Entropy Minimality of Certain Hom-Shifts

A Thesis

submitted to
Indian Institute of Science Education and Research Pune
in partial fulfillment of the requirements for the
BS-MS Dual Degree Programme

by

Kinjal Dey

Indian Institute of Science Education and Research Pune
Dr. Homi Bhabha Road,
Pashan, Pune 411008, INDIA.

May, 2026

Supervisor: Dr. Nishant Chandgotia

Certificate

This is to certify that this dissertation entitled Entropy Minimality of Certain Hom-Shifts, towards the partial fulfilment of the BS-MS dual degree programme at the Indian Institute of Science Education and Research, Pune, represents study/work carried out by Kinjal Dey at TIFR CAM under the supervision of Dr. Nishant Chandgotia, Reader, TIFR CAM , during the academic year 2025-2026.

Nishant Chandgotia

Dr. Nishant Chandgotia

Committee:
Dr. Nishant Chandgotia
Dr. Moumanti Podder

This thesis is dedicated to Epilepsy

Declaration

I hereby declare that the matter embodied in the report entitled Entropy Minimality of Certain Hom-Shifts are the results of the work carried out by me at TIFR CAM, under the supervision of Dr. Nishant Chandgotia and the same has not been submitted elsewhere for any other degree. The use of Generative AI in this thesis did not extend beyond what is described in Section 4.6.1 of the 'Guidelines for Generative AI usage at IISER Pune', and accordingly, no attribution is required.

Kinjal Dey

Kinjal Dey

Acknowledgments

I am deeply grateful to my advisor, Dr. Nishant Chandgotia, for his patient guidance, constant encouragement, and insight throughout this work. This thesis would not have been possible without his support. I also thank Dr. Moumanti Podder for serving as the expert on my thesis. Many of the ideas in this thesis were shaped through long conversations with Aditya Thorat, which helped clarify my understanding throughout.

My parents and my brother have been a constant source of support and belief. I am grateful to my friends at IISER Pune and TIFR CAM - Aayush, Aditya, Ahan, Aishika, Akash, Akshara, Arjun, Authisha, Christopher, Devesh, Eshwar, Kunal, Madhukar, Medha, Moksh, Nikhil, Parasuram, Rahitya, Rajalaxmi, Rajat, Rakshaditya, Sagar, Sanchit, Saranya, Sasidhar, Satyen, Shashank, Sourajit, Subhankar, Sumiran, Sushant, Vishnu, Yash and many more - for the discussions, ideas, and camaraderie that made this journey more meaningful.

Finally, I acknowledge Epilepsy, to whom this thesis is dedicated.

Abstract

In this thesis, we study structural and entropic properties of hom-shifts arising from finite connected graphs. Our main goal is to understand how the topology of the underlying graph influences rigidity phenomena in the associated shift space. We prove that whenever the square cover of a graph is a tree, the corresponding hom-shift is entropy minimal, meaning that forbidding any admissible pattern strictly decreases entropy.

To develop the necessary tools, we investigate Lipschitz extension problems for graph homomorphisms and establish a bipartite version of the Kirszbraun–Helly equivalence, adapted to parity-preserving maps. This extension framework plays a key role in enabling global constructions from local constraints.

A central component of our approach is the use of height functions obtained by lifting configurations using the square cover of the graph. This cover provides a natural setting in which lifts are well defined and distances encode global structural information. Using the ergodic theorems, we analyze the asymptotic growth of these height functions and show that maximal directional slope forces strong rigidity.

A final chapter discusses a related direction for certain tiling systems, where we prove a necessary and sufficient condition for tileability.

Contents

Chapter 1

Introduction

In this thesis we study shift spaces arising from graph homomorphisms from the lattice $\mathbb{Z}^d$ into a finite graph. Let G be a finite undirected graph without multiple edges. The space

$$X_G := \mathrm{Hom}(\mathbb{Z}^d, G)$$

consists of all graph homomorphisms from $\mathbb{Z}^d$ to G. Such spaces, called *hom-shifts*, form an important class of multidimensional shifts of finite type. They arise naturally in several models in statistical mechanics and combinatorics like the hard-core model, the Widom-Rowlinson model and proper k-colorings of $\mathbb{Z}^d$.

A fundamental invariant associated with a shift space X is its *topological entropy*, denoted by $h_{\mathrm{top}}(X)$, which measures the exponential growth rate of admissible patterns as the size of the underlying region increases. Computing the entropy of multidimensional shifts of finite type is generally a difficult problem. Instead of attempting to compute entropy explicitly, we study a structural property of shift spaces known as *entropy minimality*. A shift space X is said to be entropy minimal if for every proper subshift $Y \subsetneq X$ we have

$$h_{\mathrm{top}}(Y) < h_{\mathrm{top}}(X).$$

For $d = 1$, irreducible nearest-neighbour shifts of finite type are entropy minimal (see [LM21]). In higher dimensions, however, much less is known. Certain strong mixing conditions imply entropy minimality ([Sch10]), but nearest-neighbour shifts of finite type with weaker mixing properties need not be entropy minimal ([BPS10]). Hom-shifts provide a natural class of systems in which these questions can be studied.

The goal of this thesis is to prove entropy minimality for a class of hom-shifts by analyzing the structure of measures of maximal entropy. A key observation from thermodynamic formalism is that to prove entropy minimality it suffices to show that every measure of maximal entropy is fully supported. Moreover, it is enough to consider ergodic measures.

The approach developed here relies on associating height functions to configurations in X_G. These height functions arise from lifting configurations to a covering graph of G known as the *square cover*, which is obtained as a quotient of the universal cover. Tools from algebraic topology on graphs provide a framework for studying these lifts and for ensuring their existence. Once a lift of a configuration is fixed, it naturally induces a height function whose differences capture rigidity properties of configurations.

The asymptotic behaviour of these height functions is studied using the ergodic theorems. In particular, one can define *slopes* describing the average rate of change of the height function in each coordinate direction. Measures whose slopes are maximal in some direction correspond to configurations with rigid behaviour and lead to zero entropy. For measures whose slopes are submaximal in every direction, we obtain uniform bounds on the growth of height differences. More precisely, for sufficiently large boxes, the difference between the height at the origin and the height along the boundary is bounded by a constant multiple of the side length of the box.

A further ingredient in the argument is a bipartite version of the Kirszbraun–Helly equivalence for graphs. This result provides a mechanism for extending Lipschitz maps defined on subsets of the lattice. In the present setting it allows one to combine any finite admissible pattern in the interior of a large region with a configuration outside the region belonging to the support of a measure of maximal entropy. Using this extension property, we show that every admissible finite pattern appears in configurations in the support, and hence that every measure of maximal entropy is fully supported.

Combining these ingredients yields the entropy minimality of the hom-shifts considered in this thesis.

Theorem 1.0.1. Let G be a finite connected graph and let $X_G = \mathrm{Hom}(\mathbb{Z}^d, G)$ denote the associated hom-shift. If the square cover of G is a tree, then X_G is entropy minimal.

We show in Lemma 5.5.1 that the square cover of G is a tree exactly when G is square-free.

The thesis is organized as follows. Chapter 2 introduces shift spaces and the basic symbolic dynamics needed for the rest of the thesis. Chapter 3 introduces ergodic theoretic tools, and the thermodynamic formalism required to relate entropy minimality to the structure of measures of maximal entropy. Chapter 4 establishes the bipartite Kirszbraun–Helly equivalence for graphs. Chapter 5 reviews the algebraic topology of graphs, including covering graphs and the square group. Chapter 6 introduces height functions and slopes associated with configurations and studies their properties. Finally, Chapter 7 combines these tools to prove entropy minimality for the class of hom-shifts studied in this work. The thesis concludes with Chapter 8 where we consider a related problem of tilings by $1 \times m$ and $n \times 1$ rectangles.

Chapter 2

Shift Spaces

2.1 Introduction

Let $\mathcal{A}$ be a finite set, viewed as a discrete topological space. The full configuration space $\mathcal{A}^{\mathbb{Z}^d}$ is endowed with the product topology, and is therefore compact. An element $x \in \mathcal{A}^{\mathbb{Z}^d}$ will be referred to as a *configuration*, and we write $x = (x_{\vec{i}})_{\vec{i}\in\mathbb{Z}^d}$, where $x_{\vec{i}}$ denotes the symbol placed at site $\vec{i} \in \mathbb{Z}^d$. The group $\mathbb{Z}^d$ acts naturally on $\mathcal{A}^{\mathbb{Z}^d}$ by translations.

Definition 2.1.1. *For each $\vec{i} \in \mathbb{Z}^d$, define the* shift map $\sigma_{\vec{i}} : \mathcal{A}^{\mathbb{Z}^d} \longrightarrow \mathcal{A}^{\mathbb{Z}^d}$ *by*

$$(\sigma_{\vec{i}}(x))_{\vec{j}} = x_{\vec{i}+\vec{j}}.$$

The family $\{\sigma_{\vec{i}}\}_{\vec{i}\in\mathbb{Z}^d}$ defines a $\mathbb{Z}^d$-action on $\mathcal{A}^{\mathbb{Z}^d}$.

Intuitively, the shift translates the configuration by a lattice vector, reindexing the symbols accordingly.

Definition 2.1.2. *A subset $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ is called a* shift space *if it is closed (in the product topology) and invariant under the shift action.*

Shift spaces can equivalently be described through local constraints: Let $F \subset \mathbb{Z}^d$ be finite. A *pattern* supported on F is an element $a \in \mathcal{A}^F$. Given such a pattern, the associated *cylinder set* is as follows.

$$[a]_F = \{x \in \mathcal{A}^{\mathbb{Z}^d} \mid x|_F = a\}.$$

If $x|_F = a$, we may also denote the cylinder by $[x]_F$.

