by the above observation that every vertical line meets all the S -lines. Since S is co-complete, we also have that $\ell, r \neq 1$. Otherwise, if $\ell = 1$, every point in the vertical line extending from x would belong to a distinct S -line, hence $S \cap [z]$ would necessarily consist of all horizontal lines in $[z]$ (if all of the distinct points on the vertical line through x belong to distinct S -classes then, by disjointness of the distinct S -lines, for every point $(0, i) \cdot x$, the S -line through $(0, i) \cdot x$ must move horizontally, that is, the point $(1, i) \cdot x$ is on the same S -line). This contradicts that S is co-complete. Again, we conclude that $\ell \neq r$ since $p \neq q$.

Let $m \in \mathbb{Z}^+$ be an arbitrary multiple of all $f(i)$ for $i \in [a, b]$. Let $x' = (0, m) \cdot x$ and $y' = (0, m) \cdot y$. Let L be the segment from x to x' and R be the segment from y to y' . The restriction of S to the vertical strip $H = \{(u, v) \cdot x : 0 \leq u \leq b - a, v \in \mathbb{Z}\}$ is periodic with period dividing m . Let t be a positive integer, and consider the vertical strip $H_t = \{(u, v) \cdot x : 0 \leq u \leq b - a, 0 \leq v \leq tm\}$. Let L_t be the left edge of this strip and let R_t be the right edge. By periodicity, the number of distinct S -classes meeting L_t is ℓtm , and the number meeting the R_t is rtm . So, on the one hand we have that the difference between the number of distinct S -lines intersecting the left-edge L_t of H_t and the right-edge R_t of H_t would be $|\ell - r|tm$. On the other hand, this difference is bounded by a constant independent of t using the fact that the lines in S have a slope bounded by M . This is a contradiction for t large enough. \square

In view of Theorem 1.6.3 it is natural to ask to what extent we can remove the requirement that the line section of $F(2^{\mathbb{Z}^2})$ consists of CBS lines as defined in Definition 1.6.2. In the forthcoming [13] we will prove that there does exist a Borel line section of $F(2^{\mathbb{Z}^2})$ such that every class is covered by a single line. Such a line section is called a *lining* of the whole space. This result has also been obtained independently by N. Chandgotia and S. Unger [4] by a different argument.

In the following, we rule out a clopen line section S of $F(2^{\mathbb{Z}^2})$ with the property that each class $[x]$ contains exactly one S -line. We call such sections *single line sections*.

THEOREM 1.6.4. *There does not exist any clopen single line section of $F(2^{\mathbb{Z}^2})$. In particular, there does not exist any clopen lining of $F(2^{\mathbb{Z}^2})$.*

PROOF. Let S be a clopen single line section. Let $x \in F(2^{\mathbb{Z}^2})$ be a hyperaperiodic element (for this argument, no other special properties of x are required).

Let $K = \overline{[x]}$, so $K \subseteq F(2^{\mathbb{Z}^2})$ is compact. The compactness of K and the clopenness of S give that there is a $d_0 \in \mathbb{Z}^+$ such that for any $y \in K$, $y \upharpoonright [-d_0, d_0]^2$ determines the membership of y in $V(S)$ and in $\{w : (w, e_i \cdot w) \in S\}$, $i = 1, 2$. The compactness of K and the fact that S is a complete section on $F(2^{\mathbb{Z}^2})$ also give a d_1 such that every $y \in K$ is within ρ distance d_1 of a point in $V(S)$. Finally, the compactness of K and the fact that S is a single line on each class give that there is a d_2 such that for all $y \in K$ and all points z, w within ρ distance $5d_1$ of y , if $z, w \in V(S)$ then $\rho_S(z, w) \leq d_2$. This is because the quantity $\rho_S(z, w)$ is continuously determined from y (if $\rho_S(z, w) = d$ then the S -path of length d joining

them is contained in a ball of bounded radius around y , and the behavior of S in that ball is determined by the restriction of y to some larger, but still finite, ball).

Fix $x_0 \in [x] \cap V(S)$. Consider the horizontal line L through x_0 , that is, the points of the form $(a, 0) \cdot x_0$ for all $a \in \mathbb{Z}$. From the definition of d_1 , for every $n \in \mathbb{Z}$ we can fix a point $x_n \in V(S)$ with $\rho(x_n, (3nd_1, 0) \cdot x_0) \leq d_1$. Observe that $x_{n+1} = (a, b) \cdot x_n$ with $d_1 \leq a \leq 5d_1$ and $x_n = (a', b') \cdot x_0$ with $|b'| \leq d_1$. That is, the points x_n are horizontally spaced with sequential distance between d_1 and $5d_1$, and are within vertical distance d_1 of L . This is illustrated in Figure 2. From the definition of d_2 , the path along S between consecutive points x_n and x_{n+1} has length at most d_2 . In particular, all of these paths stay within the horizontal strip centered about L of width $2(d_2 + d_1)$ (see Figure 2).




FIGURE 2. Analyzing $S \cap [x]$

Let $S' \subseteq S \upharpoonright [x]$ be the union of these paths connecting x_n and x_{n+1} for all $n \in \mathbb{Z}$. Arbitrarily fix an orientation of $V(S) \cap [x]$ as a line. This orientation gives rise to a linear order $<_S$ on $V(S) \cap [x]$. If $z, w \in V(S) \cap [x]$, then z and w lie on the same S -line; define $t(z, w) = \rho_S(z, w)$ if $z <_S w$, and $t(z, w) = -\rho_S(z, w)$ if $w <_S z$. Let $T = \{t \in \mathbb{Z} : \exists y \in V(S') t(x_0, y) = t\}$. This is the set of “times” t for which, starting at the point x_0 , the element on the S -line that is t steps away from x_0 is in the set $V(S')$. Since S' contains the union of connected subsets of S connecting x_n to x_{n+1} , it follows that T is an infinite interval in \mathbb{Z} . We claim that $T = \mathbb{Z}$. Suppose, for example, $T = [a, +\infty)$. Then obviously $a < 0$ since $0 \in T$. Consider x_{a-1} and $x_{|a|+1}$. By our construction, we have both $\rho(x_{a-1}, x_0) \geq |a-1| = |a|+1$ and $\rho(x_{|a|+1}, x_0) \geq |a|+1$, and therefore $|t(x_{a-1}, x_0)|, |t(x_{|a|+1}, x_0)| \geq |a|+1$. By construction the S -line segment between x_{a-1} and $x_{|a|+1}$ passes through x_0 , thus $t(x_{a-1}, x_0)t(x_{|a|+1}, x_0) < 0$. This implies that one of $t(x_{a-1}, x_0)$ and $t(x_{|a|+1}, x_0)$ must be smaller than a , a contradiction.

Thus, $T = \mathbb{Z}$ and so $S' = S$. Thus, S is confined to the horizontal strip of width $2(d_2 + d_1)$ centered about L . This contradicts the fact that every point of $[x]$ is within d_1 of a point in $V(S)$. \square

The following questions remain unsolved.

QUESTION 1.6.5. Does there exist a clopen line section of $F(2^{\mathbb{Z}^2})$ that is both complete and co-complete?

QUESTION 1.6.6. Does there exist a clopen line section S of $F(2^{\mathbb{Z}^2})$ such that there are finitely many S -lines on each class?

1.7. Continuous cocycles

The potential function $P(x, y)$ we defined in the proof of Theorem 1.5.1 is in fact a cocycle from $F(2^{\mathbb{Z}^2})$ to \mathbb{Z} . The essential part of the proof was an analysis

of how this cocycle works. In this subsection we give a more general analysis for any continuous cocycle from $F(2^{\mathbb{Z}^2})$ to \mathbb{Z} . The main result is that they can be approximated by homomorphisms. We recall the definition of a cocycle in a general context.

DEFINITION 1.7.1. For groups G and H and an action $G \curvearrowright X$, recall that a *cocycle* is a map $\sigma: G \times X \rightarrow H$ satisfying the following cycle identity: for all $g_1, g_2 \in G$ and $x \in X$ we have

$$\sigma(g_1 g_2, x) = \sigma(g_1, g_2 \cdot x) \cdot \sigma(g_2, x).$$

In fact, we show this cocycle analysis in the more general context of residually finite groups, which is stated in Theorem 1.7.3. As an immediate consequence we obtain the analysis of \mathbb{Z} -cocycles for \mathbb{Z}^2 actions, which we state as Theorem 1.7.4. Recall that a countable group G is *residually finite* if it admits a decreasing sequence of finite-index normal subgroups G_n with $\bigcap_n G_n = \{1_G\}$. Finitely generated abelian groups are residually finite.

LEMMA 1.7.2. *Let G be a countable residually finite group, and let $(G_n)_{n \in \mathbb{N}}$ be a decreasing sequence of finite-index normal subgroups with $\bigcap_n G_n = \{1_G\}$. Then there is a hyper-aperiodic element $x \in F(2^G)$ with the following property. If \mathcal{P} is any clopen partition of $F(2^G)$ then there is a finite set $A \subseteq G$ so that for every $n \in \mathbb{N}$ the restriction of \mathcal{P} to $(G \setminus AG_n) \cdot x$ is G_n -invariant.*

PROOF. This is trivial if G is finite, so we assume G is infinite. By passing to a subsequence, we may assume G_{n+1} is a proper subgroup of G_n for every n . Define $x \in 2^G$ by setting $x(1_G) = 0$ and for $g \neq 1_G$ setting

$$x(g) = \min\{n \in \mathbb{N} : g \notin G_n\} \pmod{2}.$$

We check that x is hyper-aperiodic. Fix $1_G \neq s \in G$ and let k be least with $s \notin G_k$. Let T be a finite set satisfying $T^{-1}(G_k \setminus G_{k+1}) = G$. Such a T exists since $G_k \setminus G_{k+1}$ contains a coset of G_{k+1} and G_{k+1} has finite index in G . Now fix $g \in G$. Pick $t \in T$ with $tg \in G_k \setminus G_{k+1}$. Then $x(tg) = k+1 \pmod{2}$ while $x(tsg) = k \pmod{2}$ since, by normality of the G_n 's, $tsg = (tst^{-1})tg \in (G_{k-1} \setminus G_k)G_k = G_{k-1} \setminus G_k$. This proves x is hyper-aperiodic.

We claim that x is constant on $G_n h$ when $h \notin G_n$. Indeed, if $g \in G_n$ and $k < n$ satisfies $h \in G_k \setminus G_{k+1}$ then $gh \in G_k \setminus G_{k+1}$ as well and thus both $x(h)$ and $x(gh)$ are equal to $k+1 \pmod{2}$.

Now let \mathcal{P} be a clopen partition of $F(2^G)$. Since $\overline{[x]}$ is compact there is a finite $W \subseteq G$ such that for every $y \in \overline{[x]}$, the piece of \mathcal{P} to which y belongs is determined by $y \upharpoonright W$. We will check the lemma is satisfied with $A = W^{-1}$. So fix $n \in \mathbb{N}$, and for non-triviality assume $G \setminus W^{-1}G_n \neq \emptyset$. Fix $y \in (G \setminus W^{-1}G_n) \cdot x$, say $y = s \cdot x$ with $s \notin W^{-1}G_n$. Note that $ws \notin G_n$ for every $w \in W$ and so the claim of the previous paragraph applies. Using normality of G_n , we have that for every $g \in G_n$ and $w \in W$

$$(g \cdot y)(w) = (gs \cdot x)(w) = x(wgs) = x((wgw^{-1})ws) = x(ws) = (s \cdot x)(w) = y(w).$$

By letting $w \in W$ vary, we find that $g \cdot y$ and y belong to the same piece of \mathcal{P} . Since $g \in G_n$ was arbitrary we conclude that the restriction of \mathcal{P} to $(G \setminus W^{-1}G_n) \cdot x$ is G_n -invariant. \square

Recall that a finitely generated group G is *one-ended* if for every finite set $A \subseteq G$ the graph on $G \setminus A$ induced from the Cayley graph of G has the property that precisely one of its connected components is infinite (this property does not depend upon the choice of finite generating set for G used in constructing the Cayley graph).

THEOREM 1.7.3. *Let G be a finitely-generated, one-ended, residually finite group, let H be a countable group, and let $\delta: G \times F(2^G) \rightarrow H$ be a continuous cocycle. Let $(G_n)_{n \in \mathbb{N}}$ be a decreasing sequence of finite-index normal subgroups of G with $\bigcap_n G_n = \{1_G\}$. Then there are $x \in F(2^G)$, $n \in \mathbb{N}$, a group homomorphism $\phi: G_n \rightarrow H$, and finite sets $T \subseteq G$, $W \subseteq H$ such that*

- (i) $\delta(g, x) = \phi(g)$ for all $g \in G_n$,
- (ii) $TG_n = G$ and if $g \in G_n$ and $t \in T$ then $\delta(tg, x) \in W \cdot \phi(g)$.

PROOF. Let $y \in F(2^G)$ be the hyper-aperiodic element constructed in Lemma 1.7.2. Fix a finite generating set S for G . The map $z \in F(2^G) \mapsto (\delta(s, z))_{s \in S \cup S^{-1}} \in H^{S \cup S^{-1}}$ is continuous by assumption and thus produces a clopen partition \mathcal{P} of $F(2^G)$. Thus z_1 and z_2 belong to the same set in the partition \mathcal{P} if and only if for all $s \in S \cup S^{-1}$, $\delta(s, z_1) = \delta(s, z_2)$. Let $A \subseteq G$ be the finite set given by Lemma 1.7.2. By one-endedness, removing the vertices A from $\Gamma_{G,S}$ leaves only one infinite connected component and possibly a finite number of finite connected components. By enlarging A if necessary, we may assume that $\Gamma_{G,S} \setminus A$ is connected. For each pair of elements $g_1, g_2 \in G \setminus A$ which are adjacent to A , fix a path in $\Gamma_{G,S} \setminus A$ from g_1 to g_2 . Let B be a finite set containing all such paths with $B \cap A = \emptyset$. Now pick $t \in G$ and $n \in \mathbb{N}$ so that G_n is disjoint with $At^{-1} \cup ((tA^{-1}At^{-1} \cup tB^{-1}Bt^{-1}) \setminus \{1_G\})$. It follows from this that $G_n \cap At^{-1}G_n = \emptyset$ and that $At^{-1}g_1 \cap At^{-1}g_2 = Bt^{-1}g_1 \cap Bt^{-1}g_2 = \emptyset$ for all $g_1 \neq g_2 \in G_n$.

Fix a finite symmetric generating set V for G_n (this exists since G_n is finite-index in G). Fix $v \in V$ and consider a word w in $S \cup S^{-1}$ representing a path from 1_G to v in $\Gamma_{G,S}$. If this path does not intersect $At^{-1}G_n$ then set $w(v) = w$. Otherwise we will modify the path $w(v)$, without changing its endpoints, so that it does not intersect $At^{-1}G_n$. We can do this by applying the definition of B : if w meets $At^{-1}g$ then by considering the vertex it first enters $At^{-1}g$ from and the last vertex it exits $At^{-1}g$ to, we can replace the path in w between those two vertices by a path in $Bt^{-1}g$. Since both families of sets $At^{-1}g$ and $Bt^{-1}g$, $g \in G_n$, are pairwise disjoint as $g \in G_n$ varies, we can make all such changes to w in a single step. We let $w(v)$ be the newly obtained path. In summary, for each generator $v \in V$ of G_n , we have a path $w(v)$ in $\Gamma_{G,S}$ from 1_G to v which does not intersect $At^{-1}G_n$. From this it follows that, for every $g \in G_n$ and $v \in V$, $w(v)$ provides a path in $\Gamma_{G,S}$ from g to vg which does not intersect $At^{-1}G_n$.

Set $x = t \cdot y$. By Lemma 1.7.2 the restriction of \mathcal{P} to $(G \setminus At^{-1}G_n) \cdot x = (G \setminus AG_n) \cdot y$ is G_n -invariant. This means that if $h \notin At^{-1}G_n$ and $g \in G_n$ then $\delta(s, h \cdot x) = \delta(s, (hgh^{-1})h \cdot x) = \delta(s, hg \cdot x)$ for all $s \in S \cup S^{-1}$. Now consider $g \in G_n$ and $v \in V$. Write the word $w(v)$ as $s_{\ell(v)}(v) \cdots s_1(v)$ where each $s_i(v) \in S \cup S^{-1}$, and set $w_i(v) = s_i(v) \cdots s_1(v)$. Then

$$\begin{aligned} \delta(v, g \cdot x) &= \delta(s_{\ell(v)}(v), w_{\ell(v)-1}(v)g \cdot x) \cdots \delta(s_1(v), g \cdot x) \\ &= \delta(s_{\ell(v)}(v), w_{\ell(v)-1}(v) \cdot x) \cdots \delta(s_1(v), x) \\ &= \delta(v, x). \end{aligned}$$

Define $\phi : G_n \rightarrow H$ by $\phi(g) = \delta(g, x)$. The above computation shows that for $g \in G_n$ and $v \in V$

$$\phi(vg) = \delta(vg, x) = \delta(v, g \cdot x)\delta(g, x) = \delta(v, x)\delta(g, x) = \phi(v)\phi(g).$$

Thus ϕ is a group homomorphism.

Finally, let T be a finite set satisfying $TG_n = G$. Since $\overline{[x]}$ is compact and δ is continuous, there is a finite set $W \subseteq H$ with $\delta(t, z) \in W$ for all $t \in T$ and all $z \in \overline{[x]}$. Then for $t \in T$ and $g \in G_n$ we have

$$\delta(tg, x) = \delta(t, g \cdot x)\delta(g, x) \in W\phi(g). \quad \square$$

THEOREM 1.7.4. *Let $\delta : \mathbb{Z}^2 \times F(2^{\mathbb{Z}^2}) \rightarrow \mathbb{Z}$ be a continuous cocycle. Then there are $x \in F(2^{\mathbb{Z}^2})$, $\alpha, \beta \in \mathbb{Q}$, and $C \in \mathbb{Z}^+$ such that for any $a, b \in \mathbb{Z}$,*

$$\delta((a, b), x) = \alpha a + \beta b + \epsilon,$$

where $|\epsilon| \leq C$.

PROOF. Apply Theorem 1.7.3 with $G_n = (3^n\mathbb{Z}) \times (3^n\mathbb{Z})$ to obtain x , a homomorphism $\phi : (3^n\mathbb{Z}) \times (3^n\mathbb{Z}) \rightarrow \mathbb{Z}$, and finite sets $T \subseteq \mathbb{Z}^2$, $W \subseteq \mathbb{Z}$. Define $\alpha = \phi((3^n, 0))/3^n$ and $\beta = \phi((0, 3^n))/3^n$. Let C be the maximum value of $|-t_1\alpha - t_2\beta + w|$ for $(t_1, t_2) \in T$ and $w \in W$. Now fix $(a, b) \in \mathbb{Z}^2$ and pick $(p, q) \in \mathbb{Z}^2$ and $(t_1, t_2) \in T$ with $(a, b) = (t_1, t_2) + (p3^n, q3^n)$. Then there is $w \in W$ with

$$\begin{aligned} \delta((a, b), x) &= \delta((t_1, t_2) + (p3^n, q3^n), x) \\ &= w + p3^n\alpha + q3^n\beta \\ &= (a - t_1)\alpha + (b - t_2)\beta + w \\ &= a\alpha + b\beta - t_1\alpha - t_2\beta + w. \end{aligned}$$

This completes the proof since $|-t_1\alpha - t_2\beta + w| \leq C$. \square

1.8. Marker distortion

In this section we present a result which serves to illustrate some of the methods we will be using, but in a much simpler setting. Namely, we will present a result concerning continuous actions of \mathbb{Z} . In this one-dimensional case, the arguments reduce to simpler, but not trivial, arguments. In particular, the proof of Theorem 1.8.4 will use a one-dimensional version of a “tile” and will serve as a warm-up for the more general tile construction given later. Specifically, in Theorem 2.2.2 we will present a general “tiles theorem” for $F(2^{\mathbb{Z}})$ and in Theorem 2.3.5 a corresponding result for $2^{\mathbb{Z}^2}$.

The result we present has to do with the notion of *marker distortion* which we now introduce. We first recall the following well-known fact (see [11]).

FACT 1.8.1. Let X be a zero-dimensional Polish space, and let $\mathbb{Z} \curvearrowright X$ be a continuous, free action of \mathbb{Z} on X . Then for any integer $d > 1$ there is a clopen complete section $M_d \subseteq X$ such that for any $x \in M_d$, the first point of M_d to the right (or left) of x has distance d or $d + 1$ from x .

In the above, by “to the right of x ” we mean that $y = n \cdot x$ where $n \in \mathbb{Z}^+$, and similarly for “left”. By “distance” from x to y we mean $\rho(x, y) = |n|$ where $n \in \mathbb{Z}$ is the unique integer such that $n \cdot x = y$.

A natural question, which occurs in many contexts, is how “regular” of a marker structure can we put on an equivalence relation. The notion of regular can be measured in different ways, such as size, geometric shape, etc. In the one-dimensional case, since we are able to use marker distances of d and $d + 1$, we might ask how close we can keep the average marker distance to $d + \frac{1}{2}$. We make this precise in the following definition.

DEFINITION 1.8.2. Let $d > 1$ be a positive real number, and M a complete section for the action of \mathbb{Z} on the Polish space X . The d -distortion function for M is the function $f_d: X \times \mathbb{Z}^+ \rightarrow \mathbb{R}$ is defined by:

$$f_d(x, n) = \rho(x, x_M^n) - nd,$$

where x_M^n is the n th point of M to the right of x (if this is not defined we leave $f_d(x, n)$ undefined). When d is understood we simply write $f(x, n)$.

Even more generally, we can allow d to depend on the equivalence class of x . If $d: X \rightarrow \mathbb{R}$ is an invariant Borel function (that is, $f(x) = f(y)$ whenever $[x] = [y]$), then we write $f_{d(x)}(x, n)$ for $f_d(x, n)$ where $d = d(x)$. This measures the marker distortion with respect to the “expected” average distance $d(x)$ between marker points in $[x]$. One could use a prescribed function $d(x)$, or given the complete section M , use the function $\Delta(x)$ associated to M by: $\Delta(x) = \limsup_n \frac{\rho(x, x_M^n)}{n}$.

A particular case of this is when $M = M_d$ is a complete section such that the ρ distance between consecutive points of M_d is in $\{d, d + 1\}$. In this case, we can interpret the distortion function as follows. As we move from a point to the next marker point to the right, we record a charge of $+\frac{1}{2}$ if the distance is $d + 1$, and a charge of $-\frac{1}{2}$ if the distance is d . The distortion function $f_{d+\frac{1}{2}}(x, n)$ then measures the total charge accumulated as we move from x to the n th marker point to the right.

It is natural to ask if we can keep the distortion function bounded, or how slow-growing can we keep it? Can we do better if we allow Borel marker sets M_d instead of clopen ones? We next state three theorems which answer some of these questions.

The first result says that bounded distortion is impossible to attain for $F(2^{\mathbb{Z}})$.

THEOREM 1.8.3. *For any invariant Borel function $d: F(2^{\mathbb{Z}}) \rightarrow \mathbb{R}$ with $d(x) > 1$ for all x , there does not exist a Borel complete section $M \subseteq F(2^{\mathbb{Z}})$ with bounded d -distortion, that is, having the property that for some $B \in \mathbb{R}$, $|f_{d(x)}(x, n)| \leq B$ for all $x \in F(2^{\mathbb{Z}})$ and all $n \in \omega$.*

The second result says that using clopen marker sets in $F(2^{\mathbb{Z}})$ we cannot do better than linear distortion. Note that the basic Fact 1.8.1 shows that linear distortion is possible with clopen markers, in fact we can get $f_d(x, n) \leq n$.

THEOREM 1.8.4. *Suppose $d \in \mathbb{R}$ with $d > 1$ and M is a clopen complete section in $F(2^{\mathbb{Z}})$. Then there are $C > 0$ and $x \in F(2^{\mathbb{Z}})$ such that for all $g \in \mathbb{Z}$ and all large enough n we have $|f_d(g \cdot x, n)| > Cn$.*

The third result says that that using Borel marker sets, we can improve the distortion from linear growth to arbitrarily slow growth.

THEOREM 1.8.5. *Let X be a Polish space and $\mathbb{Z} \curvearrowright X$ be a free, Borel action. Let $d > 1$ be an integer. Let $f: \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$ be monotonically increasing with*

$\lim_n f(n) = +\infty$ and $f(1) > 4(2d + 1)$. Then there is a Borel complete section $M_d \subseteq X$ such that for all $x \in X$ and $n \in \mathbb{Z}^+$ we have $|f_d(x, n)| \leq f(n)$.

The proof of Theorem 1.8.3 is a simple category argument which we give below for the sake of completeness. The proof of Theorem 1.8.4 given below will be a special version of the more general “tile” argument for $F(2^{\mathbb{Z}})$ to be given later in Theorem 2.2.2. The proof of Theorem 1.8.5 uses different techniques, specifically the method of orthogonal markers, and will be presented in [13].

PROOF OF THEOREM 1.8.3. Suppose $M \subseteq F(2^{\mathbb{Z}})$ is a Borel complete section with d -distortion bounded by B . The orbit equivalence relation on $F(2^{\mathbb{Z}})$ is generically ergodic (see [10] Proposition 6.1.9), which implies (by [10] Proposition 10.1.2) that the function d is constant on an invariant comeager subset of $F(2^{\mathbb{Z}})$. Denote this constant value by d_0 . Let $C \subseteq F(2^{\mathbb{Z}})$ be an invariant comeager set such that $d(x) = d_0$ for all $x \in C$ and such that $M \cap C$ is relatively clopen in C (which we can obtain as M is Borel so has the Baire property). Let $\{D_n\}_{n \in \omega}$ be dense open sets in $2^{\mathbb{Z}}$ with $\bigcap_n D_n \subseteq C$.

We construct a particular point $x \in C$. Write K for the set of integers $k \in \mathbb{Z}$ satisfying $\inf_{n \in \mathbb{Z}} |k - nd_0| \geq 1/3$. Note that K is infinite since $d_0 > 1$. Let $\{(I_n, J_n)\}$ enumerate all pairs of intervals in \mathbb{Z} with $I_n \subseteq J_n$, such that each possible pair (I, J) occurs infinitely often in the enumeration. We construct a decreasing sequence of basic open sets $U_0 \supseteq U_1 \supseteq \dots$ in $2^{\mathbb{Z}}$. In the following argument, if $U \subseteq 2^{\mathbb{Z}}$ is a basic open set, then by $\text{dom}(U)$ we mean the interval $[i, j] \subseteq \mathbb{Z}$ such that for some $p \in 2^{[i, j]}$ we have $U = \{y \in 2^{\mathbb{Z}} : y \upharpoonright [i, j] \in p\}$. We say that U is the basic open set determined by $[i, j]$ and p .

Let $U_0 = 2^{\mathbb{Z}}$. Given U_n , first get a basic open $U'_n \subseteq U_n$ with $U'_n \subseteq D_n$. If $J_n \not\subseteq \text{dom}(U'_n)$, then set $U_{n+1} = U'_n$. If $J_n \subseteq \text{dom}(U'_n)$, then we take $U_{n+1} \subseteq U'_n$ such that $\text{dom}(U_{n+1})$ contains a disjoint copy J'_n (that is, J'_n has the same length as J_n and $J_n \cap J'_n = \emptyset$) and such that if a is the left endpoint of I_n , and a' the left endpoint of I'_n (the corresponding subinterval of J'_n), then $a' > a$ and $a' - a \in K$, and finally if U_{n+1} is determined by the interval $\text{dom}(U_{n+1}) = [i, j] \subseteq \mathbb{Z}$ and $p \in 2^{[i, j]}$, then $p \upharpoonright J_n \cong p \upharpoonright J'_n$ (meaning precisely that if $J_n = [c, d]$, $J'_n = [c', d']$, then for all $0 \leq i < d - c$ we have $p(c + i) = p(c' + i)$). Let $\text{dom}(U_n) = [i_n, j_n]$, and we may assume that $\lim_{n \rightarrow \infty} i_n = -\infty$, $\lim_{n \rightarrow \infty} j_n = \infty$. Let $\{x\} = \bigcap_n U_n \subseteq C$. See Figure 14 for an illustration of the construction of x .




FIGURE 3. The construction of U_{n+1} . We have $x \upharpoonright J_n \cong x \upharpoonright J'_n$ for any $x \in U_{n+1}$.

Since $x \in C$, we have for all $m, n > 0$ that $|f_{d_0}(m \cdot x, n)| \leq B$. Let $\alpha \in \mathbb{R}$ be the infimum of the values of $f_{d_0}(m \cdot x, n)$ as m ranges over all integers such that $m \cdot x \in M$ and n ranges over all positive integers. Note that, while α may be negative, it exists by the boundedness of the distortion. Likewise, let $\beta \in \mathbb{R}$ be

the supremum value of $f_{d_0}(m \cdot x, n)$ as m ranges within the set of values for which $m \cdot x \in M$, $\inf_r f_{d_0}(m \cdot x, r) < \alpha + 1/3$, and n ranges over all positive integers. Fix now $a \in \mathbb{Z}$ and $n_0 > 0$ such that $a \cdot x \in M$, and $f_{d_0}(a \cdot x, r)$ achieves values below $\alpha + 1/3$ and above $\beta - 1/3$ for $r \in [0, n_0]$. From our definitions, we have that $\alpha \leq f_{d_0}(a \cdot x, n) \leq \beta$ for all $n \in \mathbb{N}$. Let I be the interval $[a, b]$ where $b \cdot x$ is the n_0 th element of M to the right of $a \cdot x$.

Let $J \supseteq I$ be an interval such that if $U_{x \upharpoonright J}$ denotes the basic open set determined by J and $x \upharpoonright J$, then for all $r \in [a, b]$ we have that $r \cdot U_{x \upharpoonright J} \subseteq M$ or $r \cdot U_{x \upharpoonright J} \subseteq 2^{\mathbb{Z}} \setminus M$. This is possible as $M \cap C$ is relatively clopen in C . Let n be such that $(I_n, J_n) = (I, J)$ and $J \subseteq \text{dom}(U_n)$. By definition of U_{n+1} , we have that $\text{dom}(U_{n+1})$ includes a disjoint copy J'_n of J_n , and if U_{n+1} is determined by $p \in 2^{\text{dom}(U_{n+1})}$, then $p \upharpoonright J_n \cong p \upharpoonright J'_n$.

Let $I'_n = [a', b']$ be the corresponding subinterval of J'_n . Recall that $I_n = I = [a, b]$ and note $b - a = b' - a'$. We have that $a \cdot x \in M$ and $b \cdot x \in M$. Also, from the definition of J and the fact that $p \upharpoonright J \cong p \upharpoonright J'$ we have that for all $0 \leq i < b - a$ that $(a + i) \cdot x \in M$ if and only if $(a' + i) \cdot x \in M$. It follows that the maps $r \mapsto f_{d_0}(a \cdot x, r)$ and $r \mapsto f_{d_0}(a' \cdot x, r)$ for $r \in [0, n_0]$ are identical. In particular, $f_{d_0}(a' \cdot x, r)$ achieves values below $\alpha + 1/3$ and above $\beta - 1/3$ for $r \in [0, n_0]$.

Suppose $a' \cdot x$ is the q th point of M to the right of $a \cdot x$. Then $f_{d_0}(a \cdot x, q) = \rho(a \cdot x, a' \cdot x) - qd_0$. Since $\rho(a \cdot x, a' \cdot x) = a' - a \in K$, we have that $|f_{d_0}(a \cdot x, q)| \geq 1/3$. Notice that

$$f_{d_0}(a \cdot x, q + r) = f_{d_0}(a \cdot x, q) + f_{d_0}(a' \cdot x, r).$$

By construction, $f_{d_0}(a' \cdot x, r)$ achieves values below $\alpha + 1/3$ and above $\beta - 1/3$ for $r \in [0, n_0]$. If $f_{d_0}(a \cdot x, q) \leq -1/3$, then for some $r \leq n_0$ we have that $f_{d_0}(a \cdot x, q + r) < -1/3 + \alpha + 1/3 = \alpha$. If $f_{d_0}(a \cdot x, q) \geq 1/3$, then we likewise get that $f_{d_0}(a \cdot x, q + r) > 1/3 + \beta - 1/3 = \beta$ for some $r \leq n_0$. In either case, this contradicts $\alpha \leq f_{d_0}(a \cdot x, r) \leq \beta$ for all $r \in \mathbb{N}$. \square

The proof of Theorem 1.8.4 will involve constructing a special hyper-aperiodic element $x \in F(2^{\mathbb{Z}})$ which contains copies of certain ‘‘tiles’’ $T_{w,p}$ for $w, p \in \mathbb{Z}^+$. The existence of a hyper-aperiodic element with these tile structures will allow us to deduce the theorem.

PROOF OF THEOREM 1.8.4. Let $d > 1$ be a real number and $M \subseteq F(2^{\mathbb{Z}})$ be a clopen complete section. Again let K be the set of integers k satisfying $\inf_{n \in \mathbb{Z}} |k - nd| \geq 1/3$. Note that K is infinite. For $w, p \in \mathbb{Z}^+$, a $T_{w,p}$ tile is a function $t: [a, b] \rightarrow 2$ with $b - a = 2w + p$ such that $t \upharpoonright [a, a + 2w] \cong t \upharpoonright [b - 2w, b]$. Thus, a $T_{w,p}$ tile has a block R of size $2w + 1$ (meaning $|\text{dom}(R)| = 2w + 1$), followed by a block S of size $p - 2w - 1$, and then followed by the block R again. This is illustrated in Figure 4. Notice that the domain of a $T_{w,p}$ tile has a central interval $[a + w, a + w + p]$, which is of length p (and size $p + 1$), surrounded on each side by a ‘‘buffer’’ interval of size w , as shown in Figure 4.




FIGURE 4. A $T_{w,p}$ tile

We now construct a hyper-aperiodic element $x \in F(2^{\mathbb{Z}})$ with the property that for arbitrarily large w and p , with $p \in K$, there is an interval $[a, b]$ such that $x \upharpoonright [a, b]$ is a $T_{w,p}$ tile. Fix an arbitrary hyper-aperiodic element $y \in M$. Such y exists since M is a complete section. Let $\{(w_n, p_n)\}$ be any sequence with $\lim_n w_n = \lim_n (p_n - 2w_n - 2) = \infty$, with $p_n - 2w_n - 2 \geq 0$, and with $p_n \in K$. Let $I_n = [a_n, b_n]$, with $0 < a_n < b_n < a_{n+1}$, be a pairwise disjoint sequence of intervals in \mathbb{Z} such that $b_n - a_n = 2w_n + p_n$, and the gap $a_{n+1} - b_n$ between I_n and I_{n+1} also tends to ∞ with n . Each interval I_n can be viewed as the disjoint union of three subintervals A_n, B_n, C_n , where $A_n = [a_n, a_n + 2w_n]$, $B_n = [a_n + 2w_n + 1, a_n + p_n - 1]$, and $C_n = [a_n + p_n, a_n + 2w_n + p_n]$. Fix any $x \in 2^{\mathbb{Z}}$ satisfying:

- (1) For $i \notin \bigcup_n I_n$, $x(i) = y(i)$.
- (2) For all $n \in \mathbb{N}$, $x \upharpoonright B_n \cong y \upharpoonright [-\lfloor p_n/2 \rfloor + w_n + 1, \lfloor p_n/2 \rfloor - w_n - 1]$.
- (3) For all $n \in \mathbb{N}$, $x \upharpoonright A_n \cong x \upharpoonright C_n \cong y \upharpoonright [-w_n, w_n]$.

Now we check that x is hyper-aperiodic. Fix $s \in \mathbb{Z} \setminus \{0\}$, and let T' be a finite set witnessing the hyper-aperiodicity of s for y , that is, for every $g \in \mathbb{Z}$ there is $t \in T'$ with $y(g+t) \neq y(g+s+t)$. Consider the partition $Q = \{\mathbb{Z} \setminus \bigcup_n I_n\} \cup \bigcup_n \{A_n, B_n, C_n\}$ of \mathbb{Z} . By construction there is a finite set $U \subseteq \mathbb{Z}$ such that for all $g \in \mathbb{Z}$ there is $u \in U$ with $(g+u+T') \cup (g+s+u+T')$ contained in a single piece of Q . Set $T = U + T'$. Now fix $g \in G$. Pick $u \in U$ with $L = (g+u+T') \cup (g+s+u+T')$ contained in a single piece of Q . If L is contained in $\mathbb{Z} \setminus \bigcup_n I_n$ then $x \upharpoonright L = y \upharpoonright L$ and thus we are done since there is $t \in T'$ with

$$x(g+u+t) = y((g+u)+t) \neq y((g+u)+s+t) = y(g+s+u+t) = x(g+s+u+t).$$

On the other hand, if L is contained in some A_n, B_n , or C_n , then there is $h \in \mathbb{Z}$ with $x \upharpoonright L = (h \cdot y) \upharpoonright L$. Again we are done because we can find $t \in T'$ with

$$x(g+u+t) = y(h+g+u+t) \neq y(h+g+s+u+t) = x(g+s+u+t).$$

Consider now the compact set $C = \overline{[x]} \subseteq F(2^{\mathbb{Z}})$. Since M is a clopen subset of $F(2^{\mathbb{Z}})$, for each $\alpha \in F(2^{\mathbb{Z}})$ there is an $m \in \mathbb{Z}^+$ such that, letting $U_{\alpha \upharpoonright [-m, m]}$ be the basic open set of elements extending $\alpha \upharpoonright [-m, m]$, we have either $U_{\alpha \upharpoonright [-m, m]} \cap F(2^{\mathbb{Z}}) \subseteq M$ or $U_{\alpha \upharpoonright [-m, m]} \cap F(2^{\mathbb{Z}}) \cap M = \emptyset$. For each $\alpha \in F(2^{\mathbb{Z}})$ let $m(\alpha)$ be the least such $m \in \mathbb{Z}^+$. Then the function $\alpha \mapsto m(\alpha)$ is continuous on $F(2^{\mathbb{Z}})$. In particular, its restriction on C is continuous. The compactness of C now gives that there is an m_0 such that $m(\alpha) \leq m_0$ for all $\alpha \in C$. In particular, for all $\alpha \in C$, either $U_{\alpha \upharpoonright [-m_0, m_0]} \cap F(2^{\mathbb{Z}}) \subseteq M$ or $U_{\alpha \upharpoonright [-m_0, m_0]} \cap F(2^{\mathbb{Z}}) \cap M = \emptyset$. Applying this to elements of $[x]$, we get that for any $a \in \mathbb{Z}$, $x \upharpoonright [a - m_0, a + m_0]$ determines whether $\alpha \in M$ for any $\alpha \in F(2^{\mathbb{Z}})$ with $\alpha \in U_{x \upharpoonright [a - m_0, a + m_0]}$.

Now consider a large enough n in the construction of x where $w_n > m_0$ and $p_n > 2w_n + 2m_0 + 2$. Then $x \upharpoonright B_n \cong y \upharpoonright I$ for an interval I with $[-m_0, m_0] \subseteq I$. Since $y \in M$ we deduce that there is $i \in B_n$ such that $i \cdot x \in M$. For notational simplicity we denote $w = w_n, p = p_n$, and let τ denote the $T_{w,p}$ tile $x \upharpoonright I_n$.

We construct an element $\tilde{x} \in F(2^{\mathbb{Z}})$ as follows. First let z be a periodic element of $2^{\mathbb{Z}}$ which is defined as an overlapping concatenation of the $T_{w,p}$ tile τ . To be precise, letting $\tau = s \circ t \circ s$ where $|\text{dom}(s)| = 2w + 1$ and $|\text{dom}(t)| = p - 2w - 1$, then z is a bi-infinite concatenation of the form $\dots s \circ t \circ s \circ t \circ s \dots$. Without loss of generality, assume $z \upharpoonright [0, p - 2w - 2] \cong t$. Next for all $k \geq 0$ let $L_k = [2^{k+1}p, 2^{k+1}p + p - 2w - 2]$ and $N_k = [2^{k+1}p + p - 2w - 1, 2^{k+2}p - 1]$. Then

$z \upharpoonright L_k \cong t$, N_k is the interval in between L_k and L_{k+1} , and $z \upharpoonright N_k \cong s \frown t \frown \dots \frown t \frown s$ with $2^{k+1} - 1$ many occurrences of t . Define \tilde{x} by

$$\tilde{x}(i) = \begin{cases} z(i) & \text{if } i \notin \{2^{k+1}p : k \geq 0\}, \\ z(i) & \text{if } i = 2^{k+1}p \text{ and } z(i) \neq z(i - k - 1), \\ 1 - z(i) & \text{if } i = 2^{k+1}p \text{ and } z(i) = z(i - k - 1). \end{cases}$$

Then $\tilde{x} \in F(2^{\mathbb{Z}})$ because the definition of \tilde{x} makes it explicit that any $k + 1$ is not a period. Note also that \tilde{x} and z agree on all intervals except $\bigcup_{k \geq 0} L_k$. In particular, $\tilde{x} \upharpoonright N_k = z \upharpoonright N_k$ for all $k \geq 0$. The construction of \tilde{x} is illustrated in Figure 5.



FIGURE 5. The element \tilde{x}

It follows from the construction of τ that for each interval J for which $\tilde{x} \upharpoonright J \cong \tau$, there is $i \in J$ such that $i \cdot \tilde{x} \in M$. Moreover, for each $k \geq 0$, $i \mapsto \chi_M(i \cdot \tilde{x})$ has period p for $i \in N_k$. The same is true for $i \in (-\infty, 2p - 1]$. Consider any fixed $k \geq 0$. Let $i_0 \in N_k$ be the least such that $i_0 \cdot \tilde{x} \in M$. Let $r_0 \in \mathbb{Z}^+$ be such that $(i_0 + p) \cdot \tilde{x}$ is the r_0 -th element of M to the right of $i_0 \cdot \tilde{x}$. Let $c_0 = f_d(i_0 \cdot \tilde{x}, r_0)$ and note that $|c_0| \geq 1/3$ since $p \in K$. It follows that $f_d(i_0 \cdot \tilde{x}, \ell r_0) = \ell c_0$ for $\ell < 2^{k+1} - 1$. Thus there is a constant $\gamma > 0$ such that for all $q < (2^{k+1} - 1)r_0$,

$$\frac{q}{r_0} c_0 - \gamma \leq f_d(i_0 \cdot \tilde{x}, q) \leq \frac{q}{r_0} c_0 + \gamma.$$

Now let $j_0 \in N_k$ be the largest such that $j_0 \cdot \tilde{x} \in M$ and $i_1 \in N_{k+1}$ be the least such that $i_1 \cdot \tilde{x} \in M$. Then the absolute value of the total d -distortion from $j_0 \cdot \tilde{x}$ to $i_1 \cdot \tilde{x}$ is bounded as k varies through all natural numbers. Let $\beta > 0$ be such a bound. By a direct computation we get that for all $g \in \mathbb{Z}$ and for all ℓ ,

$$\frac{c_0}{r_0} \ell - \gamma - \beta \log \ell \leq f_d(g \cdot \tilde{x}, \ell) \leq \frac{c_0}{r_0} \ell + \gamma + \beta \log \ell.$$

Thus for large enough ℓ we have $|f_d(g \cdot \tilde{x}, \ell)| > \frac{|c_0|}{2r_0} \ell$. □

The Twelve Tiles Theorem

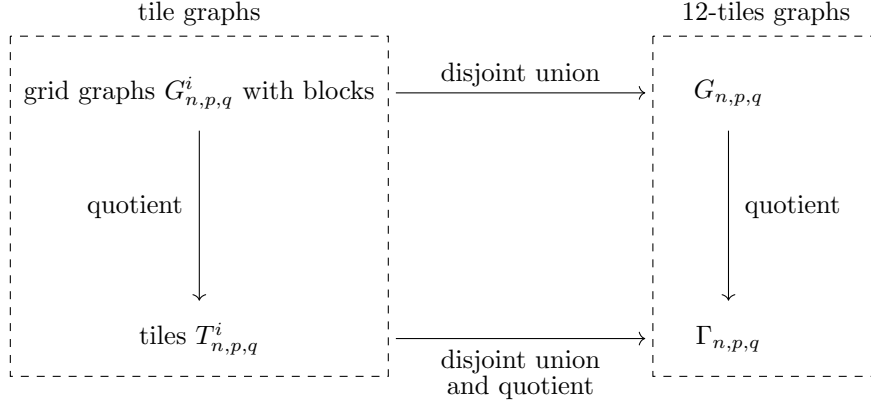
In this chapter we prove a theorem which completely answers the question of when there exists a continuous equivariant map from $F(2^{\mathbb{Z}^2})$ into a given subshift of finite type. This type of question can also be formulated through the notion of a “continuous structuring” on $F(2^{\mathbb{Z}^2})$ (to be defined in §2.6) which means a continuous way of associating a first-order structure to the equivalence classes of $F(2^{\mathbb{Z}^2})$. We prove our main theorem, Theorem 2.3.5, first in the terminology of subshifts of finite type and then in §2.6 we present the theorem in terms of continuous structurings. Although these versions of the theorem are equivalent, we find both points of view useful.