Let $\mathcal{F}$ be a collection of patterns, interpreted as *forbidden configurations*. We define

$$X_{\mathcal{F}} = \left\{ x \in \mathcal{A}^{\mathbb{Z}^d} \;\middle|\; \text{no translate of } x \text{ contains a pattern from } \mathcal{F} \right\}.$$

More explicitly,

$$x \in X_{\mathcal{F}} \quad \Longleftrightarrow \quad \sigma_{\vec{i}}(x)|_F \notin \mathcal{F} \text{ for all } \vec{i} \in \mathbb{Z}^d.$$

It is a standard fact that a subset $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ is a shift space if and only if it can be described as $X = X_{\mathcal{F}}$ for some collection $\mathcal{F}$ of forbidden patterns. Different collections of patterns may define the same shift space. Readers can refer to [LM21] for proofs in the case $d = 1$. For higher dimensions the same arguments still hold.

Definition 2.1.3. *A shift space X is called a* shift of finite type *if there exists a finite collection of forbidden patterns $\mathcal{F}$ such that $X = X_{\mathcal{F}}$.*

Definition 2.1.4. *A shift of finite type is called a* nearest neighbour shift of finite type *if the forbidden patterns are supported only on single vertices and edges of $\mathbb{Z}^d$ (more precisely the Cayley graph of $\mathbb{Z}^d$).*

Any shift of finite type can be recoded so as to become a nearest neighbour shift of finite type, although the alphabet may become larger under this recoding.

2.2 Hom-Shifts

We now restrict attention to configurations taking values in a graph. Let G be a graph. A configuration $x \in G^{\mathbb{Z}^d}$ may be viewed as assigning to each lattice site a vertex of G.

Definition 2.2.1. *Define*

$$X_G^d = \left\{ x \in G^{\mathbb{Z}^d} \;\middle|\; x_{\vec{i}} \sim_G x_{\vec{j}} \text{ whenever } \vec{i} \sim \vec{j} \right\}.$$

Thus X_G^d consists precisely of graph homomorphisms from the lattice $\mathbb{Z}^d$ to G. We refer to X_G^d as the *hom-shift* associated to G.

When G is finite, X_G^d is a nearest neighbour shift of finite type. Indeed, letting

$$\mathcal{F}_G = \left\{ [v, w]_{\vec{0}, \vec{e}_j} \;\middle|\; v \not\sim_G w,\ 1 \leq j \leq d \right\},$$

we obtain

$$X_G^d = X_{\mathcal{F}_G}.$$

In this sense, hom-shifts are precisely those nearest neighbour shifts whose local constraints enforce adjacency in a fixed graph, uniformly across all coordinate directions.

Figure 2.1: The constraint graph of the hard-core shift and an example configuration on a finite region of $\mathbb{Z}^2$.

2.3 Topological Entropy

Definition 2.3.1 (Globally Allowed Patterns)**.** *Let $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ be a shift space and let $A \subset \mathbb{Z}^d$ be finite. The set of* globally allowed patterns of X on A *is defined by*

$$\mathcal{L}_A(X) := \left\{ a \in \mathcal{A}^A \;\middle|\; \exists\, x \in X \text{ such that } x|_A = a \right\}.$$

Definition 2.3.2 (Language of a Shift Space)**.** *The* language *of X is the collection* *of all finite patterns that appear in configurations of X, namely*

$$\mathcal{L}(X) := \bigcup_{\substack{A \subset \mathbb{Z}^d \\ A \text{ finite}}} \mathcal{L}_A(X).$$

Remark 2.3.3 (Locally vs. Globally Allowed Patterns)**.** The notion of global admissibility differs from that of *local admissibility*. Suppose X is defined by a forbidden set of patterns $\mathcal{F}$. Given a finite set $A \subset \mathbb{Z}^d$, a pattern $a \in \mathcal{A}^A$ is said to be *locally allowed* if no translate of any pattern from $\mathcal{F}$ appears inside a.

In other words, a is locally allowed whenever it does not violate any of the defining local constraints, even if it does not extend to a full configuration in X.

Definition 2.3.4 (Topological Entropy)**.** *Let $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ be a shift space. The* topological entropy *of X is defined as the exponential growth rate of the number of globally allowed patterns on large finite regions. More precisely,*

$$h_{\text{top}}(X) := \lim_{n\to\infty} \frac{\log\left|\mathcal{L}_{B_n}(X)\right|}{|B_n|},$$

where $B_n \subset \mathbb{Z}^d$ denotes the box $[-n, n]^d$ and $|B_n|$ is its cardinality.

The existence of the limit follows from standard subadditivity arguments.

Note that if the graph G is a single edge or a single vertex with a self-loop, then the number of globally allowed patterns are 2 and 1 respectively. In both cases, $h_{\text{top}}(X_G) = 0$. We will now see that in all graphs containing either of these two as a strict subgraph, the topological entropy is positive.

Proposition 2.3.5. Let H be a finite graph with distinct vertices a, b, c such that $a \sim_H b$ and $b \sim_H c$. Then

$$h_{\text{top}}(X_H) \geq \frac{\log 2}{2}.$$

Proof. It suffices to prove the claim for the graph H consisting of exactly three vertices a, b, c with edges $a \sim_H b$ and $b \sim_H c$. Any larger graph satisfying the hypothesis contains this graph as a subgraph, and hence the corresponding homshift has entropy at least as large.

For this graph, any configuration in X_H assigns the symbol b to one partite class of $\mathbb{Z}^d$, while vertices in the other partite class may be assigned either a or c independently. Thus the values on half of the sites are forced, while the other half can be chosen freely between two symbols.

Consequently we obtain

$$|\mathcal{L}_{B_n}(X_H)| = 2^{\lfloor (2n+1)^d/2 \rfloor} + 2^{\lceil (2n+1)^d/2 \rceil}.$$

Taking logarithms and dividing by $|B_n| = (2n+1)^d$ gives

$$h_{\text{top}}(X_H) = \frac{\log 2}{2}.$$

□

Proposition 2.3.6. Let $X_{\mathcal{HC}} \subset \{0,1\}^{\mathbb{Z}^d}$ be the hard-core shift. Then

$$h_{\text{top}}(X_{\mathcal{HC}}) \geq \frac{\log 2}{2}.$$

Proof. Fix one partite class of $\mathbb{Z}^d$ and assign the value 0 to all its vertices. The hard-core constraint is then automatically satisfied, and the vertices in the other partite class may be assigned 0 or 1 independently.

Thus we have

$$|\mathcal{L}_{B_n}(X_{\mathcal{HC}})| \geq 2^{\lfloor (2n+1)^d/2 \rfloor}.$$

Hence

$$\frac{1}{|B_n|} \log |\mathcal{L}_{B_n}(X_{\mathcal{HC}})| \geq \frac{\lfloor (2n+1)^d/2 \rfloor}{(2n+1)^d} \log 2.$$

Taking $n \to \infty$ gives

$$h_{\text{top}}(X_{\mathcal{HC}}) \geq \frac{\log 2}{2}.$$

□

Definition 2.3.7. *A shift space X is said to be* entropy minimal *if for every non-empty proper subshift $Y \subsetneq X$ we have*

$$h_{\text{top}}(Y) < h_{\text{top}}(X).$$

Equivalently, X is entropy minimal if forbidding any globally allowed pattern results in a strict decrease in topological entropy.

From before, it is easy to see that if G is a single edge or a single vertex with a self-loop, then X_G is entropy minimal, since forbidding any globally-allowed amounts to getting a shift space which is empty.

Chapter 3

Invariant Measures and Entropy

3.1 Ergodicity

Ergodic theory studies the long-term statistical behavior of dynamical systems through the lens of measure theory. Rather than analyzing individual trajectories, one investigates how measurable sets and observables evolve under repeated application of a transformation. In this, and the upcoming sections, we introduce the notion of ergodicity and some of the Ergodic theorems, which will be useful in proving limit theorems for heights in Chapter 6 and Chapter 7. For further details and proofs of the results in Section 3.1, Section 3.2, and Section 3.3, refer to [Sar19, EW11, Wal82].

Let $(X, \mathcal{B}, \mu)$ be a measure space, where X is a set, $\mathcal{B}$ is a σ-algebra on X, and μ is a measure.

A measurable map $T : X \to X$ is called a *measure-preserving transformation (m.p.t.)* if for every $A \in \mathcal{B}$ we have

$$\mu(T^{-1}A) = \mu(A).$$

In this case the quadruple $(X, \mathcal{B}, \mu, T)$ is called a *measure-preserving system*. If the measure μ is a probability measure then, T is called a *probability-preserving transformation (p.p.t.)*.

A set $A \in \mathcal{B}$ is said to be T*-invariant* if $T^{-1}A = A$.

Definition 3.1.1. *Let $(X, \mathcal{B}, \mu, T)$ be a measure-preserving system. The transformation T is said to be* ergodic *if for every T-invariant set $A \in \mathcal{B}$, $\mu(A) = 0$ or $\mu(A) = 1$.*

Thus, in an ergodic system, the only measurable sets that remain unchanged under the dynamics (up to null sets) are trivial from the measure-theoretic point of view.

3.2 Ergodic Decomposition

The ergodic decomposition theorem asserts that every invariant probability measure can be represented as an integral of ergodic invariant measures. As a consequence, statements about invariant measures can often be reduced to the ergodic case.

Theorem 3.2.1 (Regular Conditional Probabilities)**.** Let μ be a Borel probability measure on a standard probability space $(X, \mathcal{B}, \mu)$, and let $\mathcal{F} \subset \mathcal{B}$ be a σ-algebra. Then there exists a family of Borel probability measures $\{\mu_x\}_{x \in X}$ such that:

1. For every $E \in \mathcal{B}$, the map
$$x \longmapsto \mu_x(E)$$
is $\mathcal{F}$-measurable;

2. If f is μ-integrable, then
$$\int f \, d\mu = \int_X \left(\int f \, d\mu_x \right) d\mu(x);$$

3. If f is μ-integrable, then
$$\int f \, d\mu_x = \mathbb{E}(f \mid \mathcal{F})(x) \quad \text{for } \mu\text{-almost every } x.$$

Theorem 3.2.2 (Ergodic Decomposition)**.** Let $(X, \mathcal{B}, \mu)$ be a standard probability space, and let $T : X \to X$ be a Borel measurable map preserving μ. Let $\{\mu_x\}_{x \in X}$ be the family of regular conditional probabilities with respect to $\mathfrak{Inv}(T)$.