The main theorem generalizes to $F(2^{\mathbb{Z}^n})$, but the case $n = 2$ illustrates the general arguments and suffices for the applications we have. As a warm-up, we prove an analogous result in §2.2 for $F(2^{\mathbb{Z}})$. This illustrates some of the main ideas, but the argument is significantly easier, and the answer involves a graph built from only two tiles instead of 12 as in the case for $F(2^{\mathbb{Z}^2})$. This “two tiles” theorem, Theorem 2.2.2, gives a complete answer as to when there is a continuous equivariant map from $F(2^{\mathbb{Z}})$ to a subshift of finite type $Y \subseteq \mathfrak{b}^{\mathbb{Z}}$, just as the “twelve tiles” theorem, Theorem 2.3.5, provides a complete answer for $F(2^{\mathbb{Z}^2})$ (considering subshifts of finite type $Y \subseteq \mathfrak{b}^{\mathbb{Z}^2}$). We will also need the simpler Theorem 2.2.2 for the arguments of §4.2.

The twelve tiles theorem for $F(2^{\mathbb{Z}^2})$ is given in terms of certain finite graphs which will be denoted $\Gamma_{n,p,q}$. The graphs $\Gamma_{n,p,q}$ will have a uniform definition from the three parameters n, p, q which will be positive integers. The theorem will state that a map from $F(2^{\mathbb{Z}^2})$ into a subshift of finite type (or a continuous \mathcal{L} -structuring) exists if and only if there is a triple n, p, q with $\gcd(p, q) = 1$ such that there is a corresponding map from $\Gamma_{n,p,q}$ into the subshift (or \mathcal{L} -structuring). These notions will be made precise in the following sections. The theorem will also state that this existence criterion is equivalent to the condition that for *all* sufficiently large n, p, q with $\gcd(p, q) = 1$ we have a map from $\Gamma_{n,p,q}$ into the subshift (or \mathcal{L} -structures), a result which is not obvious directly. Again, §2.2 will present an easier version of this result for $F(2^{\mathbb{Z}})$, using a simpler family of graphs $\Gamma_{n,p,q}^{(1)}$.

The graphs $\Gamma_{n,p,q}$, as we said above, will all be defined in the same manner. We will start with a collection of 12 rectangular “tiles.” These will be rectangular grid graphs $G_{n,p,q}^i$ for $i = 1, \dots, 12$, each of which will be divided into certain “blocks” of vertices, with “block labels” given to some of them. Blocks with the same labels will eventually be identified with each other to form quotient graphs $T_{n,p,q}^i$. We will refer to both collections $G_{n,p,q}^i$ and $T_{n,p,q}^i$ as “tile graphs.” To construct $\Gamma_{n,p,q}$ we will form some disjoint unions of the twelve tiles $G_{n,p,q}^i$ which will be denoted $G_{n,p,q}$, and finally $\Gamma_{n,p,q}$ will be a quotient graph of $G_{n,p,q}$ by identifying the vertices in

similarly labeled blocks. Alternatively, one can also obtain $\Gamma_{n,p,q}$ by first taking the union of all $T_{n,p,q}^i$ and then again taking the quotient by identifying the vertices in similarly labeled blocks. We refer to $G_{n,p,q}$ and $\Gamma_{n,p,q}$ as the “12-tiles graphs.” The following diagram illustrates the constructions and the terminology. The simpler graphs $\Gamma_{n,p,q}^{(1)}$ of §2.2 will be a simplified version of this construction, starting from just two tiles.



With the main theorem proved in this chapter we will be able to answer many questions about $F(2^{\mathbb{Z}^2})$ concerning continuous functions on or clopen sets in $F(2^{\mathbb{Z}^2})$. In the rest of the paper we will see that many natural questions such as proper coloring questions, perfect matching questions, tiling problems, and questions about continuous graph homomorphisms can be formulated in terms of equivariant maps into subshifts of finite type or \mathcal{L} -structures, and thus solved by an (often simple) argument about the finite graphs $\Gamma_{n,p,q}$. In particular, we will be able to derive $\chi_c(F(2^{\mathbb{Z}^2})) = 4$ as a quick corollary of our main theorem, thus giving another proof of Theorem 1.5.1.

It is somewhat amusing to note that the twelve tiles theorem makes it possible to approach some problems in the descriptive set theory of countable Borel equivalence relations through a computer analysis. Indeed, some of the arguments of the following chapters were motivated by computer studies.

2.1. From proper colorings to subshifts of finite type

In the last chapter we already studied continuous proper colorings on $F(2^{\mathbb{Z}^n})$. In order to study other related continuous combinatorial questions about $F(2^{\mathbb{Z}^n})$, we note that the proper coloring problem is a special case of the more general question on subshifts of finite type. In this section we also fix some definitions and notations that will be used throughout the rest of the paper.

First, we will be working with a notion of a G -graph for a group G , which is a notion weaker than the Cayley graph or the Schreier graph of a free action.

DEFINITION 2.1.1. Let G be a group, S be a generating set of G , and Γ be a directed graph. We say that Γ is a (G, S) -graph (or in short a G -graph if there is no danger of confusion) if each element of $E(\Gamma)$ is uniquely labeled with an element of S .

Later in the chapter we will introduce “block labels” for vertices of G -graphs. To properly distinguish different kinds of labels, we refer to the above labels for the edges of a G -graph as G -labels.

Recall our convention from §1.1 on page 14 that S is always finite and minimal in this paper. In particular, the identity is not an element of S . Also, since S is minimal, i.e., no proper subset of S is a generating set of G , we have that if $s \in S$ then $s^{-1} \notin S$. In the case of \mathbb{Z}^n , our standard generating set is $\{e_1, \dots, e_n\}$, where for $1 \leq i \leq n$, $e_i = (a_1, \dots, a_n)$ with $a_i = 1$ and $a_j = 0$ for $j \neq i$. Under this convention, the Cayley graphs under our consideration will never have both an edge (x, y) and its reverse edge (y, x) simultaneously contained in the edge set. This sometimes causes inconvenience. To remedy this, we adopt another convention that, whenever (x, y) is an edge in a Cayley graph with label s , we consider (y, x) to have label s^{-1} .

The G -graphs Γ under our consideration will be either a part of a Cayley graph or a combination of some parts of a Cayley graph, and thus they will have the same property that if $(x, y) \in E(\Gamma)$ then $(y, x) \notin E(\Gamma)$, which causes inconvenience sometimes. We will therefore also make a convention that in a G -graph Γ , whenever $(x, y) \in E(\Gamma)$ and is labeled s , then we consider (y, x) to be labeled by s^{-1} . With this additional convention, each G -graph we consider will always be symmetric, and corresponding uniquely to an undirected graph.

We will also be working with G -homomorphisms between G -graphs, which are defined below.

DEFINITION 2.1.2. Let G be a group with a generating set S , and Γ_1, Γ_2 be G -graphs. A map $\varphi: \Gamma_1 \rightarrow \Gamma_2$ is called a G -homomorphism if whenever $(x, y) \in E(\Gamma_1)$ is labeled with $s \in S$, then $(\varphi(x), \varphi(y)) \in E(\Gamma_2)$ is also labeled with s .

Proper colorings and chromatic numbers are defined for G -graphs in the same manner as for general directed graphs. In the case that Γ is a topological graph, we also define continuous chromatic number, Borel chromatic number, etc. in a standard manner.

For $n \geq 1$, by an n -dimensional (rectangular) grid graph we mean the Cayley graph of \mathbb{Z}^n restricted to an n -dimensional rectangle $[a_1, b_1] \times \dots \times [a_n, b_n]$ where $a_i \leq b_i \in \mathbb{Z}$ for all $1 \leq i \leq n$. An n -dimensional grid graph is naturally a \mathbb{Z}^n -graph. If we do not specify the dimension n , then tacitly $n = 2$.

Next we turn to the shift space $\mathfrak{b}^{\mathbb{Z}^n}$, where $\mathfrak{b} \geq 1$ is an integer, and its subshifts.

We first formalize some notions related to the topology of this space. If $g \in \mathbb{Z}^n$ and $p: \text{dom}(p) \rightarrow \{0, 1, \dots, \mathfrak{b} - 1\}$ and $q: \text{dom}(q) \rightarrow \{0, 1, \dots, \mathfrak{b} - 1\}$ are finite partial functions from \mathbb{Z}^n to $\mathfrak{b} = \{0, 1, \dots, \mathfrak{b} - 1\}$, i.e., $\text{dom}(p), \text{dom}(q) \subseteq \mathbb{Z}^n$ are finite, we make the following definitions:

- We write $g \cdot p = q$ if $\text{dom}(p) - g = \text{dom}(q)$ and for all $h \in \text{dom}(q)$, $p(h + g) = q(h)$.
- We write $p \cong q$ if there is $h \in \mathbb{Z}^n$ such that $h \cdot p = q$.
- We say that p and q are *compatible* if for all $i \in \text{dom}(p) \cap \text{dom}(q)$, $p(i) = q(i)$; otherwise we say that p and q are *incompatible*.

For $a \leq b \in \mathbb{Z}$ we write $[a, b]$, $[a, b)$, $(a, b]$, (a, b) for the set of integers between a and b , with the obvious conventions as to whether each of the two inequalities is strict or not. For an interval $[a, b]$ in \mathbb{Z} , we will be consistently speaking of its *width*, which is $b - a$, and *size*, which is $b - a + 1$. Intuitively, the size of the interval is the

number of vertices in the 1-dimensional grid graph, and the width is the length of the path from one end to the other.

If a_1, \dots, a_n are positive integers, an (a_1, \dots, a_n) -*pattern* is a map

$$p: [0, a_1] \times \cdots \times [0, a_n] \rightarrow \{0, 1, \dots, \mathfrak{b} - 1\}.$$

We think of $\text{dom}(p)$ naturally as an n -dimensional grid graph, which is in particular a \mathbb{Z}^n -graph. The *width* of an (a_1, \dots, a_n) -pattern is defined as $\max\{a_i - 1 : 1 \leq i \leq n\}$. We say an (a_1, \dots, a_n) -pattern p *occurs* in $x \in \mathfrak{b}^{\mathbb{Z}^n}$ if

$$x \upharpoonright [c_1, c_1 + a_1] \times \cdots \times [c_n, c_n + a_n] \cong p$$

for some $(c_1, \dots, c_n) \in \mathbb{Z}^n$. Note that this is an invariant notion, that is, if p occurs in x then p occurs in any $g \cdot x$. Of course, patterns give the topology on $\mathfrak{b}^{\mathbb{Z}^n}$, the basic open sets being

$$U_p = \{x \in \mathfrak{b}^{\mathbb{Z}^n} : x \upharpoonright [-k, k]^n \cong p\}$$

for some $(2k + 1, \dots, 2k + 1)$ -pattern p .

Recall that a \mathbb{Z}^n -*subshift* Y is a closed, invariant subset of $\mathfrak{b}^{\mathbb{Z}^n}$, where Y inherits a \mathbb{Z}^n -action from $\mathfrak{b}^{\mathbb{Z}^n}$. Some \mathbb{Z}^n -subshifts are said to be of finite type, which we define below.

DEFINITION 2.1.3. A \mathbb{Z}^n -*subshift of finite type* is a $Y \subseteq \mathfrak{b}^{\mathbb{Z}^n}$ for which there is a finite set $\{p_1, \dots, p_k\}$ of patterns such that for any $x \in \mathfrak{b}^{\mathbb{Z}^n}$, $x \in Y$ if and only if none of the patterns p_1, \dots, p_k occur in x .

It is easily verified that given any patterns p_1, \dots, p_k , the set

$$Y = \{x \in \mathfrak{b}^{\mathbb{Z}^n} : \text{none of the patterns } p_1, \dots, p_k \text{ occur in } x\}$$

(which could be empty) defines a \mathbb{Z}^n -subshift of $\mathfrak{b}^{\mathbb{Z}^n}$. Thus any finite collection of patterns determines a \mathbb{Z}^n -subshift of finite type. Naturally we describe the above Y by the sequence $(\mathfrak{b}; p_1, \dots, p_k)$. We refer to the patterns p_1, \dots, p_k as *forbidden patterns*, and the *width* of the \mathbb{Z}^n -subshift of finite type Y is defined as the maximum of the widths of p_1, \dots, p_k .

Note there is no loss of generality in requiring all the forbidden patterns p_i to have the same domain. In this case, if all the p_i have domain $[0, \ell]^n$, we describe the subshift Y by the sequence $(\mathfrak{b}; \ell; p_1, \dots, p_k)$. In this case, the width of Y is $\ell - 1$.

A special case of a width 1 \mathbb{Z}^n -subshift of finite type is what we call an *edge \mathbb{Z}^n -subshift*, which we define below. By an *edge pattern* p we mean a function $p: \text{dom}(p) \rightarrow \mathfrak{b}$ where $\text{dom}(p)$ is the two point set $\{\vec{0}, e_i\}$ for some $1 \leq i \leq n$ (here $\vec{0} = (0, 0, \dots, 0)$). An *edge \mathbb{Z}^n -subshift of finite type* is a \mathbb{Z}^n -subshift of finite type determined by a finite collection of edge patterns.

EXAMPLE. The set of $Y \subseteq \mathfrak{b}^{\mathbb{Z}^n}$ which are proper \mathfrak{b} -colorings of \mathbb{Z}^n is an edge \mathbb{Z}^n -subshift of finite type. In this case, Y is described by $\mathfrak{b}n$ many forbidden edge patterns $p: \{\vec{0}, e_i\} \rightarrow \mathfrak{b}$ where $p(\vec{0}) = p(e_i)$ for $1 \leq i \leq n$.

In the case $n = 1$, the above definition of edge \mathbb{Z} -subshift of finite type is sometimes taken in the literature as the definition of a \mathbb{Z} -subshift of finite type. We note that any \mathbb{Z}^n -subshift of finite type is equivalent to an edge \mathbb{Z}^n -subshift of finite type in the following sense.

FACT 2.1.4. Let $Y \subseteq \mathfrak{b}^{\mathbb{Z}^n}$ be a \mathbb{Z}^n -subshift of finite type. Then there is a \mathfrak{b}' and an edge \mathbb{Z}^n -subshift of finite type $Y' \subseteq \mathfrak{b}'^{\mathbb{Z}^n}$ such that there is an equivariant isomorphism $\varphi: Y \rightarrow Y'$.

PROOF. Let Y be determined by the forbidden patterns p_1, \dots, p_k , and we may assume without loss of generality that all the p_i have domain $[0, \ell)^n$, and thus are $(\ell, \ell, \dots, \ell)$ -patterns. Let $\mathfrak{b}' = \mathfrak{b}^{\ell^n}$. Then \mathfrak{b}' represents also the set of all $(\ell, \ell, \dots, \ell)$ -patterns. The set of forbidden edge patterns $p': \{\vec{0}, e_i\} \rightarrow \mathfrak{b}'$ determining Y' is the p' such that the following holds for $p'(\vec{0})$ and $p'(e_i)$ as $(\ell, \ell, \dots, \ell)$ -patterns:

- either $e_i \cdot p'(\vec{0})$ is incompatible with $p'(e_i)$, or at least one of $p'(\vec{0})$,
- $p'(e_i)$ occurs in the list p_1, \dots, p_k .

For $x \in Y \subseteq \mathfrak{b}^{\mathbb{Z}^n}$, let $\varphi(x) \in \mathfrak{b}'^{\mathbb{Z}^n}$ be defined so that for every $\vec{c} \in \mathbb{Z}^n$, $\varphi(x)(\vec{c})$ is the unique $(\ell, \ell, \dots, \ell)$ -pattern such that $\varphi(x)(\vec{c}) \cong x \upharpoonright (\vec{c} + [0, \ell)^n)$. This is easily an equivariant isomorphism. \square

2.2. The two tiles theorem in one dimension

In this section we prove a two tiles theorem for dimension one as a warm-up for the higher dimensional cases.

For all $n, p, q \geq 1$ with $p, q \geq n$ we define the 2-tiles graph $\Gamma_{n,p,q}^{(1)}$. First consider the 1-dimensional grid graph of size $p + n$, which can be identified as the Cayley graph of \mathbb{Z} restricted to $[0, p + n)$. Decompose the vertex set of the graph into three blocks: $[0, n)$, $[n, p)$, and $[p, p + n)$. Assign a block label $R_\times^{(1)}$ to the vertices of the blocks $[0, n)$ and $[p, p + n)$, and leave the vertices of the block $[n, p)$ unlabeled. (The notation for the label $R_\times^{(1)}$ is chosen to be consistent with the notation for higher dimensions in the next section.) We give two notations to the resulting tile graph $G_{n,p,q}^{(1),1} = T_1(n, p, q)$, the first to be consistent with our notation of higher dimensional cases, and the second for notational simplicity in the rest of this section. Similarly, we define the second tile graph $G_{n,p,q}^{(1),2} = T_2(n, p, q)$, which will be the Cayley graph of \mathbb{Z} restricted to $[0, q + n)$, with block label $R_\times^{(1)}$ assigned to blocks $[0, n)$ and $[q, q + n)$.




FIGURE 6. The 2-tiles graphs $G_{3,7,9}^{(1)}$ and $\Gamma_{3,7,9}^{(1)}$.

Note that each of the $R_\times^{(1)}$ -labeled block is a 1-dimensional grid graph of size n . When taking quotient graphs the corresponding vertices in these $R_\times^{(1)}$ -labeled blocks will be identified. The 2-tiles graph $G_{n,p,q}^{(1)}$ is defined as the disjoint union of $T_1(n,p,q)$ and $T_2(n,p,q)$. To obtain $\Gamma_{n,p,q}^{(1)}$, we take the quotient of $G_{n,p,q}^{(1)}$ by identifying the corresponding vertices in all $R_\times^{(1)}$ -labeled blocks. This finishes the definition of the graph $\Gamma_{n,p,q}^{(1)}$. Note that $\Gamma_{n,p,q}^{(1)}$ is a \mathbb{Z} -graph. Figure 6 illustrates the 2-tiles graphs $G_{n,p,q}^{(1)}$ and $\Gamma_{n,p,q}^{(1)}$ for $n = 3$, $p = 7$, and $q = 9$. In the figure the top graph shows $G_{3,7,9}^{(1)}$, which is a disjoint union of $T_1(3,7,9)$ and $T_2(3,7,9)$, and the bottom graph shows $\Gamma_{3,7,9}^{(1)}$, a graph with 13 vertices.

DEFINITION 2.2.1. Let $Y \subseteq \mathfrak{b}^{\mathbb{Z}}$ be a subshift of finite type described by the sequence $(\mathfrak{b}; p_1, \dots, p_k)$ and let $g: \Gamma_{n,p,q}^{(1)} \rightarrow \mathfrak{b}$ be a map. We say that g respects Y if for all $1 \leq i \leq k$ and for any \mathbb{Z} -homomorphism $\varphi: \text{dom}(p_i) \rightarrow \Gamma_{n,p,q}^{(1)}$, $g \circ \varphi: \text{dom}(p_i) \rightarrow \mathfrak{b}$ is not equal to p_i .

THEOREM 2.2.2 (Two tiles theorem for $F(2^{\mathbb{Z}})$). *Let $Y \subseteq \mathfrak{b}^{\mathbb{Z}}$ be a subshift of finite type. Then the following are equivalent.*

- (1) *There is a continuous, equivariant map $f: F(2^{\mathbb{Z}}) \rightarrow Y$.*
- (2) *There are n, p, q with $n < p, q$, $\gcd(p, q) = 1$, and n greater than or equal to the width of Y , and there is $g: \Gamma_{n,p,q}^{(1)} \rightarrow \mathfrak{b}$ which respects Y .*
- (3) *For all n greater than or equal to the width of Y and all sufficiently large p, q , there is $g: \Gamma_{n,p,q}^{(1)} \rightarrow \mathfrak{b}$ which respects Y .*

PROOF. The implication (3) \Rightarrow (2) is obvious. We show (1) \Rightarrow (3) and (2) \Rightarrow (1).

Suppose first that there is a continuous, equivariant map $f: F(2^{\mathbb{Z}}) \rightarrow Y$. We construct a hyper-aperiodic element $x \in 2^{\mathbb{Z}}$ as follows. Let $y \in 2^{\mathbb{Z}}$ be any hyper-aperiodic element. Enumerate without repetition all tuples (n_i, p_i, q_i, w_i) of positive integers where $p_i, q_i > n_i + 2w_i$. On \mathbb{Z} mark off a collection of disjoint intervals

$$\mathcal{I} = \{I_1^{n_i, p_i, q_i, w_i}, I_2^{n_i, p_i, q_i, w_i} : i \geq 1\}$$

such that for each i , $I_1^{n_i, p_i, q_i, w_i}, I_2^{n_i, p_i, q_i, w_i}$ are of sizes $p_i + n_i + 2w_i, q_i + n_i + 2w_i$ respectively, and the distance between successive intervals in \mathcal{I} goes to infinity. We identify each $I_1^{n_i, p_i, q_i, w_i}$ with an interval of size w_i followed by a copy of $T_1(n_i, p_i, q_i)$ and then followed by another interval of size w_i . That is, it is a copy of $T_1(n_i, p_i, q_i)$ surrounded by two intervals of size w_i . We likewise identify each $I_2^{n_i, p_i, q_i, w_i}$ with a copy of $T_2(n_i, p_i, q_i)$ surrounded by two intervals of size w_i .

For $m \notin \bigcup_i (I_1^{n_i, p_i, q_i, w_i} \cup I_2^{n_i, p_i, q_i, w_i})$, we set $x(m) = y(m)$. Next, assuming $I_1^{n_i, p_i, q_i, w_i} = [a_i, a_i + p_i + n_i + 2w_i)$, we set $x(k_i + m) = y(m)$ for each $m \in [0, p_i)$ and $x(k_i + p_i + m) = y(m)$ for each $m \in [0, n_i + 2w_i)$. We define x on the intervals $I_2^{n_i, p_i, q_i, w_i}$ similarly, using q_i instead of p_i . The definition of x on these intervals is shown in Figure 7.

For each i there are a total of four subintervals of $I_1^{n_i, p_i, q_i, w_i}$ and $I_2^{n_i, p_i, q_i, w_i}$ that are corresponding to the $R_\times^{(1)}$ -labeled blocks of $T_1(n_i, p_i, q_i)$ and $T_2(n_i, p_i, q_i)$. Extending each of these intervals by an interval of size w_i to the left and another interval of size w_i to the right, we obtain four intervals of size $n_i + 2w_i$ within $I_1^{n_i, p_i, q_i, w_i}$ and $I_2^{n_i, p_i, q_i, w_i}$, which happen to be also the initial segment and the final segment of each of $I_1^{n_i, p_i, q_i, w_i}$ and $I_2^{n_i, p_i, q_i, w_i}$ of size $n_i + 2w_i$. The key observation is that x looks the same on all these four intervals.




FIGURE 7. The definition of x on intervals $I_1^{n_i, p_i, q_i, w_i}$ and $I_2^{n_i, p_i, q_i, w_i}$.

Since y is hyper-aperiodic, the enumeration (n_i, p_i, q_i, w_i) is non-repetitive, and the distance between successive intervals in the collection \mathcal{I} tends to infinity, it is straightforward to check that x is also hyper-aperiodic. Let $K = \overline{[x]}$, so $K \subseteq F(2^{\mathbb{Z}})$ is compact (and invariant), and define $f_0 : K \rightarrow \mathfrak{b}$ by $f_0(z) = f(z)(0)$. Since K is compact and f_0 is continuous on K , there is an integer $w > 0$ such that for all $z \in K$, $f_0(z)$ is determined by $z \upharpoonright [-w, w]$.

Now let n be greater than or equal to the width of Y , and $p, q > n + 2w$. Let i be such that $n_i = n$, $p_i = p$, $q_i = q$, and $w_i = w$. Consider the copy of $T_1(n_i, p_i, q_i)$ as a subinterval of $I_1^{n_i, p_i, q_i, w_i}$ and the copy of $T_2(n_i, p_i, q_i)$ as a subinterval of $I_2^{n_i, p_i, q_i, w_i}$. For any $a \in T_1(n_i, p_i, q_i) \cup T_2(n_i, p_i, q_i)$, $f_0(a \cdot x)$ is determined by $x \upharpoonright [a-w, a+w]$. By the key observation, $f_0(a \cdot x)$, as a function of a , looks the same on all four $R_{\times}^{(1)}$ -labeled blocks within $T_1(n_i, p_i, q_i)$ and $T_2(n_i, p_i, q_i)$. Thus, if we define $g' : G_{n_i, p_i, q_i}^{(1)} \rightarrow \mathfrak{b}$ by $g'(a) = f_0(a \cdot x)$, then g' gives rise to a well-defined map $g : \Gamma_{n_i, p_i, q_i}^{(1)} \rightarrow \mathfrak{b}$. Since f is an equivariant map into the subshift Y , it follows that g respects Y . This proves (1) \Rightarrow (3).

Finally, we show that (2) implies (1). Fix n, p, q with $n < p, q$, $\gcd(p, q) = 1$, and n greater than or equal to the width of Y , and assume there is $g : \Gamma_{n, p, q}^{(1)} \rightarrow \mathfrak{b}$ which respects Y . Let N be a positive integer large enough so that every $k \geq N$ is a non-negative integral linear combination of p and q . Arbitrarily fix a point c in the $R_{\times}^{(1)}$ -labeled block in $\Gamma_{n, p, q}^{(1)}$. Then for every $k \geq N$ we can fix a positively-oriented path $(v_0^k, v_1^k, \dots, v_k^k)$ in $\Gamma_{n, p, q}^{(1)}$ (positively-oriented meaning that the directed path consists of only edges with \mathbb{Z} -label $e_1 = +1$) that begins and ends at $v_0^k = v_k^k = c$. From Lemma 0.0.1, fix a clopen set $M_N \subseteq F(2^{\mathbb{Z}})$ which is an N -marker set for $F(2^{\mathbb{Z}})$ (i.e., the conclusions of Lemma 0.0.1 hold). We define $f_0 : F(2^{\mathbb{Z}}) \rightarrow \mathfrak{b}$ as follows. For each pair $x, y \in M_N$ of consecutive points of M_N with x preceding y (meaning there is $k > 0$ with $k \cdot x = y$ but no $\ell \in (0, k)$ with $\ell \cdot x \in M_N$), let $k > 0$ satisfy $k \cdot x = y$ and define $f_0(t \cdot x) = g(v_t^k)$ for each $0 \leq t \leq k$. The map $f_0 : F(2^{\mathbb{Z}}) \rightarrow \mathfrak{b}$ we have just defined is easily seen to be continuous since the set M_N is clopen in $F(2^{\mathbb{Z}})$. Now define $f : F(2^{\mathbb{Z}}) \rightarrow \mathfrak{b}^{\mathbb{Z}}$ by extending f_0 equivariantly: $f(x)(a) = f_0(a \cdot x)$. The map f takes values in Y since g respects Y and since we have the following simple fact: if $T_a, T_b \in \{T_1(n, p, q), T_2(n, p, q)\}$ and T_a, T_b are laid down so that the final $R_{\times}^{(1)}$ -labeled block in T_a coincides with the initial $R_{\times}^{(1)}$ -labeled block in T_b to form a single interval I , then any subinterval J of I of width n lies entirely in the copy of T_a or in the copy of T_b . \square

2.3. Statement of the twelve tiles theorem

In this section we define the 12-tiles graphs $\Gamma_{n,p,q}$, and state the twelve tiles theorem, Theorem 2.3.5. This theorem is stated in terms of continuous, equivariant maps into subshifts of finite type, but in §2.6 we give the alternate formulation in terms of continuous structurings. The proof of Theorem 2.3.5 will be given in the next two sections.

We define the 12-tiles graph $\Gamma_{n,p,q}$ for any fixed n, p, q with $p, q > n$ by following the outline given in the preamble of this chapter. We begin with defining 12 (2-dimensional) rectangular grid graphs with blocks. They will be denoted as $G_{n,p,q}^i$ for $1 \leq i \leq 12$. Since a grid graph is isomorphic to the Cayley graph of \mathbb{Z}^2 restricted to a rectangle $[0, k) \times [0, h)$ for some $k, h > 0$, we only need to specify the “dimensions” (in this case $k \times h$) of the rectangle in order to describe a grid graph. In Table 1 we give a summary of the dimensions of the 12 grid graphs $G_{n,p,q}^i$ for $1 \leq i \leq 12$.

Each of the grid graphs $G_{n,p,q}^i$ will be decomposed into blocks of vertices, and some blocks (in fact all but one for each grid graph) will be assigned block labels. Each block, being a grid graph itself, also has dimensions. We will be working with five block labels $R_\times, R_a, R_b, R_c, R_d$, and their dimensions are as follows:

$$\begin{aligned} R_\times &: && n \times n \\ R_a &: && n \times (p-n) \\ R_b &: && n \times (q-n) \\ R_c &: && (p-n) \times n \\ R_d &: && (q-n) \times n \end{aligned}$$

The specifics of the construction of $G_{n,p,q}^i$ will be given in the next few pages and mostly done by illustrations rather than formal definitions. Here we recall that the purpose of assigning the block labels is so that we can form quotient graphs by identifying all vertices that are in the same relative positions within similarly labeled blocks. We refer to these quotient graphs as obtained *modulo labeled blocks*.

To finish the construction of $\Gamma_{n,p,q}$, we take

$$G_{n,p,q} = \bigcup_{i=1}^{12} G_{n,p,q}^i$$

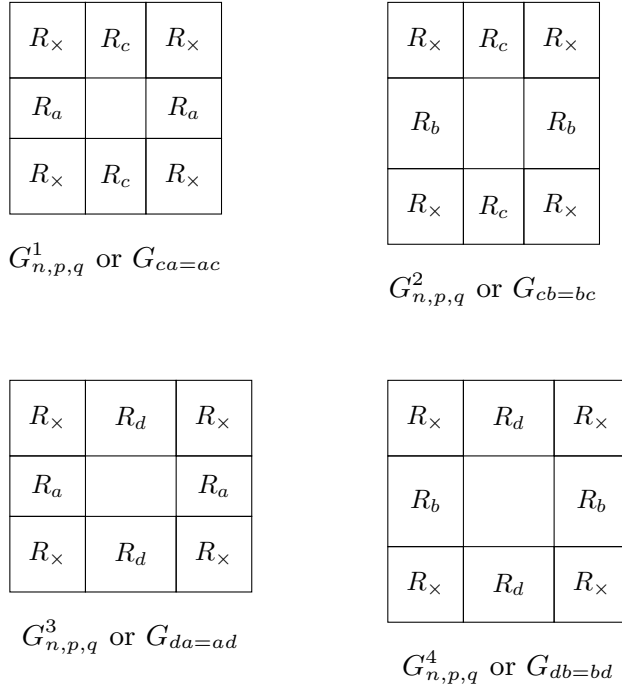
as a disjoint union of the 12 grid graphs with blocks, and obtain $\Gamma_{n,p,q}$ as the quotient graph modulo labeled blocks.

As outlined before, an alternative construction of $\Gamma_{n,p,q}$ is as follows. From each of $G_{n,p,q}^i$, we obtain a tile graph $T_{n,p,q}^i$ as a quotient graph of $G_{n,p,q}^i$ modulo labeled blocks. Then $\Gamma_{n,p,q}$ can be obtained by taking again the disjoint union of the tiles $T_{n,p,q}^i$ for $1 \leq i \leq 12$, and forming the quotient graph of the union again modulo labeled blocks.

The individual tiles $T_{n,p,q}^i$ will be convenient to use in some of the later proofs. An additional forewarning is that the grid graphs $G_{n,p,q}^i$ and the tile graphs $T_{n,p,q}^i$ will also be named in view of the homotopy they realize. These alternative names will be more convenient to use in later proofs. We summarize all these graphs and their names in Table 1.

We now turn to the specifics of the definition of $G_{n,p,q}^i$ for $1 \leq i \leq 12$.

The construction of the first four of the 12 tile graphs are illustrated in Figure 8. We use one example to elaborate the construction, the others being similar. We first consider the grid graph $G_{n,p,q}^1$ of dimensions $(p+n) \times (p+n)$. Assume the vertex set of $G_{n,p,q}^1$ to be $[0, p+n] \times [0, q+n]$. We then decompose the vertex set into nine blocks by dividing the horizontal interval $[0, p+n]$ into three intervals $[0, n)$, $[n, p)$, $[p, p+n)$ and similarly dividing the vertical interval. Then we assign the block label R_\times to the four ‘‘corner’’ blocks $I \times J$ where $I, J \in \{[0, n), [p, p+n)\}$, assign the block label R_a to the two blocks $I \times [n, p)$ where $I \in \{[0, n), [p, p+n)\}$, assign the block label R_c to the two blocks $[n, p) \times J$ where $J \in \{[0, n), [p, p+n)\}$, and leave the remaining ‘‘interior’’ block $[n, p) \times [n, p)$ unlabeled. This completes the definition of $G_{n,p,q}^1$. $G_{n,p,q}^1$ is also named $G_{ca=ac}$ to signify the statement that the two (undirected) paths along the boundaries of the grid graph from the upper-left corner to the lower-right corner are homotopic to each other (in a sense that will be made precise later in the next chapter). Furthermore, the tile graph $T_{n,p,q}^1$ is obtained as a quotient graph of $G_{n,p,q}^1$ modulo labeled blocks. To be precise, for any $i, j \in [0, n)$, the four vertices (i, j) , $(p+i, j)$, $(i, p+j)$ and $(p+i, p+j)$ in $G_{n,p,q}^1$ are identified as one vertex in $T_{n,p,q}^1$. Similarly, for $i \in [0, n)$ and $k \in [n, p)$, the two vertices (i, k) and $(p+i, k)$ are identified, and the two vertices (k, i) and $(k, p+i)$ are identified. $T_{n,p,q}^1$ can be viewed as a grid graph on a 2-torus, and for this reason we refer to it as a *torus tile*. The other three grid graphs and torus tiles are constructed in a similar manner, the only difference being their dimensions.

FIGURE 8. The torus tiles in $\Gamma_{n,p,q}$.

The definition of the remaining 8 grid graphs are illustrated in Figures 9, 10, and 11. Again, tile graphs are obtained as quotient graphs modulo labeled blocks. Each

of the 12 tiles contains a distinct interior block which is unlabeled. This completes the definition of $G_{n,p,q}^i$ and $T_{n,p,q}^i$ for $1 \leq i \leq 12$. The grid graphs are also named according to the homotopy between the two paths along their boundaries from the upper-left corner to the lower-right corner. The tile graphs are correspondingly named, and are classified as “commutativity tiles”, “long horizontal tiles”, and “long vertical tiles” according to their unique features.

R_\times	R_d	R_\times	R_c	R_\times
R_a				R_a
R_\times	R_c	R_\times	R_d	R_\times

$G_{n,p,q}^5$ or $G_{dca=acd}$

R_\times	R_c	R_\times
R_a		R_b
R_\times		R_\times
R_b		R_a
R_\times	R_c	R_\times

$G_{n,p,q}^6$ or $G_{cba=abc}$

R_\times	R_c	R_\times	R_d	R_\times
R_a				R_a
R_\times	R_d	R_\times	R_c	R_\times

$G_{n,p,q}^7$ or $G_{cda=adc}$

R_\times	R_c	R_\times
R_b		R_a
R_\times		R_\times
R_a		R_b
R_\times	R_c	R_\times


$G_{n,p,q}^8$ or $G_{cab=bac}$

FIGURE 9. The commutativity tiles in $\Gamma_{n,p,q}$. Tiles $G_{cab=bac}$ and $G_{cba=abc}$ commute R_a with R_b and tiles $G_{cda=adc}$ and $G_{dca=acd}$ commute R_c with R_d .

We summarize the notation for and the basic features of the 12 tiles and their corresponding grid graphs in Table 1.

This completes the definitions of the tile graphs $G_{n,p,q}^i$ and $T_{n,p,q}^i$ for $1 \leq i \leq 12$ and the 12-tiles graphs $G_{n,p,q}$ and $\Gamma_{n,p,q}$. Each vertex in a quotient graph inherits


q copies of R_c , $q + 1$ copies of R_x



p copies of R_d , $p + 1$ copies of R_x

$$G_{n,p,q}^9 \text{ or } G_{c^q a = ad^p}$$

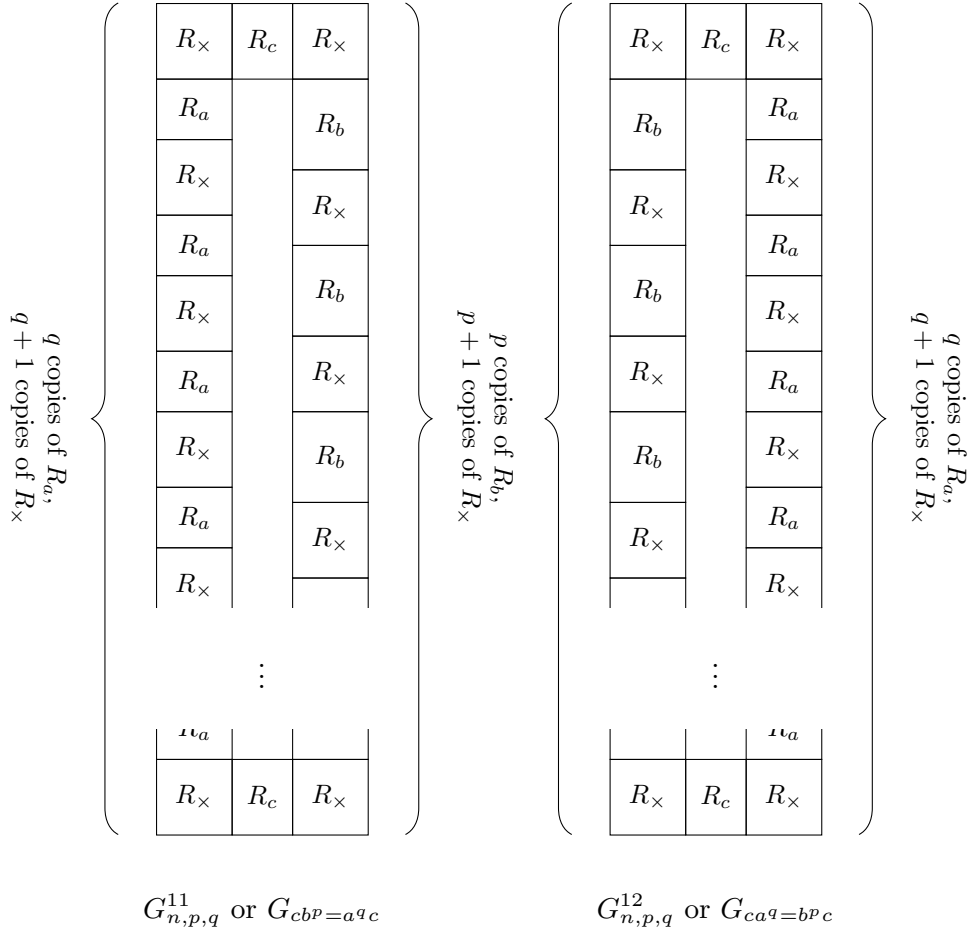
p copies of R_d , $p + 1$ copies of R_x



q copies of R_c , $q + 1$ copies of R_x

$$G_{n,p,q}^{10} \text{ or } G_{d^p a = ac^q}$$

FIGURE 10. The long horizontal tiles in $\Gamma_{n,p,q}$.

FIGURE 11. The long vertical tiles in $\Gamma_{n,p,q}$.

the block label from the rectangular grid graphs. We make the trivial observation that the $\Gamma_{n,p,q}$ contains exactly one copy of a labeled block for each label type, which is isomorphic to the grid graph of appropriate dimensions.

We note that all of the graphs constructed are naturally \mathbb{Z}^2 -graphs. In fact, considering $\Gamma_{n,p,q}$, we note that every directed edge in $\Gamma_{n,p,q}$ is labeled with a unique generator of \mathbb{Z}^2 . For example, suppose (x, y) is an edge in $\Gamma_{n,p,q}$ and has \mathbb{Z}^2 -label $e_1 = (1, 0)$ (the other cases being similar). If x, y are in the same block R labeled with R_x, R_a, R_b, R_c , or R_d , then e_1 is the unique \mathbb{Z}^2 -label for (x, y) , as all edges (x', y') in one of the 12 grid graphs in the same relative position within a copy of R have the same label. If at least one of x, y is in an interior block of a grid graph, then there is a unique pair (x', y') forming an edge in the union of the grid graphs which represent (x, y) in the quotient $\Gamma_{n,p,q}$, so the \mathbb{Z}^2 -label is unique. Finally, if x, y are in different, non-interior, blocks then in any of the 12 grid graphs in which an edge (x', y') appears which represents (x, y) in the quotient, this edge has the same \mathbb{Z}^2 -label. This is because, by inspection, if $x \in R$ then two different generators

Category	Notation	Notation	Corresponding grid graphs		
			Notation	Notation	Dimensions
Torus tiles	$T_{ca=ac}$	$T_{n,p,q}^1$	$G_{ca=ac}$	$G_{n,p,q}^1$	$(p+n) \times (p+n)$
	$T_{cb=bc}$	$T_{n,p,q}^2$	$G_{cb=bc}$	$G_{n,p,q}^2$	$(p+n) \times (q+n)$
	$T_{da=ad}$	$T_{n,p,q}^3$	$G_{da=ad}$	$G_{n,p,q}^3$	$(q+n) \times (p+n)$
	$T_{db=bd}$	$T_{n,p,q}^4$	$G_{db=bd}$	$G_{n,p,q}^4$	$(q+n) \times (q+n)$
Commutativity tiles	$T_{dca=acd}$	$T_{n,p,q}^5$	$G_{dca=acd}$	$G_{n,p,q}^5$	$(p+q+n) \times (p+n)$
	$T_{cba=abc}$	$T_{n,p,q}^6$	$G_{cba=abc}$	$G_{n,p,q}^6$	$(p+n) \times (p+q+n)$
	$T_{cda=adc}$	$T_{n,p,q}^7$	$G_{cda=adc}$	$G_{n,p,q}^7$	$(p+q+n) \times (p+n)$
	$T_{cab=bac}$	$T_{n,p,q}^8$	$G_{cab=bac}$	$G_{n,p,q}^8$	$(p+n) \times (p+q+n)$
Long horizontal tiles	$T_{c^q a=ad^p}$	$T_{n,p,q}^9$	$G_{c^q a=ad^p}$	$G_{n,p,q}^9$	$(pq+n) \times (p+n)$
	$T_{d^p a=ac^q}$	$T_{n,p,q}^{10}$	$G_{d^p a=ac^q}$	$G_{n,p,q}^{10}$	$(pq+n) \times (p+n)$
Long vertical tiles	$T_{cb^p=a^q c}$	$T_{n,p,q}^{11}$	$G_{cb^p=a^q c}$	$G_{n,p,q}^{11}$	$(p+n) \times (pq+n)$
	$T_{ca^q=b^p c}$	$T_{n,p,q}^{12}$	$G_{ca^q=b^p c}$	$G_{n,p,q}^{12}$	$(p+n) \times (pq+n)$

TABLE 1. The notation for the 12 tiles and their corresponding grid graphs.

cannot both move x to the same block R' where $R' \neq R$ either has a different label or is unlabeled.

The next lemma explains the significance of the parameter n in $\Gamma_{n,p,q}$, relating it to a possible width of a subshift of finite type.