Then:

1. For every $f \in L^1(X, \mathcal{B}, \mu)$,

$$\int f \, d\mu = \int_X \left(\int f \, d\mu_x \right) d\mu(x);$$

2. For μ-almost every $x \in X$, the measure μ_x is T-invariant;

3. For μ-almost every $x \in X$, the measure μ_x is ergodic.

3.3 Ergodic Theorems

A fundamental question in the study of dynamical systems is whether long-term time averages along a single orbit reflect the global statistical properties of the system. Given a probability-preserving transformation $(X, \mathcal{B}, \mu, T)$ and an integrable function f, we may interpret $f(T^k x)$ as the observation of a quantity f at time k along the trajectory of x. One may then ask whether the averages

$$\frac{1}{N} \sum_{k=0}^{N-1} f(T^k x)$$

stabilize as $N \to \infty$, and if so, what their limit represents. The ergodic theorem provides a precise answer: time averages always converge almost everywhere. In the special case of an ergodic transformation, the limit is constant and equals the global average $\int f \, d\mu$, expressing the principle that long-term temporal behavior reflects spatial averages.

Definition 3.3.1. *Let $(X, \mathcal{B}, \mu, T)$ be a probability-preserving transformation. The* T-invariant σ-algebra *is defined by*

$$\mathfrak{Inv}(T) = \{E \in \mathcal{B} : T^{-1}E = E\}.$$

Equivalently, $\mathfrak{Inv}(T)$ is the σ-algebra generated by all T-invariant measurable functions i.e. all f satisfying $f \circ T = f$ almost everywhere.

Theorem 3.3.2 (Birkhoff's Ergodic Theorem)**.** Let $(X, \mathcal{B}, \mu, T)$ be a probability-preserving transformation, and let $f \in L^1(X, \mathcal{B}, \mu)$. Then the limit

$$\tilde{f}(x) := \lim_{N\to\infty} \frac{1}{N} \sum_{k=0}^{N-1} f(T^k x)$$

exists for μ-almost every x, and

$$\frac{1}{N} \sum_{k=0}^{N-1} f \circ T^k \longrightarrow \tilde{f} \quad \text{in } L^1.$$

Moreover,

$$\tilde{f} = \mathbb{E}(f \mid \mathfrak{Inv}(T)) \quad \text{almost everywhere,}$$

so that $\tilde{f}$ is T-invariant, belongs to L^1, and satisfies

$$\int \tilde{f}\, d\mu = \int f\, d\mu.$$

If T is ergodic, then

$$\tilde{f} = \int f\, d\mu \quad \text{almost everywhere.}$$

In many dynamical systems one is interested in quantities that measure growth along an orbit and satisfy a subadditivity property.

Definition 3.3.3 (Subadditive cocycle)**.** *Let $(X, \mathcal{B}, \mu, T)$ be a measure-preserving dynamical system. A sequence of measurable functions $\{g_n\}_{n\geq 1}$ with $g_n : X \to \mathbb{R}$ is called a* subadditive cocycle *if $g_n \in L^1(X, \mathcal{B}, \mu)$ for all n and for all $m, n \geq 1$,*

$$g_{m+n}(x) \leq g_m(x) + g_n(T^m x) \quad \textit{for } \mu\textit{-almost every } x.$$

Theorem 3.3.4 (Kingman's Subadditive Ergodic Theorem)**.** Let $(X, \mathcal{B}, \mu, T)$ be a probability-preserving transformation, and let $\{g_n\}_{n\geq 1}$ be a subadditive cocycle over T with $g_1 \in L^1(X, \mathcal{B}, \mu)$.

Then the limit

$$g(x) := \lim_{n\to\infty} \frac{1}{n} g_n(x)$$

exists for μ-almost every $x \in X$, and the function g is T-invariant. Moreover,

$$\int g\,d\mu = \inf_{n\geq 1} \frac{1}{n} \int g_n\,d\mu.$$

If, in addition, T is ergodic, then g is constant almost everywhere and

$$g(x) = \inf_{n\geq 1} \frac{1}{n} \int g_n\,d\mu \quad \text{for } \mu\text{-almost every } x.$$

3.4 Support of Invariant Measures

Throughout this thesis, μ will denote a shift-invariant Borel probability measure on a shift space X. The *support* of μ, denoted by $\operatorname{supp}(\mu)$, is defined as the intersection of all closed subsets $Y \subseteq X$ satisfying $\mu(Y) = 1$. Equivalently, it is the smallest closed subset of full measure.

Proposition 3.4.1. $\operatorname{supp}(\mu)$ is a shift space.

Proof.

$$\operatorname{supp}(\mu) = \bigcap \{F \subset X : F \text{ is closed and } \mu(F) = 1\}.$$

Since $\operatorname{supp}(\mu)$ is an intersection of closed sets, it is closed.

If F is closed and $\mu(F) = 1$, then by shift-invariance of μ,

$$\mu(\sigma^{\vec{i}} F) = \mu(F) = 1.$$

Moreover, $\sigma^{\vec{i}} F$ is closed because $\sigma^{\vec{i}}$ is a homeomorphism.

Since $\operatorname{supp}(\mu)$ is the intersection of all closed full-measure sets, it is contained in every such set, and in particular in its translations. Therefore

$$\operatorname{supp}(\mu) \subseteq \sigma^{\vec{i}}(\operatorname{supp}(\mu)).$$

Applying the same argument to $\sigma^{-\vec{i}}$ gives the reverse inclusion, hence

$$\sigma^{\vec{i}}(\operatorname{supp}(\mu)) = \operatorname{supp}(\mu).$$

Since $\operatorname{supp}(\mu)$ is closed and shift-invariant, it is a shift space. $\square$

Proposition 3.4.2. The support of μ may equivalently be defined as

$$\operatorname{supp}(\mu) = \{x \in X : \mu(U) > 0 \text{ for every open set } U \text{ containing } x\}.$$

Proof. Let

$$S := \{x \in X : \mu(U) > 0 \text{ for every open set } U \ni x\}.$$

If $x \notin S$, then there exists an open neighborhood U of x with $\mu(U) = 0$. Hence $X \setminus U$ is a closed set with $\mu(X \setminus U) = 1$ that does not contain x. Therefore $x \notin \operatorname{supp}(\mu)$, so $S \subseteq \operatorname{supp}(\mu)$.

Conversely, if $x \notin \operatorname{supp}(\mu)$, then there exists a closed set F with $\mu(F) = 1$ and $x \notin F$. Then $U := X \setminus F$ is an open neighborhood of x with $\mu(U) = 0$, so $x \notin S$. Thus $\operatorname{supp}(\mu) \subseteq S$, proving the claim. $\square$

3.5 Measure-Theoretic Entropy

We begin with the classical notion of Shannon entropy for a discrete random variable, which quantifies uncertainty in a probability distribution.

Definition 3.5.1 (Shannon entropy)**.** *Let X be a discrete random variable with distribution (p_i). The Shannon entropy of X is*

$$H(X) := -\sum_i p_i \log p_i,$$

with the convention that $0 \log 0 = 0$.

This notion extends naturally to finite measurable partitions of a probability space.

Definition 3.5.2. *Let $(X, \mathcal{B}, \mu)$ be a probability space and let $\mathcal{P} = \{A_1, \dots, A_k\}$ be a finite measurable partition. The entropy of $\mathcal{P}$ is*

$$H_\mu(\mathcal{P}) := -\sum_{i=1}^{k} \mu(A_i) \log \mu(A_i).$$

We now specialise to shift-invariant measures on $\mathbb{Z}^d$. Let μ be such a measure. For $n \geq 1$, define

$$B_n := [-n, n]^d,$$

and let $\mathcal{P}_{B_n}$ denote the partition into cylinder sets determined by configurations on B_n.

The entropy of μ is obtained by normalising the entropy of these finite regions.

Definition 3.5.3 (Measure-theoretic Entropy)**.** *The* measure-theoretic entropy *of μ is*

$$h_\mu := \lim_{n\to\infty} \frac{1}{|B_n|} H_\mu(\mathcal{P}_{B_n}).$$

The existence of the limit follows from subadditivity arguments.
The reader can refer to [Sar19, Wal82] for more a more general theory of entropy for dynamical systems.

Definition 3.5.4 (Measure of maximal entropy)**.** *A shift-invariant probability measure μ on X is called a* measure of maximal entropy *if it maximizes entropy among all invariant measures, i.e.*

$$h_\mu = \sup_\nu h_\nu,$$

where the supremum is taken over all shift-invariant probability measures ν on X.

The existence of measures of maximal entropy follows from the upper semicontinuity of the entropy map

$$\nu \longmapsto h_\nu$$

with respect to the weak-$*$ topology on the space of invariant probability measures (see [Wal82]).

The Variational Principle gives a relation between topological entropy and measure-theoretic entropy. Refer to [Din70, Goo71, Wal82] for a proof.

Theorem 3.5.5 (Variational Principle)**.** Let X be a $\mathbb{Z}^d$ shift space. Then the topological entropy of X satisfies

$$h_{\text{top}}(X) = \sup_\nu h_\nu$$

where the supremum is taken over all shift-invariant probability measures on X. In particular, if μ is a measure of maximal entropy, then

$$h_{\text{top}}(X) = h_\mu.$$

Note that from Proposition 2.3.5 and Proposition 2.3.6, h_μ is positive for any measure of maximal entropy on a hom-shift which is either the hard-core shift or where the graph has distinct vertices a, b, c such that $a \sim_H b$ and $b \sim_H c$.

We will now rephrase the problem of entropy minimality in terms of these measures of maximal entropy ([QT00]).

Proposition 3.5.6. A shift space X is entropy minimal if and only if every measure of maximal entropy for X is fully supported.