LEMMA 2.3.1. *Consider $\Gamma_{n,p,q}$ where $n < p, q$, and let $n_0 \leq n+1$. Let A be the rectangular grid graph on $[0, n_0) \times [0, n_0)$. Suppose φ is a \mathbb{Z}^2 -homomorphism from A to $\Gamma_{n,p,q}$. Then there is a \mathbb{Z}^2 -homomorphism ψ from A to $G_{n,p,q}$ such that $\varphi = \pi \circ \psi$, where $\pi = \pi_{n,p,q}: G_{n,p,q} \rightarrow \Gamma_{n,p,q}$ is the quotient map.*

PROOF. Recall that in the construction of $\Gamma_{n,p,q}$ only the labeled blocks of the tiles $G_{n,p,q}^i$ are identified. Thus if $\text{range}(\varphi)$ contains a vertex of $\Gamma_{n,p,q}$ which is an element of an interior block of some $G_i = G_{n,p,q}^i$ ($1 \leq i \leq 12$), then this i is unique and since $n_0 \leq n+1$ there is a unique \mathbb{Z}^2 -homomorphism ψ from A to G_i such that $\varphi = \pi \circ \psi$. So we can assume that $\text{range}(\varphi)$ contains only vertices not in interior blocks of $\Gamma_{n,p,q}$. If $\text{range}(\varphi)$ consists of vertices all with a single block label, then we may take any G_i where that block appears and define ψ to map into that block of G_i in the obvious manner. The other possibility is that $\text{range}(\varphi)$ contains vertices of exactly two block labels types (by inspection, one can verify that more than two is not possible as $n_0 \leq n+1$ and $n < p, q$), one of which is R_\times . We can then take any G_i where these two block types appear and define ψ to have range in two such blocks of G_i . \square

The following related lemma analyzes \mathbb{Z}^2 -homomorphisms from the Cayley graph of \mathbb{Z}^2 to $\Gamma_{n,p,q}$. It shows that these correspond to consistent “tilings” of \mathbb{Z}^2 by grid graphs each of which is isomorphic to one of the 12 grid graphs $G_{n,p,q}^i$ with labeled blocks.

LEMMA 2.3.2. *Let $n < p, q$ be positive integers. Suppose $\varphi: \mathbb{Z}^2 \rightarrow \Gamma_{n,p,q}$ is a \mathbb{Z}^2 -homomorphism. Then there is a collection $\{C_i\}$ of rectangles in \mathbb{Z}^2 and a map $i \mapsto \alpha(i) \in [1, 12]$ such that the following hold:*

- (1) *For each i , C_i is of the same dimensions as $G_{n,p,q}^{\alpha(i)}$; thus there is a unique isomorphism between $G_{n,p,q}^{\alpha(i)}$ and C_i which makes C_i a grid graph with labeled blocks;*
- (2) $\bigcup_i C_i = \mathbb{Z}^2$;
- (3) *If $x \in \mathbb{Z}^2$ is in $C_i \cap C_j$, then there is a labeled block R of C_i which contains x as a vertex and a labeled block R' of C_j which contains x as a vertex so that $R = R'$ as labeled blocks (meaning they are the same rectangles with the same labels);*
- (4) *For any $n_0 \leq n + 1$ and any rectangle $A \subseteq \mathbb{Z}^2$ of dimensions $n_0 \times n_0$, there is i such that $A \subseteq C_i$.*

PROOF. Suppose $\varphi: \mathbb{Z}^2 \rightarrow \Gamma_{n,p,q}$ is a \mathbb{Z}^2 -homomorphism. For each of the 12 rectangular grid graphs $G_{n,p,q}^i$, $1 \leq i \leq 12$, pick a distinguished vertex in the interior block of $G_{n,p,q}^i$, say the upper-left vertex of the interior block. Let $D \subseteq \mathbb{Z}^2$ be the set of all points $x \in \mathbb{Z}^2$ such that $\varphi(x)$ is a distinguished vertex of one of the $G_{n,p,q}^i$. For each $x \in D$, let $\alpha(x) \in [1, 12]$ such that $\varphi(x)$ is a distinguished vertex of $G_{n,p,q}^{\alpha(x)}$, and let C_x be the rectangular grid subgraph of \mathbb{Z}^2 obtained by placing a copy of $G_{n,p,q}^{\alpha(x)}$ such that x corresponds with the position of the distinguished vertex of $G_{n,p,q}^{\alpha(x)}$. It is easily checked that the collection $\{C_x\}_{x \in D}$ and the map $x \mapsto \alpha(x)$ satisfy (1), (2) and (3), and it follows from Lemma 2.3.1 that it satisfies (4). \square

The following is a converse of the above lemma with a weaker condition.

LEMMA 2.3.3. *Let $n < p, q$ be positive integers. Suppose there is a collection $\{C_i\}$ of rectangles in \mathbb{Z}^2 and a map $i \mapsto \alpha(i) \in [1, 12]$ such that the following hold:*

- (1) *For each i , C_i is of the same dimensions as $G_{n,p,q}^{\alpha(i)}$; thus there is a unique isomorphism between $G_{n,p,q}^{\alpha(i)}$ and C_i which makes C_i a grid graph with labeled blocks;*
- (2) $\bigcup_i C_i = \mathbb{Z}^2$;
- (3) *If $x \in \mathbb{Z}^2$ is in $C_i \cap C_j$, then there is a labeled block R of C_i which contains x as a vertex and a labeled block R' of C_j which contains x as a vertex so that $R = R'$ as labeled blocks;*
- (4') *For any rectangle $A \subseteq \mathbb{Z}^2$ of dimensions 2×2 , there is i such that $A \subseteq C_i$.*

Then there is a \mathbb{Z}^2 -homomorphism $\varphi: \mathbb{Z}^2 \rightarrow \Gamma_{n,p,q}$.

PROOF. Suppose $\{C_i\}$ and the map $i \mapsto \alpha(i) \in [1, 12]$ are given and satisfy (1), (2), (3) and (4'). Let $\pi: G_{n,p,q} \rightarrow \Gamma_{n,p,q}$ be the quotient map. Given any $x \in \mathbb{Z}^2$, fix any i with $x \in C_i$, fix the unique isomorphism $\psi: C_i \rightarrow G_{n,p,q}^{\alpha(i)}$ as grid graphs with labeled blocks, and define $\varphi(x) = \pi \circ \psi(x)$. By (1), (2) and (3), φ is well-defined. To see φ is a \mathbb{Z}^2 -homomorphism, consider say $x, y = e_1 \cdot x \in \mathbb{Z}^2$. From

(4'), there is an i such that $\{x, y\} \subseteq C_i$. If we use this same C_i for the definition of $\varphi(x)$ and $\varphi(y)$, it is then clear that the edge $(\varphi(x), \varphi(y))$ has label e_1 in $\Gamma_{n,p,q}$. \square

DEFINITION 2.3.4. Let Y be a subshift of finite type described by $(\mathfrak{b}; p_1, \dots, p_k)$. Let $g: \Gamma_{n,p,q} \rightarrow \mathfrak{b}$. We say g *respects* Y if for any $1 \leq i \leq k$ and for any \mathbb{Z}^2 -homomorphism φ from $\text{dom}(p_i)$ to $\Gamma_{n,p,q}$, $g \circ \varphi: \text{dom}(p_i) \rightarrow \mathfrak{b}$ is not equal to p_i .

From Lemma 2.3.1 it follows that if n is greater than or equal to the width of Y , then $g: \Gamma_{n,p,q} \rightarrow \mathfrak{b}$ respects Y if and only if for any forbidden pattern p_i defining Y and for any \mathbb{Z}^2 -homomorphism $\psi: \text{dom}(p_i) \rightarrow G_{n,p,q}$, $g \circ \pi \circ \psi: A_i \rightarrow \mathfrak{b}$ is not equal to p_i , where $\pi = \pi_{n,p,q}: G_{n,p,q} \rightarrow \Gamma_{n,p,q}$ is the quotient map. Thus, for n greater than or equal to the width of Y , the condition that $g: \Gamma_{n,p,q} \rightarrow \mathfrak{b}$ respects Y is a property of $\tilde{g}: G_{n,p,q} \rightarrow \mathfrak{b}$, where $\tilde{g} = g \circ \pi$ is the map on $G_{n,p,q}$ that g induces.

We are now ready to state our main theorem.

THEOREM 2.3.5 (Twelve tiles theorem). *Let $Y \subseteq \mathfrak{b}^{\mathbb{Z}^2}$ be a subshift of finite type. Then the following are equivalent.*

- (1) *There is a continuous, equivariant map $f: F(2^{\mathbb{Z}^2}) \rightarrow Y$.*
- (2) *There are positive integers n, p, q with $n < p, q$, $\text{gcd}(p, q) = 1$, and n greater than or equal to the width of Y , and there is $g: \Gamma_{n,p,q} \rightarrow \mathfrak{b}$ which respects Y .*
- (3) *For all n greater than or equal to the width of Y and for all sufficiently large p, q , there is $g: \Gamma_{n,p,q} \rightarrow \mathfrak{b}$ which respects Y .*

The implication (3) \Rightarrow (2) is obvious. We will prove (1) \Rightarrow (3) in §2.4 and (2) \Rightarrow (1) in §2.5.

2.4. Proof of the negative direction

In this section we prove the negative direction of the twelve tiles theorem. That is, we prove the implication (1) \Rightarrow (3) for Theorem 2.3.5 (in §2.6 we give the corresponding result, Theorem 2.6.8, in terms of continuous structurings). We refer to this implication as the negative direction as it is often used to show that various continuous structurings of $F(2^{\mathbb{Z}^2})$ do not exist.

The idea of the proof is to construct a hyper-aperiodic element $x \in F(2^{\mathbb{Z}^2})$ which carries the structures $\Gamma_{n,p,q}$ embedded into it for all suitable n, p, q in a way that is similar to the proof of the two tiles theorem. We use this element for both Theorems 2.3.5 and 2.6.8. To construct x , we will use a fixed hyper-aperiodic element $y \in F(2^{\mathbb{Z}^2})$ which we now fix. The element x will be constructed by copying various regions in y to various places in x .

Let w be a positive integer and fix n, p, q with $p, q > 2n + 4w$. The integer w will be the size of a “buffer” region we will use to “pad” the construction of $G_{n,p,q}$ to produce a modified version $G_{w,n,p,q}$. Just as $G_{n,p,q}$ is the disjoint union of the 12 rectangular subgraphs $G_{n,p,q}^i$ ($1 \leq i \leq 12$), $G_{w,n,p,q}$ will be the disjoint union of the rectangular subgraphs $G_{w,n,p,q}^i$. The element x will then be built up from y and the $G_{w,n,p,q}^i$.

Consider the grid graph $G_{n,p,q}^1$ for the first torus tile. Recall this is a rectangular grid graph of dimensions $(p+n) \times (p+n)$, and is subdivided into blocks with labels R_\times , R_a , and R_c , together with an unlabeled interior block as in Figure 8. $G_{w,n,p,q}^1$

will be a grid graph of dimensions $(p + n + 2w) \times (p + n + 2w)$ which is similarly divided into blocks with new labels R'_x , R'_a , and R'_c . The definition of $G_{w,n,p,q}^1$ is shown in Figure 12. The solid lines indicate the new rectangular sub-graphs R'_x , R'_a and R'_c . The dashed lines indicate the original $G_{n,p,q}^1$ along with the original R_x , R_a , and R_c . Note that the R'_x block has dimensions $(n + 2w) \times (n + 2w)$, the R'_a block has dimensions $(n + 2w) \times (p - n - 2w)$, and the R'_c block has dimensions $(p - n - 2w) \times (n + 2w)$. The other three padded grid graphs for torus tiles, $G_{w,n,p,q}^2$, $G_{w,n,p,q}^3$, $G_{w,n,p,q}^4$, are defined in the same manner.




FIGURE 12. Construction of the padded grid graph $G_{w,n,p,q}^1$. The boundary of $G_{n,p,q}$ has been padded by w .

The remaining $G_{w,n,p,q}^i$ are defined similarly. In each case, the block R'_x is expanded in each dimension by $2w$, while the other blocks (R'_a , R'_b , R'_c , R'_d) are expanded in one direction by $2w$ and contracted in the other direction by $2w$. Figure 13 shows the construction for the grid graph $G_{w,n,p,q}^5$, which corresponds to the tile $T_{dca=acd}$. This completes the description of the rectangular grid graphs $G_{w,n,p,q}^i$.

To construct the hyper-aperiodic element x , let (w_k, n_k, p_k, q_k) enumerate, without repetition, all 4-tuples of positive integers such that $p_k, q_k > 2n_k + 4w_k$. Fix a doubly-indexed sequence (ℓ_k^i) of integers for $1 \leq i \leq 12$ and $k \in \mathbb{Z}^+$ such that for any k, i, j ,

- (i) $\ell_k^i < \ell_{k+1}^j$,
- (ii) $\ell_k^i < \ell_k^j$ if and only if $i < j$,
- (iii) $\ell_{k+1}^1 - \ell_k^{12} > k + p_k q_k + n_k + 2w_k$, and
- (iv) $\ell_k^{i+1} - \ell_k^i > k + p_k q_k + n_k + 2w_k$.

By (i) and (ii), the ℓ_k^i are in lexicographic order of the pairs (k, i) . To continue our construction, we place a copy of the rectangular grid graph G_{w_k, n_k, p_k, q_k}^i in \mathbb{Z}^2 with the lower-left corner of G_{w_k, n_k, p_k, q_k}^i at the point $(\ell_k^i, 0)$. By (iii) and (iv), the distance between any point of the copy of G_{w_k, n_k, p_k, q_k}^i and any other copy is at




FIGURE 13. Constructing the rectangular grid graph $G^5_{w,n,p,q}$ corresponding to $T_{dca=acd}$. The solid lines represent the blocks R'_x etc. of $G^5_{w,n,p,q}$, while the dashed lines indicate the original blocks of $G^5_{n,p,q}$.

least k . Formally, when we “place a copy” of $G^i_{w_k,n_k,p_k,q_k}$ in \mathbb{Z}^2 we mean we have identified a rectangular grid subgraph of \mathbb{Z}^2 and have partitioned its elements into labeled and unlabeled blocks isomorphically to $G^i_{w_k,n_k,p_k,q_k}$. Figure 14 illustrates the placement of the $G^i_{w_k,n_k,p_k,q_k}$ and the construction of x .




FIGURE 14. Constructing the hyper-aperiodic element x

We now define x . For $(a, b) \in \mathbb{Z}^2$ not in any copy of $G^i_{w_k,n_k,p_k,q_k}$, we define $x(a, b) = y(a, b)$. Suppose (a, b) is in a copy of $G^i_{w_k,n_k,p_k,q_k}$ (note that the i and k are unique). Say (a, b) is in a labeled block R' , and the lower-left corner point of block R' has coordinates (α, β) . Write $(a, b) = (\alpha, \beta) + (a', b')$. We then set $x(a, b) = y(a', b')$. In other words, we fill-in the portion of x in the block R' by a partial copy of the first quadrant of y . Finally, if (a, b) is in an unlabeled block, then set $x(a, b) = y(a, b)$.

This completes the definition of $x \in 2^{\mathbb{Z}^2}$. We next show that x is hyper-aperiodic.

LEMMA 2.4.1. *The element $x \in 2^{\mathbb{Z}^2}$ is hyper-aperiodic.*

PROOF. Let $s \neq (0, 0)$ be a non-identity element of the group \mathbb{Z}^2 . Since y is a hyper-aperiodic element, there is a finite $T_0 \subseteq \mathbb{Z}^2$ such that for any $g \in \mathbb{Z}^2$ there is a $t \in T_0$ such that $y(g + t) \neq y(g + s + t)$. Let $T_1 = \{(0, 0), s\} \cup T_0 \cup (s + T_0)$ and $M \in \mathbb{Z}^+$ be such that $T_1 \subseteq [-M, M] \times [-M, M]$.

Let S be the set of all $g \in \mathbb{Z}^2$ such that either $\{g, g+s\} \cup (g+T_0) \cup (g+s+T_0)$ does not intersect any copy of G_{w_k, n_k, p_k, q_k}^i in the construction of x or there is a labeled or unlabeled block R' of some copy of G_{w_k, n_k, p_k, q_k}^i such that $\{g, g+s\} \cup (g+T_0) \cup (g+s+T_0) \subseteq R'$. Intuitively, S is the set of $g \in \mathbb{Z}^2$ such that some $t \in T_0$ will witness the hyper-aperiodicity of x at g for s , i.e., there is $t \in T_0$ such that $x(g+t) \neq x(g+s+t)$.

To verify that x satisfies the hyper-aperiodicity condition, it suffices to show that S is *syndetic*, i.e., there is a finite set $T \subseteq \mathbb{Z}^2$ such that for any $g \in \mathbb{Z}^2$ there is $t \in T$ so that $g+t \in S$.

Let S' be the set of all $g \in \mathbb{Z}^2$ such that $g + [-M, M] \times [-M, M]$ does not intersect any copy of G_{w_k, n_k, p_k, q_k}^i . Intuitively, S' is the set of all $g \in \mathbb{Z}^2$ that are of a distance M away from any copy of G_{w_k, n_k, p_k, q_k}^i . We have $S' \subseteq S$.

Let T' be the set of all $g \in \mathbb{Z}^2$ such that g is within distance M to a copy of G_{w_k, n_k, p_k, q_k}^i for $k \leq 2M$. T' is finite. Assume that $T' \subseteq I \times \mathbb{Z}$ where I is finite. Let

$$H = \max\{p_k q_k + n_k + 2w_k : k \leq 2M\}.$$

Then H is the maximum height of the grid graphs G_{w_k, n_k, p_k, q_k}^i for $k \leq 2M$. We have that $T' \subseteq I \times [-M, H+M]$. Thus for any $g \in T'$, $g + (0, H+2M) \in S' \subseteq S$.

Suppose $g \notin S' \cup T'$, i.e., g is within distance M to a copy of G_{w_k, n_k, p_k, q_k}^i for some $k > 2M$. Note that this k is unique, since the distance between this copy of G_{w_k, n_k, p_k, q_k}^i and the neighboring copies is greater than $2M$. We consider several cases. Let $p'_k = p_k - n_k - 2w_k$, $q'_k = q_k - n_k - 2w_k$. Note that $p'_k, q'_k \geq n_k + 2w_k$.

Case 1: $n_k + 2w_k > 2M$. Since $n_k + 2w_k$ is the smallest dimension of any block within the copy of G_{w_k, n_k, p_k, q_k}^i , any interior point of a block that is of distance M away from its boundary is a point of S . It follows that g is within distance $2M$ of such a point, i.e., g is within distance $2M$ of a point of S .

Case 2: $n_k + 2w_k \leq 2M < \min\{p'_k, q'_k\}$. In this case all the labeled blocks within the copy of G_{w_k, n_k, p_k, q_k}^i has one of its dimensions $\leq 2M$, but the interior unlabeled block has dimensions $> 2M$, and any interior point of this block that is of distance M away from its boundary is a point of S . It follows that g is within distance $4M$ of such a point.

Case 3: $\max\{p'_k, q'_k\} \leq 2M$. In this case, if G_{w_k, n_k, p_k, q_k}^i is not correspondent to a long vertical tile, g is within a vertical distance of $4M$ to a point of $S' \subseteq S$. If G_{w_k, n_k, p_k, q_k}^i is correspondent to a long vertical tile, then g is within a horizontal distance of $4M$ to a point of S' .

Case 4: $\min\{p'_k, q'_k\} \leq 2M < \max\{p'_k, q'_k\}$. This is broken up into two subcases, depending on either $p'_k < q'_k$ or $q'_k < p'_k$.

Subcase 4.1: $p'_k \leq 2M < q'_k$. In this subcase, if G_{w_k, n_k, p_k, q_k}^i is not correspondent to the tile $T_{db=bd}$, then similar to Case 3, g is within either vertical or horizontal distance of $4M$ to a point of S' . If G_{w_k, n_k, p_k, q_k}^i is correspondent to the tile $T_{db=bd}$, then similar to Case 2, g is within distance $4M$ of a point of S which is an interior point of the unlabeled interior block.

Subcase 4.2: $q'_k \leq 2M < p'_k$. This is opposite to Subcase 4.1. In this subcase, if G_{w_k, n_k, p_k, q_k}^i is not correspondent to one of the tiles $T_{db=bd}$, $T_{cb=bc}$, or $T_{da=ad}$, then similar to Case 2, g is within distance $4M$ of a point of S which is an interior point of the unlabeled interior block. If G_{w_k, n_k, p_k, q_k}^i is correspondent to one of these tiles, then similar to Case 3, g is within distance $4M$ to a point of S' .

This completes the proof that S is syndetic, and hence x is hyper-aperiodic. \square

PROOF OF (1) \Rightarrow (3) OF THEOREM 2.3.5. Let $Y \subseteq \mathfrak{b}^{\mathbb{Z}^2}$ be a subshift of finite type described by $(\mathfrak{b}; p_1, \dots, p_m)$. Let $f: F(2^{\mathbb{Z}^2}) \rightarrow Y$ be a continuous, equivariant map. Fix n greater or equal to the width of Y . We need to show that for all large enough p, q there is a $g: \Gamma_{n,p,q} \rightarrow \mathfrak{b}$ which respects Y .

Let $x \in F(2^{\mathbb{Z}^2})$ be the hyper-aperiodic element produced in Lemma 2.4.1. Let $K = \overline{[x]}$. Since x is hyper-aperiodic, $K \subseteq F(2^{\mathbb{Z}^2})$, and K is compact. Thus, f is defined on K . Define $f_0: K \rightarrow \mathfrak{b}$ by $f_0(z) = f(z)(0,0)$. Since f_0 is continuous and K is compact, there is a $w \in \mathbb{Z}^+$ such that for any $z \in K$, the value of $f_0(z)$ is determined by $z \upharpoonright [-w, w]^2$. In particular, this holds for all z in the orbit of x . So, for any $(a, b) \in \mathbb{Z}^2$, $f_0((a, b) \cdot x)$ is determined by the values of x on the $(2w+1) \times (2w+1)$ rectangle centered about (a, b) .

Now consider arbitrary $p, q > 2n + 4w$. Fix a k such that $p_k = p$, $q_k = q$, $n_k = n$, and $w_k = w$. Consider the grid graphs $G_{w_k, n_k, p_k, q_k}^i = G_{w, n, p, q}^i$ along with the “unpadded” graph $G_{n, p, q}^i$ induced by $G_{w_k, n, p, q}^i$ (the region within the dotted lines in Figure 15). We first define $g': G_{n, p, q} \rightarrow \mathfrak{b}$ as follows. If $(a, b) \in \mathbb{Z}^2$ is in a copy of $G_{n, p, q}^i$ within a copy of $G_{w_k, n, p, q}^i$, we let $g'(a, b) = f_0((a, b) \cdot x)$. From the construction of x , it is easy to verify the following claims by inspection:

- If (a, b) and (a', b') are in a copy of $G_{n, p, q}^i$ within a copy of $G_{w_k, n, p, q}^i$, and if (a, b) , (a', b') are in the same positions in labeled blocks with the same labels, then $f_0((a, b) \cdot x) = f_0((a', b') \cdot x)$.
- If (a, b) is in a copy of $G_{n, p, q}^i$ within a copy of $G_{w_k, n, p, q}^i$ and (a', b') is in a copy of $G_{n, p, q}^{i'}$ within a copy of $G_{w_k, n, p, q}^{i'}$, and if (a, b) , (a', b') are in the same positions in labeled blocks with the same labels, then $f_0((a, b) \cdot x) = f_0((a', b') \cdot x)$.

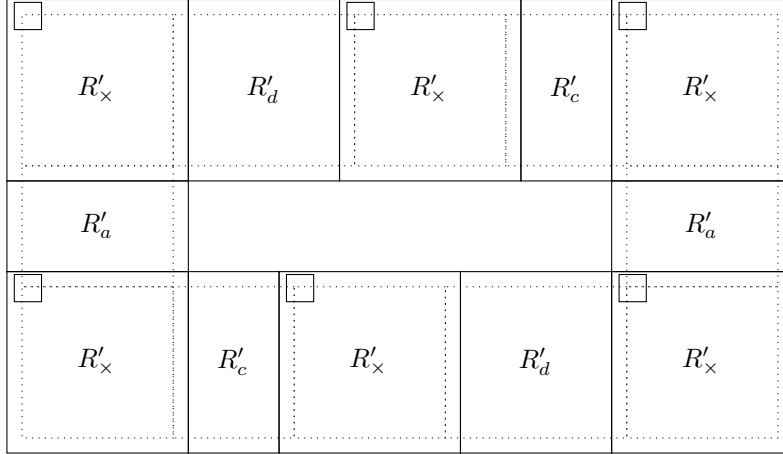


FIGURE 15. Verifying that if (a, b) and (a', b') are in the same positions in a labeled block of $G_{w_k, n, p, q}$, then $f((a, b) \cdot x) = f((a', b') \cdot x)$.

In fact, suppose, as an example, that (a, b) , (a', b') are both in blocks with the label R'_x . In this case, each of the R'_x blocks is in a strictly larger block R'_x of $G_{w_k, n, p, q}$ (see Figure 15). Since $w_k = w$, the $(2w+1) \times (2w+1)$ rectangle centered at (a, b) lies entirely within the block R'_x , and likewise for (a', b') . Since x is identical

on the various R'_\times blocks, and the $(2w+1) \times (2w+1)$ rectangle determines the value of f_0 , it follows that $f_0((a,b) \cdot x) = f_0((a',b') \cdot x)$, and so $g'(a,b) = g'(a',b')$.

Suppose next that $(a,b), (a',b')$ are in blocks of $G_{n,p,q}$ with label R_a (the other cases are identical). In this case, the R_a block containing (a,b) is not contained in the corresponding R'_a block of $G_{w_k,n,p,q}$, as the R'_a block is smaller in one direction by $2w_k$ than the R_a block. However, every time an R_a block appears in a $G_{n,p,q}^i$, the block immediately above and below have the label R'_\times . Thus, for an (a,b) in an R_a block, the $(2w+1) \times (2w+1)$ rectangle about (a,b) is contained in the union of two adjacent blocks with labels R'_\times and R'_a . Since x is identical on the various R'_\times blocks, and also on the various R'_a blocks, it again follows that $f_0((a,b) \cdot x) = f_0((a',b') \cdot x)$, and so $g'(a,b) = g'(a',b')$.

In summary, if (a,b) and (a',b') are vertices in $G_{n,p,q}$ such that they are in the same positions in blocks with the same label, then $g'(a,b) = g'(a',b')$. It follows that g' induces a map $g: \Gamma_{n,p,q} \rightarrow \mathbf{b}$. To see that g respects Y , consider a forbidden pattern p_h where $1 \leq h \leq m$, and assume $\text{dom}(p_h) = [0,a) \times [0,b)$. By our assumption $a,b \leq n+1$. Let $\varphi: \text{dom}(p_h) \rightarrow \Gamma_{n,p,q}$ be a \mathbb{Z}^2 -homomorphism. From Lemma 2.3.1, there is a \mathbb{Z}^2 -homomorphism $\psi: \text{dom}(p_h) \rightarrow G_{n,p,q}^i$ for some $1 \leq i \leq 12$ such that $g \circ \varphi = g' \circ \psi$. We need to see that $g \circ \varphi = g' \circ \psi$ is not equal to p_h . From the definition of g' , we have that

$$g' \circ \psi(a',b') = f_0(\psi(a',b') \cdot x) = f(x)(\psi(a',b'))$$

for all $a' \in [0,a)$ and $b' \in [0,b)$. Since f maps into Y , $f(x) \upharpoonright \psi(\text{dom}(p_h))$ is not equal to p_h , and hence $g' \circ \psi$ is not equal to p_h . \square

2.5. Proof of the positive direction

In this section we show the (2) \Rightarrow (1) direction of Theorem 2.3.5. The proof will follow easily from the main technical result below. In the following the finite graph $\Gamma_{n,p,q}$ is viewed as having the discrete topology.

THEOREM 2.5.1. *Let $n < p, q$ with $\text{gcd}(p, q) = 1$. Then there is a continuous \mathbb{Z}^2 -homomorphism $\varphi: F(2^{\mathbb{Z}^2}) \rightarrow \Gamma_{n,p,q}$.*

PROOF. Fix $n < p, q$ with $\text{gcd}(p, q) = 1$. We will need the following lemma which is a simple special case of the ‘‘orthogonal marker construction’’ of [11] (see the proof of Theorem 3.1 in [11]). For $s = (u, v) \in \mathbb{Z}^2$ let $\|s\|_1 = |u| + |v|$ be the ℓ_1 -norm. For $x, y \in F(2^{\mathbb{Z}^2})$ which are equivalent, define $\rho_1(x, y) = \|s\|_1$, where $s \in \mathbb{Z}^2$ is the unique element such that $s \cdot x = y$.

LEMMA 2.5.2. *Let $d > 1$ be an integer. Then there is a clopen subequivalence relation E_d of $F(2^{\mathbb{Z}^2})$ satisfying:*

- (1) *Each equivalence class of E_d is rectangular, that is, isomorphic to a grid subgraph of \mathbb{Z}^2 .*
- (2) *If x, y are $F(2^{\mathbb{Z}^2})$ equivalent and both are corner points of their respective E_d classes, then either $\rho_1(x, y) \leq 1$ or $\rho_1(x, y) > d$.*

Fix d large compared to p, q , say $d > p^{10}q^{10}$, and fix the clopen subequivalence relation E_d as in Lemma 2.5.2. We say $x \in F(2^{\mathbb{Z}^2})$ is a *corner point* if it is a corner point of the rectangular grid graph isomorphic to its E_d class. Since the rectangular E_d classes partition each $F(2^{\mathbb{Z}^2})$ class, it is easy to see that the corner points occur in groups of 2 points, such that the points in a group are within ρ_1

distance 1 of each other. Consider the set $C \subseteq F(2^{\mathbb{Z}^2})$ of *canonical corner points*, which are the lower-left corner points of the rectangular grid-graph isomorphic to its E_d class. Each corner point occurs in a group of 4 points, which we call a *corner group*, consisting of a 2×2 square, of which two points are corner points and the other two are boundary points of an E_d class. Figure 16 illustrates the possible arrangements involving corner points, with the canonical corner points and corner groups shown.




FIGURE 16. Arrangements involving corner points. The canonical corner points are shown in black.

For each $x \in F(2^{\mathbb{Z}^2})$, consider a rectangular E_d class $R \subseteq [x]$. Say R is an $a \times b$ rectangle with lower-left corner at $y \in [x]$. By the *enlargement* R' of R we mean the subset of $[x]$ isomorphic to the rectangular grid graph of dimensions $(a + 1) \times (b + 1)$ with lower-left corner point at y . Thus, R' is obtained from R by extending its right and top edges by one. The enlarged rectangles are no longer disjoint, but have overlapping boundaries.

DEFINITION 2.5.3. The *scaffolding* $S_x \subseteq [x]$, for $x \in F(2^{\mathbb{Z}^2})$, is the set of points in the boundaries of the enlarged E_d classes in $[x]$.

By considering the cases shown in Figure 16, it is easy to see that the scaffolding can also be obtained in the following manner. From each canonical corner-point $y \in [x]$, draw a horizontal line to the right, and a vertical line in the upwards direction, up to and including the first point which is not E_d equivalent to y . The scaffolding is the set of points in $[x]$ which lie on these lines. By definition, every canonical corner point is in the scaffolding. The Cayley graph structure on \mathbb{Z}^2 induces a graph on the scaffolding. Each point of the scaffolding has degree 2 or 3, with the points of degree 3 being the upper-right points of the corner groups. The connected components of $[x] - S_x$ are the E_d classes minus the left column and bottom row of each E_d class. Figure 17 gives a local view of the scaffolding in relation to the corner groups.




FIGURE 17. A local view of the scaffolding S_x . Points in S_x are in orange.

According to Lemmas 2.3.2 and 2.3.3, defining a continuous \mathbb{Z}^2 -homomorphism $\varphi: F(2^{\mathbb{Z}^2}) \rightarrow \Gamma_{n,p,q}$ is the equivalent to describing a tiling $\{C_k\}$ of each equivalence class $[x]$ by the grid graphs $G_{n,p,q}^i$ (as in Lemma 2.3.3) which is clopen in the sense that for each $i \in \{1, \dots, 12\}$, the set of $x \in F(2^{\mathbb{Z}^2})$ which are lower-left corners of a C_k isomorphic to $G_{n,p,q}^i$ is clopen in $F(2^{\mathbb{Z}^2})$.

Let $x \in F(2^{\mathbb{Z}^2})$ and we define the tiling of $[x]$ by the grid graphs $G_{n,p,q}^i$. First we use the scaffolding S_x as follows. For each point $y \in [x]$ which is of degree 3 in S_x , we put down a copy of the R_\times block “centered” at y , meaning we place the lower-left corner of the R_\times block at $(-\frac{n}{2}, -\frac{n}{2}) \cdot y$ if n is even and at $(-\frac{n-1}{2}, -\frac{n-1}{2}) \cdot y$ if n is odd.

Second, we tile the scaffolding between the R_\times copies around the degree three points. Suppose y and z in $[x]$ are degree three points of S_x such that there is a horizontal or vertical segment of S_x connecting y and z with no degree three points of S_x on this segment other than y and z . Suppose that the segment of S_x connecting y and z is vertical. We then tile the vertical column between the two copies of R_\times around y and z by alternating copies of an R_\times tile with a R_a or R_b tile. Recall that $\rho(y, z) > p^{10}q^{10}$. Since $\gcd(p, q) = 1$ and

$$\begin{aligned} \text{height}(R_\times) + \text{height}(R_a) &= p, \\ \text{height}(R_\times) + \text{height}(R_b) &= q, \end{aligned}$$

it is possible to tile the vertical column in this manner. Furthermore, we require that the copies of R_b are sparse in the following sense: the distance between any two copies of R_b block is at least p^5q^5 , and the distance from any R_b copy to the starting and ending R_\times blocks is also at least p^5q^5 . This is possible since at most p many of the R_b blocks are needed for the tiling, because a consideration of sizes gives that every p many R_b blocks can be replaced by q many R_a blocks. Similarly, if the segment of S_x connecting y and z is horizontal, we tile the horizontal regions between them in the same manner, using R_\times alternating with R_c or R_d blocks, and have the same sparsity condition on the R_d . Figure 18 illustrates this step of the construction.

Let $S'_x \subseteq [x]$ denote the points in $[x]$ which have been covered by tiles at this point in the construction. Note that $[x] - S'_x$ is a collection of pairwise disjoint rectangular grid graphs, surrounded by alternating copies of R_\times and R_a/R_b on the left and right edges, and alternating R_\times with R_c/R_d on the top and bottom edges. The corners of the boundary tiling are R_\times copies, and the edges satisfy the sparsity conditions on R_b and R_d . The remainder of the construction is entirely finitary. We show that any such rectangular region (including the boundary tiles) can be tiled with the grid graphs $G_{n,p,q}^i$ in a manner consistent with the boundary tiles.

Let \mathcal{M} denote a finite rectangular grid graph whose boundary has been tiled with sub-blocks R_\times and $R_* \in \{R_a, R_b, R_c, R_d\}$ as described above. We will describe three tiling algorithms we will use in combination to fill the interior of \mathcal{M} . Each algorithm is used to fill a rectangular grid graph with certain part of the boundary already specified.

The first algorithm, which we call algorithm (I), deals with a rectangular grid graph \mathcal{N} with its boundary already tiled with alternating R_\times blocks and R_* blocks, where $R_* \in \{R_a, R_b, R_c, R_d\}$, such that only one kind of labeled blocks appear in the left and right columns and only one kind of blocks appear in the top and bottom rows. For instance, suppose \mathcal{N} has a boundary devoid of copies of R_b or R_d . In




FIGURE 18. Edge boundaries for marker regions. R_\times blocks alternate with R_* blocks with $R_* \in \{R_a, R_b, R_c, R_d\}$. Instances of R_b and R_d are labeled. This figure differs from the actual construction in that instances of R_b and R_d in the actual construction are much less common than depicted: we require that no two copies of R_b or R_d blocks be within distance p^5q^5 of one another.

this case we simply fill the entire grid graph with overlapping copies of tile $G_{ca=ac}$. This is depicted in Figure 19. The other instances of algorithm (I) are similar, with the $G_{ca=ac}$ tiles replaced by the appropriate tiles $G_{da=ad}$, $G_{cb=bc}$, or $G_{db=bd}$.

$T_{ca=ac}$	$T_{ca=ac}$	$T_{ca=ac}$	$T_{ca=ac}$	$T_{ca=ac}$	$T_{ca=ac}$		
$T_{ca=ac}$	$T_{ca=ac}$	$T_{ca=ac}$	$T_{ca=ac}$	$T_{ca=ac}$	$T_{ca=ac}$		
$T_{ca=ac}$	$T_{ca=ac}$	$T_{ca=ac}$	$T_{ca=ac}$	$T_{ca=ac}$	$T_{ca=ac}$		

FIGURE 19. An instance of algorithm (I) with boundaries consisting of only R_\times , R_a and R_c blocks.

The next two algorithms deal with rectangular grid graphs where at least one of the R_* blocks are missing. For the sake of definiteness, let us consider a rectangular grid graph \mathcal{N} where R_b blocks are missing, i.e., the left and right boundaries are filled with R_\times blocks alternating with R_a blocks. Moreover, we assume that only the top boundary is specified. In this case, algorithm (I) can still be used to achieve a propagation of the arrangement of the top boundary down through the grid graph. Figure 20 illustrates this situation.

It is obvious that a similar algorithm can achieve horizontal propagation of the left boundary tiling if one of R_c or R_d is missing from the top and bottom boundary.

$T_{ca=ac}$	$T_{ca=ac}$	$T_{da=ad}$	$T_{ca=ac}$	$T_{ca=ac}$
$T_{ca=ac}$	$T_{ca=ac}$	$T_{da=ad}$	$T_{ca=ac}$	$T_{ca=ac}$
$T_{ca=ac}$	$T_{ca=ac}$	$T_{da=ad}$	$T_{ca=ac}$	$T_{ca=ac}$

FIGURE 20. Using only algorithm (I), we achieve a downward propagation of an instance of R_d using tile $G_{da=ad}$ through a region mostly tiled with $G_{ca=ac}$.

Our next algorithm, algorithm (II), seeks to propagate the top edge tiling downward while moving the positions of the R_d tiles to the left or right. To do this we need to use the $G_{dca=acd}$ tiles as shown in Figure 21.

		$T_{da=ad}$					
		$T_{dca=acd}$					
		$T_{dca=acd}$					
		$T_{dca=acd}$					
		$T_{da=ad}$					

FIGURE 21. Algorithm (II) achieves “diagonal” propagation of an instance of R_d using tile $G_{dca=acd}$, through a region mostly tiled with $G_{ca=ac}$ (labels omitted to reduce clutter).

A similar version of algorithm (II) can achieve diagonal propagation of an instance of R_b though a region mostly tiled with $G_{ca=ac}$.

We next describe the final algorithm, algorithm (III). Again the algorithm has various versions. Here we concentrate on the version where the rectangular grid graph \mathcal{N} has exactly $q + 3$ many R_a blocks on the left and right boundaries and no R_b blocks, but we assume that the top boundary of \mathcal{N} has been tiled with R_c and R_d blocks with the scarcity condition on R_d . The objective of this algorithm is to tile \mathcal{N} so that the bottom boundary will be first tiled with less than p many R_d blocks and then with R_c blocks for the rest of the boundary. Note that the number of R_d blocks is uniquely determined by the top boundary. Suppose the top boundary contains j many R_c blocks and k many R_d blocks. Then the size of the top boundary is $jp + kq + n$. The number of R_d blocks on the bottom boundary is




FIGURE 22. An illustration of algorithm (III). The smallest tiles are $G_{ca=ac}$. The snaking paths of slightly larger tiles are $G_{da=ad}$ and $G_{dca=acd}$ as in Figure 21. The widest tiles are $G_{c^q a=ad^p}$ and $G_{d^p a=ac^q}$, with the former on the row above the latter.

then projected to be

$$k - p \left\lfloor \frac{k}{p} \right\rfloor,$$

while there would be

$$j + q \left\lfloor \frac{k}{p} \right\rfloor$$

many copies of R_c on the bottom boundary. The reader should consult Figure 22 in the following description of algorithm (III).

We fill \mathcal{N} entirely with $G_{ca=ac}$, $G_{da=ad}$, $G_{dca=acd}$, $G_{c^q a=ad^p}$, and $G_{d^p a=ac^q}$. Each such tile is of the height of $G_{ca=ac}$, namely $p + n$, so we must specify $q + 1$ many overlapping rows of tiles in the interior of \mathcal{N} , since the top and bottom boundaries have been specified.

Beginning with the top row of \mathcal{N} , we propagate the $G_{da=ad}$ tiles downward using algorithms (I) and (II), and at each step either place a $G_{da=ad}$ tile immediately below a $G_{da=ad}$ tile from the row above, or offset it to the right by a $G_{dca=acd}$ tile as shown in Figure 22. We do this so that the bottom of the $q + 1$ st row of tiles consists of groups of alternating R_\times and R_c blocks followed by a single R_d block. We require that each group of alternating R_\times and R_c blocks has a number of R_c blocks which is a multiple of q , except for the last group of alternating sequence of R_\times and R_c blocks which follows the last copy of R_d in this row. We can meet this requirement since there are $q + 1$ rows as we move downward, and the distance between copies of R_d in the top row is greater than pq .

The last two rows of \mathcal{N} are tiled as follows. For the next to last row (i.e., the $q + 2$ nd row) of tiles, we use copies of $G_{c^q a=ad^p}$ to convert each group of q many alternating blocks of R_\times and R_c at the bottom of the $q + 1$ st row into an alternating sequence of R_\times and R_d blocks (with p many R_d blocks). For the last alternating sequence (whose number of R_c blocks is not necessarily a multiple of q), we might have an alternating sequence of R_\times and R_c left over, with $< q$ many copies of R_c . Thus, the bottom of the $q + 2$ nd row consists of an alternating sequence of R_\times and R_d , followed by a final alternating sequence of R_\times and R_c .

For the final row in the tiling of \mathcal{N} , we first copy down the final alternating sequence of R_\times and R_c from the row above. Then, starting from the right and working left, we use the $G_{d^p a=ac^q}$ tiles to convert each block of R_\times/R_d alternations

with p many copies of R_d , into an alternating sequence of R_\times/R_c with q many copies of R_c . As we approach the left edge, we will have a number, less than p , of R_d blocks (alternating with R_\times) left over, which we copy down. This completes our description of algorithm (III) in this setup.

Intuitively, what the construction accomplished is to gather copies of R_d blocks on the top and pass them to the left of the bottom boundary. This construction is readily adapted to gather copies of R_d on the bottom as well, and gather copies of R_b on the left and right.

We are now ready to describe the tiling of \mathcal{M} . Recall that the top and bottom boundary of \mathcal{M} are tiled with alternating copies of R_\times and R_c or R_d tiles, with the sparsity condition for R_d tiles, and the left and right boundary of \mathcal{M} are tiled with alternating copies of R_\times and R_a or R_b tiles, with the sparsity condition for R_b tiles. The reader should consult Figure 23 in the following description of the tiling.




FIGURE 23. Tiling of any finite rectangular grid graph satisfying sparsity conditions for R_b and R_d .

We first introduce two full columns and two full rows in the \mathcal{M} region. These are placed so that they define four corner subregions whose dimensions are exactly $(q+3)p+n$. Note that the sparsity conditions guarantee that the specified boundaries of the four corner regions consists of only R_a and R_c blocks and no R_b or R_d blocks. In Figure 23 we put these labels next to the relevant part of the boundary

to indicate this fact. The full columns and rows introduced also create four regions along the boundaries of \mathcal{M} , and we denote them by \mathcal{R}_{top} , $\mathcal{R}_{\text{bottom}}$, $\mathcal{R}_{\text{left}}$, and $\mathcal{R}_{\text{right}}$, respectively. We now apply algorithm (I) to each of the corner regions and the relevant version of algorithm (III) to each of the regions \mathcal{R}_{top} , $\mathcal{R}_{\text{bottom}}$, $\mathcal{R}_{\text{left}}$, and $\mathcal{R}_{\text{right}}$. This gives rise to an inner rectangular grid graph whose boundaries are tiled with R_a, R_b, R_c, R_d tiles in a fashion specified in Figure 23. Note that the number of R_b blocks and the number of R_d blocks on these boundaries are completely determined by the dimensions of \mathcal{M} , and therefore the numbers of R_b blocks on the left and right parts of the boundaries of the inner grid graph are the same, and the numbers of R_d blocks on the top and bottom parts of the boundaries of the inner grid graph are also the same. At this point we can introduce a column and a row of the inner grid graph and a complete tiling of the graph with algorithm (I).