Proof. Suppose first that X is entropy minimal and let μ be a measure of maximal entropy on X. By the variational principle applied to X,

$$h_{\text{top}}(X) = h_\mu.$$

Since $\operatorname{supp}(\mu)$ is a shift space contained in X, the variational principle applied to $\operatorname{supp}(\mu)$ gives

$$h_\mu \leq h_{\text{top}}(\operatorname{supp}(\mu)).$$

On the other hand,

$$h_{\text{top}}(\operatorname{supp}(\mu)) \leq h_{\text{top}}(X),$$

because $\operatorname{supp}(\mu) \subseteq X$. Combining these inequalities yields

$$h_{\text{top}}(X) \leq h_{\text{top}}(\operatorname{supp}(\mu)) \leq h_{\text{top}}(X),$$

and hence

$$h_{\text{top}}(\operatorname{supp}(\mu)) = h_{\text{top}}(X).$$

Since X is entropy minimal, this forces $\operatorname{supp}(\mu) = X$. Thus every measure of maximal entropy is fully supported.

Conversely, assume that every measure of maximal entropy for X is fully supported and suppose, for the sake of a contradiction, that X is not entropy minimal. Then there exists a proper subshift $Y \subsetneq X$ such that

$$h_{\mathrm{top}}(Y) = h_{\mathrm{top}}(X).$$

By the variational principle applied to Y, there exists a shift-invariant probability measure μ supported on Y satisfying

$$h_\mu = h_{\mathrm{top}}(Y) = h_{\mathrm{top}}(X).$$

Thus μ is a measure of maximal entropy for X, but its support is contained in Y and hence is not equal to X, contradicting the assumption. Therefore X must be entropy minimal. □

3.6 Markov Random Fields

Definition 3.6.1 (Markov random field)**.** *A Borel probability measure μ on $\mathcal{A}^{\mathbb{Z}^d}$ is called a Markov random field if for all finite sets $A, B \subset \mathbb{Z}^d$ satisfying*

$$\partial A \subseteq B \subseteq A^c,$$

and for all patterns $a \in \mathcal{L}_A$ and $b \in \mathcal{L}_B$ with $\mu([b]_B) > 0$, we have

$$\mu([a]_A \,|\, [b]_B) = \mu([a]_A \,|\, [b]_{\partial A})\,.$$

In words, the conditional distribution inside a finite region depends only on the configuration along its boundary.

Definition 3.6.2 (Uniform Markov random field)**.** *A Markov random field μ is said to be* uniform *if, for every finite $A \subset \mathbb{Z}^d$ and every boundary pattern $b \in \mathcal{L}_{\partial A}$ with $\mu([b]_{\partial A}) > 0$, the conditional distribution on A is uniform. That is,*

$$\mu([a]_A \,|\, [b]_{\partial A}) = \frac{1}{N_{A,b}},$$

where

$$N_{A,b} := \left|\left\{a \in \mathcal{L}_A \;\middle|\; \mu([a]_A \cap [b]_{\partial A}) > 0\right\}\right|.$$

Definition 3.6.3 (Homoclinic equivalence [PS97, Sch97])**.** *Let X be a shift space. The* homoclinic equivalence relation *on X, denoted by Δ_X, is defined by*

$$\Delta_X := \left\{(x, y) \in X \times X \;\middle|\; x_i = y_i \textit{ for all but finitely many } i \in \mathbb{Z}^d\right\}.$$

Definition 3.6.4 (Adapted measure)**.** *Let X be a shift space. A probability measure μ is said to be* adapted *to X if*

$$\mathrm{supp}(\mu) \subset X,$$

and the support of μ is closed under homoclinic equivalence; that is, whenever $x \in \mathrm{supp}(\mu)$ and $y \in X$ satisfy $(x, y) \in \Delta_X$, it follows that $y \in \mathrm{supp}(\mu)$.

Theorem 3.6.5. Let X be a nearest–neighbour shift of finite type. Then every measure of maximal entropy on X is a shift-invariant uniform Markov random field μ which is adapted to X.

The above result is a special case of the classical Lanford–Ruelle theorem (see [LR69, Rue04]).

The following fact is implied by Theorem 14.15 in [Geo88] and Theorem 4.3.7 in [Kel98].

Theorem 3.6.6. Let X be a shift space and let μ be a shift-invariant uniform Markov random field adapted to X. Suppose the ergodic decomposition (see Theorem 3.2.2) of μ is given by a measurable map $x \mapsto \mu_x$ such that

$$\mu = \int_X \mu_x \, d\mu(x).$$

Then, for μ-almost every $x \in X$, the measure μ_x is itself a shift-invariant uniform Markov random field adapted to X, and moreover

$$\mathrm{supp}(\mu_x) \subset \mathrm{supp}(\mu).$$

In addition, the entropy decomposes as

$$h_\mu = \int_X h_{\mu_x}\, d\mu(x).$$

As a consequence of this theorem, it will suffice for us to check certain claims only for the ergodic measures.

Chapter 4

Kirszbraun-Helly equivalence for Graphs

4.1 Kirszbraun-type theorems amd Helly Property

A central question in metric geometry is whether Lipschitz maps defined on a subset of a space can be extended to the whole space without increasing their Lipschitz constant. Such extension problems arise naturally in analysis, probability, and combinatorics, and are closely related to convexity and intersection properties of metric balls. The classical result addressing this question in Euclidean spaces is due to Kirszbraun, McShane and Whitney ([Kir34, McS34, Whi34]).

Notation 4.1.1. For $1 \leq p \leq \infty$ and $x = (x_1, \ldots, x_d) \in \mathbb{R}^d$, the ℓ_p norm is defined by

$$\|x\|_p := \begin{cases} \left(\sum_{i=1}^{d} |x_i|^p\right)^{1/p}, & 1 \leq p < \infty, \\ \max_{1\leq i\leq d} |x_i|, & p = \infty. \end{cases}$$

Theorem 4.1.2 (Kirszbraun Extension Theorem)**.** Let $A \subseteq \mathbb{R}^n$ and let $f : A \to \mathbb{R}^m$ be a Lipschitz map with Lipschitz constant $L \geq 0$, i.e.,

$$\|f(x) - f(y)\|_2 \leq L\|x - y\|_2 \quad \text{for all } x, y \in A.$$

Then there exists an extension

$$\tilde{f} : \mathbb{R}^n \to \mathbb{R}^m$$

such that

1. $\tilde{f}|_A = f$, and
2. $\tilde{f}$ is Lipschitz with the same constant L, i.e.,

$$\|\tilde{f}(x) - \tilde{f}(y)\|_2 \leq L\|x - y\|_2 \quad \text{for all } x, y \in \mathbb{R}^n.$$

The Helly property is a fundamental intersection principle for families of subsets of a metric or combinatorial space (see [Hel23]). Roughly speaking, it asserts that global intersection behaviour can be detected from small subfamilies. In geometric settings, this phenomenon is classically associated with convexity: for instance, in Euclidean spaces, finite collections of convex sets intersect whenever all sufficiently small subcollections do.

Theorem 4.1.3 (Helly's Theorem in $\mathbb{R}^m$)**.** Let $\mathcal{F}$ be a family of convex subsets of $\mathbb{R}^m$ such that every subfamily of at most $m + 1$ sets has nonempty intersection.

If $\mathcal{F}$ is finite, then

$$\bigcap_{C \in \mathcal{F}} C \neq \varnothing.$$

The same result holds if $\mathcal{F}$ is an infinite family of convex compact sets.

A classical connection between Helly-type intersection results and Lipschitz extension theorem was established by Valentine [Val45], who showed that Helly's theorem can be used to obtain the Kirszbraun extension theorem.

An analogous phenomenon in the setting of graphs was explored more recently by Chandgotia, Pak, and Tassy in [CPT18]. In their work, the authors showed that Helly-type intersection properties admit an equivalent formulation in terms of Lipschitz extension phenomena for graph metrics. More precisely, they proved that a graph satisfies an appropriate Helly property for metric balls if and only if it enjoys a Kirszbraun-type extension property for Lipschitz maps defined on its subsets. For a more extensive survey of the Helly property in graphs refer to [BC08].

Definition 4.1.4 (X-Kirszbraun Property)**.** *Let (X, d_X) and (Y, d_Y) be metric spaces. We say that Y is X-*Kirszbraun *if every 1-Lipschitz map defined on an arbitrary subset $A \subseteq X$ can be extended to a 1-Lipschitz map on the whole space X.*

Under this terminology, Kirszbraun's theorem asserts that $(\mathbb{R}^m, \ell_2)$ is $(\mathbb{R}^n, \ell_2)$-Kirszbraun.

Definition 4.1.5 ((n, m)-Helly Property)**.** *Let $m \in \mathbb{N}$ and let $n \in \mathbb{N} \cup \{\infty\}$ with $n > m$. A metric space (X, d) is said to satisfy the (n, m)-*Helly property *if the following holds:*

For every collection of closed balls

$$B_1, B_2, \ldots, B_n \subseteq X$$

each of radius at least 1, whenever every subcollection of at most m balls has nonempty intersection, we have

$$\bigcap_{i=1}^{n} B_i \neq \varnothing.$$

Theorem 4.1.6. [CPT18] A graph H is $\mathbb{Z}^d$-Kirszbraun (where both H and $\mathbb{Z}^d$ are equipped with the graph metric) if and only if it is $(2d, 2)$-Helly.

4.2 The Bipartite case

In many combinatorial settings, particularly when working with graph homomorphisms from $\mathbb{Z}^d$, an additional parity constraint naturally appears. Since $\mathbb{Z}^d$ is bipartite, any graph homomorphism must respect the bipartite structure of the target graph. Consequently, extension problems in this discrete framework require preserving parity in addition to Lipschitz continuity. An underlying Helly property in the target graphs ensures the existence of extensions. This leads to a bipartite analogue of the Kirszbraun and Helly properties, tailored to parity-preserving maps.

Definition 4.2.1 (Bipartite (n,m)-Helly Property)**.** *Let H be a bipartite graph with partite classes H_0 and H_1. Let $m \in \mathbb{N}$ and let $n \in \mathbb{N} \cup \{\infty\}$ with $n > m$.*

We say that H is bipartite (n,m)-Helly *if the following holds:*

For every collection of balls

$$B_1, B_2, \ldots, B_n$$

in H (finite if $n \neq \infty$, and any finite collection otherwise), whenever every subcollection of size m among

$$B_1 \cap H_k,\ B_2 \cap H_k,\ \ldots,\ B_n \cap H_k$$

has nonempty intersection, it follows that

$$\bigcap_{i=1}^{n} (B_i \cap H_k) \neq \varnothing$$

for $k = 0, 1$.