This completes the construction of the tilings of the equivalence classes $[x]$ for $x \in F(2^{\mathbb{Z}^2})$ by the $G_{n,p,q}^i$, and by Lemma 2.3.3 thus completes the definition of $\varphi: F(2^{\mathbb{Z}^2}) \rightarrow \Gamma_{n,p,q}$. Since the subequivalence relation E_d was clopen in $F(2^{\mathbb{Z}^2})$, the local nature of the constructions of the S_x and the tilings of the \mathcal{M} show that φ is continuous. This completes the proof of Theorem 2.5.1. \square

PROOF OF (2) \Rightarrow (1) OF THEOREM 2.3.5. Let $Y \subseteq \mathfrak{b}^{\mathbb{Z}^2}$ be a subshift of finite type, n greater than or equal to the width of Y , $p, q > n$, and $g: \Gamma_{n,p,q} \rightarrow \mathfrak{b}$ which respects Y . Let $\varphi: F(2^{\mathbb{Z}^2}) \rightarrow \Gamma_{n,p,q}$ be the continuous \mathbb{Z}^2 -homomorphism in Theorem 2.5.1. Define $f: F(2^{\mathbb{Z}^2}) \rightarrow Y$ by letting $f(x)(a, b) = g \circ \varphi((a, b) \cdot x)$ for all $x \in F(2^{\mathbb{Z}^2})$ and $(a, b) \in \mathbb{Z}^2$. Then f is a continuous, equivariant map. \square

2.6. Continuous structurings

In this section we present a reformulation of the main theorem, Theorem 2.3.5, in terms of *continuous structurings* (defined below in Definition 2.6.4) instead of subshifts of finite type. This makes precise the notion of a continuous assignment of a first-order structure to an equivalence class. We show that the existence of continuous structurings can be reformulated as the existence of continuous, equivariant maps into certain subshifts of finite type and conversely. We then state a version of Theorem 2.3.5 in terms of continuous structurings. This alternate viewpoint is useful as some problems are naturally presented as continuous structuring questions. In fact, many questions of the form “Does there exist a continuous assignment of a certain type of structure to the classes of $F(2^{\mathbb{Z}^n})$?” can be directly rephrased in the terminology of continuous structurings.

We next make these notions precise by introducing the notions of a continuous structuring and Π_1 formulas.

DEFINITION 2.6.1. Let $\mathcal{L} = (F_1, \dots, F_n, F_{n+1}, \dots, F_{2n}, R_1, \dots, R_k)$ be a finite first-order language, where $F_1, \dots, F_n, F_{n+1}, \dots, F_{2n}$ are unary function symbols and R_1, \dots, R_k are relation symbols of arity $\alpha_1, \dots, \alpha_k$ respectively. By a \mathbb{Z}^n -structure for \mathcal{L} we mean a structure of the form

$$\mathfrak{A} = (A; F_1^{\mathfrak{A}}, \dots, F_n^{\mathfrak{A}}, F_{n+1}^{\mathfrak{A}}, \dots, F_{2n}^{\mathfrak{A}}, R_1^{\mathfrak{A}}, \dots, R_k^{\mathfrak{A}})$$

where A is a set with an action of \mathbb{Z}^n that is transitive and free, and for $1 \leq i \leq n$,

$$F_i^{\mathfrak{A}}(a) = e_i \cdot a, \quad F_{n+i}^{\mathfrak{A}}(a) = (-e_i) \cdot a$$

for $a \in A$, and for $1 \leq j \leq k$, $R_j^{\mathfrak{A}} \subseteq A^{\alpha_j}$.

The interpretations of the function symbols in a \mathbb{Z}^n -structure for \mathcal{L} are always fixed and are given by some transitive, free action of \mathbb{Z}^n . An important case is $A = \mathbb{Z}^n$, where the action is understood to be the addition, and we denote the standard interpretations of the function symbols by $f_1, \dots, f_n, f_{n+1}, \dots, f_{2n}$. Similarly, if $A = [x]$ is an equivalence class in $F(2\mathbb{Z}^n)$, we also denote the standard interpretation of the function symbols by $f_1, \dots, f_n, f_{n+1}, \dots, f_{2n}$.

We will adopt the following definition of a Π_1 sentence. This comes from the standard definition of Π_1 formulas in logic, except we restrict the quantifiers to a single variable.

DEFINITION 2.6.2. Let $\mathcal{L} = (F_1, \dots, F_n, F_{n+1}, \dots, F_{2n}, R_1, \dots, R_k)$ and φ be a sentence in \mathcal{L} . We say that φ is Π_1 if it is of the form

$$\varphi = \forall v \psi(v),$$

where ψ is a boolean combination of atomic formulas in the language \mathcal{L} .

Note that if $\mathcal{L} = (F_1, \dots, F_n, F_{n+1}, \dots, F_{2n}, R_1, \dots, R_k)$ with R_j of arity α_j for $1 \leq j \leq k$, the atomic formulas in the language of \mathcal{L} are of the form

$$R_j(t_1, \dots, t_{\alpha_j})$$

for some $1 \leq j \leq k$, where t_1, \dots, t_{α_j} are terms in the language \mathcal{L} , or of the form $t_1 = t_2$, where t_1 and t_2 are terms. Each term in the language \mathcal{L} is of the form

$$F_{i_1} \dots F_{i_\ell}(v)$$

for a variable v , where $1 \leq i_1, \dots, i_\ell \leq 2n$.

DEFINITION. Let $\mathcal{L} = (F_1, \dots, F_n, F_{n+1}, \dots, F_{2n}, R_1, \dots, R_k)$, t be a term in \mathcal{L} , and ψ be a formula in \mathcal{L} . The *depth* of t is defined as the number of function symbols appearing in t (counted with repetitions). The *depth* of ψ is defined as the maximum depth of terms appearing in ψ .

We will also consider *partial \mathbb{Z}^n -structures* as defined in the following. These are not technically \mathcal{L} -structures in the sense of first-order logic. Nevertheless, some instantiations of Π_1 formulas can be interpreted in such structures and have a truth value.

DEFINITION 2.6.3. Let $\mathcal{L} = (F_1, \dots, F_n, F_{n+1}, \dots, F_{2n}, R_1, \dots, R_k)$ be a language as in Definition 2.6.1. Let ℓ be a non-negative integer. An ℓ -structure for \mathcal{L} is a tuple

$$\mathfrak{A} = (D_\ell; p_1, \dots, p_n, p_{n+1}, \dots, p_{2n}, R_1^{\mathfrak{A}}, \dots, R_k^{\mathfrak{A}})$$

where $D_\ell = [-\ell, \ell]^n$, $p_1, \dots, p_n, p_{n+1}, \dots, p_{2n}$ are partial functions on D_ℓ defined as $p_i(a) = a + e_i$ whenever $a, a + e_i \in D_\ell$, and $p_{n+i}(a) = a - e_i$ whenever $a, a - e_i \in D_\ell$ for $1 \leq i \leq n$, and $R_j^{\mathfrak{A}} \subseteq (D_\ell)^{\alpha_j}$ for all $1 \leq j \leq k$.

Since the p_i are partial, the ℓ -structure just defined is not technically an \mathcal{L} -structure in the sense of first-order logic. Nevertheless, if $t(v)$ is a term (with only one variable v) in the language \mathcal{L} of depth d and $a \in D_\ell$ is of distance at least d from the boundary of D_ℓ , then $t(a)$ is well-defined in \mathfrak{A} with all function symbols interpreted by p_i . Moreover, if $\psi(v)$ is a boolean combination of atomic formulas (with only one variable v) in the language \mathcal{L} of depth d and $a \in D_\ell$ is of distance at least d from the boundary of D_ℓ , then $\psi(a)$ can be interpreted in \mathfrak{A} with all relation

symbols interpreted by $R_j^{\mathfrak{A}}$ and has a truth value. In this case, we say \mathfrak{A} *satisfies* $\psi(a)$ if $\psi(a)$ holds under the intended interpretations.

We next introduce the notion of an \mathcal{L} -structuring of $F(2^{\mathbb{Z}^n})$.

DEFINITION 2.6.4. Let \mathcal{L} be a language as in Definition 2.6.1. An \mathcal{L} -*structuring* of $F(2^{\mathbb{Z}^n})$ is a function Φ with domain $F(2^{\mathbb{Z}^n})$ which assigns to each $x \in F(2^{\mathbb{Z}^n})$ and each relation symbol R_j , $1 \leq j \leq k$, a set $\Phi(x, R_j) \subseteq [x]^{\alpha_j}$ in an invariant manner, meaning that $\Phi(x, R_j) = \Phi(y, R_j)$ whenever $[x] = [y]$. For each $x \in F(2^{\mathbb{Z}^n})$, this defines a \mathbb{Z}^n -structure for \mathcal{L}

$$\mathfrak{A}_x = ([x]; f_1, \dots, f_n, f_{n+1}, \dots, f_{2n}, \Phi(x, R_1), \dots, \Phi(x, R_k)).$$

We say the \mathcal{L} -structuring Φ is *continuous* if for each $1 \leq j \leq k$ and each $(g_1, \dots, g_{\alpha_j}) \in (\mathbb{Z}^n)^{\alpha_j}$, the map $x \in F(2^{\mathbb{Z}^n}) \mapsto \chi_{\Phi(x, R_j)}(g_1 \cdot x, \dots, g_{\alpha_j} \cdot x)$ is continuous.

The next theorem says that the question of the existence of a continuous structuring of $F(2^{\mathbb{Z}^n})$ can be reduced to the existence of a continuous equivariant map to a certain \mathbb{Z}^n -subshift of finite type.

THEOREM 2.6.5. *For every language \mathcal{L} as in Definition 2.6.1 and every Π_1 sentence φ in \mathcal{L} , there is a canonically constructed \mathbb{Z}^n -subshift of finite type $Y = Y(\varphi, \mathcal{L})$ such that the following are equivalent:*

- (1) *There is a continuous structuring Φ of $F(2^{\mathbb{Z}^n})$ such that $\mathfrak{A}_x \models \varphi$ for all $x \in F(2^{\mathbb{Z}^n})$.*
- (2) *There is a continuous, equivariant map π from $F(2^{\mathbb{Z}^n})$ to Y .*

PROOF. Let $\mathcal{L} = (F_1, \dots, F_n, F_{n+1}, \dots, F_{2n}, R_1, \dots, R_k)$ and $\varphi = \forall v \psi(v)$ be a Π_1 sentence. Let d denote the depth of ψ . Let $\mathfrak{b} = \prod_{j=1}^k 2^{(2d+1)^{n\alpha_j}}$ where again α_j is the arity of R_j . Consider the set of all d -structures for \mathcal{L} as defined in Definition 2.6.3. This set is naturally identified with the integers in $\mathfrak{b} = \{0, 1, \dots, \mathfrak{b}-1\}$. We view each element of \mathfrak{b} as a *code* for a d -structure.

Note that, since ψ has depth d , all the terms in $\psi(\vec{0})$ for $\vec{0} = (0, \dots, 0) \in \mathbb{Z}^n$ are well-defined, and thus $\psi(\vec{0})$ has a truth value in any d -structure \mathfrak{A} . We need the following definition before defining $Y(\varphi, \mathcal{L})$.

If \mathfrak{A} and \mathfrak{B} are two d -structures and $1 \leq i \leq n$, we say that \mathfrak{A} is *e_i -consistent* with \mathfrak{B} if for each relation symbol R_j , $1 \leq j \leq k$, and all $a_1, \dots, a_{\alpha_j} \in D_d$ with also $a_1 + e_i, \dots, a_{\alpha_j} + e_i \in D_d$ we have that $R_j^{\mathfrak{A}}(a_1 + e_i, \dots, a_{\alpha_j} + e_i)$ if and only if $R_j^{\mathfrak{B}}(a_1, \dots, a_{\alpha_j})$.

Let $Y = Y(\varphi, \mathcal{L})$ be the set of all $y \in \mathfrak{b}^{\mathbb{Z}^n}$ satisfying:

- (i) For each $a \in \mathbb{Z}^n$ and $1 \leq i \leq n$, the d -structure \mathfrak{A} coded by $y(a)$ is e_i -consistent with the d -structure \mathfrak{B} coded by $y(a + e_i)$;
- (ii) For each $a \in \mathbb{Z}^n$, the d -structure \mathfrak{A} coded by $y(a)$ satisfies $\psi(\vec{0})$.

Note Y is in fact a \mathbb{Z}^n -subshift of finite type of width 1.

It is now straightforward to check that there is a continuous \mathcal{L} -structuring Φ such that $\mathfrak{A}_x \models \varphi$ for all $x \in F(2^{\mathbb{Z}^n})$ if and only if there is a continuous, equivariant map into Y . In fact, suppose there is such a continuous \mathcal{L} -structuring Φ . Let $\pi: F(2^{\mathbb{Z}^n}) \rightarrow Y$ be defined as follows. For $x \in F(2^{\mathbb{Z}^n})$ and $a \in \mathbb{Z}^n$, $\pi(x)(a)$ is the integer coding the d -structure $\mathfrak{A}(x, a)$ which is given by

$$R_j^{\mathfrak{A}(x, a)}(a_1, \dots, a_{\alpha_j}) \leftrightarrow \Phi(x, R_j)((a + a_1) \cdot x, \dots, (a + a_{\alpha_j}) \cdot x)$$

for any $1 \leq j \leq k$ and $a_1, \dots, a_{\alpha_j} \in D_d$. The structures $\mathfrak{A}(x, a)$ satisfy the consistency condition (i) for all $1 \leq i \leq n$ since $\Phi(x, R_1), \dots, \Phi(x, R_k)$ defines a single structure \mathfrak{A}_x , and (ii) holds since $\mathfrak{A}_x \models \varphi$ and thus $\psi(\vec{a} \cdot x)$. The map π is clearly continuous and equivariant. Conversely, suppose π is a continuous, equivariant map from $F(2^{\mathbb{Z}^n})$ to Y . Then the consistency condition (i) will allow us to assemble a single structure \mathfrak{A}_x , and condition (ii) guarantees that $\mathfrak{A}_x \models \varphi$. \square

From now on we consider continuous \mathcal{L} -structurings of $F(2^{\mathbb{Z}^2})$. To state a version of the main theorem 2.3.5 for continuous \mathcal{L} -structurings of $F(2^{\mathbb{Z}^2})$, we need to extend the notion of an \mathcal{L} -structuring to the graphs $\Gamma_{n,p,q}$. In the following definitions we will see that the interpretations of the relation symbols in $\Gamma_{n,p,q}$ is straightforward, but to interpret the function symbols, we need to work in $G_{n,p,q}$ and consider only instantiations that are sufficiently far away from the boundary of $G_{n,p,q}$.

DEFINITION 2.6.6. Let $\mathcal{L} = (F_1, \dots, F_4, R_1, \dots, R_k)$ be a language as in Definition 2.6.1, where R_j has arity α_j for $1 \leq j \leq k$. An \mathcal{L} -structuring of $\Gamma_{n,p,q}$ is a map Φ which assigns to each relation symbol R_j a set $\Phi(R_j) \in (\Gamma_{n,p,q})^{\alpha_j}$.

Note that the relations $\Phi(R_j)$ on $\Gamma_{n,p,q}$ lift to relations $\Phi'(R_j)$ on $G_{n,p,q}$ via the quotient map $\pi : G_{n,p,q} \rightarrow \Gamma_{n,p,q}$.

DEFINITION 2.6.7. Let Φ be a \mathcal{L} -structuring of $\Gamma_{n,p,q}$, and let $\varphi = \forall v \psi(v)$ be a Π_1 sentence in the language \mathcal{L} of depth d . We say $(\Gamma_{n,p,q}, \Phi)$ satisfies φ if for every $x \in \Gamma_{n,p,q}$ and every $a \in G_{n,p,q}$ with $\pi(a) = x$, say $a \in G_{n,p,q}^i$ for some $1 \leq i \leq 12$, if a is of distance at least d from the boundary of $G_{n,p,q}^i$ then the formula $\psi(\vec{0})$ holds for $\vec{0} = (0, 0) \in \mathbb{Z}^2$ in the d -structure

$$\mathfrak{A} = ([-d, d]^2; p_1, \dots, p_4, R_1^{\mathfrak{A}}, \dots, R_k^{\mathfrak{A}})$$

where for all $1 \leq j \leq k$ and $a_1, \dots, a_{\alpha_j} \in \mathbb{Z}^2$, $R_j^{\mathfrak{A}}(a_1, \dots, a_{\alpha_j})$ holds if and only if $\Phi'(R_j)(a + a_1, \dots, a + a_{\alpha_j})$ holds in $G_{n,p,q}$.

The following is the \mathcal{L} -structuring version of the main theorem.

THEOREM 2.6.8. Let $\mathcal{L} = (F_1, \dots, F_4, R_1, \dots, R_k)$ be a language as in Definition 2.6.1, and let $\varphi = \forall v \psi$ be a Π_1 sentence of depth d . Then the following are equivalent.

- (1) There is a continuous \mathcal{L} -structuring of $F(2^{\mathbb{Z}^2})$ which satisfies φ .
- (2) There are $n < p, q$ with $n \geq 2d$ and $\gcd(p, q) = 1$, and there is an \mathcal{L} -structuring of $\Gamma_{n,p,q}$ which satisfies φ .
- (3) For any $n \geq 2d$ and all large enough $p, q > n$, there is an \mathcal{L} -structuring of $\Gamma_{n,p,q}$ which satisfies φ .

PROOF. The proof of (1) \Rightarrow (3) is very similar to that of the negative direction of Theorem 2.3.5 in §2.4. Here we sketch the changes. Let x again be the hyperaperiodic element of Lemma 2.4.1. Given the language \mathcal{L} and the Π_1 formula $\varphi = \forall x \psi$ of depth d , fix $n \geq 2d$. Let $(x, j) \mapsto \Phi(x, R_j)$ be a continuous \mathcal{L} -structuring of $F(2^{\mathbb{Z}^2})$ which satisfies φ . From the compactness of $K = \overline{[x]}$ and the continuity of Φ , there is a w such that for all $y \in K$ and terms $t_1(v), \dots, t_{\alpha_j}(v)$ of depth at most d , whether $\Phi(y, R_j)(t_1(y), \dots, t_{\alpha_j}(y))$ holds is determined by the value of y restricted to the $(2w + 1) \times (2w + 1)$ rectangle centered about $(0, 0)$.

This implies that, for any $(a, b) \in \mathbb{Z}^2$, whether $\psi((a, b) \cdot x)$ holds in $\mathfrak{A}_x = \mathfrak{A}_{(a,b) \cdot x}$ is determined by the value of x restricted to the $(2w+1) \times (2w+1)$ rectangle centered about (a, b) .

As in the proof in §2.4, let $p, q > n + 2w$ and k be such that $n_k = n$, $p_k = q$, $q_k = q$, and $w_k = w$. To define the \mathcal{L} -structuring Ψ of $\Gamma_{n,p,q}$, consider the α_j -arity relation symbol R_j and points y_1, \dots, y_{α_j} in $\Gamma_{n,p,q}$, and we need to define the truth value of $\Psi(R_j)(y_1, \dots, y_{\alpha_j})$. For this we first stipulate that $\Psi(R_j)(y_1, \dots, y_{\alpha_j})$ will not hold unless there are $1 \leq i \leq 12$ and $z, z_1, \dots, z_{\alpha_j} \in G_{n,p,q}^i$ such that $\pi(z_1) = y_1, \dots, \pi(z_{\alpha_j}) = y_{\alpha_j}$ and $z_1, \dots, z_{\alpha_j} \in [-d, d]^2 \cdot z$. In case these witnesses exist, we consider the grid graph $G_{w_k, n, p, q}^i$ in the construction of x , and let z, z_1, \dots, z_n continue to denote the corresponding points in $G_{w_k, n, p, q}^i$. We define that $\Psi(R_j)(y_1, \dots, y_{\alpha_j})$ holds if and only if $\Phi(x, R_j)(z_1 \cdot x, \dots, z_n \cdot x)$ holds.

Since $n \geq 2d$, Lemma 2.3.1 and the proof of Theorem 2.3.5 ensure that this definition does not depend on the particular choice of $z, z_1, \dots, z_{\alpha_j}$ when such witnesses exist. Since Φ satisfies φ and φ has depth d , it follows that the \mathcal{L} -structuring Ψ of $\Gamma_{n,p,q}$ satisfies φ .

For the proof of (2) \Rightarrow (1), let Ψ be an \mathcal{L} -structuring of $\Gamma_{n,p,q}$ as in (2). Exactly as in the corresponding proof in §2.5 of Theorem 2.3.5, we first construct a continuous tiling of $F(2\mathbb{Z}^2)$ by partially overlapping copies of the grid graphs $G_{n,p,q}^i$, $1 \leq i \leq 12$. Then Ψ induces a continuous \mathcal{L} -structuring Φ of $F(2\mathbb{Z}^2)$ in the natural way: $\Phi(x, R_j)(a_1 \cdot x, \dots, a_{\alpha_j} \cdot x)$ if and only if there is a copy of some $G_{n,p,q}^i$ in the tiling of $[x]$ which contains the points $a_1 \cdot x, \dots, a_{\alpha_j} \cdot x$ and, letting y_1, \dots, y_{α_j} be the points in $G_{n,p,q}^i$ corresponding to $a_1 \cdot x, \dots, a_{\alpha_j} \cdot x$, then $\Psi(R_j)(y_1, \dots, y_{\alpha_j})$ holds. To see that $\mathfrak{A}_x \models \varphi$, it is enough to fix $(a, b) \in \mathbb{Z}^2$ and show that $\mathfrak{A}_x \models \psi((a, b) \cdot x)$. Fix a copy of some $G_{n,p,q}^i$ in the tiling of $[x]$ which contains $([-d, d]^2 + (a, b)) \cdot x$. This can be done as $n \geq 2d$. Let y be the point in $G_{n,p,q}^i$ corresponding to $(a, b) \cdot x$. Since Ψ satisfies φ and y is of distance at least d from the boundary of $G_{n,p,q}^i$, we have Ψ satisfies $\psi(y)$. It follows that $\mathfrak{A}_x \models \psi((a, b) \cdot x)$. \square

REMARK 2.6.9. It is possible to deduce Theorem 2.6.8 from Theorems 2.3.5 and 2.6.5. For the reader willing to check a few details, we give an outline of this proof below. Suppose $\mathcal{L} = (F_1, \dots, F_4, R_1, \dots, R_k)$ is a language and $\varphi = \forall v \psi(v)$ is a Π_1 formula. Let $Y \subseteq \mathfrak{b}^{\mathbb{Z}^n}$ be the subshift of finite type from Theorem 2.6.5. Suppose first that there is a continuous structuring Φ with associated structures \mathfrak{A}_x which satisfies φ (that is, $\mathfrak{A}_x \models \varphi$ for all $x \in F(2\mathbb{Z}^n)$). Let $\pi: F(2\mathbb{Z}^n) \rightarrow \mathfrak{b}^{\mathbb{Z}^n}$ be the equivariant map given by Theorem 2.6.5. By Theorem 2.3.5 there is a map $g: \Gamma_{n,p,q} \rightarrow \mathfrak{b}$ which respect Y . Let $g': G_{n,p,q} \rightarrow \mathfrak{b}$ be the lift of g to the graph $G_{n,p,q}$. Define the relations R'_j on $G_{n,p,q}^{\alpha_j}$ as follows. If x_1, \dots, x_{α_j} are in one of the $G_{n,p,q}^i$ and least one of these points is an interior point of this graph, then we set $R'_j(x_1, \dots, x_{\alpha_j})$ if there is a D_d square contained in $G_{n,p,q}^i$, say centered at the point z , for which $g'(z)$ declares the points in D_d corresponding to x_1, \dots, x_{α_j} to be in the relation R_j . Since interior points of the 12 tiles are never identified, there are no other points y_1, \dots, y_{α_j} in $G_{n,p,q}$ which are identified with the x_1, \dots, x_{α_j} in $\Gamma_{n,p,q}$ and which occur in the same relative positions within some D_d square in $G_{n,p,q}$. If none of the points x_1, \dots, x_{α_j} is an interior point of $G_{n,p,q}$, then we set $R'_j(x_1, \dots, x_{\alpha_j})$ if there is a D_d square contained in $G_{n,p,q}^i$, say centered at the point z , with D_d contained in the interior of some $G_{n,p,q}^i$, say centered say at

z (necessarily z is not an interior point), where $g'(z)$ declares the corresponding points to be R_j related. If y_1, \dots, y_{α_j} are equivalent to the x_i points in $\Gamma_{n,p,q}$, then it is easy to check that $R'_j(x_1, \dots, x_{\alpha_j})$ iff $R'_j(y_1, \dots, y_{\alpha_j})$. So, the relation R'_j corresponds to a relation R_j on $\Gamma_{n,p,q}$. It is straightforward to check that this structuring of $\Gamma_{n,p,q}$ satisfies φ . The other direction is similar to the corresponding direction of the proof presented above.

REMARK 2.6.10. The analog of Theorem 2.6.8 holds for $F(2^{\mathbb{Z}})$ as well. The proof is similar (but easier) to that of Theorem 2.6.8, except we follow the proof of the two tiles theorem (Theorem 2.2.2) instead of Theorem 2.3.5.

REMARK 2.6.11. In the definition of Π_1 formulas it is important that we restrict the formula $\psi(v)$ to having a single free variable. The statement of Theorem 2.6.8 (or the corresponding theorem for $F(2^{\mathbb{Z}})$) fails for formulas of the form $\varphi = \forall u \forall v \psi(u, v)$ where ψ is quantifier free. Consider, for example, the case of $F(2^{\mathbb{Z}})$. There is a formula φ of the above form which asserts that the unary relations R_1, R_2, R_3 describe a proper 3-coloring of the \mathbb{Z} -graph, and that at most one vertex satisfies R_3 . This formula can easily be seen to hold in the two-tiles graphs $\Gamma_{n,p,q}^1$, but there is no continuous structuring of $F(2^{\mathbb{Z}})$ which satisfies φ .

2.7. Higher dimensions

The twelve tiles theorem has generalizations from the two-dimensional case $F(2^{\mathbb{Z}^2})$ to the general n -dimensional case $F(2^{\mathbb{Z}^n})$. For the applications of the current paper, however, it suffices to establish a weaker form of the negative direction of the theorem. Namely, we show that a continuous, equivariant map from $F(2^{\mathbb{Z}^n})$ to a subshift of finite type implies that there is a corresponding map from the n -dimensional $q_1 \times \dots \times q_n$ torus tile to the subshift for all sufficiently large q_1, \dots, q_n . This could be proved by the method of §2.4, but here we use instead Lemma 1.7.2.

Recall Definition 2.1.3 for the notion of a subshift of finite type $Y \subseteq \mathfrak{b}^{\mathbb{Z}^n}$. As before, a subshift Y can be described by a sequence $(\mathfrak{b}; p_1, \dots, p_k)$, where $p_i: [0, a_{i,1}) \times \dots \times [0, a_{i,n}) \rightarrow \mathfrak{b}$ are the forbidden patterns for the subshift.

By the $q_1 \times \dots \times q_n$ torus tile T_{q_1, \dots, q_n} we mean the quotient of \mathbb{Z}^n by the relation $(a_1, \dots, a_n) \sim (b_1, \dots, b_n)$ if and only if $a_1 \equiv b_1 \pmod{q_1}, \dots, a_n \equiv b_n \pmod{q_n}$. The tile inherits the \mathbb{Z}^n -graph structure from the \mathbb{Z}^n -graph structure of \mathbb{Z}^n . The tile has size $q_1 \dots q_n$, and every vertex has degree $2n$.

If $Y \subseteq \mathfrak{b}^{\mathbb{Z}^n}$ is a subshift of finite type described by $(\mathfrak{b}; p_1, \dots, p_k)$, then we say that a map $g: T_{q_1, \dots, q_n} \rightarrow \mathfrak{b}$ respects Y just as in Definition 2.3.4, using the \mathbb{Z}^n -graph T_{q_1, \dots, q_n} instead of $\Gamma_{n,p,q}$. This is equivalent to saying that the map $g \circ \pi: \mathbb{Z}^n \rightarrow \mathfrak{b}$ (where $\pi: \mathbb{Z}^n \rightarrow T_{q_1, \dots, q_n}$ is the quotient map) respect Y in that for any rectangle $R = [0, a_{i,1}) \times \dots \times [0, a_{i,n}) \subseteq \mathbb{Z}^n$, $g \circ \pi \upharpoonright R \neq p_i$.

THEOREM 2.7.1. *Let $Y \subseteq \mathfrak{b}^{\mathbb{Z}^n}$ be a subshift of finite type. If there is a continuous, equivariant map $f: F(2^{\mathbb{Z}^n}) \rightarrow Y$, then for all $q_1, \dots, q_n \geq 2$ and all sufficiently large k there is $g: T_{q_1^k, \dots, q_n^k} \rightarrow \mathfrak{b}$ which respects Y .*

PROOF. Let $f: F(2^{\mathbb{Z}^n}) \rightarrow Y$ be a continuous and equivariant map into the subshift Y , and define the map $f_0: F(2^{\mathbb{Z}^n}) \rightarrow \mathfrak{b}$ by $f_0(x) = f(x)(\vec{0})$, where $\vec{0} = (0, \dots, 0) \in \mathbb{Z}^n$. Let $x \in F(2^{\mathbb{Z}^n})$ be the hyper-aperiodic element of Lemma 1.7.2 for $G = \mathbb{Z}^n$, constructed with respect to the subgroups $G_k = q_1^k \mathbb{Z} \times \dots \times q_n^k \mathbb{Z}$. Let $\mathcal{P} = \{P_1, \dots, P_{\mathfrak{b}}\}$ be the clopen partition of $F(2^{\mathbb{Z}^n})$ determined by f_0 , that is, for

all $1 \leq m \leq \mathfrak{b}$, $y \in P_m$ if and only if $f_0(y) = m$. Let $A \subseteq \mathbb{Z}^n$ be the finite set given by Lemma 1.7.2. Let $w \geq 1$ be larger than the width of Y and let k be large enough so that the diameter of A (with respect to the ℓ_1 -norm on \mathbb{Z}^n) plus $2w$ is less than $\frac{q_i^k}{2}$ for all i . Set $R_1 = [w, q_1^k) \times \cdots \times [w, q_n^k)$ and $R_2 = [0, q_1^k + w) \times \cdots \times [0, q_n^k + w)$.

Next, our choice of k allows us to pick a point $y \in [x]$ satisfying $(R_2 \setminus R_1) \cdot y \cap AG_k = \emptyset$. From Lemma 1.7.2 we have that \mathcal{P} is G_k -invariant on $G_k(R_2 \setminus R_1) \cdot y$ since $G_k(R_2 \setminus R_1) \subseteq G \setminus AG_k$. We now define $g : T_{q_1^k, \dots, q_n^k} \rightarrow \mathfrak{b}$ as follows. For any $r + G_k \in T_{q_1^k, \dots, q_n^k} = \mathbb{Z}^n / G_k$, let $a \in (r + G_k) \cap R_2$ be arbitrary, and define $g(r + G_k) = f_0(a \cdot y)$. To see that g is well-defined, note that $(r + G_k) \cap R_2 = \{r\}$ for every $r \in R_1$, and for $a, a' \in R_2 \setminus R_1$, if $a + G_k = a' + G_k$, then $f_0(a \cdot y) = f_0(a' \cdot y)$.

Lastly, by our choice of R_2 , every copy of $[0, w) \times \cdots \times [0, w)$ in $T_{q_1^k, \dots, q_n^k}$ is the image under the quotient map of a translate of $[0, w) \times \cdots \times [0, w)$ contained in R_2 . Therefore g respects Y since the image of f is contained in Y . \square

REMARK 2.7.2. The construction of the hyper-aperiodic element of §2.4 shows a little more, namely that if $f : F(2^{\mathbb{Z}^n}) \rightarrow Y$ is continuous and equivariant, then for all sufficiently large q_1, \dots, q_n there is $g : T_{q_1, \dots, q_n} \rightarrow \mathfrak{b}$ which respects Y . Theorem 2.7.1 is sufficient for our purposes, however.

Applications of the Twelve Tiles Theorem

3.1. Basic applications of the Twelve Tiles Theorem

In this section we prove several results which follow easily from the Twelve Tiles Theorem, by which we mean either of Theorems 2.3.5 or 2.6.8.

For our first application, we give another proof of Theorem 1.5.1 stating that the continuous chromatic number of $F(2^{\mathbb{Z}^2})$ is 4. This includes a new proof of Theorem 1.5.1 for the lower bound as well as a proof different from [11] for the upper bound. These are included in the next two results.

THEOREM 3.1.1. *For any n and $p, q > n$ with $\gcd(p, q) = 1$, neither of the long tiles $T_{c^q a = ad^p}, T_{d^p a = ac^q}, T_{cb^p = a^q c}, T_{ca^q = b^p c}$ can be properly 3-colored.*

PROOF. Fix n, p, q with $p, q > n$ and $\gcd(p, q) = 1$. We only consider the long horizontal tile $T_{c^q a = ad^p}$, which is shown in Figure 24. The proof for the other long tiles is similar. Assume $\varphi: T_{c^q a = ad^p} \rightarrow \{0, 1, 2\}$ is a proper 3-coloring. We identify $\{0, 1, 2\}$ with the vertices of K_3 , the complete graph on three vertices. For any closed walk α in K_3 , let $w(\alpha)$ be its winding number in K_3 , with the convention that the 3-cycle $(0, 1, 2, 0)$ has winding number $+1$. If β is a cycle in $T_{c^q a = ad^p}$, then $\varphi(\beta)$ is a cycle in K_3 .




FIGURE 24. The long horizontal tile $T_{c^q a = ad^p}$ is not properly 3-colorable. Consider the cycle indicated by the thick line.

Consider the cycle in $T_{c^q a = ad^p}$ illustrated in Figure 24. The cycle is constructed as follows. First identify the upper-left corner vertex of each R_x block. We name three cycles that start and end at this vertex. The cycle along the left edges of R_x and R_a blocks, going downward in the figure, is denoted α . The cycle along the top edges of R_x and R_c blocks, going to the right, is denoted γ . The cycle along the top edges of R_x and R_d blocks, going to the right, is denoted δ . Note that α and γ have length p and δ has length q . For any cycle β we also use β^{-1} to denote the cycle obtained from β by reversing its direction, that is, if we list out the nodes in β as $v_1 \dots v_n$, then β^{-1} can be represented as $v_n \dots v_1$. In this notation, the cycle

we are considering is $\gamma^q \alpha \delta^{-p} \alpha^{-1}$. This cycle is depicted in Figure 24 by the thick line and has clockwise orientation.

The cycle $\gamma^q \alpha \delta^{-p} \alpha^{-1}$ encompasses a region in $T_{c^q a = ad^p}$ which is filled by a grid of cycles of length 4. It follows that the winding number of $\varphi(\gamma^q \alpha \delta^{-p} \alpha^{-1})$ is the sum of the winding numbers of several 4-cycles in K_3 . However, it is easy to check that any 4-cycle in K_3 has winding number zero. Therefore

$$q \cdot w(\varphi(\gamma)) + w(\varphi(\alpha)) - p \cdot w(\varphi(\delta)) - w(\varphi(\alpha)) = w(\varphi(\gamma^q \alpha \delta^{-p} \alpha^{-1})) = 0$$

which implies that

$$q \cdot w(\varphi(\gamma)) = p \cdot w(\varphi(\delta)).$$

Since $\gcd(p, q) = 1$, at least one of p, q must be odd. An odd cycle in K_3 must have non-zero winding number, so at least one of $w(\varphi(\gamma)), w(\varphi(\delta))$ must be non-zero. The above equation then shows that both are non-zero and since $\gcd(p, q) = 1$, $p|w(\varphi(\gamma))$. But, $\varphi(\gamma)$ is a p -cycle in K_3 , and so can have winding number at most $\frac{p}{3}$, a contradiction. \square

Using either Theorem 2.3.5 or 2.6.8, Theorem 3.1.1 immediately implies that the continuous chromatic number of $F(2^{\mathbb{Z}^2})$ is greater than 3. The proof of Theorem 1.5.1 can also be adapted to give an alternative proof of Theorem 3.1.1.

We next turn to the other direction, namely showing that there is a continuous proper 4-coloring of $F(2^{\mathbb{Z}^2})$. By Theorem 2.3.5 it suffices to show that $\Gamma_{n,p,q}$ can be properly 4-colored for some $1 \leq n < p, q$ with $\gcd(p, q) = 1$. We can, in fact, take $n = 1, p = 2, q = 3$.

THEOREM 3.1.2. *There is a proper 4-coloring of $\Gamma_{1,2,3}$.*

PROOF. We describe a proper 4-coloring on $\Gamma_{1,2,3}$ by giving the coloring on the grid graphs $G_{1,2,3}^i$ for $1 \leq i \leq 12$ so that similarly labeled blocks are colored the same way. Because proper colorings remain proper after reflections and rotations, it suffices to color the tiles $T_{ca=ac}, T_{da=ad}, T_{db=bd}, T_{c^q a = ad^p}$ and $T_{cda=adc}$. We give this coloring in Figure 25. For better readability of the colors the grid graphs are represented as boards, with square boxes representing vertices and adjacency of squares representing edges. The figure also shows the coloring on the R_x, R_a, R_b, R_c blocks by themselves for reference, so that it is easy to check the consistency of the coloring across tiles. \square

For our next application we show that there is no continuous perfect matching in the Schreier graph of $F(2^{\mathbb{Z}^n})$. Here we think of the Schreier graph of $F(2^{\mathbb{Z}^n})$ as an undirected graph. Recall that a *matching* in an undirected graph is a subset M of the edges such that every vertex is incident with at most one of the edges of M . A matching M is a *perfect matching* if every vertex is incident with exactly one edge in M . For a countable group G with a minimal generating set S , a perfect matching M in the Schreier graph $\Gamma(F(2^G), G, S)$ (see the Definition in §1.1) is given by a map $\eta: F(2^G) \rightarrow S \cup S^{-1}$ such that, for every $x \in F(2^G)$, $(x, \eta(x) \cdot x) \in M$. In order to give rise to a perfect matching, such η must satisfy $\eta(\eta(x) \cdot x) = \eta(x)^{-1}$ for all $x \in F(2^G)$. When η is continuous, we say that M is continuous.

THEOREM 3.1.3. *There is no continuous perfect matching in $F(2^{\mathbb{Z}^n})$ for any $n \geq 1$.*




FIGURE 25. A proper 4-coloring of $\Gamma_{1,2,3}$. Vertices are represented by square boxes and edges are represented by adjacency. Thick lines separate labeled blocks.

PROOF. Assume a continuous perfect matching in $F(2^{\mathbb{Z}^n})$ exists and is given by a continuous function

$$\eta : F(2^{\mathbb{Z}^n}) \rightarrow \{\pm e_1, \pm e_2, \dots, \pm e_n\}.$$

The condition $\eta(\eta(x) \cdot x) = \eta(x)^{-1} = -\eta(x)$ can be described by a number of forbidden patterns of width 1, and thus gives rise to a subshift of finite type $Y \subseteq (2n)^{\mathbb{Z}^n}$. Let q_1, \dots, q_n be odd. From Theorem 2.7.1 there is k such that there is $g: T_{q_1^k, \dots, q_n^k} \rightarrow 2n$ which respects Y . As g respects Y , it gives a perfect matching in the torus $T_{q_1^k, \dots, q_n^k}$. This is a contradiction as this torus has an odd number of vertices. \square

Next we compute the continuous edge chromatic number of $F(2^{\mathbb{Z}^2})$. Again we think of the Schreier graph of $F(2^{\mathbb{Z}^2})$ as an undirected graph. Recall that in an undirected graph two edges are *adjacent* if they are incident to a common vertex. An *edge coloring* is an assignment of colors to the edges so that adjacent edges are assigned different colors. The *edge chromatic number* of a graph is the smallest number of colors one can use in an edge coloring. In the Cayley graph of \mathbb{Z}^2 each edge is adjacent to 6 other edges, and it is easy to see that the edge chromatic number is 4. In the next result we again need only consider the torus tiles.

THEOREM 3.1.4. *The continuous edge chromatic number of $F(2^{\mathbb{Z}^2})$ is 5.*

PROOF. We first show that it is at least 5. This, in fact, follows immediately from Theorem 3.1.3. For if there were a 4-edge-coloring of $F(2^{\mathbb{Z}^2})$, then every vertex

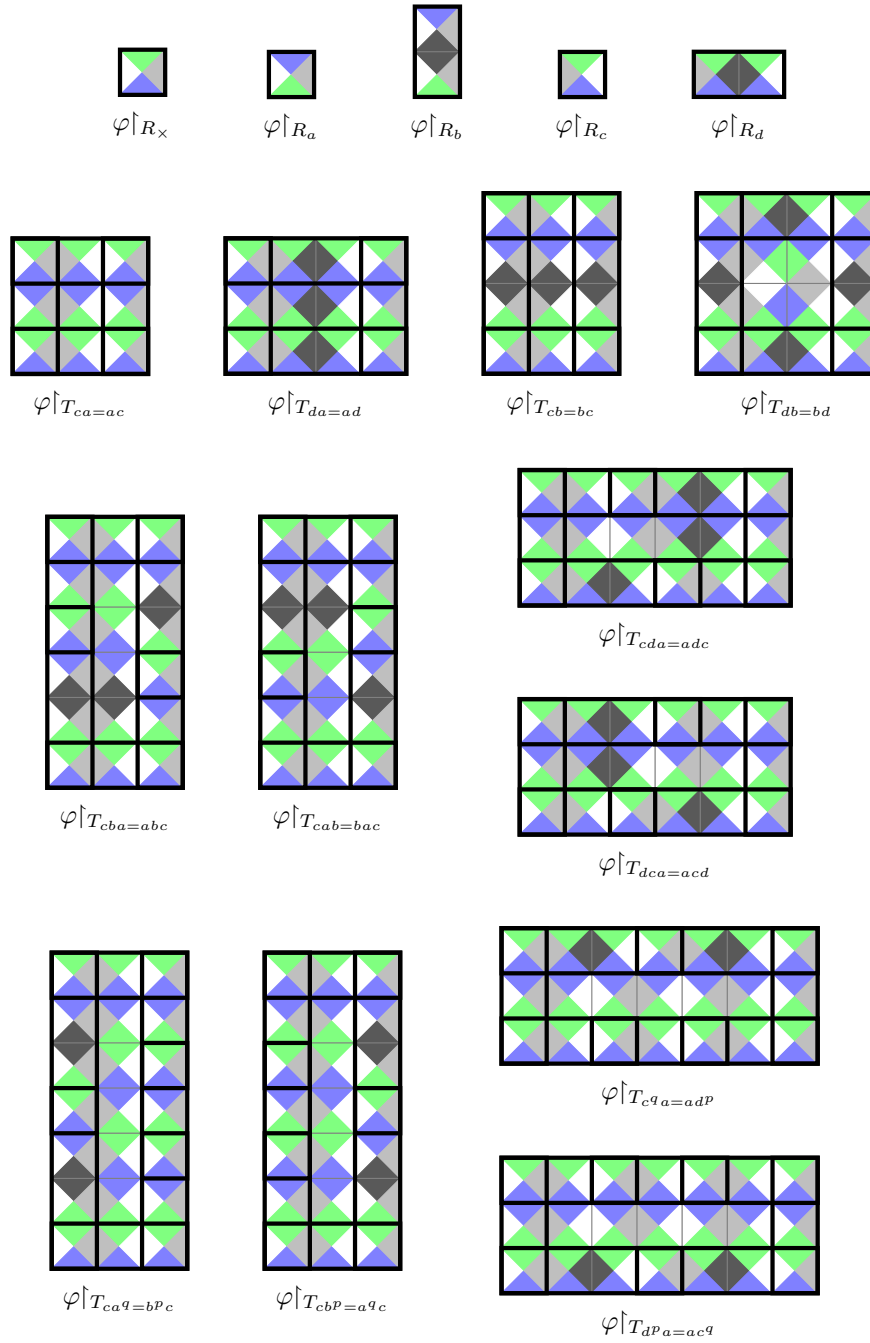


FIGURE 26. A 5-edge-coloring of $\Gamma_{1,2,3}$. Each node is assigned four colors representing the colors of its adjacent edges. Also observe that the boundary nodes of each tile are colored in a consistent way, e.g., R_d is colored the same way everywhere it appears in the figure.

of the Schreier graph has exactly one edge of each color incident with it. Thus, for any fixed color, the set of edges of that color gives a perfect matching in $F(2^{\mathbb{Z}^2})$.

We next show there is a continuous 5-edge-coloring of $F(2^{\mathbb{Z}^2})$. Note that the condition for an edge coloring can be described by forbidden patterns of width 1, thus we may apply Theorem 2.3.5 with $n = 1$. It suffices to find $p, q > 1$ with $\gcd(p, q) = 1$ such that $\Gamma_{1,p,q}$ can be 5-edge-colored. We take $p = 2, q = 3$. A 5-edge-coloring of $\Gamma_{1,2,3}$ is shown in Figure 26. \square

The Borel edge chromatic number for $F(2^{\mathbb{Z}^2})$ is either 4 or 5 (as the Borel number is less than or equal to the continuous number). We do not know the actual answer, so we state this as a question.