Definition 4.2.2 (Bipartite G-Kirszbraun Property)**.** *Let G and H be bipartite graphs with partite classes G_0, G_1 and H_0, H_1, respectively.*

We say that H is bipartite G-Kirszbraun *if the following holds:*

For every subset $A \subseteq G$ and every 1-Lipschitz map

$$f : A \to H$$

such that

$$f(A \cap G_0) \subseteq H_0 \quad \textit{and} \quad f(A \cap G_1) \subseteq H_1,$$

there exists a graph homomorphism

$$\tilde{f} \in \mathrm{Hom}(G, H)$$

extending f, i.e.,

$$\tilde{f}(x) = f(x) \quad \textit{for all } x \in A.$$

Theorem 4.2.3. A bipartite graph H is bipartite $\mathbb{Z}^d$-Kirszbraun if and only if it is bipartite $(2d, 2)$-Helly.

The proof of the bipartite case proceeds along the same lines as the proof of the non-bipartite theorem given in [CPT18], and we adapt their arguments to the present setting.

Definition 4.2.4 (Geodesic Extension)**.** *Let G be a graph, let $A \subseteq G$, and let $b \in G \setminus A$. The* geodesic extension *of A with respect to b is defined as*

$$\mathrm{Ext}(A, b) := \left\{ a \in A \;\middle|\; \begin{array}{l} \textit{there is no } a' \in A \setminus \{a\} \\ \textit{such that there exists a geodesic from } a \textit{ to } b \textit{ passing through } a' \end{array} \right\}.$$

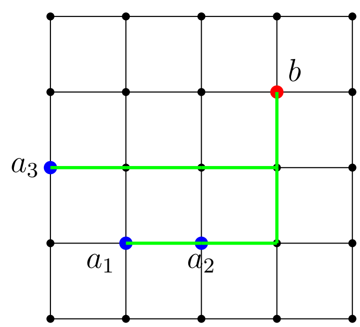

Figure 4.1: Illustration of $\mathrm{Ext}(A, b)$ where $A = \{a_1, a_2, a_3\}$ in the lattice graph. There is a geodesic from a_1 to b which passes through a_2, so $a_1 \notin \mathrm{Ext}(A, b)$. Thus $\mathrm{Ext}(A, b) = \{a_2, a_3\}$.

Remark 4.2.5. When G and H are bipartite graphs, parity-preserving 1-Lipschitz maps coincide with graph homomorphisms. Indeed, if $f : G \to H$ is 1-Lipschitz and preserves the bipartitions, then adjacent vertices in G have opposite parity, and hence their images lie in opposite partite classes of H. Since a 1-Lipschitz map sends adjacent vertices to vertices at distance at most one, the images must be adjacent, so f is a graph homomorphism. The converse is immediate.

Proposition 4.2.6. Suppose G and H are bipartite graphs. Let $A \subseteq G$ and let $f : A \to H$ be a 1-Lipschitz map which preserves parity. Let $b \in G \setminus A$. Then f admits a parity-preserving 1-Lipschitz extension to $A \cup \{b\}$ if and only if the restricted map

$$f\big|_{\mathrm{Ext}(A,b)}$$

admits a parity-preserving 1-Lipschitz extension to

$$\mathrm{Ext}(A,b) \cup \{b\}.$$

.

Proof. The forward direction is immediate since $\mathrm{Ext}(A,b) \subseteq A$.

For the converse, suppose $f|_{\mathrm{Ext}(A,b)}$ admits a parity-preserving 1-Lipschitz extension

$$\tilde{f} : \mathrm{Ext}(A,b) \cup \{b\} \to H.$$

Define $\hat{f} : A \cup \{b\} \to H$ by

$$\hat{f}(a) := \begin{cases} f(a) & a \in A, \\ \tilde{f}(b) & a = b. \end{cases}$$

Clearly $\hat{f}$ preserves parity and extends f. It remains to show that $\hat{f}$ is 1-Lipschitz. Thus we must verify that for every $a \in A$,

$$d_H(\hat{f}(a), \hat{f}(b)) \leq d_G(a,b).$$

If $a \in \mathrm{Ext}(A,b)$ this follows immediately since $\tilde{f}$ is 1-Lipschitz. Otherwise $a \in A \setminus \mathrm{Ext}(A,b)$. Then there exists $a' \in \mathrm{Ext}(A,b)$ lying on a geodesic from a to b, so

$$d_G(a,b) = d_G(a,a') + d_G(a',b).$$

Since f and $\tilde{f}$ are 1-Lipschitz we have

$$d_G(a,a') \geq d_H(f(a), f(a')), \qquad d_G(a',b) \geq d_H(\tilde{f}(a'), \tilde{f}(b)).$$

Applying the triangle inequality gives

$$d_H(\hat{f}(a), \hat{f}(b)) \leq d_H(f(a), f(a')) + d_H(f(a'), \tilde{f}(b)) \leq d_G(a, a') + d_G(a', b) = d_G(a, b),$$

which completes the proof. □

Notation 4.2.7. Denote the parity classes of a bipartite graph G by G_0 and G_1. If $g \in G_i$ then define parity$(g) := i$. We can relabel the indices for parity classes so that parity-preserving maps become index-preserving.

Proposition 4.2.8. Let H be a graph satisfying the bipartite $(n, 2)$-Helly property. Let $f : A \to H$ be a parity-preserving 1-Lipschitz map with $A \subseteq G$, and let $b \in G \backslash A$. If

$$|\mathrm{Ext}(A, b)| \leq n,$$

then f admits a parity-preserving 1-Lipschitz extension to $A \cup \{b\}$.

Proof. For all $a, a' \in \mathrm{Ext}(A, b)$, we have

$$d_H\big(f(a), f(a')\big) \;\leq\; d_G(a, a') \;\leq\; d_G(a, b) + d_G(b, a').$$

Consequently,

$$B^H_{d_G(a,b)}\big(f(a)\big) \;\cap\; B^H_{d_G(b,a')}\big(f(a')\big) \neq \varnothing.$$

Now let

$$x \in B^H_{d_G(a,b)}\big(f(a)\big) \;\cap\; B^H_{d_G(a',b)}\big(f(a')\big).$$

Then

$$d_H\big(x, f(a)\big) \leq d_G(a, b), \qquad d_H\big(x, f(a')\big) \leq d_G(a', b).$$

Since H is bipartite, distances determine parity classes modulo 2. Because f is parity-preserving, we have

$$\mathrm{parity}\big(f(a)\big) = \mathrm{parity}(a).$$

Hence

$$\mathrm{parity}(x) \equiv \mathrm{parity}\big(f(a)\big) + d_H\big(x, f(a)\big) \pmod 2.$$

If $B^H_{d_G(a,b)}\big(f(a)\big) \ \cap \ B^H_{d_G(b,a')}\big(f(a')\big)$ contains two adjacent vertices, then we can make a choice of x such that parity(x) = parity(b). Otherwise, $d_H\big(f(a), f(a')\big) = d_G(a,b) + d_G(b,a')$. So, we can choose x such that it is at $d_G(a,b)$ distance from $f(a)$ along a geodesic from $f(a)$ to $f(a')$ in H. Therefore

$$\text{parity}(x) \equiv \text{parity}(a) + d_G(a,b) \equiv \text{parity}(b) \pmod 2.$$

Thus x lies in the same parity class as b i.e. $x \in H_{\text{parity}(b)}$.

We have that

$$H_{\text{parity}(b)} \ \cap \ B^H_{d_G(a,b)}\big(f(a)\big) \ \cap \ B^H_{d_G(b,a')}\big(f(a')\big) \neq \varnothing$$

for all $a, a' \in \text{Ext}(A,b)$.

Since $|\text{Ext}(A,b)| \leq n$ and H satisfies the bipartite $(n,2)$-Helly property, it follows that

$$H_{\text{parity}(b)} \ \cap \bigcap_{a \in \text{Ext}(A,b)} B^H_{d_G(a,b)}\big(f(a)\big) \neq \varnothing.$$

Let y be a vertex in this intersection. Defining $f(b) = y$ yields a parity-preserving 1-Lipschitz extension of f to $\text{Ext}(A,b) \cup \{b\}$ and by Proposition 4.2.6 we have the result.

□

Notation 4.2.9 ($T_{\mathbf{r}}$ graph)**.** Let $\mathbf{r} = (r_1, r_2, \ldots, r_n) \in \mathbb{N}^n$. Define $T_{\mathbf{r}}$ to be the tree formed by taking a vertex b_0 and attaching to it n pairwise vertex-disjoint paths of lengths $r_1, \ldots, r_n$. Denote the terminal vertex of the i-th path by b_i for $1 \leq i \leq n$.

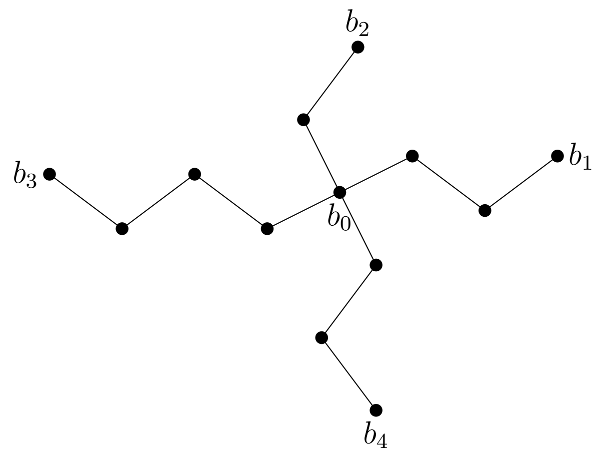

Figure 4.2: The tree $T_{\mathbf{r}}$ where $\mathbf{r} = (3, 2, 4, 3)$

Proof of Theorem 4.2.3. ($\implies$) Let H be a graph that is bipartite $\mathbb{Z}^d$-Kirszbraun. Given any $\mathbf{r} \in \mathbb{N}^{2d}$, one can construct an isometry from $T_{\mathbf{r}}$ to $\mathbb{Z}^d$ under which the geodesic rays starting at the central vertex correspond to walks on the coordinate axes. Consequently, H is bipartite $T_{\mathbf{r}}$-Kirszbraun for every $\mathbf{r} \in \mathbb{N}^{2d}$.