QUESTION 3.1.5. Is the Borel edge chromatic number for $F(2^{\mathbb{Z}^2})$ equal to 4 or 5?

Remark. Since the completion of an earlier version of this manuscript the above question has been answered independently by Grebík–Rozhoň [16], Bencs–Hrušková–Tóth [3] and Weilacher [28]. The Borel edge chromatic number for $F(2^{\mathbb{Z}^2})$ is 4.

We also do not know the continuous chromatic edge number of $F(2^{\mathbb{Z}^n})$ for $n \geq 2$. Using Theorem 3.1.3, a similar argument as above shows that it is at least $2n + 1$.

In recent work Marks [24] has shown, using techniques from games and determinacy, that every Borel function $\varphi: F(\omega^{F_2}) \rightarrow 2$, where F_2 is the free group with two generators, has an infinitely large monochromatic graph component. In the case of $F(2^{\mathbb{Z}^2})$, the situation is quite different. The proof of the following result uses just the positive direction of Theorem 2.3.5.

THEOREM 3.1.6. *There exists a continuous function $\varphi: F(2^{\mathbb{Z}^2}) \rightarrow 2$ such that each monochromatic component is not only finite, it has size at most 3.*

PROOF. We build the function $\varphi: F(2^{\mathbb{Z}^2}) \rightarrow 2$ by constructing a function $\psi: \Gamma_{n,p,q} \rightarrow 2$ for $n = 1, p = 5, q = 2$. The function ψ is constructed as in Figure 27. Theorem 2.5.1 produces a continuous \mathbb{Z}^2 -homomorphism $\varphi': F(2^{\mathbb{Z}^2}) \rightarrow \Gamma_{1,5,2}$. Let $\varphi = \psi \circ \varphi': F(2^{\mathbb{Z}^2}) \rightarrow 2$.

We need to observe that each monochromatic component of φ has size at most 3. From Figure 27 we see that each boundary point of one of the $G_{1,5,2}^i$ (which are illustrated by boundary boxes in the figure) is adjacent to interior points (i.e., points in unlabeled blocks) of a different color. Also, the largest monochromatic component of the boundary of any $G_{1,5,2}^i$ is at most 2 by inspection. Thus, if $\varphi'(x)$ is within some labeled block, then the monochromatic component containing x has size at most 2. If $\varphi'(x)$ is an interior point of one of the $G_{1,5,2}^i$, then the monochromatic component of x must consist of points y with $\varphi'(y)$ in the interior of the same $G_{1,5,2}^i$. By inspection of Figure 27 we see that this component has size at most 3. \square

Next we consider tiling questions about $F(2^{\mathbb{Z}^2})$. The general problem can be formulated as follows. Let T_1, T_2, \dots be a finite or countably infinite collection of finite subsets of \mathbb{Z}^2 (which we call the “tiles”).

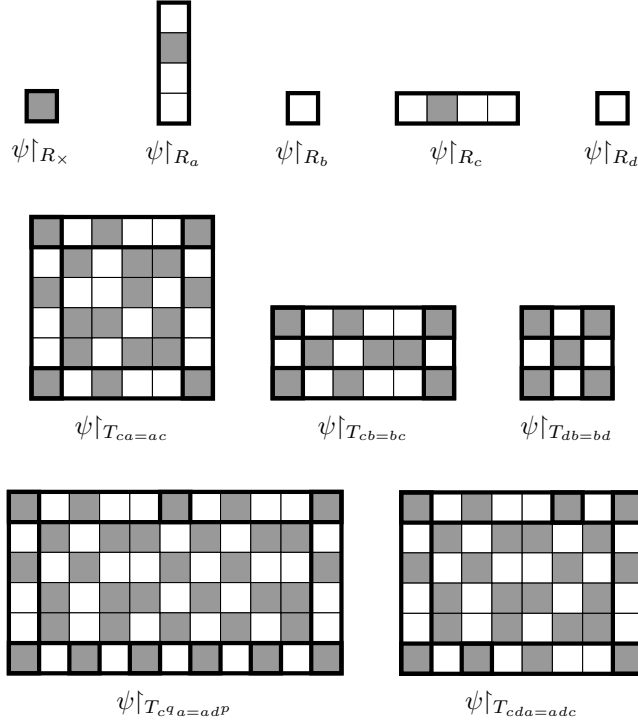


FIGURE 27. A function $\psi: \Gamma_{1,5,2} \rightarrow 2$ inducing a continuous function $\varphi: F(2^{\mathbb{Z}^2}) \rightarrow 2$ which has no infinite monochromatic graph component. Five of the twelve tiles are given; the other seven are rotations and/or reflections of these five.

DEFINITION 3.1.7. We say $F(2^{\mathbb{Z}^2})$ has a continuous (or Borel) tiling by $\{T_i\}$ if there is a clopen (or Borel) subequivalence relation E of $F(2^{\mathbb{Z}^2})$ such that every E class is isomorphic to one of the T_i . More precisely, for every $x \in F(2^{\mathbb{Z}^2})$, $\{g \in \mathbb{Z}^2: g \cdot x E x\}$ is a translate of some T_i .

This leads to the following general question.

QUESTION 3.1.8. For which sets of tiles $\{T_i\}$ is there a continuous tiling of $F(2^{\mathbb{Z}^2})$? For which sets of tiles is there a Borel tiling?

For finite sets of tiles, Theorems 2.3.5, 2.6.8 give a theoretical answer to the continuous version of the question. Namely, if we take n larger than the maximum diameter of a tile T_i , then there is a continuous tiling of $F(2^{\mathbb{Z}^2})$ by the $\{T_i\}$ if and only if there are $p, q > n$ with $\gcd(p, q) = 1$ such that there is a tiling of $\Gamma_{n,p,q}$ by the $\{T_i\}$ (as n is greater than any tile diameter, the notion of a tiling of $\Gamma_{n,p,q}$ by the T_i is naturally defined). This does not give a decision procedure for the continuous tiling question, however, and even in some specific fairly simple cases we do not have an answer.

Let us consider one specific instance of the continuous tiling question which is connected with the theory of marker structures. In [11] it is shown that for every $d > 1$ there is a tiling of $F(2^{\mathbb{Z}^n})$ by rectangular regions, each of which has side

lengths in $\{d, d+1\}$. In particular, $F(2^{\mathbb{Z}^2})$ has a continuous tiling using the four rectangular tiles of dimensions $\{d \times d, d \times (d+1), (d+1) \times d, (d+1) \times (d+1)\}$. We call this set of four tiles the $d, d+1$ tiles. A natural question is to ask for which subsets of this set can we get a continuous tiling of $F(2^{\mathbb{Z}^2})$.

The next result shows that if we omit the “small” $d \times d$ tile or the “large” $(d+1) \times (d+1)$ tile, then a continuous tiling is impossible.

THEOREM 3.1.9. *There is no continuous tiling of $F(2^{\mathbb{Z}^2})$ by the $d, d+1$ tiles such that every equivalence class omits either the $d \times d$ tile or the $(d+1) \times (d+1)$ tile. In particular, if we omit the $d \times d$ or the $(d+1) \times (d+1)$ tile from the set of $d, d+1$ tiles, then there is no continuous tiling of $F(2^{\mathbb{Z}^2})$ by this set of tiles.*

PROOF. Suppose there is a continuous tiling of $F(2^{\mathbb{Z}^2})$ by the $d, d+1$ tiles, with each class omitting either the small or large tile. The proof of Theorem 2.3.5 shows that the tiling of $\Gamma_{n,p,q}$ (we assume $n > d+1$) produced comes from a single equivalence class, that of the hyper-aperiodic element x constructed. Thus, we get a tiling of $\Gamma_{n,p,q}$ for all sufficiently large p, q without using the small tile, or without using the large tile. Say we tile $\Gamma_{n,p,q}$ without using the large tile. Since $n > d+1$, this gives a tiling of the torus tile $T_{ca=ac}$ without the large tile. Note that the available tiles all have a number of vertices which is divisible by d . However, if we take p to be relatively prime to d , then $T_{ca=ac}$ has p^2 vertices which is also relatively prime to d , so not divisible by d , a contradiction. If, on the other hand, the small tile is omitted, then the number of vertices in each available tile is divisible by $d+1$. This leads to a contradiction if we similarly consider $T_{ca=ac}$ and take p to be relatively prime to $d+1$. \square

3.2. Continuous subactions and regular tilings

In this section we provide a negative answer to the continuous version of Problem 10.6 of [11]. Specifically, we will show in Corollary 3.2.4 that there does not exist a clopen rectangular marker structure on $F(2^{\mathbb{Z}^2})$ which is too regular. We will make this notion of regularity precise later in the section. To prove this we will first demonstrate a result concerning continuous subactions of $F(2^{\mathbb{Z}^2})$ on clopen complete sections. We make this notion of subaction precise in the following definition.

DEFINITION 3.2.1. Let X be a space with a group action (G, \cdot) . Let $Y \subseteq X$, and (H, \star) an action of H on Y . We say (Y, H, \star) is a *subaction* of (X, G, \cdot) if for all $y_1, y_2 \in Y$, there is an $h \in H$ with $h \star y_1 = y_2$ if and only if there is a $g \in G$ such that $g \cdot y_1 = y_2$. If X is a topological space and G, H are topological groups, then we say the subaction is *continuous* if there is a continuous $\varphi: H \times Y \rightarrow G$, where the topology on Y is the subspace topology inherited from X , such that $h \star y = \varphi(h, y) \cdot y$ for all $y \in Y$ and $h \in H$.

Note that if (X, G, \cdot) is a continuous action and the subaction (Y, H, \star) is continuous then as an action (Y, H, \star) is continuous. If $H = G$ in Definition 3.2.1, we just write (Y, \star) for the continuous subaction. The next theorem says that there are no proper continuous subactions on clopen complete, co-complete sections of $F(2^{\mathbb{Z}^2})$.

THEOREM 3.2.2. *There does not exist a clopen, complete, co-complete section $M \subseteq F(2^{\mathbb{Z}^2})$ and a continuous free subaction (M, \star) of $F(2^{\mathbb{Z}^2})$.*

PROOF. Suppose towards a contradiction that $M \subseteq F(2^{\mathbb{Z}^2})$ is a clopen complete, co-complete section and \star is a continuous free subaction of \mathbb{Z}^2 on M . Let $x \in F(2^{\mathbb{Z}^2})$ be the hyper-aperiodic element constructed in §2.4. Let $K = \overline{\{x\}} \subseteq F(2^{\mathbb{Z}^2})$. Then K is compact. As M is a complete section, $[y] \cap M \neq \emptyset$ for any $y \in K$. As M is co-complete, we also have $[y] \cap (2^{\mathbb{Z}^2} - M) \neq \emptyset$ for any $y \in K$.

In the rest of the proof we will use the following terminology. Given $y \in F(2^{\mathbb{Z}^2})$ and $n > 0$, the n -ball about y is the set $\{z \in [y] : \rho(y, z) \leq n\}$, where ρ is the distance function defined in §1.1. Given any $y \in F(2^{\mathbb{Z}^2})$, the *first quadrant* from y is the set of all points $(a_1, a_2) \cdot y$ where $a_1, a_2 > 0$. Similarly one can define the second, third, and fourth quadrants from a point $y \in F(2^{\mathbb{Z}^2})$.

By the compactness of K and the clopen-ness of M , there is an $n_0 \in \mathbb{Z}^+$ such that for any $y \in K$, whether $y \in M$ is completely determined by $y \upharpoonright [-n_0, n_0]^2$. Suppose $\varphi : \mathbb{Z}^2 \times M \rightarrow \mathbb{Z}^2$ witnesses the continuity of the subaction (M, \star) . Since $M \cap K$ is compact and φ is continuous, the set $E = \{\varphi(\pm e_1, y), \varphi(\pm e_2, y) : y \in M \cap K\}$ is finite. This means that there is an $n_1 \in \mathbb{Z}^+$ such that for any $y \in M \cap K$, $\rho(y, \pm e_1 \star y), \rho(y, \pm e_2 \star y) \leq n_1$; in other words, for all $y \in M \cap K$ and all $\zeta \in \{\pm e_1, \pm e_2\}$, $\zeta \star y$ is within the n_1 -ball about y . In addition, we claim there is an $n_2 \in \mathbb{Z}^+$ such that

- (1) the values $\varphi(\zeta, y)$ for all $\zeta \in \{\pm e_1, \pm e_2\}$ and $y \in M \cap K$ are completely determined by $y \upharpoonright [-n_2, n_2]^2$, and
- (2) for any $\zeta \in \{\pm e_1, \pm e_2\}$, $y \in M \cap K$ and $\bar{y} \in M$, if $\bar{y} \upharpoonright [-n_2, n_2]^2 = y \upharpoonright [-n_2, n_2]^2$, then $\varphi(\zeta, \bar{y}) = \varphi(\zeta, y)$.

To see this claim, we fix without loss of generality a $\zeta \in \{\pm e_1, \pm e_2\}$ and an element $g \in E$, and consider the set $W = \{z \in M : \varphi(\zeta, z) = g\}$. W is clopen in $F(2^{\mathbb{Z}^2})$. We write W as a disjoint union of basic open neighborhoods

$$W = \bigcup_{i \in \mathbb{Z}^+} N_{s_i},$$

where each $s_i : [-m_i, m_i]^2 \rightarrow \{0, 1\}$ for some $m_i \in \mathbb{Z}^+$, and $N_{s_i} = \{u \in F(2^{\mathbb{Z}^2}) : u \upharpoonright [-m_i, m_i]^2 = s_i\}$. Note that $W \cap K$ is compact, and therefore there is $k \in \mathbb{Z}^+$ such that

$$W \cap K = \bigcup_{i \leq k} (N_{s_i} \cap K).$$

Without loss of generality, assume $N_{s_i} \cap K \neq \emptyset$ for all $i \leq k$. Then for all $y \in M \cap K$, $\varphi(\zeta, y) = g$ if and only if $y \in \bigcup_{i \leq k} N_{s_i}$. Moreover, for any $\bar{y} \in M$, if $\bar{y} \in \bigcup_{i \leq k} N_{s_i}$, we also have $\varphi(\zeta, \bar{y}) = g$. This proves the claim.

Since M is a complete section, we have that

$$\bigcup_{n > 0} [-n, n]^2 \cdot M = F(2^{\mathbb{Z}^2}).$$

By the openness of M and the compactness of K , there is an $n_3 \in \mathbb{Z}^+$ such that $[-n_3, n_3]^2 \cdot (M \cap K) = K$. In other words, for any $y \in K$ there is a $u \in M$ with $\rho(y, u) \leq n_3$, or the n_3 -ball about y contains a point $u \in M$. Similarly, since M is also a co-complete section and M is closed, there is an $n_4 \in \mathbb{Z}^+$ such that for any $y \in K$ the n_4 -ball about y contains a point $v \notin M$. Let

$$N_0 = 5 \max\{n_0, n_1, n_2, n_3, n_4\}.$$

Since $N_0 > 2 \max\{n_3, n_4\} + 1$, it follows that the N_0 -ball about any $y \in K$ contains points $u \in M$, $v \notin M$ in each of the four quadrants from y .

Let

$$A = \{(u, g \cdot u) : u \in M, g \in [-2N_0, 2N_0]^2, g \cdot u \in M\}$$

and for each $n > 0$

$$B_n = \{(u, h \star u) : u \in M, h \in [-n, n]^2\}.$$

Note that $A \subseteq \bigcup_{n>0} B_n$. Moreover $A \cap K^2$ is compact and, since \star is a continuous action, each B_n is open. Thus there is an $n_5 \in \mathbb{Z}^+$ such that $A \cap K^2 \subseteq \bigcup_{n \leq n_5} B_n$. Since \star is also a free action, we get that for each pair $(u, v) \in (M \cap K)^2$ with v being in the $2N_0$ -ball about u , $\varphi(g, u) \in [-n_5, n_5]^2$, where $g \in [-2N_0, 2N_0]^2$ is the unique group element with $g \cdot u = v$. Using the above finite set E , we get an $n_6 \in \mathbb{Z}^+$ such that for each pair $(u, v) \in (M \cap K)^2$ with v being in the $2N_0$ -ball about u , there is a \star -path from u to v which stays inside the n_6 -ball about u . By a \star -path we mean a sequence $u = w_0, w_1, \dots, w_l = v$ such that for each $0 \leq i < l$, $w_{i+1} = \zeta \star w_i$ for some $\zeta \in \{\pm e_1, \pm e_2\}$. Let

$$N_1 = N_0 + n_6.$$

Then for any $y \in K$ and any $u, v \in M$ within the N_0 -ball about y , there is a \star -path from u to v which stays inside the N_1 -ball about y .

Consider now the hyper-aperiodic element $x \in K$. Recall that for any n and sufficiently large p, q , there is a copy of $G_{n,p,q}^i$ for each $1 \leq i \leq 12$ in the domain of x . We fix $n \geq 2N_1 + 4N_0$ and sufficiently large distinct primes $p, q > n$, and consider a particular copy of $G_{n,p,q}^1$ in the domain of x . The alternative notation for this copy of grid graph is $G_{ca=ac}$. We think of $G_{ca=ac}$ both as a sub-grid-graph of \mathbb{Z}^2 and as a subset of $[x]$ that corresponds to the points $G_{ca=ac} \cdot x$. Consider the R_\times block in the upper-left corner of this copy of $G_{ca=ac}$. Let C be the $N_0 \times N_0$ rectangle at the center of this R_\times block. From the definition of N_0 , there is a $g \in C$ such that $g \cdot x \in M$. We fix this g and let $m = g \cdot x \in M$. Let $m' = (p, 0) \cdot m$ and $m'' = (0, -p) \cdot m$. Then m' corresponds to the point in the same position within the R_\times block in the upper-right corner of this copy of $G_{ca=ac}$, and m'' corresponds to the point in the same position within the R_\times block in the lower-left corner of this copy of $G_{ca=ac}$. The top part of this copy of $G_{ca=ac}$ along with the elements m and m' are illustrated in Figure 28.




FIGURE 28. The top part of $G_{ca=ac}$. The outer rectangle has dimensions $n \times (p+n)$ with $n > 2N_1 + 4N_0$; the inner rectangle is positioned in the center and has dimensions $(p+2N_0) \times 2N_0$. The \star -path from m to m' stays inside the outer rectangle.

In Figure 28, the outer rectangle shows the top part of $G_{ca=ac}$ and has dimensions $(p+n) \times n$. We also consider an inner rectangle of dimensions $(p+2N_0) \times 2N_0$

that is positioned in the center of the outer rectangle. This region is denoted as R . We see that $m, m' \in R$.

Consider the point m . By the definition of N_0 , in the first or fourth quadrant from m there is another point m_1 in M within the N_0 -ball of m , and with x -coordinate strictly greater than that of m . From the definition of N_1 , there is a \star -path from m to m_1 that stays within the N_1 -ball of m , and therefore stays within the outer rectangle. Continue this construction with m_1 replacing m , and we find $m_2 \in R$ to the right of m_1 and a \star -path from m_1 to m_2 that stays within the outer rectangle. Note that the intersection of the N_0 -ball about m_1 with either the first quadrant from m_1 or the fourth quadrant from m_1 is entirely contained in the region R , thus we can guarantee that $m_2 \in R$. This construction can be repeated until the \star -path extends to m' . We have thus found a \star -path from m to m' with intermediate endpoints in R and the entire path staying within the outer rectangle. This is illustrated in Figure 28.

Let h be the unique element of \mathbb{Z}^2 such that $h \star m = m'$. Thus, $h = \zeta_k \cdot \zeta_{k-1} \cdots \zeta_0$, where each $\zeta_i \in \{\pm e_1, \pm e_2\}$ and $m, \zeta_0 \star m, \zeta_1 \star \zeta_0 \star m, \dots, \zeta_k \star \cdots \star \zeta_0 \star m = m'$ is the \star -path from m to m' which stays inside the outer rectangle. Likewise, there is an element $h' \in \mathbb{Z}^2$ such that $h' \star m = m''$ and a \star -path from m to m'' that also stays inside the left part of $G_{ca=ac}$. Let $H = \langle h, h' \rangle \leq \mathbb{Z}^2$, and let $G = \mathbb{Z}^2/H$.

Let $T = ([0, p-1] \times [-p+1, 0]) \cdot m$ and $M_T = M \cap T$. Then T corresponds to the $p \times p$ rectangular region with m at the upper-left corner, and M_T is the set of all points of M in T . In the following we will define an action $*$ of G on M_T .

Before defining $*$ we define another action \star' of \mathbb{Z}^2 on M_T . We first define this action for the standard generators of \mathbb{Z}^2 . Suppose $\zeta \in \{\pm e_1, \pm e_2\}$. Let $u \in M_T$. If $\zeta \star u \in M_T$ then let $\zeta \star' u = \zeta \star u$. Assume $\zeta \star u \notin M_T$. Then by the definition of N_0 , both $u, \zeta \star u$ must be within ρ distance N_0 to the boundary of T , and therefore each of $u, \zeta \star u$ must be a point in one of the labeled blocks of $G_{ca=ac}$. Consider another block of $G_{ca=ac}$ with the same label and let u' be the point within this block in the same position as u . Then $u' = (a_1, a_2) \cdot u$ with some $a_1, a_2 \equiv 0 \pmod{p}$. By the definition of N_0 , we have that $\zeta \star u' = (a_1, a_2) \cdot (\zeta \star u)$. Note that exactly one of such $\zeta \star u'$ is contained in T . We let $\zeta \star' u$ be this unique $\zeta \star u' \in M_T$. This finishes the definition of $\zeta \star' u$ for all $u \in M_T$.

In general, suppose $g \in \mathbb{Z}^2$ and $u \in M_T$, we write $g = \zeta_k \zeta_{k-1} \cdots \zeta_0$ where $\zeta_i \in \{\pm e_1, \pm e_2\}$, and define

$$g \star' u = \zeta_k \star' \zeta_{k-1} \cdots \star' \zeta_0 \star' u.$$

We must verify that the definition does not depend on the decomposition of g into a product of standard generators. For this we show that if ζ_1, \dots, ζ_k is a sequence of standard generators of \mathbb{Z}^2 such that $\sum \zeta_i = (0, 0)$, then $\zeta_k \star' \zeta_{k-1} \cdots \star' \zeta_1 \star' u = u$ for all $u \in M_T$. Fix such a sequence of standard generators and $u \in M_T$. We construct an element $\bar{x} \in F(2\mathbb{Z}^2)$ as follows. In the domain of \bar{x} first lay down a copy of $G_{ca=ac}$ in exactly the same position as the copy of $G_{ca=ac}$ in the domain of x . Then we put down a tessellation of copies of $G_{ca=ac}$ with overlapping boundaries as illustrated in Figure 19, with the first copy of $G_{ca=ac}$ in the center of the tessellation, and with $2k+1$ many copies of $G_{ca=ac}$ horizontally and $2k+1$ many copies vertically. So the total number of copies of $G_{ca=ac}$ is $(2k+1)^2$. We then define \bar{x} so that the values of \bar{x} on any copy of $G_{ca=ac}$ is the same as the values of x on $G_{ca=ac}$, and outside the tessellation of copies of $G_{ca=ac}$ the values of \bar{x} are defined in any way

to make it an element of $F(2\mathbb{Z}^2)$. Let $\bar{u} \in [\bar{x}]$ be the corresponding point to $u \in [x]$. By the definition of N_0 , we have that $\bar{u} \in M$ and that for any standard generator ζ , $\varphi(\zeta, u) = \varphi(\zeta, \bar{u})$ since $u \uparrow [-N_0, N_0]^2 = \bar{u} \uparrow [-N_0, N_0]^2$. By the same argument we can see that for all $v \in [x] \cap M$ and $\bar{v} \in [\bar{x}]$, if v is within $G_{ca=ac}$ in the definition of x , \bar{v} is within any one of the $(2k-1)^2$ many copies of $G_{ca=ac}$ in the definition of \bar{x} that are not on the boundary of the tessellation, and v and \bar{v} are in corresponding positions, then $\bar{v} \in M$ and $\varphi(\zeta, v) = \varphi(\zeta, \bar{v})$. As a consequence, if we compare the two sequences

$$u, \zeta_1 \star' u, \zeta_2 \star' \zeta_1 \star' u, \dots, \zeta_k \star' \dots \star' \zeta_1 \star' u$$

and

$$\bar{u}, \zeta_1 \star \bar{u}, \zeta_2 \star \zeta_1 \star \bar{u}, \dots, \zeta_k \star \dots \star \zeta_1 \star \bar{u},$$

we conclude that the corresponding terms are points in corresponding positions in some copies of $G_{ca=ac}$. Since each sequence has $k+1$ terms, the second sequence of points is entirely contained in the tessellation of copies of $G_{ca=ac}$ in the construction of \bar{x} . Since \star is a free action and $\sum \zeta_i = (0, 0)$, we see that $\zeta_k \star \dots \star \zeta_1 \star \bar{u} = \bar{u}$. By the above comparison, we have $\zeta_k \star' \dots \star' \zeta_1 \star' u = u$ as required.

We have thus defined an action \star' of \mathbb{Z}^2 on M_T . Next we show that this action induces an action $*$ of G on M_T . First we claim that $h \star' m = m$ and $h' \star' m = m$. To see this, consider the \star -path from m to m' that stays within the top part of $G_{ca=ac}$. We now make a stronger observation that every point on this \star -path is in fact of ρ -distance at least N_0 away from the boundary of the top part of $G_{ca=ac}$. This is because $n > 2N_1 + 4N_0$ and that the intermediate endpoints used in the construction of the \star -path are all within the $2N_0 \times p$ region R . Write the sequence of points on the \star -path as

$$m, \zeta_1 \star m, \dots, \zeta_k \star \dots \star \zeta_1 \star m,$$

where $h = \sum_{i \leq k} \zeta_i$. Then the N_0 -ball about each of these points are still contained in the top part of $G_{ca=ac}$. Compare it with another sequence

$$m, \zeta_1 \star' m, \dots, \zeta_k \star' \dots \star' \zeta_1 \star' m,$$

which by definition is contained in M_T . We conclude that the corresponding terms are in corresponding positions within some labeled blocks of $G_{ca=ac}$. Since $h \star m = m'$, this comparison gives that $h \star' m = m$. Similarly $h' \star' m = m$ because $h' \star m = m''$. Next we claim that for any $u \in M_T$, $h \star' u = u$ and $h' \star' u = u$. Fix $u \in M_T$. We construct a \star -path from m to u which stays entirely within $G_{ca=ac}$ and so that every point on the \star -path is of ρ distance N_0 away from the boundary of $G_{ca=ac}$. This \star -path is constructed in a similar fashion as the \star -path from m to m' , making use of the fact that from any point v of M there are other points of M within the N_0 -ball of v in each of the four quadrants from v . Let $t \in \mathbb{Z}^2$ be the unique element such that $t \star m = u$. Then the \star -path corresponds to a decomposition of t as $\sum_{i \leq k} \zeta_i$ for some sequence ζ_1, \dots, ζ_k of standard generators of \mathbb{Z}^2 . A similar comparison argument as above gives that $t \star' m = u$. Now commutativity of \mathbb{Z}^2 gives

$$h \star' u = (tht^{-1}) \star' u = t \star' h \star' t^{-1} \star' u = u.$$

Similarly $h' \star' u = u$ as well. These claims guarantee that for any $g_1, g_2 \in \mathbb{Z}^2$ with $g_1 - g_2 \in H$, and for any $u \in M_T$, $g_1 \star' u = g_2 \star' u$. In other words, if we define $(g + H) \star u = g \star' u$, we obtain an action of G on M_T .

We next claim that this action \star of G on M_T is free. For this suppose $u \in M_T$ and $g \in \mathbb{Z}^2$ are such that $g \star u = u$. We need to show that $g \in H$. Note that we have $g \star m = m$. Write g as $\sum_{i \leq k} \zeta_i$ where ζ_1, \dots, ζ_k are standard generators of \mathbb{Z}^2 . Consider the element $\bar{x} \in F(2^{\mathbb{Z}^2})$ constructed above, with $(2k+1)^2$ many copies of $G_{ca=ac}$ in its domain. Let $\bar{m} \in [\bar{x}]$ correspond to $m \in [x]$. Then the comparison argument gives that $g \star \bar{m}$ and $g \star m$ are in corresponding positions in some copies of $G_{ca=ac}$. This implies that there are integers a_1, a_2 such that $g \star \bar{m} = (a_1 p, -a_2 p) \cdot \bar{m}$. The comparison argument also gives that $h \star \bar{m} = (p, 0) \cdot \bar{m}$ and $h' \star \bar{m} = (0, -p) \cdot \bar{m}$. This can be applied to all copies of $G_{ca=ac}$ to conclude that, if $\bar{m} \in [\bar{x}]$ is in the position corresponding to m within that copy of $G_{ca=ac}$, then $h \star \bar{m} = (p, 0) \cdot \bar{m}$ and $h' \star \bar{m} = (0, -p) \cdot \bar{m}$. Repeatedly applying this fact, we get that $g \star \bar{m} = (a_1 h + a_2 h') \star \bar{m}$. Since \star is a free action, we conclude that $g = a_1 h + a_2 h' \in H$ as required.

Thus we have shown that \star is a free action of G on M_T . This in particular implies that G is finite, and that $|M_T| = |G|$. Also, since $G = \mathbb{Z}^2 / \langle h, h' \rangle$, $|G|$ is also the absolute value of the determinant of the 2×2 matrix P whose rows are the components of h and h' , as this represents the volume of the fundamental region of the lattice in \mathbb{Z}^2 generated by the vectors h and h' . Finally, we show that the components of h and h' are all divisible by p , which shows that $|\det(P)| \geq p^2$. Since $|M_T| \leq |T| = p^2$, it follows that $M_T = T$, a contradiction to the fact that T contains points not in M .

To see that the components of h and h' are divisible by p , we use the long tiles. Again we concentrate on h and assume $h = (h_1, h_2)$, the argument for h' being similar. Consider the long tile grid graph $G_{c^q a = ad^p}$ (see Figure 10). Let $m_0 \in M$ be near the center (to be exact, within the N_0 -ball of the center) of the upper-left copy of R_\times block in $G_{c^q a = ad^p}$. Let $m_1 = (p, 0) \cdot m_0$, $m_q = (pq, 0) \cdot m_0$, $u_0 = (0, -p) \cdot m_0$, $u_1 = (q, 0) \cdot u_0$, and $u_p = (pq, 0) \cdot u_0$. Then $m_0, m_1, m_q, u_0, u_1, u_p$ are all in the same position within an R_\times block. From the definition of $G_{c^q a = ad^p}$ we have that $h \star m_0 = m_1$, $qh \star m_0 = m_q$, $h' \star m_0 = u_0$, and $h' \star m_q = u_p$ because all of these relations are witnessed by points in $G_{ca=ac}$. Now since $n > 2N_1 + 4N_0$, we can find a \star -path from u_0 to u_1 that stays within the bottom portion of $G_{c^q a = ad^p}$ consisting of only the first and second R_\times blocks and the R_d block in between them. This gives rise to an $\ell \in \mathbb{Z}^2$ such that $\ell \star u_0 = u_1$ and furthermore $p\ell \star u_0 = u_p$. Thus, $h' \star (qh) \star m_0 = u_p = (p\ell) \star h' \star m_0$. Since \star is a free action, we must have $h' + qh = p\ell + h'$ in \mathbb{Z}^2 , so $qh = p\ell$. Since $\gcd(p, q) = 1$, this shows that h_1, h_2 are divisible by p . \square

As a consequence of Theorem 3.2.2 we can provide a partial answer to Problem 10.6 of [11] concerning “almost lined-up” marker regions. Recall that by a “rectangular tiling” of $F(2^{\mathbb{Z}^2})$ we mean a finite subequivalence relation E of $F(2^{\mathbb{Z}^2})$ such that each E -class R is a finite subset of a $F(2^{\mathbb{Z}^2})$ class $[x]$, and corresponds to a rectangle in that it is equal to $R = ([a, b] \times [c, d]) \cdot x$ for some $a < b, c < d$ in \mathbb{Z} . We have the natural notions of the “top-edge” of R , etc. If R_1, R_2 are rectangular E -classes in the same $F(2^{\mathbb{Z}^2})$ class $[x]$, we say R_1, R_2 are *adjacent* if there are $y \in R_1$ and $z \in R_2$ such that $\zeta \cdot y = z$ for some $\zeta \in \{\pm e_1, \pm e_2\}$. Note that if R_1, R_2 are adjacent, then from the ordered pair (R_1, R_2) the generator is uniquely determined, and so it makes sense to say “ R_2 is above R_1 ,” etc. If R_1 and R_2 are adjacent, it

also makes sense to define the “overlap” of R_1 and R_2 . For example, if R_2 is above R_1 , this is the number of points $y \in R_1$ such that $e_2 \cdot y \in R_2$.

DEFINITION 3.2.3. Suppose E is a rectangular tiling of $F(2^{\mathbb{Z}^2})$ and suppose all the rectangles in the tiling E have side lengths bounded above by M . We say the rectangular tiling E is *almost lined-up* if for every rectangle R_1 in E , there is a rectangle R_2 in E which is adjacent to R_1 , and to the top of R_1 , with the overlap of R_1 and R_2 being strictly greater than $\frac{M}{2}$. Similarly if “top” is replaced with “bottom”, “left”, or “right.”

Problem 10.6 of [11] asks if there exist Borel rectangular tilings of $F(2^{\mathbb{Z}^2})$ that are almost lined-up. The next result rules out the existence of clopen, almost lined-up rectangular tilings.

COROLLARY 3.2.4. *There does not exist a clopen rectangular tiling of $F(2^{\mathbb{Z}^2})$ which is almost lined-up.*

PROOF. Suppose E were a clopen, almost lined-up, rectangular tiling of $F(2^{\mathbb{Z}^2})$. Since the E classes are finite, and E is a clopen tiling, there is a clopen set $S \subseteq F(2^{\mathbb{Z}^2})$ which is a selector for E . We define a free action \star of \mathbb{Z}^2 on S as follows. If $s \in R_1$ (R_1 a class of E), then $(0, 1) \star s$ is the unique point of $S \cap R_2$, where R_2 is the class of E which is adjacent to R_1 , to the top of R_1 , and has overlap with R_1 which is $> \frac{M}{2}$. By assumption, such a class R_2 exists, and it is easy to see it is unique. The action of the other generators is defined similarly. Since E is clopen, it is easy to check that \star is a continuous action on S , and it is also clearly a free action. This contradicts Theorem 3.2.2. \square

3.3. Undirected graph homomorphisms and the reduced homotopy group

In the rest of this chapter we apply the Twelve Tiles Theorem to study the existence of continuous graph homomorphisms $\varphi : F(2^{\mathbb{Z}^2}) \rightarrow \Gamma$ for various graphs Γ . For this purpose, all our graphs will be assumed to be undirected and without any assigned \mathbb{Z}^2 -labeling of their edges.

In the setting of continuous graph homomorphisms, our main theorem, Theorem 2.3.5, takes the following form.

COROLLARY 3.3.1. *Let Γ be a graph. Then the following are equivalent.*

- (1) *There is a continuous graph homomorphism $\varphi : F(2^{\mathbb{Z}^2}) \rightarrow \Gamma$.*
- (2) *There is a continuous graph homomorphism $\varphi : F(2^{\mathbb{Z}^2}) \rightarrow \Gamma_0$ where Γ_0 is a connected component of Γ .*
- (3) *There are positive integers n, p, q with $n < p, q$ and $\gcd(p, q) = 1$ such that $\Gamma_{n,p,q}$ admits a graph homomorphism to Γ .*
- (4) *For all $n \geq 1$ and for all sufficiently large p, q , $\Gamma_{n,p,q}$ admits a graph homomorphism to Γ .*

PROOF. Let S be the set of all functions $x \in V(\Gamma)^{\mathbb{Z}^2}$ such that $x : \mathbb{Z}^2 \rightarrow V(\Gamma)$ is a graph homomorphism. It is easily seen that S is a subshift of finite type. In fact S is defined by a condition on the edges and so has width 1 (the notion of width is defined before Definition 2.1.3). The existence of a continuous graph homomorphism from $F(2^{\mathbb{Z}^2})$ to Γ is equivalent to the existence of a continuous equivariant

map from $F(2^{\mathbb{Z}^2})$ to S . Similarly, the existence of a graph homomorphism from $\Gamma_{n,p,q}$ to Γ is equivalent to the existence of a map $g: \Gamma_{n,p,q} \rightarrow V(\Gamma)$ which respects S . Having noted these equivalences, the equivalence of (1), (3) and (4) follows immediately from Theorem 2.3.5.

We next argue that (1) and (2) are equivalent. It is obvious that (2) implies (1). For the other direction, assume (1) holds. By the above equivalence, we have (3) holds, and consider a graph homomorphism φ from $\Gamma_{n,p,q}$ to Γ , where $n < p, q$ and $\gcd(p, q) = 1$. Since $\Gamma_{n,p,q}$ is connected, we know that the image of φ is contained in a connected component Γ_0 of Γ . Thus we have (3) for Γ_0 instead of Γ . Applying the equivalence between (3) and (1) for Γ_0 instead of Γ , we get that (2) holds. \square

By the above corollary, the continuous homomorphism problem is essentially equivalent to the instances of the problem in which the graph Γ is connected. Thus in the remainder of this chapter, we tacitly assume that Γ is connected unless explicitly stated otherwise.

In some sense the above corollary has provided a solution to the continuous graph homomorphism problem, but the conditions are not easy to apply in practice, and so we still seek to improve them.

We prove two main theorems on this topic. The first one, Theorem 3.4.1, gives a sufficient condition on the graph Γ to establish that there does not exist a continuous graph homomorphism from $F(2^{\mathbb{Z}^2})$ to Γ . Similarly, the second theorem, Theorem 3.5.1, gives a sufficient condition on Γ to establish the existence of a continuous homomorphism from $F(2^{\mathbb{Z}^2})$ to Γ . We refer to these theorems as the “negative” and “positive” theorems, and to the conditions in these two theorems as the “negative” and “positive” conditions. Both of these conditions are phrased in terms of a certain “reduced homotopy group” $\pi_1^*(\Gamma)$ associated to the graph Γ . Roughly speaking, $\pi_1^*(\Gamma)$ will be the quotient of the first homotopy group $\pi_1(\Gamma)$ of Γ by the normal subgroup generated by the 4-cycles in Γ (we give the precise definitions below). The negative theorem has a number of corollaries which give easier-to-apply, but less general, conditions sufficient to guarantee the non-existence of a continuous graph homomorphism from $F(2^{\mathbb{Z}^2})$ to Γ .

We will illustrate each of these theorems by considering a number of specific examples of graphs for which these theorems provide an answer. We note that even in the case of simple graphs such as K_3 or K_4 (the complete graphs on 3 and 4 vertices respectively), the question is not trivial (Theorems 1.5.1 and 3.1.1 show there is no continuous homomorphism for $\Gamma = K_3$, while Theorem 3.1.2 shows there is a continuous homomorphism for $\Gamma = K_4$). We will give some simpler versions of the negative conditions, as mentioned above, but also give some examples of graphs Γ where these simpler conditions do not suffice to answer the continuous homomorphism question for Γ , but nevertheless we can show a negative answer.

In §4.4 we will return to the graph homomorphism problem. Building on the results of this chapter, we will show in Theorem 4.4.1 that the continuous graph homomorphism problem for $F(2^{\mathbb{Z}^2})$ is not decidable. That is, the set of finite graphs Γ for which there exists a continuous graph homomorphism from $F(2^{\mathbb{Z}^2})$ to Γ is not a computable set (we will observe that it is a computably enumerable set, however).

The arguments in the next two sections will be based on homotopy theory. In the rest of this section, we review the homotopy notions we will need within the specific setting we will use them.

Fix an undirected graph $\Gamma = (V(\Gamma), E(\Gamma))$. For simplicity we still write edges of Γ as pairs $(u, v) \in E(\Gamma)$ for $u, v \in V(\Gamma)$. A path is a finite sequence of vertices $\sigma = (v_0, \dots, v_n)$ where $(v_i, v_{i+1}) \in E(\Gamma)$ for every $0 \leq i < n$ (we allow paths to visit the same vertex multiple times). The length of a path is the length of the sequence minus one (i.e. the number of edges traversed). A trivial path consists of one vertex and has length 0. A *back-tracking* path is a path $\sigma = (v_0, \dots, v_{2n})$ such that $v_{n-i} = v_{n+i}$ for all $0 \leq i \leq n$. If σ is a path from u to v , then we write σ^{-1} for the reverse path from v to u . If σ ends at v and τ begins at v then we write $\sigma\tau$ for the concatenated path which first follows σ and then follows τ . Concatenation is done without removing instances of back-tracking, so the length of $\sigma\tau$ is the sum of the lengths of σ and τ . Removing instances of back-tracking is done by passing to the homotopy equivalence class of a path. *Homotopy equivalence* is an equivalence relation on paths defined as follows: it is the smallest equivalence relation with the property that if σ is a path ending at v , τ is a path beginning at v , and θ is a back-tracking path beginning and ending at v , then $\sigma\theta\tau$ is homotopy equivalent to $\sigma\tau$. Notice that $\sigma\sigma^{-1}$ is homotopy equivalent to the trivial path.

A path which begins and ends at the same vertex will be called a *cycle*, and when it has length p it will be called a p -*cycle*. If σ is a 4-cycle beginning at u and τ is any path ending at u , then $\tau\sigma\tau^{-1}$ is a cycle and we call it a *square-cycle*. Notice that, for any fixed vertex $v \in V(\Gamma)$ serving as a base point, the set of homotopy equivalence classes of cycles beginning at v forms a group under concatenation. This is called the *first fundamental group of Γ* and is denoted $\pi_1(\Gamma)$. It is a basic exercise to show that up to isomorphism the group $\pi_1(\Gamma)$ does not depend upon the choice of the base point v , assuming Γ is connected. We will always implicitly assume without mention that a fixed base point v has been chosen, and that all cycles being discussed are cycles which begin and end at v . For a cycle γ we write $\pi_1(\gamma)$ for its homotopy equivalence class. The set of π_1 -images of square-cycles is conjugation invariant and thus generates a normal subgroup N . We write $\pi_1^*(\Gamma) = \pi_1(\Gamma)/N$ for the quotient of $\pi_1(\Gamma)$ by the subgroup N . We refer to $\pi_1^*(\Gamma)$ as the *reduced homotopy group of Γ* , and to N as the *kernel* of π_1^* . For a cycle γ we write $\pi_1^*(\gamma)$ for the image of γ in this group. If $\pi_1^*(\gamma) = \pi_1^*(\eta)$ we say that γ and η are *reduced homotopy equivalent*.

Notice that if two paths are homotopy equivalent then their lengths differ by an even number. Since square-cycles also have even length, the lengths of reduced homotopy equivalent cycles also differ by an even number. It follows that $\pi_1^*(\gamma)$ is non-trivial whenever γ is an odd-length cycle.

Note that if γ is a cycle in Γ then it is well-defined, without mentioning the base point, to say that $\pi_1^*(\gamma)$ satisfies a particular group-theoretic statement φ in the group $\pi_1^*(\Gamma)$. This is because choosing a different base point for computing the homotopy groups results in considering a conjugate $\gamma' = \delta\gamma\delta^{-1}$ of γ , and since conjugation is an isomorphism we have that $\pi_1^*(\Gamma)$ satisfies $\varphi(\gamma)$ if and only if $\delta\pi_1^*(\Gamma)\delta^{-1}$ satisfies $\varphi(\gamma')$. For example, if γ is a cycle in Γ , it is well-defined to say that $\pi_1^*(\gamma)$ has order p in $\pi_1^*(\Gamma)$.

We mention an alternative view of $\pi_1^*(\Gamma)$ which requires a bit of homotopy theory and fundamental groups in a slightly more general setting. One could view Γ as a topological space by viewing each edge as a copy of $[0, 1]$. We can form a larger topological space Γ^* by attaching a copy of the unit-square $[0, 1] \times [0, 1]$ to each 4-cycle in Γ , identifying the boundary of the square with the edges and vertices

used by the 4-cycle. This is an instance of a so-called CW complex (see [17] for the definition and discussion. In [23] it is shown how a CW complex is associated to a group presentation). In this setting, a path is any continuous function from $[0, 1]$ to Γ^* , and two paths are homotopy equivalent if they have the same endpoints and, while maintaining the same endpoints, one path can be continuously morphed into the other. Again, one can fix a point v and consider the set of homotopy equivalence classes of paths which begin and end at v . This forms a group denoted $\pi_1(\Gamma^*)$. It is not difficult to check that $\pi_1^*(\Gamma) = \pi_1(\Gamma^*)$. We will only briefly use this identity in one example, but the reader may find it intuitively helpful.

We point out that any graph homomorphism $\Gamma \rightarrow \Gamma'$ induces a group homomorphism $\pi_1^*(\Gamma) \rightarrow \pi_1^*(\Gamma')$ (and similarly with each π_1^* replaced by π_1).

As an example, we calculate the reduced homotopy group for the 12-tiles graph $\Gamma_{n,p,q}$ below. We will not need this result in the rest of the paper.

PROPOSITION 3.3.2. *Let $p, q \geq n$ with $\gcd(p, q) = 1$. Then $\pi_1^*(\Gamma_{n,p,q}) \cong \mathbb{Z}^2$.*

PROOF. We simply sketch the argument as we will not need this fact, and the argument is similar to one given later in Lemma 4.4.3. In computing $\pi_1^*(\Gamma_{n,p,q})$ we choose our base point v to lie in the R_\times block. Any cycle σ beginning at v can be viewed as a collection of path segments on the 12 tiles. Each of the 12 tiles forming $\Gamma_{n,p,q}$ consists of an interior region (these points are not identified with any others in $\Gamma_{n,p,q}$) surrounded by rectangular blocks labeled by $R_\times, R_a, R_b, R_c,$ or R_d . For the purposes of this proof we refer to these surrounding blocks as the boundary blocks of the tile. Since every block on the boundary either is an R_\times block or has copies of R_\times blocks adjacent to it, it is possible to adjust σ modulo the kernel of π_1^* so that σ always stays in the boundary blocks and enters and leaves tiles through a copy of R_\times (in fact through v). With this adjustment, each path segment within the blocks of a given tile represents a cycle and can be expressed as a product in $\alpha, \beta, \gamma, \delta$, where α is the cycle at v passing vertically through R_a once, and similarly $\beta, \gamma,$ and δ pass through $R_b, R_c,$ and R_d once, respectively. This implies $\pi_1^*(\Gamma_{n,p,q})$ is generated by $\pi_1^*(\{\alpha, \beta, \gamma, \delta\})$. In fact, we have named each of the 12 tiles with an equation in its subscript, and this equation describes a relation which is satisfied in $\pi_1^*(\Gamma_{n,p,q})$. Let R be the set of these 12 relations. So, in the group $\pi_1^*(\Gamma_{n,p,q})$ we have that each of α, β commutes with each of γ, δ , $\alpha\beta = \beta\alpha$, $\gamma\delta = \delta\gamma$, and $\alpha^q = \beta^p$, $\gamma^q = \delta^p$. Let m, k be integers such that $mp + kq = 1$. Let $u = \alpha^m \beta^k$ and $v = \gamma^m \delta^k$. Then $u^q = \alpha^{mq} \beta^{kq} = \beta^{mp+kq} = \beta$ in $\pi_1^*(\Gamma_{n,p,q})$, and similarly $\alpha = u^p$, $\gamma = v^p$, $\delta = v^q$. So, $\pi_1^*(\Gamma_{n,p,q})$ is generated by the commuting elements u, v . It is clear then that the group $\langle \alpha, \beta, \gamma, \delta \mid R \rangle$ is isomorphic to \mathbb{Z}^2 , the free abelian group generated by u and v .

It remains to show that in the group $\pi_1^*(\Gamma_{n,p,q})$ there are no further relations than those generated by R , that is, the kernel of the map $\varphi: \langle \alpha, \beta, \gamma, \delta \rangle \rightarrow \pi_1^*(\Gamma_{n,p,q})$ is the normal subgroup generated by the words in R . Suppose $\sigma = \eta_1^{\epsilon_1} \cdots \eta_\ell^{\epsilon_\ell}$ is null-homotopic in $\Gamma_{n,p,q}$, where each $\eta_i \in \{\alpha, \beta, \gamma, \delta\}$ and $\epsilon_i \in \{\pm 1\}$. Let $\sigma_0 = \sigma, \sigma_1, \dots, \sigma_k$ be a sequence of cycles witnessing this, i.e., each σ_{i+1} differs from σ_i by the insertion/deletion of a four-cycle in $\Gamma_{n,p,q}$, or a two-cycle, and σ_k is the trivial cycle at v . Let σ_i^* be the result of applying the above algorithm of adjustment to express σ_i as a word in $\{\alpha, \beta, \gamma, \delta\}$. Thus, every part u_0, u_1, \dots, u_m of the path σ_i for which u_0, u_m are in the boundary blocks of one of the 12 tiles, but u_1, \dots, u_{m-1} are in the interior of this tile, is replaced by a path in this tile which stays in the boundary blocks. σ_i^* is the path in $\langle \alpha, \beta, \gamma, \delta \rangle$ which is equivalent

to this boundary-block path in the natural manner. Since σ_i^* differs from σ_i by a combination of four-cycles (or two-cycles), and since it is easy to check that all four-cycles in $\Gamma_{n,p,q}$ are part of a single tile, it is straightforward to check that σ_{i+1}^* and σ_i^* differ by either the insertion/deletion of a canceling pair such as $\alpha\alpha^{-1}$, or the insertion/deletion of one of the words in R . Since σ_k is the trivial word, we must have that σ_0 is in the normal subgroup generated by R . \square

3.4. Negative conditions for graph homomorphisms

The following result is the general “negative condition” sufficient to give the non-existence of a continuous graph homomorphism from $F(2^{\mathbb{Z}^2})$ to the graph Γ .

THEOREM 3.4.1 (Homotopy-based negative condition for graph homomorphisms). *Let Γ be a graph. Suppose that for every $N > 0$ there are relatively prime integers $p, q > N$ such that for every p -cycle γ of Γ , $\pi_1^*(\gamma^q)$ is not a p^{th} power in $\pi_1^*(\Gamma)$. Then there does not exist a continuous graph homomorphism $\varphi: F(2^{\mathbb{Z}^2}) \rightarrow \Gamma$.*

PROOF. We prove the contrapositive. So let $\varphi: F(2^{\mathbb{Z}^2}) \rightarrow \Gamma$ be a continuous graph homomorphism. Fix any $n \geq 1$. By Corollary 3.3.1 there is N so that for all integers $p, q > N$ there is a graph homomorphism $\theta: \Gamma_{n,p,q} \rightarrow \Gamma$. Fix such a pair (p, q) with p, q relatively prime. Consider the rectangular grid graph given by the long tile $T_{c^q a = ad^p}$ (see Figure 10). Recall that this tile has q many copies of R_c blocks on its top boundary with each R_c block occurring between two copies of R_x blocks, a single R_a block on its left and right boundaries, and has p many copies of R_d blocks along its bottom boundary with each R_d block occurring between two copies of R_x blocks. Let σ be the cycle in $T_{c^q a = ad^p}$ which begins and ends at the vertex v which is the upper-left vertex of the upper-left copy of R_x , and goes clockwise through all of the upper left corners of the labeled blocks, and which traces out a rectangle (see Figure 24). Let $\sigma' = \theta \circ \sigma$ be the corresponding cycle in Γ .

The arrangement of the boundary blocks tells us that σ' has the form $\gamma^q \alpha \delta^{-p} \alpha^{-1}$, where γ and δ are the images of the horizontal paths from (the upper-left vertices of) R_x to R_x going through R_c and R_d , respectively, and α is the image of the vertical path from R_x to R_x going through R_a (see Figure 24). In particular, γ is a p -cycle. Clearly $\pi_1^*(\sigma) = 1$ so $\pi_1^*(\sigma') = 1$ as well and thus $\pi_1^*(\gamma^q) = \pi_1^*(\alpha \delta^p \alpha^{-1}) = \pi_1^*(\alpha \delta \alpha^{-1})^p$. \square

The next theorem, which is a consequence of Theorem 3.4.1, also gives a “negative condition” for the non-existence of continuous graph homomorphisms into Γ . This result is frequently easier to apply than Theorem 3.4.1, and has the advantage of referring directly to the graph Γ and not the homotopy group $\pi_1^*(\Gamma)$. A *weighting* of Γ , or a *weight function* for Γ , is a function $w: \tilde{E}(\Gamma) \rightarrow \mathbb{Z}$, where

$$\tilde{E}(\Gamma) = \{(u, v): \{u, v\} \in E(\Gamma)\},$$

such that for all $(u, v) \in \tilde{E}(\Gamma)$, $w(u, v) + w(v, u) = 0$. For a weighting $w: \tilde{E}(\Gamma) \rightarrow \mathbb{Z}$ and a path $\gamma = (v_0, \dots, v_n)$ in Γ , we can extend the function w by setting

$$w(\gamma) = \sum_{i=0}^{n-1} w(v_i, v_{i+1}).$$

THEOREM 3.4.2 (Weighting-based negative condition for graph homomorphisms). *Let Γ be a graph. Assume that for infinitely many integers p and every p -cycle*

γ there exists a weighting w such that $w(\gamma)$ is not divisible by p and $w(\sigma) = 0$ for all 4-cycles σ . Then there does not exist a continuous graph homomorphism $\varphi: F(2^{\mathbb{Z}^2}) \rightarrow \Gamma$.

PROOF. We again prove the contrapositive. So suppose there is a continuous graph homomorphism. By Theorem 3.4.1, there is $N > 0$ such that for all relatively prime $p, q > N$ there is a p -cycle γ in Γ such that $\pi_1^*(\gamma^q)$ is a p^{th} power in $\pi_1^*(\Gamma)$. Fix such a pair (p, q) , let γ be a p -cycle, and let τ be a cycle satisfying $\pi_1^*(\gamma^q) = \pi_1^*(\tau^p)$. Fix any weight w such that $w(\sigma) = 0$ for all 4-cycles σ . Then $w(\sigma) = 0$ for all square-cycles. This implies that w induces a group homomorphism $w: \pi_1^*(\Gamma) \rightarrow \mathbb{Z}$. Therefore $q \cdot w(\gamma) = w(\gamma^q) = w(\tau^p) = p \cdot w(\tau)$. Since $\gcd(p, q) = 1$, we conclude $w(\gamma)$ is divisible by p . \square

REMARK 3.4.3. Theorem 3.4.2 generalizes to weight functions $w: \tilde{E}(\Gamma) \rightarrow A$ for any abelian group A . In applying Theorem 3.4.2 we frequently use a fixed weighting w on the graph Γ (that is, w does not depend on p and the p -cycle γ).

The next example shows that we may use Theorem 3.4.2 to give yet another proof that the continuous chromatic number of $F(2^{\mathbb{Z}^2})$ is greater than 3. Note that a proper 3-coloring of $F(2^{\mathbb{Z}^2})$ is equivalent to a graph homomorphism from $F(2^{\mathbb{Z}^2})$ to K_3 .

EXAMPLE 3.4.4. There is no continuous graph homomorphism $\varphi: F(2^{\mathbb{Z}^2}) \rightarrow K_3$, where K_3 is the complete graph of size 3.




FIGURE 29. The complete graph of order 3, K_3 , with a weighting function.

PROOF. We apply Theorem 3.4.2 with the weight function w given in Figure 29. Since K_3 contains no non-trivial 4-cycles, we trivially have that w assigns 0 to every 4-cycle. Now consider an odd integer $p > 1$ and consider a p -cycle γ . Between any two consecutive instances of the weight 1 edge, there must be an even number of weight 0 edges. Since p is odd, it follows $w(\gamma) \neq 0$. Clearly also $|w(\gamma)| < p$ since γ has length p and must use at least one edge of weight 0. Thus $w(\gamma)$ is not divisible by p . We conclude there is no continuous homomorphism to K_3 . \square

The next theorem, extending the above argument slightly, gives a simple (in particular, computable) sufficient condition on the graph Γ for there not to exist a continuous graph homomorphism from $F(2^{\mathbb{Z}^2})$ to Γ .

THEOREM 3.4.5 (Simple negative condition). *There does not exist a continuous graph homomorphism from $F(2^{\mathbb{Z}^2})$ to the graph Γ if Γ satisfies either of the following:*

- (1) Γ is properly 3-colorable.
- (2) Γ has no non-trivial 4-cycles.

PROOF. The first statement follows immediately from Example 3.4.4. Namely, note that Γ is properly 3-colorable if and only if there is a graph homomorphism from Γ to K_3 . If Γ is properly 3-colorable, then any continuous graph homomorphism from $F(2^{\mathbb{Z}^2})$ to Γ would immediately give a continuous graph homomorphism from $F(2^{\mathbb{Z}^2})$ to K_3 , a contradiction to Example 3.4.4. Suppose now Γ has no non-trivial 4-cycles. Fix an arbitrary direction for each edge in the graph Γ . Let $p > 1$ be odd, and let γ be any p -cycle in Γ , say beginning and ending at v . There must be some edge in Γ which occurs in γ an odd number of times, as otherwise the total length of γ , which is p , would be even. Say the edge $e_0 = (u, v)$ occurs an odd number of times in γ . Let $w: \tilde{E}(\Gamma) \rightarrow \mathbb{Z}$ be the weighting function where $w(u, v) = 1$ (so $w(v, u) = -1$) and $w(e) = 0$ for all other edges e of Γ . Clearly $w(\gamma) \neq 0$, as e_0 is traversed an odd number of times by γ . Also, $|w(\gamma)| < p$ as γ contains an edge e with $w(e) = 0$ (since p is odd, γ must at some point use an edge other than e_0). So, $w(\gamma)$ is not divisible by p . As there are no non-trivial 4-cycles in Γ , the hypotheses of Theorem 3.4.2 are satisfied. \square

The simple conditions of Theorem 3.4.5 suffice to rule out the existence of continuous graph homomorphisms from $F(2^{\mathbb{Z}^2})$ to Γ for many graphs Γ . For example, the next example considers the *Petersen* graph from graph theory.




FIGURE 30. The Petersen graph Γ_P , with a proper 3-coloring of its nodes.

EXAMPLE 3.4.6. There is no continuous graph homomorphism $\varphi: F(2^{\mathbb{Z}^2}) \rightarrow \Gamma_P$, where Γ_P is the Petersen graph.

PROOF. The Petersen graph is given in Figure 30, along with a proper 3-coloring (i.e a homomorphism into K_3). There can be no continuous homomorphism into Γ_P by (1) of Theorem 3.4.5. \square

A less immediate example is given by the following “clamshell” graph Γ_C depicted in Figure 31. This graph was chosen so as to have chromatic number strictly greater than 3 and to have several nontrivial 4-cycles.

EXAMPLE 3.4.7. There is no continuous graph homomorphism $\varphi: F(2^{\mathbb{Z}^2}) \rightarrow \Gamma_C$.




FIGURE 31. The “Clamshell” graph Γ_C , together with a weight function. Observe the subtlety that the edges $x \rightarrow u_0$ and $x' \rightarrow u_0$ are labeled *positive* 1, unlike the edges from x and x' to the other u_i 's.

PROOF. A weighting w of the graph is shown in Figure 31. We first check that $w(\sigma) = 0$ for all 4-cycles. Note that there is only one non-trivial 4-cycle not containing x , namely (x', u_3, v_3, v_4) , which has weight 0. All other non-trivial 4-cycles must contain x and thus must contain two vertices adjacent to x . We quickly see the 4-cycles (x, u_m, x', u_k) for $0 \leq m \neq k \leq 4$, all of which have weight 0. The only edges leaving x we have not considered are (x, v_0) and (x, x') . The only non-trivial 4-cycle containing $\{x, v_0\}$ and not $\{x, x'\}$ is (x, v_0, v_1, u_1) , and the only non-trivial 4-cycles containing $\{x, x'\}$ are (x, x', u_0, v_0) and (x, x', v_4, u_4) . All of these have weight 0, proving the claim.

Now let p be odd and let γ be a p -cycle. Since all edges have weight ± 1 and p is odd, $w(\gamma) \neq 0$. We claim that $|w(\gamma)| < p$. Towards a contradiction suppose not. Since $|w(e)| = 1$ for all edges e in Γ_C , we have that $|w(\gamma)| = p$. By replacing

γ with γ^{-1} if necessary, we can assume $w(\gamma) = p$, meaning each edge traversed by γ contributes $+1$ to $w(\gamma)$. Notice that since γ is a cycle, every vertex used by γ must have a $+1$ incoming edge and a $+1$ outgoing edge. So γ cannot use the vertex v_2 as no edge entering v_2 has weight $+1$. Similarly, it cannot use v_1 or v_3 since the only edges with weight $+1$ entering these vertices emanate from v_2 . Similarly, it cannot use v_0 or v_4 , as their only $+1$ entering edges come from v_1 and v_3 . It also cannot use any u_i for $1 \leq i \leq 4$ since the only $+1$ edge entering u_i comes from v_i . Similarly, γ cannot use x' because every $+1$ edge entering x' comes from a u_i with $1 \leq i \leq 4$ or from v_4 . Lastly, γ cannot use x because every $+1$ edge entering x comes from a u_i with $1 \leq i \leq 4$ or from v_0 or from x' . This means that u_0 is the only vertex used by γ , a contradiction. We conclude that $w(\gamma)$ is not divisible by p , so by Theorem 3.4.2 there is no continuous graph homomorphism into Γ_C . \square

While simpler to use, Theorem 3.4.2 is strictly weaker than Theorem 3.4.1. This will follow from considering the “Klein bottle” graph Γ_K which is a quotient of the graph shown in Figure 32. In that figure we have labeled several vertices. Γ_K is obtained by identifying vertices of the same label (and identifying edges whose endpoints have identical labels). Unlabeled vertices are not identified with one another. Notice that Γ_K would be a 5×5 torus graph, i.e., the Cayley graph of $(\mathbb{Z}/5\mathbb{Z})^2$, except that the progression of nodes a_i is downward ascending on the left side of the graph and upward ascending on the right side of the graph.




FIGURE 32. Γ_K is obtained as a quotient of the above graph by identifying labeled vertices having the same label (and by identifying edges whose endpoints have the same labels). Unlabeled vertices are not identified. The 2-dimensional cell complex Γ_K^* is homeomorphic to the Klein bottle.

EXAMPLE 3.4.8. There is no continuous graph homomorphism $\varphi: F(2^{\mathbb{Z}^2}) \rightarrow \Gamma_K$. This follows from Theorem 3.4.1 but not Theorem 3.4.2.

PROOF. Define the cycles $\alpha = (x, a_1, a_2, a_3, a_4, x)$ and $\beta = (x, c_1, c_2, c_3, c_4, x)$. Let $\pi_1^*(\Gamma_K) = \pi_1(\Gamma_K)/N$ be the reduced homotopy group of Γ_K with respect to the base point x . We will use the identification $\pi_1^*(\Gamma_K) = \pi_1(\Gamma_K^*)$. The 2-dimensional complex Γ_K^* can be obtained by adding a 2-cell for every 4-cycle in the graph in Figure 32 and then taking the quotient. It follows from this that Γ_K^* is homeomorphic to the Klein bottle. So $\pi_1^*(\Gamma_K) \cong \pi_1(\Gamma_K^*) \cong G$ where $G =$

$\langle g, h | gh = hg^{-1} \rangle$ and where the isomorphism identifies $\pi_1^*(\alpha)$ with g and $\pi_1^*(\beta)$ with h . Notice that g and h have infinite order, that $\langle g \rangle$ is normal in G , and that $G/\langle g \rangle = \{h^i \langle g \rangle : i \in \mathbb{Z}\} \cong \mathbb{Z}$. It is also possible to show directly, without appeal to Γ_K^* , that $\pi_1^*(\Gamma_K) \cong G$, by using an argument similar to that of Proposition 3.3.2 to show that there are no more relations in $\pi_1^*(\Gamma_K)$ than those of G .

We will apply Theorem 3.4.1. Fix any pair of relatively prime integers (p, q) with p odd, and let γ be a p -cycle in Γ_K , say starting and ending at v . Note that any node in Γ_K is no more than four steps away from the node x . Let δ be path of length $|\delta| \leq 4$ from v to x . Consider $\delta^{-1}\gamma\delta \in \pi_1(\Gamma_K)$ with $|\delta^{-1}\gamma\delta| \leq |\gamma| + 8$.

Notice that in the quotient graph Γ_K , the edge set can be unambiguously partitioned into horizontal edges and vertical edges, and every 4-cycle has two horizontal edges and two vertical edges. Also, it is straightforward to check that starting at any vertex $v \in \Gamma_K$, following a vertical edge and then a horizontal edge is equivalent modulo a 4-cycle to following a horizontal edge and then a vertical edge. It follows that there is a cycle γ' starting and ending at x with $|\gamma'| = |\delta^{-1}\gamma\delta|$, $\pi_1^*(\gamma') = \pi_1^*(\delta^{-1}\gamma\delta)$ (recall every square-cycle is in N), and with $\gamma' = \rho_1\rho_2$ where ρ_1 consists of just horizontal edge movements and ρ_2 just vertical edge movements. Each of ρ_1, ρ_2 must be a cycle, since if ρ_1 did not end at x , $\rho_1\rho_2$ would not be a cycle. It follows that there are $i, j \in \mathbb{Z}$ and $\gamma'' = \beta^i\alpha^j$ with $|\gamma''| \leq |\gamma'|$ and $\pi_1(\gamma'') = \pi_1(\gamma')$. Note that in going from γ' to γ'' we eliminate backtracking which may reduce the cycle length. Since $p > 2$ and $5|i| + 5|j| = |\gamma''| \leq |\gamma'| + 8 = p + 8$, we must have $|i|, |j| < p$.

Towards a contradiction, suppose there is a cycle τ , starting and ending with v , with γ^q equal to τ^p in the reduced homotopy group of Γ_K with base vertex v . It follows that $\pi_1^*(\delta^{-1}\gamma^q\delta) = \pi_1^*(\delta^{-1}\tau^p\delta) = \pi_1^*((\delta^{-1}\tau\delta)^p)$. Let i', j' be such that $\pi_1^*(\delta^{-1}\tau\delta) = \pi_1^*(\beta^{i'}\alpha^{j'})$. Thus we have $(h^i g^j)^q = (h^{i'} g^{j'})^p$ in G . In $G/\langle g \rangle$ we have $h^{qi} \langle g \rangle = h^{p i'} \langle g \rangle$. Since h has infinite order in $G/\langle g \rangle$, this gives $p|(qi)$, and so $p|i$. Since $|i| < p$, this implies $i = 0$, which then gives $i' = 0$ as well. So, $g^{qj} = g^{p j'}$ in G . Since g has infinite order in G , the same argument now gives that $j = j' = 0$. Thus, $\pi_1^*(\delta^{-1}\gamma\delta)$ is the identity in $\pi_1^*(\Gamma_K)$. This is a contradiction as $\delta^{-1}\gamma\delta$ is an odd-length cycle, and N consists of only even length cycles. We conclude, via Theorem 3.4.1, that there is no continuous graph homomorphism from $F(2^{\mathbb{Z}^2})$ into Γ_K .

Finally, we claim that Theorem 3.4.2 cannot be applied to Γ_K . Let w be any weight which assigns weight 0 to all 4-cycles. Then w induces a group homomorphism $w : \pi_1^*(\Gamma_K) \rightarrow \mathbb{Z}$. Since \mathbb{Z} is abelian, w must nullify the commutator subgroup. As $\pi_1^*(\beta^{-1}\alpha\beta\alpha^{-1}) = \pi_1^*(\alpha^{-2})$ (i.e. $h^{-1}ghg^{-1} = g^{-2}$ in G), we must have $-2w(\alpha) = w(\alpha^{-2}) = 0$ and thus $w(\alpha) = 0$. Now consider any $p \geq 5$. If p is even then choose a p -cycle γ which simply alternates between two adjacent vertices. Then $w(\gamma) = 0$ is divisible by p . If p is odd then let γ be the p -cycle which first follows α once and then alternates between two adjacent vertices. Then $w(\gamma) = w(\alpha) = 0$ is divisible by p . \square

3.5. A positive condition for graph homomorphisms

In this section we present a ‘‘positive condition’’ on the graph Γ which is sufficient to guarantee the existence of a continuous graph homomorphism from $F(2^{\mathbb{Z}^2})$ to Γ . The positive condition, as was the negative condition, is stated in terms of the reduced homotopy group $\pi_1^*(\Gamma)$.

THEOREM 3.5.1. *Let Γ be a graph. Suppose Γ contains an odd-length cycle γ such that $\pi_1^*(\gamma)$ has finite order in $\pi_1^*(\Gamma)$. Then there is a continuous graph homomorphism $\varphi: F(2^{\mathbb{Z}^2}) \rightarrow \Gamma$.*

PROOF. By Corollary 3.3.1, it will suffice to find a graph homomorphism $\varphi: \Gamma_{1,p,q} \rightarrow \Gamma$ for some $p, q > 1$ with $\gcd(p, q) = 1$. Note that φ can be viewed as a ‘‘coloring’’ of the vertices of $\Gamma_{1,p,q}$ by colors in $V(\Gamma)$ (with adjacent vertices in $\Gamma_{1,p,q}$ getting colors which are adjacent in Γ). Let ℓ be the length of γ , and let $m = |\pi_1^*(\gamma)|$ be the order of $\pi_1^*(\gamma)$ in $\pi_1^*(\Gamma)$. Let $p \gg m\ell$ be prime, and let $q = p + m$. Because $p > m$, p does not divide m , so p does not divide q . Thus, $\gcd(p, q) = 1$. This gives us our parameters p, q .

Before we proceed to the coloring of $\Gamma_{1,p,q}$, we make the additional observation that q must be odd. Indeed, $\pi_1^*(\Gamma)$ is just $\pi_1(\Gamma)$ modded out by conjugates of four-cycles in Γ , and each such conjugate is a cycle of even length. Thus, no odd length cycle is null-homotopic. Since γ has odd length and γ^m (with length $m\ell$) is null-homotopic, m must be even, hence q is odd. So we have that both p, q are odd.

Let $v_0, v_1 \in V(\Gamma)$ be the first two vertices in γ . Extend γ with alternating instances of v_0, v_1 until it is length p , and call this extension α ; we can do this because p is odd and γ is of odd length. Similarly, extend γ with alternating instances of v_0, v_1 until it is of length q , and call this extension β . In our coloring, v_0 will correspond to the R_\times block (note that now the R_\times block has dimensions 1×1), α to R_a, R_c blocks, and β to R_b, R_d blocks. More precisely, we define the coloring first on the boundaries of the 12 tiles comprising $\Gamma_{1,p,q}$ as follows. Each occurrence of an R_\times block is labeled with the vertex v_0 . Each occurrence of adjacent blocks of the form $R_\times R_a R_\times$ is colored by the vertices in the cycle α , with the order of the vertices occurring in α matching the upward order of the blocks $R_\times R_a R_\times$. Similarly, the adjacent $R_\times R_b R_\times$ blocks are also colored by vertices of β in the same order. Likewise, the adjacent $R_\times R_c R_\times$ and $R_\times R_d R_\times$ blocks are colored by the vertices in α, β respectively, in the left-to-right order. This defines the coloring φ on the boundaries of the 12 tiles in $\Gamma_{1,p,q}$, and we have clearly respected the vertex identifications in $\Gamma_{1,p,q}$.

It is also instructive to note the following fact. Consider, for the moment, the boundary of the long tile $T_{c^q a = ad^p}$ in $\Gamma_{1,p,q}$ (see Figure 10). Moving counter-clockwise from the bottom-left corner and disregarding the null-homotopic ‘‘filler’’ between copies of γ consisting of two-cycles (v_0, v_1, v_0) , we encounter p copies of γ corresponding to R_d , one copy of γ corresponding to R_a , q backwards copies of γ corresponding to R_c , and one final backwards copy of γ corresponding to R_a again. Thus, the boundary of this tile is homotopic to γ^{p-q} . But we have arranged that $m = p - q$, so the boundary is actually null-homotopic. Similarly, all twelve tiles have null-homotopic boundaries, so our scheme has, so far, avoided homotopy obstructions.

Having colored the boundaries of the tiles, it suffices to color each tile in isolation. The general strategy to get started will be to connect pairs of forwards and backwards copies of γ at the boundary of a tile with special paths, and then fill the remaining space with alternating v_0, v_1 . This process is illustrated in Figure 33. Note that in order for this strategy to work, we need the occurrences of the coloring by γ to be of the same parity, as illustrated in Figure 33. More precisely, for each occurrence of the coloring by γ , whether it is a forwards or backwards copy, we

concentrate on the position of the leading v_0 (one adjacent to v_1) and regard its parity to be the parity of the occurrence of γ . To apply this strategy, we assign the label of v_0 to the points in some chosen diagonally-moving path (the diagonal movement is why parity is important) that connects the v_0 's of the two instances of γ . This diagonally-moving path divides the tile into two pieces. We can fill one side by alternating v_0 and v_1 . For the other side, the region consisting of points of distance at most ℓ to the diagonally-moving path can be filled in using γ , and the remaining portion is filled in by alternating v_0 and v_1 . Thus the process in Figure 33 works whenever the two occurrences of γ are of the same parity but of the opposite orientation. Aside from that constraint, this process is quite flexible, as shown in Figure 33; a path of width $\ell + 1$ and almost arbitrary shape can connect two copies of γ satisfying the above conditions. For future reference we refer to this connection algorithm as Algorithm (C) and the trivial algorithm to fill in alternating v_0 and v_1 as Algorithm (T).

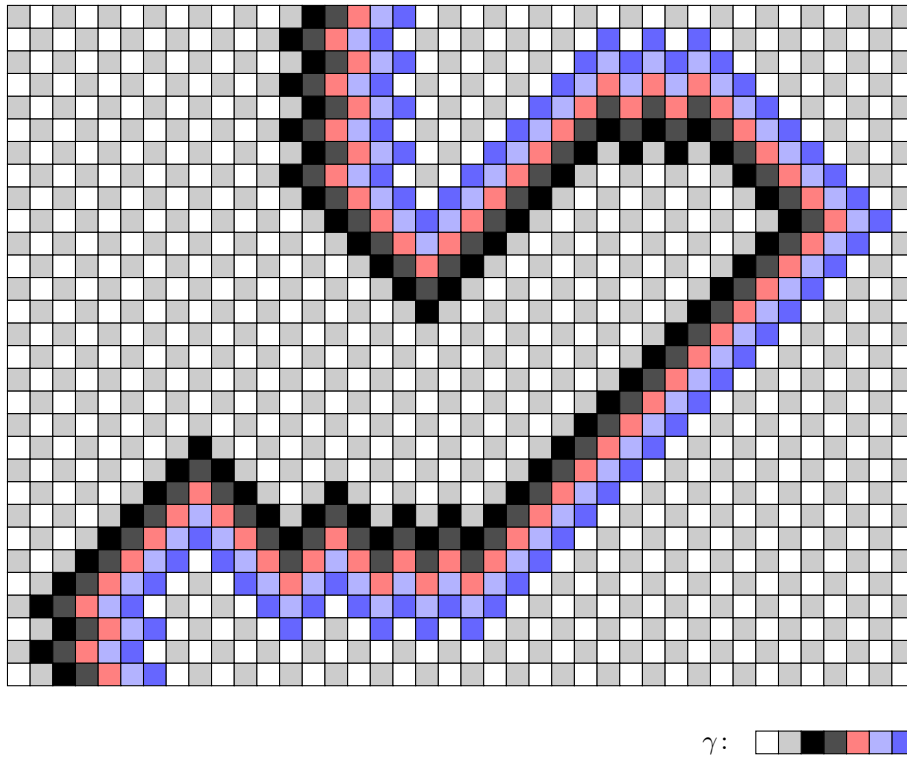


FIGURE 33. An illustration of how to “connect” two copies of γ using a winding path.

We now analyze the parity of occurrences of γ on the boundaries of the 12 tiles. For each tile, if we set the bottom-left corner to be the origin, then going around the boundaries in the counter-clockwise direction, we will see forwards copies of γ along the bottom and right boundaries and then backwards copies of γ along the top and left boundaries. As p, q are odd, the parity of the forwards copies of γ alternate, starting with the even parity for the first copy. The parity of the backwards copies

of γ also alternate, ending with the even parity for the last copy. Thus, if we pair up the first forwards copy of γ with the last backwards copy of γ , and then the second forwards copy of γ with the penultimate backwards copy of γ , etc., we will be able to apply Algorithm (C) to all these pairs. Obviously, we have to make sure that the connecting paths will not intersect each other in order for the algorithm to work, but this can be guaranteed by the fact that $p, q \gg m\ell > \ell + 1$, which guarantees enough room in between these connecting paths.

For the first eight tiles, this will result in all copies of γ being accounted for, and therefore these tiles will be completely colored with just nodes on γ . However, for the long tiles $T_{c^q a = ad^p}$, $T_{d^p a = ac^q}$, $T_{ca^q = b^p c}$, and $T_{cb^p = a^q c}$, there will be m unmatched copies of γ in the same orientation, with alternating parities. We will describe a different algorithm to handle these copies of γ .

The algorithm starts by collecting the remaining copies of γ and applying Algorithm (C) to make a single copy of γ^m . Next, we identify a rectangular region in the uncolored part of the long tile and connect this single copy of γ^m to a part of its boundary. It will be clear once we finish describing the algorithm that the dimensions of the rectangular region needed to perform the algorithm are much smaller than p, q and therefore such a region can be easily accommodated in the uncolored part of the long tile. We will call this rectangular region the *hub*. To be clear, the boundary of the hub consists of a copy of γ^m together with alternating v_0, v_1 . The hub will be completely filled to concretely manifest the fact that γ^m is null-homotopic.

It remains to describe the coloring within the hub. For definiteness we assume that the copy of γ^m occurs on the top boundary of the hub in the backwards orientation.

To describe the algorithm for filling the hub we need to define two auxiliary algorithms. The first one is called the *spur elimination* algorithm, or Algorithm (S), and the second one is the *4-cycle elimination* algorithm, or Algorithm (F). A *spur* is just a 2-cycle. Suppose that (u_0, u_1, u_0) is a 2-cycle in Γ , and three adjacent vertices $x_1, x_2 = (1, 0) \cdot x_1, x_3 = (2, 0) \cdot x_1$ in a grid graph have been colored by u_0, u_1, u_0 , in this order. Then to “eliminate” the 2-cycle (u_0, u_1, u_0) we just color the vertex above x_2 by u_0 . In general, if a cycle η is equivalent to the identity cycle in $\pi_1(\Gamma)$, then we can obtain the identity cycle from η by successively eliminating the spurs that occur in η . As an example, suppose a row of the grid graph has been colored by the cycle $\sigma = (u_0, u_1, u_2, u_3, u_2, u_4, u_2, u_1, u_0)$, which is equivalent to the identity cycle in $\pi_1(\Gamma)$. We can repeat the spur elimination scheme to accomplish the following coloring:

$$\begin{array}{cccccccc}
 & & & & & & & u_0 \\
 & & & & & & & (u_0 \ u_1 \ u_0) \\
 & & & & & & & u_0 \ (u_1 \ u_2 \ u_1) \ u_0 \\
 & & & & & & & u_0 \ u_1 \ (u_2 \ u_4 \ u_2) \ u_1 \ u_0 \\
 & & & & & & & u_0 \ u_1 \ (u_2 \ u_3 \ u_2) \ u_4 \ u_2 \ u_1 \ u_0
 \end{array}$$

In each step we marked out the spur that was eliminated. In general, suppose we have a rectangular region whose boundary consists of such an η plus alternating

is in fact the product of a conjugate of a 4-cycle with η . Now consider the cycle γ^m . Because γ^m is null-homotopic in $\pi_1^*(\Gamma)$, there are conjugates of 4-cycles $\lambda_1, \dots, \lambda_k$, all begin and end at v_0 , such that $\lambda_k \dots \lambda_1 \gamma^m$ is equivalent to the identity cycle in $\pi_1(\Gamma)$.

Inside the hub, we will color a row with nodes in $\lambda_k \dots \lambda_2 \lambda_1 \gamma^m$, along with fillers. Note that since m is even, the total parity change along $\lambda_k \dots \lambda_2 \lambda_1 \gamma^m$ is 0. Then above the row we perform the 4-cycle elimination algorithm for each of $\lambda_k, \dots, \lambda_1$, and connect γ^m to the boundary. Below the row we perform the spur elimination algorithm. The result of the coloring is depicted in Figure 35.




FIGURE 35. Coloring within a hub. We use conjugates of the four-cycles $\lambda_1, \lambda_2, \dots, \lambda_k$ to eliminate γ^m when $\lambda_k \dots \lambda_1 \gamma^m$ is equivalent to the identity in $\pi_1(\Gamma)$. The algorithm performed in each region is shown.

After performing all of these algorithms we fill the rest of the tiles by alternating v_0, v_1 . It is clear that we have described a complete coloring of the 12 tiles by nodes in Γ in a way that adjacent vertices are colored by adjacent nodes in Γ . This finishes the definition of a homomorphism $\Gamma_{1,p,q} \rightarrow \Gamma$. \square

The sufficient condition of Theorem 3.5.1 for the existence of a continuous homomorphism from $F(2^{\mathbb{Z}^2})$ to a graph Γ is a search problem suitable for a computer. For example, Theorem 3.5.1 applies to two well-studied and well-known graphs, the Chvátal graph and the Grötsch graph.

EXAMPLE 3.5.2. Let Γ_C be the Chvátal graph and Γ_G the Grötsch graph. There exist continuous graph homomorphisms $\varphi_C: F(2^{\mathbb{Z}^2}) \rightarrow \Gamma_C$ and $\varphi_G: F(2^{\mathbb{Z}^2}) \rightarrow \Gamma_G$.

PROOF. We can establish that each graph $\Gamma \in \{\Gamma_C, \Gamma_G\}$ has a cycle γ with $\pi_1^*(\gamma)$ of order 2 by coloring the nodes of a grid graph with nodes of Γ , placing γ at the top of the grid graph and γ^{-1} at the bottom, and leaving the sides alternating between two colors. The alternation between two colors will contribute nothing to

homotopy class of the boundary of the box, whence γ^2 will be null-homotopic in $\pi_1^*(\Gamma)$. The Chvátal and Grötzsch graphs are given along with their solutions in Figures 36 and 37 respectively. \square




FIGURE 36. The Chvátal graph, given at left. The odd cycle $\gamma = (0, 11, 7, 2, 1, 0)$ has order 2 in this graph, because there is a (necessarily null-homotopic) graph homomorphism from a grid graph to the Chvátal graph with boundary homotopic to γ^2 , given at right.




FIGURE 37. The Grötzsch graph. The odd cycle $\gamma = (0, 1, 2, 3, 9, 0)$ has order 2 for similar reasons as with the Chvátal graph.

It follows immediately from Theorem 3.1.2 that a simple condition sufficient to guarantee existence of continuous homomorphisms from $F(2^{2^2})$ to Γ is that Γ contains a copy of K_4 , the complete graph with 4-vertices. It is also easy to check directly with the same method as above that any 3-cycle in K_4 has order 2. Thus Theorem 3.5.1 applies also in this case. Note, however, that none of the above graphs contains a copy of K_4 . Therefore the positive condition in Theorem 3.5.1 is strictly weaker.

REMARK. According to Theorems 3.4.1 and 3.5.1, the negative and positive conditions for a graph Γ must necessarily be mutually exclusive. It may be reassuring to the reader that this mutual exclusiveness can be established by a simple direct argument, similar to what we noted in the proof of Theorem 3.5.1.

To see this, suppose Γ satisfies the positive condition of Theorem 3.5.1, and let γ be an odd-length cycle in Γ such that γ has finite order m in $\pi_1^*(\Gamma)$. Note that m must be even, as an odd-length cycle cannot be the identity element in $\pi_1^*(\Gamma)$. By raising γ to a suitable odd power, we may assume that $m = 2^{m_0}$. Let $\gamma = (v_0, v_1, \dots, v_\ell)$. We violate the negative condition of Theorem 3.4.1 in a strong way by showing that for all sufficiently large $p, q > m\ell$ with $\gcd(p, q) = 1$ there is a p -cycle α and a q -cycle β such that $\pi_1^*(\alpha^q) = \pi_1^*(\beta^p)$ in $\pi_1^*(\Gamma)$. We can assume without loss of generality that p is odd.

Suppose first that q is also odd. Let $0 < a, b < m$ be such that $a \equiv p \pmod{m}$ and $b \equiv q \pmod{m}$. Then a, b are odd and $aq - bp \equiv 0 \pmod{m}$. Let α be the cycle obtained by following γ^a with alternations of v_0, v_1 so that α has length p . This is possible since $p > a\ell$ and both are odd. Let β be obtained in a similar manner, extending γ^b by alternations of v_0, v_1 to have length q . Then $\pi_1^*(\alpha^q) = \pi_1^*(\gamma^{aq}) = \pi_1^*(\gamma^{bp}) = \pi_1^*(\beta^p)$ since $m|(aq - bp)$.

Next assume p is odd and q is even. Let 2^k be the highest power of 2 which divides q . If $k \geq m_0$ then we can take α to be the extension of γ by alternations of v_0, v_1 so that α has length p . Then $\pi_1^*(\alpha^q)$ is the identity, and the result is immediate if we let β be a q -cycle consisting only of alternations of v_0, v_1 . So assume $1 \leq k < m_0$. Let a, b be such that $0 < a < 2^{m_0-k}$ is odd, $0 < b < 2^{m_0}$ is an odd multiple of 2^k , and $a\frac{q}{2^k} - \frac{b}{2^k}p \equiv 0 \pmod{2^{m_0-k}}$. Then let α be the extension of γ^a by alternations of v_0, v_1 to have length p , and β the extension of γ^b to have length q . Then $\pi_1^*(\alpha^q) = \pi_1^*(\gamma^{aq}) = \pi_1^*(\gamma^{bp}) = \pi_1^*(\beta^p)$ since $2^{m_0}|(aq - bp)$.

Applications to Decidability Problems

4.1. Three general problems

We formulate three general problems for $F(2^{\mathbb{Z}^n})$. These problems can be considered at both the continuous and Borel levels, and are interesting in both contexts. We will, however, confine our attention to the continuous versions in this paper.

Recall from Definition 2.1.3 that a \mathbb{Z}^n -subshift of finite type is a closed invariant set $Y \subseteq \mathfrak{b}^{\mathbb{Z}^n}$ consisting of all $x \in \mathfrak{b}^{\mathbb{Z}^n}$ avoiding a prescribed set of patterns $\{p_1, \dots, p_k\}$. Thus, a subshift is determined by the tuple $\vec{p} = (\mathfrak{b}; p_1, \dots, p_k)$ which can be viewed as a tuple of integers via some coding. More explicitly, each pattern p as an n -dimensional matrix of integers can be coded by a single integer, which we still denote as p , via a computable procedure. Then the entire tuple $\vec{p} = (\mathfrak{b}; p_1, \dots, p_k)$ can be encoded as a single integer $m = \langle \mathfrak{b}, p_1, \dots, p_k \rangle$, where $\langle \cdot, \dots, \cdot \rangle : \omega^{<\omega} \rightarrow \omega$ is a fixed computable bijection. We let $Y_{\vec{p}}$ denote the subshift determined by \vec{p} . We also use Y_m to denote this subshift, where m is the integer $m = \langle \mathfrak{b}, p_1, \dots, p_k \rangle$ coding the subshift. The general subshift problem is the following.

PROBLEM (Subshift problem). For which $(\mathfrak{b}; p_1, \dots, p_k)$ does there exist a continuous (or Borel) equivariant map $\pi : F(2^{\mathbb{Z}^n}) \rightarrow Y_{\vec{p}}$? What is the complexity of the set of integers $m = \langle \mathfrak{b}, p_1, \dots, p_k \rangle$ coding subshifts Y_m for which there is such a continuous (or Borel) equivariant map? In particular, is this set computable?

A finite graph Γ can be coded by an integer, say by viewing the vertex set as $\{0, \dots, k\}$ for some k . This then suggests the general graph homomorphism problem.

PROBLEM (Graph homomorphism problem). For which finite graphs Γ does there exist a continuous (or Borel) graph homomorphism from $F(2^{\mathbb{Z}^n})$ to Γ ? What is the complexity of the set of integers m coding finite graphs Γ_m for which there is such a continuous (or Borel) graph homomorphism? In particular, is this set computable?

Recall that in Definition 3.1.7 we gave the definition of a clopen (or Borel) tiling of $F(2^{\mathbb{Z}^n})$ by a set of tiles $\{T_1, \dots, T_k\}$, where each T_i is a finite subset of \mathbb{Z}^n . We assume here that k is finite, and so the sequence of tiles can again be coded by an integer m . Partly restating Question 3.1.8 we have the following tiling problem.