Suppose that

$$B^H_{r_1}(v_1),\ B^H_{r_2}(v_2),\ \ldots,\ B^H_{r_{2d}}(v_{2d})$$

are balls in H such that

$$I^0_{ij} := H_0 \cap B^H_{r_i}(v_i)\ \cap\ B^H_{r_j}(v_j) \neq \varnothing \quad \text{for all } 1 \leq i, j \leq 2d.$$

Also define

$$I_{ij} := B^H_{r_i}(v_i)\ \cap\ B^H_{r_j}(v_j).$$

Choose the partite class labels such that $\text{parity}(b_0) = 0$. Denote by b'_i the unique vertices situated at unit distance from b_i in $T_{\mathbf{r}}$.

Consider the map

$$f : \{\bar{b}_1, \ldots, \bar{b}_{2d}\} \subset T_{\mathbf{r}} \longrightarrow H, \qquad f(\bar{b}_i) = v_i\ \ (1 \leq i \leq 2d),$$

where the vertices $\overline{b}_i$ are defined according to parity by

$$\overline{b}_i = \begin{cases} b_i, & \text{if the parity of } r_i \text{ and } v_i \text{ agree,} \\ b'_i, & \text{if the parity of } r_i \text{ and } v_i \text{ differ.} \end{cases}$$

Here parity(r_i) is simply the value of r_i modulo 2. This definition makes the map f parity-preserving. Now we check that f is $1-$Lipschitz.

Case 1 ($\overline{b_i} = b_i$, $\overline{b_j} = b_j$)**:** In this case $d(f(b_i), f(b_j)) = d(v_i, v_j) \leq r_i + r_j = d(b_i, b_j)$

Case 2 ($\overline{b_i} = b'_i$, $\overline{b_j} = b_j$)**:** This decomposes into two different cases.

> **Case** $2a$ **(There are no adjacent vertices in** I_{ij}**.):** Here, $d(v_i, v_j) = r_i + r_j$. Let $x \in I^0_{ij}$. Since v_i and r_i have different parities, but v_j and r_j have the same parity, x will have different parities when calculating using the former and latter values. This is a contradiction.
>
> **Case** $2b$ **(There are adjacent vertices in** I_{ij}.)**:** In this case, $d(f(b'_i), f(b_j)) = d(v_i, v_j) \leq r_i + r_j - 1 = d(b'_i, b_j)$.

Case 3 ($\overline{b_i} = b'_i$, $\overline{b_j} = b'_j$)**:** This decomposes into three cases.

> **Case** $3a$ **(There are no adjacent vertices in** I_{ij}**.):** Here, $d(v_i, v_j) = r_i + r_j$. This case is not possible since we can take $x \in I^0_{ij}$, but the parities of v_i and r_i are different, so they add up to 1.
>
> **Case** $3b$ **(There are adjacent vertices in** I_{ij}**, but none of the vertices adjacent to a pair of adjacent vertices are in** I_{ij}**.):** Take a pair of such adjacent vertices and consider the one nearer to v_i. It will have parity, parity(v_i) + parity$(r_i) + 1 \equiv 0$. But it should also have parity, parity(v_j) + parity$(r_j) \equiv 1$. This is a contradiction.
>
> **Case** $3c$ **(There is a path** $u_1u_2u_3$ **contained in** I_{ij}**):** In this case, $d(f(b'_i), f(b'_j)) = d(v_i, v_j) \leq r_i + r_j - 2 = d(b'_i, b'_j)$.

By the $T_{\mathbf{r}}$-Kirszbraun property, the map f admits an extension

$$\tilde{f} \in \mathrm{Hom}(T_{\mathbf{r}}, H).$$

By construction of f, the image of the central vertex b_0 satisfies

$$\tilde{f}(b_0) \ \in \ H_0 \ \cap \ \bigcap_{i=1}^{2d} B^H_{\overline{r}_i}(v_i),$$

where the radii $\overline{r}_i$ are defined by

$$\overline{r}_i = \begin{cases} r_i, & \text{if } \overline{b}_i = b_i, \\ r_i - 1, & \text{if } \overline{b}_i = b_i'. \end{cases}$$

Note that in the case where $r_i = 1$ and $\overline{b}_i = b_i'$, we simply obtain $\overline{r}_i = 0$, and hence

$$v_i \in H_0 \ \cap \ \bigcap_{i=1}^{2d} B^H_{\overline{r}_i}(v_i).$$

It follows that

$$H_0 \ \cap \ \bigcap_{i=1}^{2d} B^H_{\overline{r}_i}(v_i) \neq \varnothing.$$

($\Longleftarrow$) Suppose that H satisfies the $(2d, 2)$-Helly property. We aim to show that for every subset $A \subseteq \mathbb{Z}^d$, any parity-preserving 1-Lipschitz map $f : A \to H$ admits a parity-preserving 1-Lipschitz extension to $\mathbb{Z}^d$. It suffices to establish this when A is finite.

We argue by induction on $|A|$. For $n \in \mathbb{N}$, consider the statement:

$\mathbf{St}(n)$: *Let $A \subseteq \mathbb{Z}^d$ with $|A| = n$, and let $f : A \to H$ be parity-preserving 1-Lipschitz. Then for every $b \in \mathbb{Z}^d \setminus A$, the map f extends to a parity-preserving 1-Lipschitz map on $A \cup \{b\}$.*

From Proposition 4.2.8, $\mathbf{St}(n)$ holds for all $n \leq 2d$.

Now assume $\mathbf{St}(n)$ for some $n \geq 2d$, and let us prove $\mathbf{St}(n+1)$. Let $A \subseteq \mathbb{Z}^d$ with $|A| = n+1$, let $f : A \to H$ be parity-preserving 1-Lipschitz, and let $b \in \mathbb{Z}^d \setminus A$.

WLOG we may assume that $b = \mathbf{0}$ (shift all the vectors involved by $-b$ otherwise). Moreover, we may assume that

$$\mathrm{Ext}(A, \mathbf{0}) = A,$$

since otherwise the induction hypothesis together with Proposition 4.2.6 yields a 1-Lipschitz extension of f to $A \cup \{\mathbf{0}\}$.

We will show that there exists a subset

$$\tilde{A} \subseteq \mathbb{Z}^d$$

and a parity-preserving 1-Lipschitz map

$$\tilde{f} : \tilde{A} \to H,$$

such that the following properties hold:

1. $\tilde{f}$ admits a parity-preserving 1-Lipschitz extension to $\tilde{A} \cup \{\mathbf{0}\} \implies f$ admits a parity-preserving 1-Lipschitz extension to $A \cup \{\mathbf{0}\}$.

2. Either $\tilde{A}$ is contained in the coordinate axes of $\mathbb{Z}^d$, or $|\tilde{A}| \leq 2d$.

Since H has the $(2d, 2)$-Helly property, by Proposition 4.2.8 it follows that $\tilde{f}$ has an extension to

$$\tilde{A} \cup \{\mathbf{0}\},$$

which completes the proof.

Since $|A| \geq n + 1 > 2d$, we can find vertices $\vec{i}, \vec{j} \in A$ and a coordinate index $1 \leq k \leq d$ for which the components i_k and j_k are both nonzero and share the same sign. Without loss of generality, assume that $i_k \leq j_k$.

Consider the vertex $\vec{i} - i_k \vec{e}_k$. There is a geodesic path from $\vec{j}$ to this vertex which passes through $\vec{i}$.

$A = \mathrm{Ext}(A, \mathbf{0})$ implies that no vertex of A lies on a geodesic from another vertex of A to $\mathbf{0}$. Then we have,

$$\vec{i} - i_k \vec{e}_k \notin \{\mathbf{0}\} \cup A.$$

since $\vec{i}-i_k\vec{e}_k \in A$ would result in a geodesic from $\mathbf{0}$ to $\vec{i}$ which passes through $\vec{i}-i_k\vec{e}_k$ and $\vec{i}-i_k\vec{e}_k = \mathbf{0}$ would result in a geodesic from $\mathbf{0}$ to $\vec{j}$ which passes through $\vec{i}$.

Also,

$$\vec{j} \notin \mathrm{Ext}\big(A, \vec{i}-i_k\vec{e}_k\big),$$

and therefore

$$\big|\mathrm{Ext}\big(A, \vec{i}-i_k\vec{e}_k\big)\big| \leq n.$$

By $\mathbf{St}(n)$, the restriction

$$f\big|_{\mathrm{Ext}(A,\vec{i}-i_k\vec{e}_k)}$$

admits a parity-preserving 1-Lipschitz extension to

$$\mathrm{Ext}(A, \vec{i}-i_k\vec{e}_k) \;\cup\; \{\vec{i}-i_k\vec{e}_k\}.$$

Applying Proposition 4.2.6, this yields a parity-preserving 1-Lipschitz extension

$$f' : A \cup \{\vec{i}-i_k\vec{e}_k\} \to H$$

of the original map f.

Finally, observe that there exists a geodesic from $\vec{i}$ to $\mathbf{0}$ passing through $\vec{i}-i_k\vec{e}_k$. Consequently,

$$\mathrm{Ext}\big(A \cup \{\vec{i}-i_k\vec{e}_k\}, \mathbf{0}\big) \subseteq \big(A \setminus \{\vec{i}\}\big) \cup \{\vec{i}-i_k\vec{e}_k\}.$$

Set

$$A' := \big(A \setminus \{\vec{i}\}\big) \;\cup\; \{\vec{i}-i_k\vec{e}_k\}.$$

By Proposition 4.2.6, the map f' admits a parity-preserving 1-Lipschitz extension to

$$A \cup \{\vec{i}-i_k\vec{e}_k\} \cup \{\mathbf{0}\}$$

if and only if the restricted map

$$f'\big|_{A'}$$

admits a parity-preserving 1-Lipschitz extension to

$$A' \cup \{\mathbf{0}\}.$$

Thus we have constructed a subset $A' \subseteq \mathbb{Z}^d$ and a parity-preserving 1-Lipschitz map

$$f' : A' \to H$$

satisfying the following properties:

1. If f' extends to $A' \cup \{\mathbf{0}\}$, then f extends to $A \cup \{\mathbf{0}\}$;
2. The total number of nonzero coordinates among the vertices of A' is strictly smaller than that for A.