PROBLEM (Tiling problem). For which sets of tiles $\{T_1, \dots, T_k\}$ does there exist a clopen (or Borel) tiling of $F(2^{\mathbb{Z}^n})$? What is the complexity of the set of integers m which code finite sets of tiles for which there is a clopen (or Borel) tiling of $F(2^{\mathbb{Z}^n})$? In particular, is this set computable?

In each of these three problems, when discussing the complexity of the set of integers m we of course must use a “lightface” notion of complexity as we are dealing with a set of integers. In each of the three cases, it is not immediately clear from the definitions that the complexity of the relevant set of integers is even Δ_1^1 (hyperarithmetical). For example, the set of integers $m = \langle \mathfrak{b}, p_1, \dots, p_k \rangle$ coding subshifts Y_m for which there is a continuous equivariant map from $F(2^{\mathbb{Z}^n})$ to Y_m seems to involve quantifying over reals coding continuous functions from $F(2^{\mathbb{Z}^n})$ to $\mathfrak{b}^{\mathbb{Z}^n}$. Superficially at least, this seems to compute the set as a Σ_2^1 set of integers.

Our first observation is that Theorem 2.3.5 shows that in all of the three problems, the relevant set of integers is actually Σ_1^0 (or computably enumerable, or semi-computable).

THEOREM 4.1.1. *For each of the subshift, graph homomorphism, and tiling problems for $F(2^{\mathbb{Z}^n})$, the set of integers m for which the continuous version of the problem has a positive answer is Σ_1^0 .*

PROOF. In each case, Theorem 2.3.5 shows that the continuous object exists (e.g., for the subshift problem, a continuous equivariant map into the subshift) if and only if a corresponding object exists for the graph $\Gamma_{n,p,q}$ for some $n < p, q$ where $\gcd(p, q) = 1$ (e.g., for the subshift problem, there is a map $g: \Gamma_{n,p,q} \rightarrow \mathfrak{b}$ which respects the subshift, see Definition 2.3.4). For given n, p, q , it is a finite problem to determine if such a g from $\Gamma_{n,p,q}$ exists. Thus, the set of m coding subshifts for which there is a continuous equivariant map from $F(2^{\mathbb{Z}^2})$ to the subshift is a Σ_1^0 set. \square

In the remaining sections of this chapter, we show that the one-dimensional subshift problem corresponds to a computable set, and that for the general subshift problem and the graph homomorphism problem, the corresponding sets are Σ_1^0 -complete, and in particular not computable. The exact degree of computability of the set corresponding to the tiling problem is unknown.

4.2. The subshift problem for $F(2^{\mathbb{Z}})$

We show first that the general subshift problem in dimension one, that is for $F(2^{\mathbb{Z}})$, is decidable. It follows immediately that the one-dimensional graph homomorphism problem and tiling problem are also decidable. In contrast, we will show in §4.3 that the subshift problem in dimensions two or higher is not decidable. In fact, we will show in §4.4 that even the specific instance of the graph homomorphism problem is not decidable for $F(2^{\mathbb{Z}^n})$ for $n \geq 2$. Recall a subshift of finite type $Y \subseteq \mathfrak{b}^{\mathbb{Z}}$ is determined by a sequence $(\mathfrak{b}; \ell; p_1, \dots, p_k)$ where $\ell \geq 1$ and $p_i \in \mathfrak{b}^\ell$ are the forbidden patterns of length ℓ (width $\ell - 1$).

THEOREM 4.2.1. *The set of integers m coding subshifts $Y_m \subseteq \mathfrak{b}^{\mathbb{Z}}$ for which there is a continuous, equivariant map from $F(2^{\mathbb{Z}})$ to Y_m is a Δ_1^0 (i.e., computable) set.*

The rest of this section is devoted to a proof of Theorem 4.2.1. We recall the one-dimensional version of the tile analysis, which was given in Theorem 2.2.2. Recall the graph $\Gamma_{n,p,q}^{(1)}$ from §2.2 which was used in that theorem. We use also the other notation of §2.2, for example the tile graphs $T_1(n, p, q), T_2(n, p, q)$ and the block label $R_\times^{(1)}$ associated to certain parts of $\Gamma_{n,p,q}^{(1)}$.

Given a subshift $Y \subseteq \mathfrak{b}^{\mathbb{Z}}$ determined by $(\mathfrak{b}; \ell; p_1, \dots, p_k)$, we define a directed graph $\Lambda = \Lambda(Y)$ as follows. Let the vertex set $V(\Lambda)$ of Λ be the $v \in \mathfrak{b}^{\ell}$ which are not in $\{p_1, \dots, p_k\}$. If $u, v \in V(\Lambda)$, we put a directed edge from u to v if $(v(0), \dots, v(\ell-2)) = (u(1), \dots, u(\ell-1))$. Let $(V(\Lambda), E(\Lambda))$ be the resulting directed graph. The vertices of Λ correspond to ℓ -size window states (that is possible values of $y \upharpoonright [k, k+\ell]$ for $y \in Y$) and the edge relation corresponds to a shift of the window one unit to the right.

We say $f: F(2^{\mathbb{Z}}) \rightarrow V(\Lambda)$ is a *directed graph homomorphism* if for all $x \in F(2^{\mathbb{Z}})$ we have that $(f(x), f(1 \cdot x)) \in E(\Lambda)$, where 1 is the generator for the action of \mathbb{Z} on $F(2^{\mathbb{Z}})$. We likewise define the notion of $g: \Gamma_{n,p,q}^1 \rightarrow V(\Lambda)$ being a directed graph homomorphism. The following fact is now immediate from the definitions.

LEMMA 4.2.2. *There is a continuous, equivariant map $f: F(2^{\mathbb{Z}}) \rightarrow Y$, where Y is the subshift determined by $(\mathfrak{b}; \ell; p_1, \dots, p_k)$, if and only if there is a continuous directed graph homomorphism $f': F(2^{\mathbb{Z}}) \rightarrow \Lambda(Y)$.*

A directed graph homomorphism $f': F(2^{\mathbb{Z}}) \rightarrow \Lambda(Y)$ is, in turn, a type of equivariant subshift map where the subshift has width 1. So, Lemma 4.2.2 and Theorem 2.2.2 immediately give the following.

LEMMA 4.2.3. *There is a continuous, equivariant map $f: F(2^{\mathbb{Z}}) \rightarrow Y$, where $Y \subseteq \mathfrak{b}^{\mathbb{Z}}$ is the subshift given by $(\mathfrak{b}; \ell; p_1, \dots, p_k)$, if and only if there are $p, q > 1$ with $\gcd(p, q) = 1$ and a directed graph homomorphism $g: \Gamma_{1,p,q}^{(1)} \rightarrow \Lambda(Y)$.*

It remains to show that it is computable in the integer m coding the subshift determined by $(\mathfrak{b}; \ell; p_1, \dots, p_k)$ whether there are $p, q > 1$ with $\gcd(p, q) = 1$ and a directed graph homomorphism $g: \Gamma_{1,p,q}^{(1)} \rightarrow \Lambda(Y)$.

Consider the directed graph $\Lambda = \Lambda(Y)$. Let $V = V(\Lambda)$ be the vertex set of Λ , and let $N = |V|$. For $u, v \in V$ say $u \sim v$ if there is a directed path in Λ from u to v and also a directed path from v to u . This is clearly an equivalence relation, and we call the equivalence classes C_1, \dots, C_t the *directed components* of Λ . Note that all the vertices in a directed cycle lie in the same directed component of Λ .

We say a directed cycle $\gamma = (u_0, u_1, \dots, u_k, u_0)$ in Λ is *primitive* if no subsegment $(u_i, u_{i+1}, \dots, u_j)$ is a cycle. Clearly a primitive cycle in Λ has length at most $N \leq \mathfrak{b}^{\ell}$.

LEMMA 4.2.4. *Let C be a directed component of Λ . Then the following are equivalent:*

- (1) *The gcd of the primitive cycle lengths in C is 1.*
- (2) *The gcd of the (directed) cycle lengths in C is 1.*
- (3) *The gcd of the lengths of the cycles beginning and ending at v_0 (a fixed vertex in C) is 1.*

PROOF. Clearly (1) implies (2). To see (2) implies (1), suppose $d > 1$ divides all the lengths of primitive cycles in C . We show d divides the length of all cycles in C , a contradiction. This is done by an induction on the length of the cycle. Let $\gamma = (u_0, \dots, u_k, u_0)$ be a cycle in C of length $|\gamma| = k + 1$. Either γ is primitive, and we are done, or there is a subsegment $\delta = (u_i, \dots, u_j)$, $i < j$, of γ which is a primitive cycle (so $u_j = u_i$). Then $\gamma' = (u_0, \dots, u_i, u_{j+1}, \dots, u_k, u_0)$ is a cycle of length $|\gamma| - (j - i)$. Since δ is primitive, by assumption d divides $|\delta| = j - i$. By induction d divides $|\gamma'|$, and so d divides $|\gamma|$.

Clearly (3) implies (2) and so it remains to show that (1) implies (3). Assume (1) and assume towards a contradiction that $d > 1$ divides all the lengths of the cycles which begin and end at v_0 . Let $a > 1$ be such that $d \nmid a$ and a is the length of a primitive cycle η in C , say starting and ending at $v_1 \in C$. Let α be a directed path in C from v_0 to v_1 , and β a directed path in C from v_1 to v_0 . Then the cycle $\rho = \alpha\eta\beta$ is a cycle starting and ending at v_0 with length $|\alpha| + |\eta| + |\beta| = |\eta| + |\delta|$ where δ is the cycle $\alpha\beta$ (which starts and ends at v_0). By assumption, d divides $|\delta|$ and $|\rho|$, and so d divides $|\eta| = a$, a contradiction. \square

We now have the following algorithm to determine if there is a directed graph homomorphism from $\Gamma_{1,p,q}^{(1)}$ to Λ . First, compute the directed components C_1, \dots, C_t of Λ . Next, for each of these components C , compute the gcd of the lengths of the primitive cycles in C . This is computable as these cycles have length $\leq N$. The following lemma finishes the algorithm.

LEMMA 4.2.5. *The following are equivalent:*

- (a) *There are $p, q > 1$ with $\gcd(p, q) = 1$ and a directed graph homomorphism $g : \Gamma_{1,p,q}^{(1)} \rightarrow \Lambda$.*
- (b) *There is a directed component C of Λ such that the gcd of the lengths of the primitive cycles in C is 1.*

PROOF. Suppose first that $g : \Gamma_{1,p,q}^{(1)} \rightarrow \Lambda$ is a directed graph homomorphism, where $\gcd(p, q) = 1$. Since $n = 1$, there is a unique vertex r in $\Gamma_{1,p,q}^{(1)}$ with label $R_\times^{(1)}$. Let $v = g(r)$. Let C be the directed component of v in Λ . Note that $T_1(1, p, q)$ has length $p + 1$ and $T_2(1, p, q)$ has length $q + 1$. As we move from the left to the right in $T_1(n, p, q)$, applying g gives a directed path in C of length p , starting and ending at v . Likewise, considering T_2 gives a directed path starting and ending at v of length q . Since $\gcd(p, q) = 1$, the gcd of all the cycle lengths in C is 1. By Lemma 4.2.4, the gcd of the lengths of the primitive cycles in C is 1.

Suppose next that the gcd of the lengths of the primitive cycles in the component C is 1. From Lemma 4.2.4 we may fix a vertex $v_0 \in C$ and directed cycles η_1, \dots, η_m starting and ending at v_0 such that $\gcd(|\eta_1|, \dots, |\eta_m|) = 1$. Any sufficiently large integer is a non-negative integral linear combination of the $|\eta_i|$. So, for any sufficiently large p, q (and we may take $\gcd(p, q) = 1$) there are directed cycles γ and δ of lengths p and q respectively which come from combinations of η_1, \dots, η_m and therefore start and end at v_0 . Using γ and δ we define g on T_1 and T_2 respectively in the obvious manner, namely, set $g(r) = v_0$, and follow γ to color the points of T_1 , and likewise follow δ for T_2 . This gives a directed graph homomorphism $g : \Gamma_{1,p,q}^{(1)} \rightarrow \Lambda$. \square

This completes the proof of Theorem 4.2.1.

4.3. The subshift problem for $F(2^{\mathbb{Z}^2})$

In this section we show that the general subshift problem for dimension at least 2 is Σ_1^0 -complete, and in particular undecidable. In §4.4 we will show a similar result for the graph homomorphism problem. Since the graph homomorphism problem can be viewed as a special case of the subshift problem, the main theorem of §4.4 will imply the main result of this section. However, we present a detailed proof in this section for two reasons. The first is that our proof here will be self-contained in

the sense that we will construct a computable reduction from the halting problem of Turing machines to the subshift problem. In contrast, the proof in §4.4 uses the Σ_1^0 -completeness of a special case of the word problem as a black box. The second reason is that our proof here establishes the undecidability of the subshift problem for subshifts $Y \subseteq \mathfrak{b}^{\mathbb{Z}^2}$ whenever $\mathfrak{b} \geq 4$. In contrast, the proof in §4.4 will require a much larger \mathfrak{b} .

THEOREM 4.3.1. *Let $\mathfrak{b} \geq 4$. The set of integers m coding subshifts $Y_m \subseteq \mathfrak{b}^{\mathbb{Z}^2}$ for which there is a continuous, equivariant map from $F(2^{\mathbb{Z}^2})$ to Y_m is a Σ_1^0 -complete set.*

We note that we do not have a similar result for the subshift problem for $\mathfrak{b} = 2$ or $\mathfrak{b} = 3$. In particular, it is unknown if the subshift problem for $\mathfrak{b} = 2$ is decidable.

The rest of this section is devoted to a proof of Theorem 4.3.1. Let $A \subseteq \omega$ be the set of integers $m = \langle \mathfrak{b}, p_1, \dots, p_k \rangle$ coding subshifts Y_m such that there is a continuous, equivariant map from $F(2^{\mathbb{Z}^2})$ to Y_m . We show A is Σ_1^0 -complete by defining a computable reduction from the halting problem to A .

4.3.1. Good Turing machines. We adopt some conventions regarding Turing machines. Our Turing machines operate on a bi-infinite memory tape divided into discrete cells. In a typical step, the machine positions its head over a cell and reads the symbol there. Then, according to the symbol and its present state, the machine looks up a table of instructions and takes action following the instruction applicable to the symbol read and the state. A typical action involves the machine writing a symbol in the cell followed by positioning its head either at the present cell or one cell to the left or one cell to the right. Upon completion of an action the machine changes its state to a new state following the instruction. After this the machine is ready to proceed to the next instruction or it can halt. The Turing machines we work with will have some additional features. First, they will all have the alphabet

$$\Sigma = \{B, S, E, 0, 1\}.$$

The symbol B indicates that the cell is blank. The symbol S is intended to mark the start of a region on the tape (the “data region”) where the machine will do its work. The symbol E marks the end of the working region. The symbols 0 and 1 are used by the Turing machine for its computations. Each of our Turing machines will have a set of states $Q = \{q_0, \dots, q_t\}$, with two special states q_0 and q_t . The state q_0 is the initial state of the machine at the beginning of the computation. The state q_t is the halting state. We think of each instruction of a Turing machine to be a quadruple

$$I = (q, \sigma, \alpha, q')$$

where $q, q' \in Q$, $\sigma \in \Sigma$, and α is a pair (τ, P) where $\tau \in \Sigma$ and $P \in \{-1, 0, 1\}$. In the end, a Turing machine is written as a finite set of instructions

$$M = \{(q_i, \sigma_i, \alpha_i, q'_i) : 1 \leq i \leq s\}.$$

Figure 38 illustrates a typical step in the computation of a Turing machine.

A *tape configuration* is an element of $\Sigma^{\mathbb{Z}}$ representing the content of the tape. The following definition imposes more restrictions on the Turing machines we would like to work with.




FIGURE 38. A typical step in the computation of a Turing machine

DEFINITION 4.3.2. Let M be a Turing machine. Let $p(i) \in \mathbb{Z}$ denote the position of the reading head of M on the tape at step i . Let z_i be the tape configuration of M at step i . We say that M is *good* if it has the following properties:

- (1) $p(0) = 0$; $z_0(0) = S$, $z_0(1) = E$, and $z_0(k) = B$ for all $k \neq 0, 1$.
- (2) For any $i \geq 0$, $p(i) \geq 0$; $z_i(0) = S$, $z_i(k) = B$ for all $k < 0$, and for any $k \neq 0$, $z_i(k) \neq S$.
- (3) For any $i \geq 0$, there is a unique $K > 0$ such that $z_i(K) = E$ and $z_i(k) = B$ for all $k > K$. We refer to this cell as the *ending cell* of the working region.
- (4) For any $i \geq 0$, if K is the ending cell of the working region, then $p(i) \leq K + 1$.
- (5) For any $i \geq 0$, if K is the ending cell of the working region, then $z_i(k) \in \{0, 1, E\}$ for all $0 < k < K$. Moreover, if $z_i(k) = E$ for any $0 < k < K$, then $k = K - 1$.
- (6) If M halts at step i , then $p(i) = 0$, $z_i(1) = E$, and $z_i(k) = B$ for all $k \neq 0, 1$.

In English, a good machine only uses the non-negative portion of the tape, and always maintains a working region with a starting cell with symbol S and an ending cell with symbol E . The symbol in the cell immediately preceding the ending cell could have a symbol E ; this is intended to address the situation where the machine needs to extend its working region and therefore needs to temporarily write a symbol E next to the previous ending cell. The point is that at any time during the computation, the working region should always end with an ending cell with symbol E . Except for this convention, the working region should always consist of 0s and 1s. During the computation, the reading head stays in the work region, with the only exception that it might be positioned one cell to the right of the ending cell (again, to accommodate extension of the working region). If a good machine ever halts, it halts in the state q_t with its head over the beginning cell, with the tape configuration identical to the initial configuration of the computation.

Although good Turing machines operate in a very restricted way, it is a standard fact that they are able to simulate any computation of Turing machines. For programming techniques for Turing machines the reader can consult standard references such as Chapter 8 of [18]. For example, the technique to keep all computations on a semi-infinite tape while requiring that the machine never prints the blank symbol B is presented in §8.5.1 of [18].

We use the usual Gödel numbering for Turing machines to assign a natural number to each Turing machine. Conversely, for each natural number $n \in \omega$ there is a corresponding Turing machine M_n . The standard programming techniques for Turing machines give that there is a total computable function $f: \omega \rightarrow \omega$ such that for all $n \in \omega$, $M_{f(n)}$ is a good Turing machine, and M_n halts if and only if $M_{f(n)}$ halts. We fix this computable function f for the rest of the proof. Our halting set

is

$$H = \{n: M_n \text{ halts}\} = \{n: M_{f(n)} \text{ halts}\}.$$

This is a Σ_1^0 -complete set of integers.

We will define another computable function $h: \omega \rightarrow \omega$ such that if n codes a good Turing machine M_n , then $h(n) = \langle \mathbf{b}, p_1, \dots, p_k \rangle$ codes a subshift $Y_{h(n)}$ of finite type, and M_n halts if and only if $h(n) \in A$. Then $h \circ f$ gives a computable reduction of H to A , and so A is Σ_1^0 -complete.

4.3.2. The canonical allowable pattern. For the rest of the proof we fix a good Turing machine $M = M_n$ and describe the subshift $Y = Y_{h(n)}$. It will be clear from the construction that the map h being implicitly defined is computable. We describe the subshift by giving the local rules an element $x \in \mathbf{b}^{\mathbb{Z}^2}$ must satisfy in order to be an element of Y (we are thereby defining the set of patterns p_i for the subshift).

The idea is that, for an element $x \in \mathbf{b}^{\mathbb{Z}^2}$ to be in the subshift, x must appear to be coding a valid computation sequence following the instructions of the Turing machine M . To describe the set of patterns for the subshift, we first define the *canonical allowable pattern* $\pi(M) \in \{0, 1, 3\}^{\mathbb{Z}^2}$ associated to M . Each horizontal line $\ell_j = \{(i, j) \in \mathbb{Z}^2: i \in \mathbb{Z}\}$, for an even $j \geq 0$, of $\pi(M)$ will code a valid tape configuration of M at a particular step of the computation, and the line above, ℓ_{j+1} , will contain a code for the state and the head position of M at this step. When two lines ℓ_j, ℓ_{j+1} , where $j \geq 0$ is even, in $\pi(M)$ code the information of a step during the computation, the two lines above, ℓ_{j+2}, ℓ_{j+3} , will code the information of the next step obtained by following the applicable instruction of M .

To implement this idea, we use the following coding scheme for the alphabet of M :

$$\begin{aligned} B &\mapsto 11111 \\ S &\mapsto 01110 \\ 0 &\mapsto 01000 \\ 1 &\mapsto 01100 \\ E &\mapsto 11110 \end{aligned}$$

It is easy to see that a bi-infinite string that codes a tape configuration of a good Turing machine this way is uniquely readable. To code the head position and the state of the Turing machine, we assume without loss of generality that the set of states $Q = \{q_0 = 1, q_1 = 2, \dots, q_t = |Q|\}$ consists of integers. We then use the string

$$000001^q00000$$

to code the state $q \in Q$, where 1^q denotes q many consecutive 1s. The obvious way to code the position is to line up the leading 5 consecutive 0s in the above code with the code for the symbol of the cell the head is reading. However, we will postpone this to address another important consideration.

The consideration is regarding the parities of positions of 1s in $\pi(M)$. To be precise, an element $(i, j) \in \mathbb{Z}^2$ has *even parity* if $i + j$ is even, and *odd parity* if $i + j$ is odd. The consideration is that we would like to make sure that in our definition of the canonical allowable pattern $\pi(M)$, all 1s appear in positions of the same parity. For this we use the method of *padding*, which is to append a 0 to each digit of the code, no matter whether it is 0 or 1. Equivalently, we apply the replacement

rule

$$0 \mapsto 00, \quad 1 \mapsto 10$$

to the results of above coding schemes. Now all lines of $\pi(M)$ coding tape configurations have 1s occurring only in even positions. It only remains to determine the exact position for the (padded) string coding the state of the machine. For this we first align the beginning of the string with the beginning position of the code for the cell the head is reading and then shift one unit to the right. Thus in all lines of $\pi(M)$ coding the state and the head position, all 1s occur in odd positions. Figure 39 illustrates the construction of $\pi(M)$; some spaces are added for readability.




FIGURE 39. From the computation of M to the canonical allowable pattern $\pi(M)$

Thus the canonical allowable pattern $\pi(M)$ will start with two lines ℓ_0, ℓ_1 coding the initial configuration (a single S followed by a single E), the initial state ($q_0 = 1$), and the initial head position (over the starting cell). Then $\pi(M)$ will extend upward as the computation of M proceeds, with every two lines representing a valid transition from the two lines below according to the applicable instructions of the machine. For lines below the starting configuration line, $\pi(M)$ will consist entirely of codes for the symbol B , with every other line shifted to the right by 1 to maintain the same parity of 1s. Note that so far $\pi(M)$ is a pattern of 0s and 1s, with all 1s occurring in positions of the same parity. If M does not halt, then the region of $\pi(M)$ not coding B s will extend upwards forever and will be infinite. In case M halts, this extension process will stop, at which time we will introduce a new symbol 3 (we reserve the symbol 2 for another purpose below) and add a 10×10 block of the symbol 3 directly above the 10 leading 0s in the code for the halting state q_t . For all regions that have not been covered by this definition, $\pi(M)$ will consist of alternating 0s and 1s so that the same parity of 1s is maintained. This finishes the definition of the canonical allowable pattern $\pi(M)$. Thus in case M halts, the part of $\pi(M)$ not coding B s will be a finite pattern with symbols 0, 1, and 3. Note that every line of the domain of $\pi(M)$ has a finite interval of non-blank codes, and that the non-blank portion of $\pi(M)$ is connected. Figure 40 illustrates the canonical allowable pattern $\pi(M)$ in case M halts.

Here is the intuition behind the introduction of the block of 3s for halting computations. In a general element of the subshift (to be defined), there might be different areas appearing to code valid computations of the machine. For each of these areas the 1s occur in positions with the same parity, but the parities of 1s for different areas do not necessarily match. These computation regions of (possibly) different parity will be separated by “threads” of the symbol 2. The block of 3s




FIGURE 40. A bird's eye view of the canonical allowable pattern $\pi(M)$ for a halting computation. Black areas represent codes for starting cells and ending cells. The gray area consists of 0s and 1s and the lined area consists of 3s.

is introduced to allow these parity changing threads to emanate/terminate, thus allowing any potential parity conflicts to be resolved. In case M halts, this will allow us to show, using Theorem 2.3.5, that there is a continuous, equivariant map into the subshift. Conversely, if there is a continuous equivariant map into the subshift, Theorem 2.3.5 will imply that halting configurations must exist.

4.3.3. Prototypical elements and the subshift of finite type. We now define the notion of a *prototypical element* for M . A prototypical element will be an element of $\{0, 1, 2, 3\}^{\mathbb{Z}^2} = 4^{\mathbb{Z}^2}$. These elements will be the basis for our definition of the subshift Y associated to M . The idea is that within regions separated by 2s, we appear to see a valid (coded and padded) Turing machine computation (with alternating rows giving the tape and state/position information as described previously).

DEFINITION 4.3.3. Let $\ell = \ell(M) = 100|Q|$. We say that $x \in 4^{\mathbb{Z}^2}$ is *prototypical* if for any $\ell \times \ell$ rectangle $R \subseteq \mathbb{Z}^2$, $x \upharpoonright R$ satisfies the following:

- (1) (no parity violations) There do not exist points $a, b \in R$ of different parities in \mathbb{Z}^2 with $x(a) = x(b) = 1$, and a path p from a to b in the undirected Cayley graph on R such that $x \upharpoonright p$ takes values in $\{0, 1\}$.
- (2) (no adjacent 2 colors) If $a, b \in R$ and $x(a) = x(b) = 2$, then a, b are not related by a generator of \mathbb{Z}^2 . Furthermore, if $x(a) = 2$, $j \leq 15$, and $x(a + (i, 0)) \notin \{2, 3\}$ for $i \leq j$, then the values $x(a + (1, 0)), \dots, x(a + (j, 0))$ are alternating 0s and 1s (starting with either a 0 or a 1). Similarly, if we look to the left of a up to 15 positions we see an alternating pattern of 0s and 1s until we see a 3.
- (3) (no Turing machine violations) If $D \subseteq R$ is the set of $a \in R$ such that $x(a) \in \{0, 1, 3\}$, then $x \upharpoonright D$ has no violations of the rules for the Turing machine M . Specifically, we mean the following.
 - (a) If I is a horizontal segment of $x \upharpoonright D$, then I is consistent with being a coded and padded sequence of a portion of the data line of a Turing machine, or the coded and padded sequence of a portion of the

state/position line of a Turing machine, or consistent with being a portion of a 10×10 region of 3s surrounded to the left and right by blanks.

- (b) If $I = (a, b) \times \{j\}$ is a horizontal segment of $x \upharpoonright D$ coding 0s and 1s, then $S = (a - 3, b + 3) \times [j - 3, j + 3] \cap R$ does not contain any 2s and is consistent with the rules of a coded and padded Turing Machine (including the rule about a 10×10 region of 3s after a halting line). Also, if I is an interval of 3s, then the two lines below this is a properly aligned halting configuration.

Thus, a prototypical $x \in 4^{\mathbb{Z}^2}$ has the property that in each $\ell \times \ell$ window, each of the regions bounded by the 2s appears to be a (coded and padded) instance of a computation according to M . Note that in (3) of Definition 4.3.3 we only require that x satisfies the rules of the Turing machine M , we do *not* require that it is part of the actual canonical allowable pattern $\pi(M)$.

We are now ready to define the subshift Y associated to a good Turing machine M .

DEFINITION 4.3.4. Let $\ell = \ell(M) = 100|Q|$ as before. Let W be the set of all $\ell \times \ell$ patterns which occur in some prototypical element of $4^{\mathbb{Z}^2}$. Define Y to be the subshift of $4^{\mathbb{Z}^2}$ determined by (p_1, \dots, p_k) , where the p_i are all $\ell \times \ell$ patterns not in W .

The prototypical elements $x \in \mathfrak{b}^{\mathbb{Z}^2}$ are precisely the elements of the subshift Y . It is obvious from the definition that Y has width $\ell - 1$.

In the following we prove some basic facts about prototypical elements.

LEMMA 4.3.5. *Let $x \in 4^{\mathbb{Z}^2}$ be a prototypical element for the Turing machine M and suppose that $x(a) \neq 3$ for all $a \in \mathbb{Z}^2$. Then there is a proper 3-coloring $x' \in 3^{\mathbb{Z}^2} = \{0, 1, 2\}^{\mathbb{Z}^2}$ such that if $x'(a) \neq x(a)$, then $x(a) = 0$ and $x'(a) = 1$. Moreover, for every $a \in \mathbb{Z}^2$, whether $x'(a) \neq x(a)$ is determined by the values of x in the region $\{b \in \mathbb{Z}^2 : \|a - b\| \leq 20\}$.*

PROOF. Let $D = \{a \in \mathbb{Z}^2 : x(a) \in \{0, 1\}\}$. We claim that in any connected component C of D that the positions of 1s all have the same parity. To see this claim, let $a, b \in \mathbb{Z}^2$ be in the same component C of D . Let p be a path in C connecting a and b . So, $x(c) \in \{0, 1\}$ for all $c \in p$.

We show a subclaim that for any $c \in p$ there is $d \in \mathbb{Z}^2$ with $\|c - d\| \leq 20$ such that $x(d) = 1$ and furthermore there is a path from c to d staying within $C \cap B(c, 20)$ (the points of distance ≤ 20 from c). Let c_1 be the first point to the left of c for which $x(c_1) = 2$ (if this does not exist, the point d can be taken to the left of c and within distance 11 of c since the maximum number of consecutive 0s in a horizontal line of a prototypical element is 11 and there are no 3s in x). Likewise define c_2 to the right of c . If $\|c_1 - c_2\| \geq 13$, then the above fact that the maximum number of consecutive 0s in a horizontal line of a prototypical element is 11 gives the result. If $3 \leq \|c_1 - c_2\| \leq 12$, then condition (2) of Definition 4.3.3 implies there is a 1 in the horizontal interval between c_1 and c_2 and we are done. Finally, suppose c_1 and c_2 are the points immediately to the left and right of c . The path p must move from c to the point immediately above or below c . Say, for example, p moves from c to c' , the point immediately above c . Let c'_1, c'_2 be the points immediately above c_1 and c_2 . By (2) of Definition 4.3.3, $x(c'_1)$ and $x(c'_2)$ are

not equal to 2. Repeating the argument given for c_1 and c_2 , shows that there is a 1 on the horizontal segment between c'_1 and c'_2 , which proves the subclaim.

For each point $c \in p$, let $d(c)$ be a point of C connected to c by a path in C of length ≤ 20 with $x(d(c)) = 1$. Since $\ell > 40$, an easy induction using (1) of Definition 4.3.3 shows that all of the points $d(c)$ have the same parity. In particular, a and b have the same parity. This proves the claim.

We define x' as follows. Let $x'(a) = x(a)$ if $x(a) = 1$ or 2. If $x(a) = 0$, then consider the connected component C of D containing a . If a has the same parity as the positions of 1s in C , then let $x'(a) = 1$; otherwise let $x'(a) = x(a) = 0$. Using condition (2) of Definition 4.3.3 and the above claim, it is easy to see that x' is a proper 3-coloring.

To see the “moreover” part of the lemma, note that by the subclaim, for every $a \in \mathbb{Z}^2$ such that there is b in the connected component C of a in the subgraph induced by D with $x(b) = 1$, there is such a b with $\|a - b\| \leq 20$ and a path p from a to b such that $x \upharpoonright p$ takes values in $\{0, 1\}$. Thus $x'(a) \neq x(a)$ exactly when a and b are of the same parity and $x(a) = 0$. By our claim this does not depend on the choice of the point b and the path p . Thus the values of x in the region $\{b \in \mathbb{Z}^2 : \|a - b\| \leq 20\}$ completely determine whether $x'(a) \neq x(a)$. \square

Fix for the moment n, p, q with $p, q > n \geq \ell - 1$, where ℓ is as in Definition 4.3.3 for the good Turing machine M . Let $H_1 = G_{d^p a = ac^q}$, $H_2 = G_{c^q a = ad^p}$ be the labeled grid graphs for the two horizontal long tiles used in defining $\Gamma_{n,p,q}$ as shown in Figure 10. Let $H = H_{n,p,q}$ be the quotient graph of the disjoint union of H_1 and H_2 by identifying corresponding vertices of the blocks having the same label, as in the definition of $\Gamma_{n,p,q}$ (see §2.3). Thus, H is an induced subgraph of the graph $\Gamma_{n,p,q}$. Note that we may also obtain H by first stacking H_1 on top of H_2 so that the labeled blocks along the bottom side of H_1 coincide with the corresponding labeled blocks along the top side of H_2 , and then identifying corresponding vertices as in $\Gamma_{n,p,q}$. If z is a map from the vertex set $V(H)$ of H to $4 = \{0, 1, 2, 3\}$, we say z is prototypical if for any $\ell \times \ell$ rectangle R contained in one of the grid graphs H_1 or H_2 , $z \upharpoonright R$ satisfies (1)–(3) of Definition 4.3.3.

From Lemma 4.3.5 and Theorem 3.1.1 we have the following.

LEMMA 4.3.6. *Let $p, q > n \geq \ell - 1$, where ℓ is as in Definition 4.3.3 for the good Turing machine M , and assume $\gcd(p, q) = 1$. Let $H_{n,p,q}$ be as defined above. If $z: V(H_{n,p,q}) \rightarrow 4$ is prototypical, then $3 \in \text{range}(z)$.*

PROOF. Suppose $z: V(H_{n,p,q}) \rightarrow 4$ is prototypical and $z(c) \neq 3$ for all vertices c of $H_{n,p,q}$. Let \bar{x} be the element of $3^{\mathbb{Z}^2}$ obtained by tiling \mathbb{Z}^2 with copies of H_1 and H_2 in the manner illustrated in Figure 41. That is, we stack H_1 and H_2 alternatively, with the labeled blocks along the bottom side of a copy of H_1 coinciding with the labeled blocks along the top side of the copy of H_2 below it, and similarly the labeled blocks along the bottom side of a copy of H_2 coinciding with the labeled blocks along the top side of the copy of H_1 below it. Also, neighboring columns share their labeled blocks on the sides.

Since $n \geq \ell - 1$, any $\ell \times \ell$ region is properly contained in a copy of H_1 or H_2 , and it follows that \bar{x} is a prototypical element of $3^{\mathbb{Z}^2}$. Let x' be obtained from \bar{x} as in Lemma 4.3.5. So x' is a proper 3-coloring of \mathbb{Z}^2 and x' is obtained from \bar{x} by only changing some 0 values of x to 1. Moreover, for any $a \in \mathbb{Z}^2$ whether $x'(a) \neq \bar{x}(a)$ is determined by the values of x in the region $\{b \in \mathbb{Z}^2 : \|a - b\| \leq 20\}$.




FIGURE 41. Tiling of \mathbb{Z}^2 with H_1 and H_2 .

Since $n \geq \ell - 1 > 2 \cdot 20$, from the periodic nature of \bar{x} it follows that x' also respects the identified vertices in the tiling (e.g., for two elements a, a' of \mathbb{Z}^2 in the same position of similarly labeled blocks we have $x'(a) = x'(a')$). Thus, x' induces a proper 3-coloring of the graph $H_{n,p,q}$. In particular, this gives a proper 3-coloring for each of the long horizontal tiles $H_1 = G_{d^p a = ac^q}$ and $H_2 = G_{c^q a = ad^p}$. This contradicts Theorem 3.1.1. \square

4.3.4. Proof of Theorem 4.3.1. Let M be a good Turing machine and Y the subshift associated to M as in Definition 4.3.4. To finish the proof of Theorem 4.3.1 it suffices to show that M halts if and only if there is a continuous, equivariant map from $F(2^{\mathbb{Z}^2})$ to Y .

First suppose that M halts in N steps. Fix $n > \ell$ and $p, q > \max\{2n, 100N\}$ with $\gcd(p, q) = 1$. We may also assume that n, p, q are odd, and $q = p + 2$. From Theorem 2.3.5 it suffices to find $g: \Gamma_{n,p,q} \rightarrow 4$ which respects the subshift Y . We describe such a g as a coloring of $\Gamma_{n,p,q}$ with colors $\{0, 1, 2, 3\}$.

We use the following terminology. For a subset $D \subseteq \mathbb{Z}^2$, a *checkerboard pattern* on D is a proper 2-coloring with colors $\{0, 1\}$. If D is a region connected by generators of \mathbb{Z}^2 , there are only two checkerboard patterns possible on D . For each of them the positions of 1s are of the same parity, and the parities of 1s for the two patterns are the opposite. By a *zig-zag path* in \mathbb{Z}^2 we mean a sequence u_0, u_1, \dots of points such that $u_{i+1} = (\pm 1, \pm 1) + u_i$ for all i . Note that all points on such a path have the same parity, in particular no two points of the path are connected by a generator of \mathbb{Z}^2 .

We describe the map g by first describing the values of g on the regions with labels R_\times, R_a, R_b, R_c , and R_d . For R_\times , we use a zig-zag path connecting the corner points of the region to each other as shown in Figure 42. This is possible as n is odd. The points on the path are given color 2. As demonstrated in Figure 42, the path created an interior region. This region is given a checkerboard pattern. Finally, the points on the top and right edges of R_\times are given color 0, and the points on the other two edges are given color 1.




FIGURE 42. The map g on the R_\times -labeled region.

Each of the R_a -labeled and R_b -labeled regions is given a checkerboard pattern with the upper-left corner point getting a color 1. It is easy to check that this is consistent. Because n, p, q are odd, points in the top row of the R_a -labeled region get different colors than their adjacent points in the bottom row of the R_x -labeled region, and points in the bottom row of the R_a -labeled region get different colors than their adjacent points in the top row of the R_x -labeled region. Similarly for the R_b -labeled regions.

We color the R_c -labeled and R_d -labeled regions in two steps. In the first step, each of these regions is given a background checkerboard pattern with the upper-left corner point getting a color 0. As above, the coloring of the R_x -labeled region and the fact that p, q are odd give that this is a consistent proper two-coloring of these regions. Next, we override the background coloring by inserting in each R_c -labeled and R_d -labeled region a portion of the canonical allowable pattern $\pi(M)$ that corresponds to the beginning steps of the computation of M , with the positions of 1s in the inserted piece having the same parity as the positions of 1s in the background. For example, we may insert the first two lines of $\pi(M)$ corresponding to the initial step of the computation into the top two rows of R_c and R_d . Since $p - n, q - n > n \geq \ell - 1$, there is enough room in R_c and R_d for the insertion. We will make sure that the code for the starting cell appears near the horizontal center of the region. Note that the width of the codes for the two initial steps of the computation is no more than $10N$ (since each step can access only one more cell and each symbol is coded by a 0,1-string with ten digits). By placing the code for the starting cell around the center, we ensure that the codes for the entire computation, when extended beyond the R_c -labeled or R_d -labeled regions, will stay in a region directly above the R_c -labeled or R_d -labeled region. This is illustrated in Figure 43.




FIGURE 43. Mapping $G_{c^q a=ad^p}$ into the subshift when M halts.

Note that the coloring we have defined so far respects the subshift Y , since each $\ell \times \ell$ pattern appearing in the coloring can be extended to a prototypical element of $4^{\mathbb{Z}^2}$.

Finally, it remains to color the “interior” parts of the tiles of $\Gamma_{n,p,q}$. As these points are not identified with any other points in forming the quotient $\Gamma_{n,p,q}$, we only need to make sure that the resulting coloring of each tile respects the subshift

Y separately. We consider the long tile $G_{c^q a = ad^p}$, the other cases being similar. Figure 43 shows the coloring of the long tile which defines $g \upharpoonright G_{c^q a = ad^p}$.

In $G_{c^q a = ad^p}$ we extend each copy of the codes for the entire computation from the R_d -labeled regions along the bottom side into the interior. This is possible as we noted that $p, q > 100N$, whereas the width of the codes for each step of the computation is no more than $10N$. Recall that at the top of each copy of the codes for the entire computation there is a 10×10 rectangular region consisting of color-3 points (shown by circles in Figure 43). We run zig-zag color-2 paths (shown in green in Figure 43) as shown in the figure. One end of each of these paths terminates in a corner point of an R_x -labeled region, which sets its parity. The other end terminates in one of the color-3 rectangular regions. Thus, we maintain the subshift conditions required by Y . Note that since $q = p + 2$, there are at most 4 color-2 paths needing to terminate in any given color-3 region, and so the 10×10 size of the color-3 regions is enough to accommodate them. Note that the color-2 paths isolate regions of different parity in $G_{c^q a = ad^p}$, and there is no problem filling in each of the isolated regions by a checkerboard pattern of 0s and 1s.

The key point in the construction is that the color-2 paths, which are needed to isolate regions of different parity in the tile, are free to terminate arbitrarily in the special color-3 regions. From the construction we easily see that the map $g: \Gamma_{n,p,q} \rightarrow 4$ defined above respects the subshift Y .

Next suppose that M does not halt. From Theorem 2.3.5 it is enough to fix arbitrary $n \geq \ell - 1$ and $p, q > n$ and show that there does not exist $g: \Gamma_{n,p,q} \rightarrow 4$ which respects the subshift Y . Suppose that such a g exists. Restricting G to the horizontal long tiles gives a $g': H_{n,p,q} \rightarrow 4$ which respects the subshift, that is, g' is a prototypical element in the sense of Lemma 4.3.6. From Lemma 4.3.6 we have that $3 \in \text{range}(g')$. In particular, g restricted to some horizontal long tile must attain the color 3.

Say $g \upharpoonright G_{c^q a = ad^p}$ takes the value 3 (the other case is similar). The rules of the subshift require that a point of color 3 occurs in a rectangular region of color 3 points which is immediately above a line representing a halting configuration of M . Consider the diagram shown in Figure 44.




FIGURE 44. Problem with having map into subshift when M does not halt.

This diagram consists of the tile $G_{c^q a = ad^p}$ with a copy of $G_{d^p a = ac^q}$ below, and another copy of $G_{c^q a = ad^p}$ below that. Both copies of $G_{c^q a = ad^p}$ have a rectangular grid of color 3 points at the same location, as shown in the figure. The rules of the subshift require that every other line below the top region of color 3 represents a valid tape configuration consisting of a starting cell and an ending cell. The codes for the starting cells are vertically aligned as shown in the figure. Each of these lines is obtained from two lines below by an applicable instruction of the machine M . Since M does not halt, it cannot be the case that as we move down from the top region of color 3 we reach a line corresponding to an initial configuration (as this would give a valid computation of M which halted). The rules of the subshift then imply that two lines below each line correspondent to a tape configuration is another line correspondent to a tape configuration, with starting and ending cells. This, however, contradicts that the region of color 3 points occurs in the bottom copy of $G_{c^q a = ad^p}$ in the same location, since it violates the subshift requirements.

This contradiction shows that when M does not halt that there is no continuous, equivariant map into the subshift Y , which completes the proof of Theorem 4.3.1.

4.4. The graph homomorphism problem

Note that the subshift problem (stated in §4.1) is equivalent to a problem regarding the existence of \mathbb{Z}^n -homomorphisms from $F(2^{\mathbb{Z}^n})$ to a finite \mathbb{Z}^n -graph. In fact, for each subshift of finite type $Y_{\vec{p}}$, where $\vec{p} = (\mathbf{b}; p_1, \dots, p_k)$, we can associate a finite \mathbb{Z}^n -graph Γ as follows. For a sufficiently large positive integer a (we can take any a larger than the width of Y , see Definition 2.1.3 and the preceding discussion), let $V(\Gamma)$ be the set of all patterns p with $\text{dom}(p) = [0, a]^n$ in which none of p_1, \dots, p_k occurs. For a generator s of \mathbb{Z}^n and patterns $p, q \in V(\Gamma)$, let $(p, q) \in E(\Gamma)$ with label s if $s \cdot p$ is compatible with q . This defines a \mathbb{Z}^n -graph Γ , and it is straightforward to check that there is a continuous equivariant map from $F(2^{\mathbb{Z}^n})$ to $Y_{\vec{p}}$ if and only if there is a continuous \mathbb{Z}^n -graph homomorphism from $F(2^{\mathbb{Z}^n})$ to Γ . Thus Theorem 4.3.1 implies that the existence of a continuous \mathbb{Z}^2 -homomorphism from $F(2^{\mathbb{Z}^2})$ to a finite \mathbb{Z}^2 -graph is Σ_1^0 -complete, and in particular undecidable.