Iterating this reduction procedure produces the desired subset

$$\tilde{A} \subseteq \mathbb{Z}^d$$

together with a parity-preserving 1-Lipschitz map

$$\tilde{f} : \tilde{A} \to H.$$

This completes the proof.

$\square$

Let us see why it suffices to establish the extension property for finite subsets $A \subseteq \mathbb{Z}^d$.

Indeed, extending a parity-preserving 1-Lipschitz map $f : A \to H$ to a vertex $b \in \mathbb{Z}^d \setminus A$ amounts to finding a vertex

$$x \in H_0 \ \cap \ \bigcap_{a \in A} B^H_{d_G(b,a)}\big(f(a)\big)'$$

where we assume WLOG that $\text{parity}(b) = 0$.

For each finite subset $F \subseteq A$, define

$$S_F := H_0 \cap \bigcap_{a \in F} B^H_{d_G(b,a)}\big(f(a)\big).$$

If the extension property holds for finite sets, then $S_F \neq \varnothing$ for every finite $F \subseteq A$. Moreover, if $F \subseteq F'$, then $S_{F'} \subseteq S_F$.

Fix $a_0 \in A$ and set $R := d_G(b, a_0)$. We can approximate A using an increasing sequence of finite sets $\{F_n\}$ each of which contain a_0. so the family $\{S_{F_n}\}$ is nested and decreasing. For every $F_n \subseteq A$ we have

$$S_{F_n} \subseteq B^H_R\big(f(a_0)\big).$$

Hence all sets S_{F_n} lie inside the finite set

$$K := B^H_R\big(f(a_0)\big).$$

We therefore have:

- a finite ambient set K,
- a nested family $S_F \subseteq K$,
- each $S_F \neq \varnothing$.

A decreasing family of nonempty subsets of a finite set has nonempty intersection. Consequently,

$$\bigcap_n S_{F_n} \neq \varnothing.$$

Any vertex in this intersection satisfies all Lipschitz constraints, and thus defines the required extension to b.

It remains to justify why the one–point extension property implies extension to all of $\mathbb{Z}^d$.

Since $\mathbb{Z}^d$ is countable, we may enumerate its vertices as

$$\mathbb{Z}^d = \{\vec{v}_0, \vec{v}_1, \vec{v}_2, \dots\}.$$

Reindexing if necessary, assume that the initial domain $A \subseteq \mathbb{Z}^d$ is contained in $\{\vec{v}_0, \ldots, \vec{v}_m\}$.

We construct inductively a sequence of finite sets

$$A_n := A \cup \{\vec{v}_{m+1}, \ldots, \vec{v}_n\}, \qquad n \geq m,$$

together with parity-preserving 1-Lipschitz maps

$$f_n : A_n \to H$$

satisfying

$$f_n\big|_{A_{n-1}} = f_{n-1}.$$

Set $f_m := f$. Suppose f_n has been defined. Since A_n is finite and $\vec{v}_{n+1} \notin A_n$, the one–point extension property yields a parity-preserving 1-Lipschitz extension

$$f_{n+1} : A_n \cup \{\vec{v}_{n+1}\} \to H$$

of f_n.

This inductive procedure produces a compatible family of maps $\{f_n\}_{n \geq m}$. We now define

$$F : \mathbb{Z}^d \to H$$

by setting

$$F(\vec{v}_n) := f_n(\vec{v}_n), \qquad n \geq 0.$$

This is well defined because the extensions are consistent on overlaps.

Since, F is defined pointwise by parity-preserving maps, it is also parity-preserving. Finally, we verify that F is 1-Lipschitz. Let $\vec{u}, \vec{v} \in \mathbb{Z}^d$. Then there exists N such that both vertices lie in A_N. Since f_N is parity-preserving 1-Lipschitz, we have

$$d_H\big(F(\vec{u}), F(\vec{v})\big) = d_H\big(f_N(\vec{u}), f_N(\vec{v})\big) \leq d_G(\vec{u}, \vec{v}).$$

Thus F is a parity-preserving 1-Lipschitz extension of f to all of $\mathbb{Z}^d$.

Alternatively, an application of Zorn's Lemma also gives us the same justification.

4.3 Helly for Trees

Proposition 4.3.1. Trees are $(n, 2)$-Helly for all n.

Proof. Let T be a tree. Closed balls in a tree (with the graph metric) induce connected subgraphs, in particular, they are subtrees.

Base case: Let T_1, T_2, T_3 be subtrees of T such that $T_i \cap T_j \neq \varnothing$. Suppose that $T_1 \cap T_2 \cap T_3 = \varnothing$. Then there exist vertices $v_1 \in T_2 \cap T_3$, $v_2 \in T_1 \cap T_3$ and $v_3 \in T_1 \cap T_2$ such that $v_i \notin T_i$.

Thus there exists a path $P \subseteq T_3$ joining v_1 to v_2. Along this (unique) path, choose vertices $x, y \in P$ such that $x \in T_1$ and $y \in T_2$, and no vertex lying strictly between x and y on P belongs to $T_1 \cup T_2$. By construction, x and y are distinct, although they may be adjacent.

We now construct a path Q from x to y contained entirely in $T_1 \cup T_2$. Indeed, consider the walk obtained by travelling from x to v_3 within T_1 and then from v_3 to y within T_2; let Q be a shortest path between x and y contained in this walk, and hence in $T_1 \cup T_2$.

By the choice of x and y, the path Q is internally disjoint from P. Consequently, the union $P \cup Q$ contains a cycle, contradicting the fact that T is acyclic.

Induction step: Suppose that the hypothesis holds true for some $k \geq 3$.

Consider $k+1$ subtrees, $T_1, ..., T_{k+1}$ such that $T_i \cap T_j \neq \varnothing$ for all $1 \leq i < j \leq k+1$. Applying the hypothesis for three subtrees, we get that $T_i \cap T_j \cap T_{k+1} \neq \varnothing$ i.e. $(T_i \cap T_{k+1}) \cap (T_j \cap T_{k+1}) \neq \varnothing$ for all $1 \leq i < j < k+1$. Now, $(T_1 \cap T_{k+1}), ..., (T_k \cap T_{k+1})$ is a collection of k subtrees. By the induction hypothesis,

$$\bigcap_{i=1}^{k}((T_i \cap T_{k+1})) \neq \varnothing \implies \bigcap_{i=1}^{k+1} T_i \neq \varnothing$$

□

Proposition 4.3.2. Trees are bipartite $(n, 2)$-Helly for all n.

Proof. Let T_1 and T_2 be the partite classes of a tree T and $B_1, ..., B_n$ be closed balls in T such that $(B_i \cap T_1) \cap (B_j \cap T_1) \neq \varnothing$ for all $1 \leq i < j \leq n$. Then, $B_i \cap B_j \neq \varnothing$ for all $i \neq j$ and by Proposition 4.3.1 we get that $\bigcap_{i=1}^{n} B_i \neq \varnothing$.

Pick $x \in \bigcap_{i=1}^{n} B_i$. If $x \in T_1$, then we are done.

Suppose $x \in T_2$. Fix any $1 \leq i \leq n$ and let c_i be the centre of the ball B_i. Since T is a tree, the geodesic $[c_i, x]$ is unique and the vertices alternate in parity along it. Also, by the convexity of B_i, $[c_i, x] \subseteq B_i$. Let z_i be the neighbour of x on this geodesic, then z_i lies in $B_i \cap T_1$.

Suppose z_i and z_j are distinct such that z_j does not lie in B_i and z_i does not lie in B_j. It must the case that $x \in \partial B_i \cap \partial B_j$. The removal of the vertex x splits T into two subtrees T_α and T_β such that no point in T_α lies in B_j and no point in T_β lies in B_i. But then, we would have $B_i \cap B_j \cap T_1 = \varnothing$, which is a contradiction.

If z_i and z_j are distinct such that WLOG z_j lies in B_i, then for all other z_ks which are distinct from z_i also lie in B_i, since $x \notin \partial B_i$. It follows that $x \in \bigcap_{i=1}^{n}(B_i \cap T_1)$ $\square$

Corollary 4.3.2.1. Trees are bipartite $\mathbb{Z}^d$-Kirszbraun.

Proof. Follows from Theorem 4.2.3. $\square$

Remark 4.3.3. Theorem 3.1 and Theorem 3.3 in [BC08] give general characterisations for the (n,2)-Helly and bipartite (n,2)-Helly properties respectively on graphs.

Chapter 5

Algebraic Topology on Graphs

5.1 Fundamental groups and graph covers

The fundamental group is one of the basic objects studied in algebraic topology, describing the structure of loops in a space up to deformation. For a detailed treatment we refer the reader to standard texts such as [Hat02]. In this chapter we work with a combinatorial version of this construction adapted to graphs, following the approach introduced in [CGdMO25]. While related ideas appear in [Wro17, GJKS25, KM25], the frameworks developed there do not treat self-loops in a manner suitable for applications to graph homomorphisms. This is a significant restriction, since many important examples of hom-shifts are defined using graphs with self-loops.

Notation 5.1.1. Let $p = p_0 \dots p_{\ell(p)}$ and $q = q_0 \dots q_{\ell(q)}$ be walks on a graph G, such that $p_{\ell(p)} = q_0$. We denote by $p \odot q$, the concatenated walk

$$p_0 \cdots p_{\ell(p)} q_1 \cdots q_{\ell(q)}.$$

We also denote by p^{-1}, the reverse walk

$$p_{\ell(p)} \cdots p_0.$$

Notation 5.1.2. Let φ denote the map which associates to a walk p the non-backtracking walk obtained by repeatedly removing backtracks of the form aba.