In this subsection we prove a result which answers the graph homomorphism problem, Problem 4.1. As with the subshift problem, Problem 4.1 and Theorem 4.3.1, we show that the graph homomorphism problem is not decidable. This is done via an undecidable word problem for finitely presented groups.

We view finite graphs as coded by integers in a standard manner.

THEOREM 4.4.1. *The set of finite graphs Γ such that there is a continuous graph homomorphism from $F(2^{\mathbb{Z}^2})$ to Γ is a Σ_1^0 -complete set.*

The rest of this section is devoted to a proof of Theorem 4.4.1.

We will define a computable reduction from a variation of the word problem for finitely presented groups to the graph homomorphism problem. Specifically, we use the word problem for torsion-free groups, which is known to be undecidable (see Lemma 1 of [1]). So, there is a Σ_1^0 -complete set $A \subseteq \omega$ and a computable function $f: \omega \rightarrow \omega$ such that for all n , $f(n)$ is an integer coding

- a finite group presentation $\mathcal{G}_n = \langle a_1, \dots, a_k | r_1, \dots, r_l \rangle$ for a torsion-free group G_n , and
- a distinguished word $w = w(a_1, \dots, a_k)$

(in our notation we suppress the dependence of k, l , the r_i , and w on n), such that $n \in A$ if and only if w is the identity e in $G_n = \langle a_1, \dots, a_k | r_1, \dots, r_l \rangle$. We will need one extra property of the groups G_n :

- (*) There is (a lower bound) $\alpha > 0$ such that, if the distinguished word w is not the identity in G_n , then for all integers $m \geq 1$, w^m is not equal in G_n to any word of length at most αm .

The property (*) for the groups G_n is easy to arrange. If $\mathcal{G} = \langle a_1, \dots, a_k | r_1, \dots, r_l \rangle$ is a finite presentation of a group G with distinguished word $w = w(\vec{a})$, let

$$\tilde{\mathcal{G}} = \langle a_1, \dots, a_k, b_1, \dots, b_k | r_1(\vec{a}), \dots, r_l(\vec{a}), r_1(\vec{b}), \dots, r_l(\vec{b}) \rangle,$$

and let $\tilde{w} = w(\vec{a})w(\vec{b})$. \tilde{G} is just the free product of G with itself. If $w = e$ in G , then clearly $\tilde{w} = e$ in \tilde{G} . If $w \neq e$ in G , we also have that $\tilde{w}^m = w(\vec{a})w(\vec{b}) \cdots w(\vec{a})w(\vec{b})$ is not equal to any word of length $< 2m$ by the normal form theorem for free products (see Theorem 2.6 of Chapter IV of [23]). So, (*) is satisfied with any $\alpha < 2$. Note that if G is torsion-free, then so is \tilde{G} . So, without loss of generality we henceforth assume that all of the groups G_n are torsion-free and satisfy (*).

The following technical theorem gives an outline for the rest of the proof.

THEOREM 4.4.2. *Let $\mathcal{G} = \langle a_1, \dots, a_k | r_1, \dots, r_l \rangle$ be a presentation of a torsion-free group G and let $w = w(a_1, \dots, a_k)$ be a word. Suppose that the pair (\mathcal{G}, w) satisfies (*). Then there is a finite connected graph Γ with the following properties:*

- (i) $\pi^*(\Gamma) \cong \langle a_1, \dots, a_k, z | r_1, \dots, r_l, z^2w^{-1} \rangle$, where z is a new letter distinct from a_1, \dots, a_k ;
- (ii) if $w = e$ in G , then Γ satisfies the conditions of Theorem 3.5.1, and hence there is a continuous homomorphism from $F(2^{\mathbb{Z}^2})$ to Γ ;
- (iii) if $w \neq e$ in G , then Γ satisfies the conditions of Theorem 3.4.1, and hence there is no continuous homomorphism from $F(2^{\mathbb{Z}^2})$ to Γ .

Furthermore, one can construct Γ from \mathcal{G} using a computable procedure.

4.4.1. Realizing a given reduced homotopy group. In (i) of Theorem 4.4.2 we need to construct a finite connected graph with a prescribed reduced homotopy group. In this subsection we give the procedure to do this as well as make some observations about the procedure that we need to use later in the proof. Our main objective is to show the following general fact.

LEMMA 4.4.3. *Let $\mathcal{G} = \langle b_1, \dots, b_k | s_1, \dots, s_l \rangle$ be a finite presentation for a group G . Then there is a finite connected graph $\Gamma(\mathcal{G})$ such that $\pi_1^*(\Gamma(\mathcal{G})) \cong G$. Moreover, the map $\mathcal{G} \mapsto \Gamma(\mathcal{G})$ is computable.*

PROOF. This fact can be seen as an instance of a general result about CW complexes associated to group presentations, as for example in Proposition 2.3 of [23]. For the sake of completeness, and since we need to record a few extra points, we briefly sketch a proof.

The graph $\Gamma(\mathcal{G})$ is constructed as follows. $\Gamma(\mathcal{G})$ will have a distinguished vertex v_0 . For each of the generators b_i , $1 \leq i \leq k$, we add a cycle β_i of some length $\lambda_i > 20$ that starts and ends at the vertex v_0 . Other than the common vertex v_0 , we make the vertex sets of these cycles pairwise disjoint. This gives a natural notion of length $\lambda(b_i) = \lambda_i$ which extends in the obvious manner to reduced words in the free group generated by the b_i . For each word s_j , $1 \leq j \leq l$, from the presentation

\mathcal{G} , we add to $\Gamma(\mathcal{G})$ a rectangular grid graph R_j whose length and width are both greater than 4 and whose perimeter $\text{per}(R_j)$ is equal to $\lambda(s_j)$. This is possible because each $\lambda(b_i)$ is a large enough even number. The vertices in the interiors of these grid graphs R_j are pairwise disjoint, and are disjoint from the vertices of β_i . The vertices on the sides of each R_j are assigned so that v_0 occurs as the upper-left corner of each R_j and, going clockwise, the boundary of R_j coincides with the concatenation of the paths corresponding to the generators in the word s_j . This completes the definition of the graph $\Gamma(\mathcal{G})$.

It is obvious from this definition that $\Gamma(\mathcal{G})$ is a finite connected graph and that the map $\mathcal{G} \mapsto \Gamma(\mathcal{G})$ is computable.

Next we verify that $\pi_1^*(\Gamma(\mathcal{G})) \cong G$. For this we first introduce a “normal form” for each cycle in $\Gamma(\mathcal{G})$ that starts and ends at v_0 . Consider a cycle χ in $\Gamma(\mathcal{G})$ starting and ending at v_0 . Say $\chi = (v_0, v_1, \dots, v_M)$, where $v_0 = v_M$ and the v_m are vertices of $\Gamma(\mathcal{G})$. Consider a segment $\sigma = (v_{m_0}, \dots, v_{m_1})$ of this path where v_{m_0}, v_{m_1} are vertices in the union of the β_i , and for $m_0 < m < m_1$, v_m is an interior vertex of some R_j . Note that the value of j is uniquely determined by σ . Let $\bar{\sigma}$ be the result of replacing in σ the segment $(v_{m_0+1}, \dots, v_{m_1-1})$ with a path along the perimeter of R_j (say, choose the shortest path, and if there are two of equal length, choose one randomly). Doing this on each such segment σ of χ gives a new path $\bar{\chi}$. Clearly χ is reduced homotopy equivalent to $\bar{\chi}$ in $\pi_1^*(\Gamma(\mathcal{G}))$. Note that $\bar{\chi}$ corresponds to a word in the b_i , which we still denote by $\bar{\chi}$. It is not difficult to check that if χ differs from another cycle ψ by the insertion/deletion of a 4-cycle in $\Gamma(\mathcal{G})$, then $\bar{\chi}$, as a word, is equivalent to $\bar{\psi}$ in G . We use here the fact that if all $\lambda_i > 20$ and each R_j has width and height greater than 4, then the only 4-cycles in $\Gamma(\mathcal{G})$ are the ones coming from individual R_j . It follows that, if χ is reduced homotopy equivalent to e in $\pi_1^*(\Gamma(\mathcal{G}))$, then $\bar{\chi}$, as a word in G , is either trivial or a combination of s_j .

We are now ready to argue that $\pi_1^*(\Gamma(\mathcal{G})) \cong G$. If F is the free group generated by the b_i , then there is a natural homomorphism φ from F onto $\pi_1^*(\Gamma(\mathcal{G}))$ obtained by sending b_i to the equivalence class $[\beta_i]$ of the cycle β_i in $\pi_1^*(\Gamma(\mathcal{G}))$. Clearly each s_j is in the kernel of φ , and so φ induces a homomorphism φ' from G onto $\pi_1^*(\Gamma(\mathcal{G}))$. To see that the kernel of φ' is trivial, it suffices to show that the kernel of φ is generated by s_j in F . Suppose u is a word in the b_i and $\varphi(u) = e$. Since u gives a path χ starting and ending at v_0 in $\Gamma(\mathcal{G})$, $\varphi(u) = e$ means that χ is reduced homotopy equivalent to the trivial path in $\pi_1^*(\Gamma(\mathcal{G}))$. By the argument in the preceding paragraph, this implies that $\bar{\chi}$ is either trivial or a combination of s_j . However, since χ does not contain any interior vertices of R_j , we have $\chi = \bar{\chi}$. This means that u is in the subgroup generated by s_j . This completes the proof of Lemma 4.4.3. \square

Applying Lemma 4.4.3 to the presentation demonstrated in (i) of Theorem 4.4.2 would establish that clause. However, in order to continue with our proof we need to use some details of the construction of $\Gamma(\mathcal{G})$ as well as some additional properties. We specify these additional properties below.

Given a presentation $\mathcal{G}_n = \langle a_1, \dots, a_k | r_1, \dots, r_l \rangle$ for a group G_n , let G'_n be the finitely presented group with presentation given by

$$\mathcal{G}'_n = \langle a_1, \dots, a_k, z | r_1, \dots, r_l, z^2 w^{-1} \rangle,$$

that is, we add the extra generator z and the extra relation $z^2 = w$ to the presentation for G_n . For notational convenience, we denote $a_0 = z$ and $r_0 = z^2 w^{-1}$.

Since every word in the normal subgroup of G'_n generated by the r_0, r_1, \dots, r_ℓ has an even number of occurrences of z , the parity of the number of occurrences of z in an element of G'_n is well-defined. We accordingly refer to an element $g \in G'_n$ as being either odd or even.

When associating a finite connected graph $\Gamma(\mathcal{G}'_n)$ to the presentation \mathcal{G}'_n , we deviate slightly from the construction in the above proof. The deviation is that, when assigning a cycle β_0 to the generator z , we let $\lambda_0 = \lambda(z) > 10$ be an *odd* number and ensure that the length of β_0 is λ_0 . One can easily check that this does not invalidate the proof as the only occurrences of z in the relations are in $r_0 = z^2 w^{-1}$, and in order to carry out the proof we only needed all $\lambda(r_j)$ to be sufficiently large and even.

In addition, we note that there is no loss of generality to require that the λ_i are large and that the aspect ratios of R_j are less than 2. This latter property can be achieved by choosing the width of R_j , denoted $\mathfrak{w}(R_j)$, and the height of R_j , denoted $\mathfrak{h}(R_j)$ to be close to $\frac{1}{4}\text{per}(R_j)$, subject to the requirement that $\mathfrak{w}(R_j) + \mathfrak{h}(R_j) = \frac{1}{2}\text{per}(R_j)$, which is easy to arrange. Both these requirements have no effect on the validity of the above proof.

Let the length of the word w in G_n be L . Let α be the constant given by property (*) of G_n . We henceforth assume that $\Gamma(\mathcal{G}'_n)$ has all the following properties.

- (1) $\lambda_i = \lambda(a_i) > \max(20, \frac{12}{\alpha})$ are even for all $1 \leq i \leq k$.
- (2) $\lambda_0 = \lambda(z) > \max(10, \frac{12}{\alpha})$ is odd.
- (3) For each $0 \leq j \leq l$, we have $\mathfrak{w}(R_j), \mathfrak{h}(R_j) > 4$.
- (4) For each $0 \leq j \leq l$, the perimeter length $\text{per}(R_j)$ is equal to $\lambda(r_j)$.
- (5) For each $0 \leq j \leq l$, we have $\frac{1}{2} \leq \frac{\mathfrak{w}(R_j)}{\mathfrak{h}(R_j)} \leq 2$.

Condition (5) gives the following lemma which we will need in the argument later. For any cycle χ in $\Gamma(\mathcal{G}'_n)$ let $|\chi|$ denote the length of χ .

LEMMA 4.4.4. *For any cycle χ in $\Gamma(\mathcal{G}'_n)$, there is a cycle $\bar{\chi}$ in $\Gamma(\mathcal{G}'_n)$ such that*

- (a) $\bar{\chi}$ is in the subgroup of $\pi_1(\Gamma(\mathcal{G}'_n))$ generated by the cycles corresponding to a_1, \dots, a_k and z ;
- (b) $\bar{\chi}$ is equivalent to χ in $\pi_1^*(\Gamma(\mathcal{G}'_n))$; and
- (c) $|\bar{\chi}| \leq 3|\chi|$.

PROOF. As in the proof of Lemma 4.4.3, each segment of the cycle χ with two endpoints on the boundary of a grid graph R_j can be replaced by a path which stays on the boundary of R_j . Replacing all such segments of χ in this manner gives $\bar{\chi}$. If ρ is the maximum aspect ratio of all R_j , then it is easily seen that $|\bar{\chi}| \leq (\rho + 1)|\chi|$. By condition (5) above, $\rho \leq 2$ and hence $|\bar{\chi}| \leq 3|\chi|$. \square

4.4.2. Standard forms of elements of G'_n . In this subsection we give an algebraic analysis of various standard forms of elements of G'_n . The standard forms will be instrumental in the proof of Theorem 4.4.2 in the next subsection.

The group G'_n can be described as the amalgamated free product

$$G'_n = G_n *_H K \mathbb{Z}$$

where $H = \langle w \rangle \leq G_n$ and $K = \langle z^2 \rangle \leq \langle z \rangle \cong \mathbb{Z}$. Note that $H \cong K \cong \mathbb{Z}$ as G_n is torsion-free.

DEFINITION 4.4.5. A *reduced form* for an element $g \in G'_n$ is a word of the form

$$g = g_1 \dots g_m$$

where for any $1 \leq i < m$, either $g_i = z$ and $g_{i+1} \in G_n$ or $g_i \in G_n$ and $g_{i+1} = z$, and for any $1 < i \leq m$, if $g_i \in G_n$ then $g_i \notin H$.

Note that a reduced form may start with a power of w . In general, a *word* in G'_n is a sequence $g_1 \dots g_m$, where each *term* g_i is either an element of G_n or an element of $\langle z \rangle$. For any word in G'_n that is not in reduced form, we can apply the following procedure, and repeat it if necessary, to arrive at a reduced form. First, if the word contains any term that is an odd power of z , say z^{2k+1} with $k \neq 0$, we replace the term by two terms $w^k z$. Then, if it contains any term that is an even power of z , say z^{2k} , we replace the term by the term w^k . (Note that after these two steps, in the resulting word every term that contains z is just z itself.) Next, if the resulting word contains adjacent terms that are both in G_n , we obtain a shorter word by combining the two terms. If the resulting word contains adjacent terms that are both z , we replace the two terms by a single w . If the resulting word contains three consecutive terms of the form $zw^k z$, we replace them by a single w^{k+1} . At any moment a term is obtained which is the identity of G_n , we delete the term. These steps are repeated as much as necessary. (Note that these steps only reduce the length of the word, and therefore the process must terminate.) If the last two terms of the resulting word are of the form zw^k , we replace them by $w^k z$ and repeat the length-reducing steps again as much as necessary. It is easy to see that, by following the procedure, we will eventually arrive at a reduced form for the original element. Also note that the number of occurrences of z in the original word (including those appearing in either the positive or negative powers of z) does not increase through this reduction procedure. No power of z other than z itself appears as a term in a reduced form.

We next define a normal form for elements of G'_n . Let U be a set of non-identity right coset representatives for the subgroup $H = \langle w \rangle$ of G_n . Clearly z is the unique coset representative for K in \mathbb{Z} .

DEFINITION 4.4.6. The *normal form* of an element $g \in G'_n$ is one of the following

- (1) $g = w^k u_1 z u_2 z \cdots z u_m$
- (2) $g = w^k z u_1 z u_2 \cdots z u_m$
- (3) $g = w^k u_1 z u_2 z \cdots u_m z$
- (4) $g = w^k z u_1 z u_2 \cdots u_m z$
- (5) $g = w^k z$
- (6) $g = w^k$

where $m \geq 1$, $u_i \in U$ and $k \in \mathbb{Z}$.

The normal form theorem for amalgamated free products (see Theorem 2.6 of [23]) implies that every element g of G'_n can be written uniquely in a normal form specified in the above definition. There is also a natural procedure to turn a reduced form into a normal form. For the convenience of the reader we describe this procedure below. Suppose

$$g = g_1 \cdots g_m$$

is in reduced form. The procedure is defined inductively on the length m of the reduced form for g . If $m = 1$ then either $g = z$, which is already in normal form (5), or $g \in G_n$, in which case either $g \in H$, which is in normal form (6), or $g = w^k u_1$

for some $u_1 \in U$, which is in normal form (1). Next assume $m = 2$. Then either $g = w^k z$, or $g = g_1 z$ with $g_1 \in G_n \setminus H$, or $g = z g_2$ with $g_2 \in G_n \setminus H$. The first case is already in normal form (5). In the second case, write $g_1 = w^k u_1$ for $u_1 \in U$, and $g = w^k u_1 z$ is in normal form (3). In the third case, write $g_2 = w^k u_1$, and $g = w^k z u_1$ is in normal form (2). In general, we assume $m > 2$. Then either $g_m = z$ or $g_m \in G_n \setminus H$. In case $g_m = z$ we write $g = g' z$ where $g' = g_1 \dots g_{m-1}$. Note that g' is in reduced form but of length $m - 1$. We then inductively obtain the normal form of g' , and the normal form of g is the normal form of g' followed by z . Suppose $g_m \in G_n \setminus H$. Let $g_m = w^k u_m$ for $u_m \in U$, and

$$g' = g_1 \dots g_{m-3} (g_{m-2} w^k).$$

Then g' is in reduced form and of length $m - 2$. By induction we can obtain the normal form of g' . Then the normal form of g is the normal form of g' followed by $z u_m$. This finishes the formal definition of the procedure. Informally, given a reduced form, we obtain the normal form by successively “passing” any power of w to the left, leaving behind the coset representatives in place of general elements of $G_n \setminus H$.

We next define another standard form for elements of G'_n .

DEFINITION 4.4.7. Let $g \in G'_n$. Among all representations of g as a product of the form $g = xyx^{-1}$ where y is in normal form, we let $i(g)$ be the minimum number of occurrences of z in the normal form y . A *minimal conjugate form* of g is

$$g = xyx^{-1}$$

where $x, y \in G'_n$ and y is in normal form with exactly $i(g)$ many occurrences of z .

The following fact is immediate from Definition 4.4.7.

FACT 4.4.8. Let $g \in G'_n$ and $g = xyx^{-1}$ be in minimal conjugate form. Then the normal form of g has at least $i(g)$ many occurrences of z .

PROOF. Since $g = eg'e^{-1}$, where g' is the normal form for g , it follows that the number of occurrences of z in g' is at least $i(g)$. \square

Note that if $g = xyx^{-1}$ is in minimal conjugate form, then g is odd if and only if y is odd. In particular if g is odd then $i(g) > 0$.

LEMMA 4.4.9. Let $g \in G'_n$ and $g = xyx^{-1}$ be in minimal conjugate form. Suppose $y \notin \langle z \rangle$. Then for any $N > 0$ we have that $g^N = xy^N x^{-1}$ where the normal form for y^N has exactly $N \cdot i(g)$ many occurrences of z .

PROOF. Since $y \notin \langle z \rangle$, y is not of normal form (5) or (6). Note that y cannot be of normal form (4) as then $y = z^{-1}(w^k z^2)(u_1 z \dots u_m)z$ and

$$g = (xz^{-1})(w^{k+1}u_1 z \dots u_m)(zx^{-1}).$$

This gives a conjugate form for g with a smaller number of occurrences of z in the normal form of the middle element, which violates the fact that y attained the minimum value $i(g)$.

Assume next that y is in normal form either (2) or (3). The cases are similar, so assume y is as in normal form (2), namely $y = w^k z u_1 z u_2 \dots u_m$. Then $g^N = x(w^k z u_1 \dots u_m)^N x^{-1}$. Passing each of the w^k terms to the left to put $(w^k z u_1 \dots u_m)^N$ in normal form gives

$$g^N = x(w^\ell z u'_1 z u'_2 \dots u'_m z u'_{m+1} z \dots u'_{2m} \dots z u'_{(N-1)m+1} z \dots u'_{Nm}) x^{-1}$$

for some $\ell \in \mathbb{Z}$. Here the coset representatives in U have changed (compared to the concatenation $(w^k u_1 z \dots u_m)^N$), but the number of them is still equal to Nm . So, in this case the conclusion of the lemma is immediate.

Finally, assume y is in normal form (1). Consider $y^N = (w^k u_1 z \dots u_m)^N$. To show the normal form for y^N has Nm many coset terms, it suffices to show that $u_m w^k u_1 \notin H$. Suppose $u_m w^k u_1 = h \in H$. Then

$$\begin{aligned} y &= (w^k u_1)(z u_2 \dots u_{m-1} z)(h)(w^k u_1)^{-1} \\ &= (w^k u_1 z)(u_2 \dots u_{m-1})(h w)(w^k u_1 z)^{-1} \\ &= (w^k u_1 z)(w^\ell u'_2 \dots u'_{m-1})(w^k u_1 z)^{-1} \end{aligned}$$

for some $\ell \in \mathbb{Z}$. Thus

$$g = (x w^k u_1 z)(w^\ell u'_2 \dots u'_{m-1})(x w^k u_1 z)^{-1}$$

and the normal form for the middle element contains fewer occurrences of z than y . This violates the choice of y . \square

4.4.3. Proof of Theorem 4.4.2. In this subsection we prove (ii) and (iii) of Theorem 4.4.2, and thus complete the proof of Theorem 4.4.1. Recall from §4.4.1 that we have obtained a finite connected graph $\Gamma(\mathcal{G}'_n)$ via a computable procedure so that $\pi_1^*(\Gamma(\mathcal{G}'_n)) \cong G'_n$.

To show (ii) of Theorem 4.4.2, suppose $n \in A$. Then $w = e$ in G'_n . Recall from condition (2) in §4.4.1 that the generator z in G'_n corresponds to an odd-length cycle ζ in $\Gamma(\mathcal{G}'_n)$. However $z^2 = w$ in G'_n , and so $z^2 = e$ in G'_n . Thus ζ is an odd-length cycle with $\pi_1^*(\zeta)$ having finite order in $\pi_1^*(\Gamma(\mathcal{G}'_n))$, which shows $\Gamma(\mathcal{G}'_n)$ satisfies the positive condition of Theorem 3.5.1.

It remains to verify (iii) of Theorem 4.4.2. For the rest of the argument, assume $n \notin A$ and thus $w \neq e$ in G'_n . We show that $\pi_1^*(\Gamma(\mathcal{G}'_n)) \cong G'_n$ satisfies the negative condition of Theorem 3.4.1. For this, let p, q be large odd primes. Toward a contradiction, suppose γ is a p -cycle in $\Gamma(\mathcal{G}'_n)$ and assume $\pi_1^*(\gamma^q)$ is a p^{th} power in $\pi_1^*(\Gamma(\mathcal{G}'_n))$.

To ease notation we use the isomorphism between $\pi_1^*(\Gamma(\mathcal{G}'_n))$ and G'_n tacitly and, instead of writing $\pi_1^*(\gamma)$, we simply write γ as an element of $\pi_1^*(\Gamma(\mathcal{G}'_n))$. Thus, suppose $\gamma^q = \delta^p$ in G'_n . In the group G'_n express γ and δ in their respective minimal conjugate forms, say $\gamma = x v x^{-1}$ and $\delta = y u y^{-1}$, where $i(\gamma), i(\delta)$ are the number of occurrences of z in v, u , respectively. For slight notational simplicity let $i_v = i(\gamma)$ and $i_u = i(\delta)$.

We claim that $i_v, i_u \neq 0$. Assume $i_v = 0$. Then $v \in G_n$, the subgroup of G'_n generated by a_1, \dots, a_k . By condition (1) of §4.4.1, in $\Gamma(\mathcal{G}'_n)$ the cycle corresponding to $x v x^{-1}$ has even length. But the lengths of reduced homotopy equivalent cycles differ by an even number, thus γ has even length, contradicting that p is odd. Since p, q are both odd, δ is also an odd-length cycle in $\Gamma(\mathcal{G}'_n)$. By a similar argument, $i_u \neq 0$.

Continuing our proof, assume first that $v \notin \langle z \rangle$. We have

$$\gamma^q = x v^q x^{-1} = y u^p y^{-1} = \delta^p$$

in G'_n . By Lemma 4.4.9, the normal form of v^q has exactly $q i_v$ many occurrences of z . We first claim that $u \notin \langle z \rangle$ as well. If $u = z^a$, then we would have that $x v^q x^{-1} = y z^{a p} y^{-1} = y w^{\frac{a p N}{2}} y^{-1}$ for all even N . Thus, $v^q = x^{-1} y w^{\frac{a p N}{2}} y^{-1} x$ for all even N . From Lemma 4.4.9, the normal form of v^q has exactly $q N i_v$ many occurrences of z . However, we may put $x^{-1} y w^{\frac{a p N}{2}} y^{-1} x$ into reduced form which

involves no more occurrences of z than those in the reduced forms of $x^{-1}y$ plus those in $y^{-1}x$, which is a constant that does not depend on N . It follows that the normal form of $x^{-1}yw^{\frac{apN}{2}}y^{-1}x$ has a constant number of occurrences of z that does not depend on N . Because the normal form for an element of G'_n is unique, we have a contradiction for large enough even N . So, $u \notin \langle z \rangle$. By Lemma 4.4.9 we also have that the normal form of u^p has pi_u many occurrences of z .

Now, for any N we have $xv^{qN}x^{-1} = yu^{pN}y^{-1}$, where the normal forms of v^{qN} and u^{pN} have qNi_v and pNi_u occurrences of z respectively. Next we claim that $p|i_v$. Otherwise, p is not a divisor of qi_v and so $qi_v \neq pi_u$. Thus $|qNi_v - pNi_u| \geq N$. Without loss of generality assume $qNi_v - pNi_u \geq N$. Then $v^{qN} = x^{-1}yu^{pN}y^{-1}x$. Comparing the number of occurrences of z for the normal forms of both sides, we get a contradiction when N is large. So, we have $p|i_v$.

Next, we let $\bar{\gamma}$ be a cycle in $\Gamma(\mathcal{G}'_n)$ given by Lemma 4.4.4 for the p -cycle γ . By Lemma 4.4.4 (c) we have $|\bar{\gamma}| \leq 3|\gamma| = 3p$. By Lemma 4.4.4 (b), $\bar{\gamma} = \gamma = xv x^{-1}$ as elements of G'_n . By Fact 4.4.8, the canonical form of $\bar{\gamma}$ has at least i_v many occurrences of z . Since $p|i_v$, we conclude that the canonical form of $\bar{\gamma}$ has at least p many occurrences of z . Finally, by Lemma 4.4.4 (a), $\bar{\gamma}$ is a cycle that is a combination of cycles corresponding to the generators a_1, \dots, a_k, z , which in particular gives a representation of $\bar{\gamma}$ in G'_n as a product of the generators a_1, \dots, a_k, z . For definiteness suppose

$$(7) \quad \bar{\gamma} = g_1 \dots g_m$$

is this representation. Since the normal form of $\bar{\gamma}$ contains at least p many occurrences of z , this representation also contains at least p many occurrences of z , because the procedures to produce the normal form from any representation does not increase the number of occurrences of z . Thus, as a cycle in $\Gamma(\mathcal{G}'_n)$, $\bar{\gamma}$ contains at least p many occurrences of the cycle corresponding to z . It follows that $|\bar{\gamma}| \geq p\lambda_0$ where $\lambda_0 = \lambda(z) > 10$. This contradicts $|\bar{\gamma}| \leq 3p$.

This finishes the proof for the case $v \notin \langle z \rangle$. Note that the above argument to show $u \notin \langle z \rangle$ from assuming $v \notin \langle z \rangle$ is symmetric for u and v . Thus the remaining case is in fact $u, v \in \langle z \rangle$.

Suppose $u = z^a$ and $v = z^b$. Thus $\gamma = xz^ax^{-1}$ and $\delta = yz^by^{-1}$. Note that a is odd, since otherwise the cycle corresponding to xz^ax^{-1} in $\Gamma(\mathcal{G}'_n)$ would have even length and therefore not reduced homotopy equivalent to the odd-length cycle γ . Similarly, b is also odd. So,

$$\gamma^q = xw^{\frac{aq-1}{2}}zx^{-1} = yw^{\frac{bp-1}{2}}zy^{-1} = \delta^p.$$

LEMMA 4.4.10. *For any integers c, d and $x, y \in G'_n$, if $xw^czz^{-1} = yw^dzy^{-1}$ then $c = d$ and $x^{-1}y \in \langle z \rangle$.*

PROOF. By conjugating both sides of the identity by x , the lemma is reduced to the case $x = e$ in G'_n , which we prove below. Note that $y \in \langle z \rangle$ implies $c = d$, and so the conclusion of the lemma holds. Suppose the conclusion fails. Let $y \notin \langle z \rangle$ be of minimal length in normal form such that for some c, d we have

$$w^cz = yw^dzy^{-1}.$$

We claim y cannot end with the z term. To see this, let $y = y'z$ be in normal form (so y' ends with a G_n term), then

$$w^cz = y'zw^dzz^{-1}y'^{-1} = y'w^dzy'^{-1},$$

and $y' = yz^{-1} \notin \langle z \rangle$. This violates the minimality assumption on the length of y . So, y ends with a $G_n \setminus H$ term. It follows that the normal form of yw^dzy^{-1} obtained from the procedure starting with a reduced form must contain a $G_n \setminus H$ term (that is, some $u \in U$). But w^cz is already in its normal form. This violates the uniqueness of the normal forms guaranteed by the normal form theorem. \square

From the lemma we have that $aq = bp$ and $x^{-1}y \in \langle z \rangle$. In particular, $p|a$. Let $a = pc$ for some odd integer c . Then $aq = pcq = bp$. In G'_n we have

$$\gamma = xz^{pc}x^{-1} = x(x^{-1}y)z^{pc}(y^{-1}x)x^{-1} = yz^{pc}y^{-1}$$

and $\delta = yz^{qc}y^{-1}$. Consider again a cycle $\bar{\gamma}$ in $\Gamma(G'_n)$ given by Lemma 4.4.4 for the p -cycle γ . By Lemma 4.4.4 (c), $|\bar{\gamma}| \leq 3|\gamma| = 3p$. By Lemma 4.4.4 (a), $\bar{\gamma}$ is a combination of the cycles corresponding to the generators a_1, \dots, a_k, z . This gives a representation of $\bar{\gamma}$ in G'_n as a product $g_1 \dots g_m$. Let ℓ be the minimum of $\lambda_0 = \lambda(z)$ and $\lambda(a_i)$ for all $1 \leq i \leq k$. By conditions (1) and (2) in §4.4.1, $\ell > \frac{12}{\alpha}$. We have that

$$3p \geq |\bar{\gamma}| \geq m\ell > m\frac{12}{\alpha}.$$

Thus $m \leq \frac{\alpha p}{4}$.

For any integer N we have $yz^{2pcN}y^{-1} = \gamma^{2N} = \bar{\gamma}^{2N}$, and so

$$w^{pcN} = z^{2pcN} = y^{-1}\bar{\gamma}^{2N}y.$$

Let Y be the length of y as a word in the a_i and z . Then as a word in the a_i and z , the length of the word $y^{-1}\bar{\gamma}^{2N}y$ is at most $2Y + N\frac{\alpha p}{2}$. On the other hand, from the definition of α we have that any representation of w^{pcN} in the a_i and z must have length at least $\alpha pcN \geq \alpha pN$. For N large enough, this is a contradiction.

Thus, in all cases we have a contradiction from assuming that there is a p -cycle in $\Gamma(G'_n)$ such that γ^q is a p^{th} power in $\pi_1^*(\Gamma(G'_n))$, assuming p, q are sufficiently large odd primes. So, $\Gamma(G'_n)$ satisfies the negative condition of Theorem 3.4.1. This completes the proof of Theorems 4.4.2 and 4.4.1.

4.5. The tiling problem

Recall from Definition 3.1.7 that a continuous tiling of $F(2^{\mathbb{Z}^2})$ by tiles $T_i \subseteq \mathbb{Z}^2$ is a clopen subequivalence relation E of $F(2^{\mathbb{Z}^2})$ such that every E class is isomorphic to one of the tiles T_i . We consider the continuous tiling problem for $F(2^{\mathbb{Z}^2})$ with a given finite set of finite tiles T_1, \dots, T_n . In §4.1 we stated the general tiling problem, which asks for which sets of tiles $\{T_i\}$ there is a continuous (or Borel) tiling of $F(2^{\mathbb{Z}^2})$. Of course, this problem could be asked for more general group actions, but we confine our remarks here to continuous tilings of $F(2^{\mathbb{Z}^2})$ by finite sets of finite tiles.

As we pointed out in §4.1, the continuous tiling problem for finite sets of finite tiles (which can be coded as integers) is a Σ_1^0 set. That is, the set A of integers n coding a finite set of finite tiles for which there is a continuous tiling of $F(2^{\mathbb{Z}^2})$ is a Σ_1^0 set of integers (this was Theorem 4.1.1). However, unlike the subshift and graph-homomorphism problems for which we know the corresponding sets are Σ_1^0 -complete (Theorems 4.3.1 and 4.4.1), we do not know if the tiling problem is Σ_1^0 -complete (i.e., the whether the set A is Σ_1^0 -complete). This question seems to be difficult even for specific rather simple sets of tiles. We restate the tiling problem:

PROBLEM 1. *Is the set of integers n coding a finite set of finite tiles*

$$T_1(n), \dots, T_i(n) \subseteq \mathbb{Z}^2$$

(i also depends on n) for which there is a continuous tiling of $F(2^{\mathbb{Z}^2})$ a Σ_1^0 -complete set?

In Theorem 3.1.9 and the discussion before it we showed that for the four rectangular tiles of dimensions $d \times d$, $d \times (d+1)$, $(d+1) \times d$, and $(d+1) \times (d+1)$, a continuous tiling of $F(2^{\mathbb{Z}^2})$ was possible, but if either the “small” tile T_s of dimensions $d \times d$, or the “large” tile T_ℓ of dimensions $(d+1) \times (d+1)$ is omitted from the set, then there is no continuous tiling of $F(2^{\mathbb{Z}^2})$ by the set of three remaining tiles. This argument (say in the case of omitting the large tile) proceeded by showing that, for p relatively prime to d , the $p \times p$ torus tile $T_{ca=ac}$ could not be tiled by these three tiles. This argument does not work if both the small tile and the large tile are present but one of the other tiles are omitted. In particular, we can ask:

PROBLEM 2. *Is there a continuous tiling of $F(2^{\mathbb{Z}^2})$ by the tiles $\{T_s, T_\ell\}$, where T_s is the $d \times d$ rectangle, and T_ℓ is the $(d+1) \times (d+1)$ rectangle (for $d \geq 2$)?*

It seems possible at the moment that the answer to Problem 2 could depend on the value of d . We currently do not know the answer for any $d \geq 2$.

To illustrate the nature of the problem, consider the $d = 2$ instance of Problem 2. Here we have the 2×2 tile T_s and the 3×3 tile T_ℓ . Theorem 2.3.5 gives a way, in principle, to answer the question. Namely, we must see if for all sufficiently large p, q with $\gcd(p, q) = 1$, whether we can tile (in the natural sense) the graph $\Gamma_{n,p,q}$. A necessary condition for doing this is that we be able to tile each of the 12 graphs $T_{n,p,q}^1, \dots, T_{n,p,q}^{12}$ individually.

First, unlike the case where we omit the small or large tile, we can now show that T_s and T_ℓ suffice to tile the 4 torus tiles (e.g., $T_{ca=ac}$) for $n = 1$ and p, q large enough with $\gcd(p, q) = 1$:

FACT. Let $p, q \geq 60$ with $\gcd(p, q) = 1$. Then the $p \times q$ torus $T_{ca=ac}$ can be tiled by 2×2 and 3×3 tiles.


PROOF. First of all, it is easy to see that for any positive integers a and $b > 1$, the $6a \times b$ and $b \times 6a$ grid graphs can be tiled by T_s and T_ℓ , the 2×2 and 3×3 tiles.

Let $\Lambda \subseteq \mathbb{Z}^2$ be the lattice generated by the vectors $v_1 = (3, 2)$ and $v_2 = (2, -3)$. Λ has horizontal and vertical periods of

$$\left| \det \begin{pmatrix} 3 & 2 \\ 2 & -3 \end{pmatrix} \right| = 13.$$

There is a tiling of \mathbb{Z}^2 by T_s, T_ℓ using the points of L for the centers of the T_ℓ tiles and the points of $(2, 0) + \Lambda$ for the upper-left corners of the T_s tiles. This is illustrated in Figure 45. From the periodicity, we can use restrictions of this tiling to tile any torus of dimensions $13a \times 13b$, for positive integers a, b .

Next, we define a procedure by which we can stretch any torus tiling by a multiple of 6 in either direction. Consider the horizontal direction, and suppose we have a tiling of a torus of dimensions $a \times b$. We describe a tiling of the torus of dimensions $(a + 6c) \times b$. Draw a vertical line L_0 through the torus. The line L_0 will

FIGURE 45. A periodic tiling by T_s, T_ℓ .

be partitioned into segments where it intersects various copies of the T_s, T_ℓ tiles (it may intersect the middle of a T_ℓ tile, or along one of its edges). Place another vertical line, L_1 , $6c$ units to the right of the line L_0 . We place copies of T_s and T_ℓ so that they intersect L_1 in the same way that L_0 intersected copies of these tiles. To the right of L_1 we will copy the original tiling of the $a \times b$ torus. Since 6 is a multiple of both 2 and 3, we may fill in the region between L_0 and L_1 as follows. For each segment w of L_0 which intersects a T_ℓ tile, we add $2c - 1$ copies of T_ℓ tiles horizontally to its right until we reach the corresponding tile intersecting L_1 . We likewise proceed for each segment of L_0 which intersects a T_s tile, adding $3c - 1$ copies of T_s horizontally to its right. Figure 46 illustrates the stretching procedure.




FIGURE 46. Stretching the torus tiling.

The above algorithm, together with the observation at the beginning of this proof, allow us to tile any torus of dimensions $(13a+6b) \times (13c+6d)$ for nonnegative integers a, b, c, d where $a+b, c+d > 0$. Since $\gcd(13, 6) = 1$, this shows any $p \times q$ torus tile for p, q large enough ($p, q \geq 60$) can be tiled by T_s and T_ℓ , which proves the fact. \square

We note that some smaller values of p, q can also work for different reasons. For example, the 17×19 torus can be tiled by 2×2 and 3×3 tiles as shown in Figure 47 (this tiling was discovered by a computer program), even though 17 is not of the form $13a + 6b$.

It seems likely that the tiling problem for continuous tilings of $F(2^{\mathbb{Z}^2})$ by T_s and T_ℓ hinges on whether the “long tiles” such as $T_{c^q a = ad^p}$ of Figure 10 can be tiled by T_s, T_ℓ (for large enough p, q). We do not know the answer to this question, but some computer experiments done by the authors for values of $p, q \leq 50$ were unable to find a tiling of the long tiles corresponding to these values of p, q .




FIGURE 47. A tiling of the 17×19 torus tile.

Bibliography

- [1] Gilbert Baumslag and Charles F. Miller III. The isomorphism problem for residually torsion-free nilpotent groups. *Groups, Geometry, and Dynamics*, 1(1):1–20, 2007.
- [2] Howard Becker and Alexander S. Kechris. *The descriptive set theory of Polish group actions*, volume 232 of *London Mathematical Society Lecture Notes Series*. Cambridge University Press, 1996.
- [3] Ferenc Bencs, Aranka Hrušková, and László Márton Tóth. Factor-of-iid schreier decorations of lattices in euclidean spaces. 2021. <https://arxiv.org/abs/2101.12577>.
- [4] Nishant Chandgotia and Spencer Unger. Borel factors and embeddings of systems in subshifts. 2022. <https://arxiv.org/abs/2203.09359>.
- [5] Clinton Conley, Steve Jackson, Andrew Marks, Brandon Seward, and Robin Tucker-Drob. Borel asymptotic dimension and hyperfinite equivalence relations. 2019. <https://arxiv.org/abs/2009.06721>.
- [6] Clinton Conley and Omer Tamuz. Unfriendly colorings of graphs with finite average degree. *Proceedings of the London Mathematical Society*, 122(2):229–233, 2019.
- [7] Endre Csóka, Lukasz Grabowski, András Máthé, Oleg Pikhurko, and Konstantinos Tyros. Borel version of the local lemma. 2016. <https://arxiv.org/abs/1605.04877>.
- [8] Randall Dougherty, Steve Jackson, and Alexander S. Kechris. The structure of hyperfinite Borel equivalence relations. *Transactions of the American Mathematical Society*, 341:193–225, 1994.
- [9] Jacob Feldman and Calvin Moore. Ergodic equivalence relations, cohomology, and von Neumann algebras, I. *Transactions of the American Mathematical Society*, 234(2):289–324, 1977.
- [10] Su Gao. *Invariant Descriptive Set Theory*. Tayler and Francis Group, 2009.
- [11] Su Gao and Steve Jackson. Countable abelian group actions and hyperfinite equivalence relations. *Inventiones Mathematicae*, 201(1):309–383, 2015.
- [12] Su Gao, Steve Jackson, Ed Krohne, and Brandon Seward. Forcing constructions and countable Borel equivalence relations. *The Journal of Symbolic Logic*, 87(3):873–893, 2022.
- [13] Su Gao, Steve Jackson, Ed Krohne, and Brandon Seward. Borel combinatorics of countable group actions. In preparation.
- [14] Su Gao, Steve Jackson, and Brandon Seward. A coloring property for countable groups. *Mathematical Proceedings of the Cambridge Philosophical Society*, 147(3):579–592, 2009.
- [15] Su Gao, Steve Jackson, and Brandon Seward. Group colorings and Bernoulli subflows. *Memoirs of the American Mathematical Society*, 241(1141):1–254, 2016.
- [16] Jan Grebík and Václav Rozhoň. Local problems on grids from the perspective of distributed algorithms, finitary factors, and descriptive combinatorics. 2021. <https://arxiv.org/abs/2103.08394>.
- [17] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.
- [18] John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Pearson Education, 2007.
- [19] Alexander S. Kechris. *Classical Descriptive Set Theory*. Springer-Verlag, 1995.
- [20] Alexander S. Kechris. The theory of countable Borel equivalence relations. To appear.
- [21] Alexander S. Kechris and Andrew Marks. Descriptive graph combinatorics. To appear.
- [22] Alexander S. Kechris, Slawomir Solecki, and Stevo Todorcevic. Borel chromatic numbers. *Advances in Mathematics*, 141:1–44, 1999.
- [23] R.C. Lyndon and P.E. Schupp. *Combinatorial Group Theory*. Springer, 1977.
- [24] Andrew Marks. A determinacy approach to Borel combinatorics. *Journal of the American Mathematical Society*, 29:579–600, 2016.

- [25] Scott Schneider and Brandon Seward. Locally nilpotent groups and hyperfinite equivalence relations. 2013. <https://arxiv.org/abs/1308.5853>.
- [26] Theodore A. Slaman and John R. Steel. Definable functions on degrees. In *Cabal Seminar 81–85*, number 1333 in Lecture Notes in Math., pages 37–55, Berlin, 1988. Springer.
- [27] Riley Thornton. Orienting Borel graphs. *Proceedings of the American Mathematical Society*, 150(4):1779–1793, 2022.
- [28] Felix Weilacher. Borel edge colorings for finite dimensional groups. 2021. <https://arxiv.org/abs/2104.14646v1>.
- [29] Benjamin Weiss. Measurable dynamics. In R. Beals et al., editor, *Conference in Modern Analysis and Probability*, number 26 in Contemporary Mathematics, pages 395–421, Providence RI, 1984. American Mathematical Society.