It is straightforward to verify that the resulting non-backtracking walk does not depend on the order in which backtracks are removed, and hence φ is well-defined.

For walks p and q satisfying $p_{\ell(p)} = q_0$, define the star-concatenation,

$$p \star q := \varphi(p \odot q).$$

The operation $\star$ is associative.

Notation 5.1.3. Let G be a graph and let $a \in V(G)$. We denote by $\pi_1(G, a)$ the set of all non-backtracking cycles in G that start and end at a. This set is equipped with the star-concatenation operation

$$(p, p') \longmapsto p \star p',$$

under which $\pi_1(G, a)$ forms a group. The identity element is the trivial cycle of length zero at a, and the inverse of a cycle p is given by the reverse cycle p^{-1}.

Definition 5.1.4 (Fundamental Group of a Graph)**.** *For any two vertices $a, b \in V(G)$, the groups $(\pi_1(G, a), \star)$ and $(\pi_1(G, b), \star)$ are isomorphic. Indeed, if p is any walk from b to a, the map*

$$c \longmapsto p \star c \star p^{-1}$$

defines a group isomorphism between them. The resulting isomorphism class is called the fundamental group *of G and is denoted by $\pi_1(G)$.*

In order to simplify notation, we will typically drop the choice of base vertex a and write $\pi_1(G)$ in place of $\pi_1(G, a)$. Likewise, any subsequent constructions or notations defined with respect to a base vertex will be presented without explicitly indicating it. All statements should be understood to hold independently of the chosen basepoint.

Notation 5.1.5. Let G be a graph and $a \in V(G)$. We denote by $N_G(a)$ the *neighbourhood* of a, defined by

$$N_G(a) := \{\, b \in V(G) \mid (a,b) \in E(G) \,\}.$$

In particular, if G admits self-loops, then $a \in N_G(a)$ whenever $(a,a) \in E(G)$.

Definition 5.1.6 (Covering Map of Graphs)**.** *Let G and $\widetilde{G}$ be graphs. A graph homomorphism*

$$\theta : \widetilde{G} \to G$$

is called a covering map *if it is locally bijective on neighbourhoods. More precisely, for every vertex $\tilde{a} \in V(\widetilde{G})$, the restriction*

$$\theta\big|_{N_{\widetilde{G}}(\tilde{a})} :\ N_{\widetilde{G}}(\tilde{a}) \longrightarrow N_G\big(\theta(\tilde{a})\big)$$

is a bijection. Equivalently, for each vertex $a \in V(G)$, the preimage of its neighbourhood decomposes as

$$\theta^{-1}\big(N_G(a)\big) \;=\; \bigsqcup_{\tilde{a} \in \theta^{-1}(a)} N_{\widetilde{G}}(\tilde{a}).$$

A graph $\widetilde{G}$ is said to be a *cover* of G if there exists a covering map $\theta : \widetilde{G} \to G$.

A basic feature of coverings is that walks lift uniquely. Given a walk in G and a lift of its starting vertex, there is exactly one walk in the covering graph projecting to it.

Coverings also admit symmetries called *deck transformations*. A deck transformation of a covering map $\theta : \widetilde{G} \to G$ is an automorphism η of $\widetilde{G}$ satisfying $\theta \circ \eta = \theta$. Such maps act by permuting each fibre $\theta^{-1}(a)$, for $a \in V(G)$.

5.2 Universal covers and quotient constructions

Among all covers of a graph, the universal cover plays a distinguished role, providing a canonical way to unfold all cycles while preserving local structure.

Definition 5.2.1 (Universal Cover)**.** *Let G be a graph and fix a vertex $a \in V(G)$. We define $\mathcal{U}_G[a]$ to be the graph whose vertex set consists of all non-backtracking walks in G starting at a. Two such walks are adjacent if one is obtained from the other by extending it by a single vertex. The universal cover of G, denoted $\mathcal{U}_G$, is the common isomorphism class of the graphs $\mathcal{U}_G[a]$, $a \in V(G)$.*

The isomorphism type of $\mathcal{U}_G[a]$ is independent of the choice of base vertex: if $a, b \in V(G)$ and q is a non-backtracking walk from a to b, then concatenation with q defines an isomorphism between $\mathcal{U}_G[a]$ and $\mathcal{U}_G[b]$. We therefore identify $\mathcal{U}_G$ with any representative $\mathcal{U}_G[a]$ and suppress the basepoint from the notation.

Each vertex of $\mathcal{U}_G$ is a walk in G, and hence has a well-defined endpoint. Sending a walk to its endpoint defines a graph homomorphism

$$\alpha : \mathcal{U}_G \to G,$$

which is a covering map. In particular, $\mathcal{U}_G$ is a cover of G. Moreover, $\mathcal{U}_G$ is acyclic.

The fundamental group acts naturally on the universal cover. If $g \in \pi_1(G)$ and w is a walk starting at the base vertex, then the concatenation

$$\eta_g(w) = g \star w$$

produces another vertex of $\mathcal{U}_G$. These maps preserve the projection α, and therefore act as deck transformations.

More generally, covers of G can be obtained from the universal cover by taking quotients. If Γ is a subgroup of $\pi_1(G)$, the action of Γ on $\mathcal{U}_G$ yields a quotient graph $\mathcal{U}_G/\Gamma$. The projection α then induces a covering map

$$\alpha_\Gamma : \mathcal{U}_G/\Gamma \to G.$$

In this way every covering graph of G arises from a subgroup of the fundamental group.

When Γ is normal in $\pi_1(G)$ the quotient group $\pi_1(G)/\Gamma$ acts on $\mathcal{U}_G/\Gamma$ by deck transformations. Such coverings are called regular covers. The universal cover corresponds to the trivial subgroup.

5.3 Square groups, square covers, and square lifting

Let G be a graph. A *square* in G is a non-backtracking cycle of length four. In other words it is a cycle

$$a_0 a_1 a_2 a_3 a_4$$

with $a_0 = a_4$, $a_0 \neq a_2$, and $a_1 \neq a_3$.

Squares naturally appear when studying lifts of configurations $x \in X_G^d$. If $p = (p_0, \ldots, p_k)$ is a walk in $\mathbb{Z}^d$, reading the configuration along p produces a walk

$$x_p = x_{p_0} x_{p_1} \ldots x_{p_k}$$

in G.

Suppose $\theta : \widetilde{G} \to G$ is a covering map and we try to lift a configuration x to a configuration $\tilde{x} \in X_{\widetilde{G}}^d$. Choose a lift $\tilde{v}$ of $x_{\mathbf{0}}$. For a vertex $u \in \mathbb{Z}^d$, one may attempt to define $\tilde{x}_u$ by choosing a path p from $\mathbf{0}$ to u and lifting the walk x_p starting at $\tilde{v}$. The endpoint of this lifted walk would then be taken as $\tilde{x}_u$.

For this construction to be well defined, the result must not depend on the chosen path p. The difficulty is that two lattice paths with the same endpoints typically differ by a sequence of elementary squares. Consequently the walks in G obtained by reading the configuration along these paths differ by inserting or removing cycles of length four. Understanding how such cycles lift therefore becomes essential. In particular, we would want squares in G to lift to squares in $\tilde{G}$.

This leads to the following construction. Let $\Delta(G)[a]$ denote the subgroup of $\pi_1(G)[a]$ generated by all conjugates of squares, that is, by cycles of the form

$$p \star s \star p^{-1},$$

where p is a non-backtracking walk based at a and s is a square in G. These groups are naturally isomorphic for different choices of base vertex, and we write $\Delta(G)$ for their common isomorphism class. Elements of $\Delta(G)$ will be referred to as *square-decomposable cycles*. For more references regarding this, consult Section 4 of

[CGdMO25].

Definition 5.3.1. *[CGdMO25][Square Group] The* square group *of G is defined as the quotient*

$$\pi_1^{\square}(G) := \pi_1(G)/\Delta(G).$$

Informally, this group is obtained from the fundamental group by forcing all squares to behave like trivial cycles.

As usual, subgroups of the fundamental group correspond to quotients of the universal cover. Applying this correspondence to $\Delta(G)$ produces the graph

$$\mathcal{U}_G^{\square} := \mathcal{U}_G/\Delta(G),$$

which we call the *square cover* of G. The associated covering map will be denoted

$$\alpha^{\square} : \mathcal{U}_G^{\square} \to G.$$

The relevance of this construction can be seen by examining how squares lift in covering graphs. Suppose that

$$\widetilde{G} = \mathcal{U}_G/\Gamma$$

is a quotient of the universal cover corresponding to a subgroup $\Gamma \subseteq \pi_1(G)$. One may then ask when a square in G lifts to a square in $\widetilde{G}$.

The answer is that this happens precisely when the subgroup Γ contains the square subgroup $\Delta(G)$. In other words, squares in G lift to squares in $\widetilde{G}$ exactly when

$$\Delta(G) \subseteq \Gamma.$$

In particular, the square cover

$$\mathcal{U}_G^{\square} = \mathcal{U}_G/\Delta(G)$$

is the largest cover of G in which every square of G lifts to a square.

5.4 Configuration lifting

Let $\theta : \widetilde{G} \to G$ be a covering map. A configuration $\tilde{x} \in X^d_{\widetilde{G}}$ is called a *lift* of $x \in X^d_G$ if

$$\theta(\tilde{x}_u) = x_u \qquad \text{for all } u \in \mathbb{Z}^d.$$

The following result characterizes when configurations admit lifts.

Proposition 5.4.1. [CGdMO25] Let G be a graph and let $\Gamma \subseteq \pi_1(G)$ be a subgroup. Set

$$\tilde{G} := \mathcal{U}_G/\Gamma,$$

and let $\theta : \tilde{G} \to G$ be the associated covering map. Then

$$\Delta(G) \subseteq \Gamma$$

if and only if the following lifting property holds:

For every configuration $x \in X^d_G$ and every lift w of $x_{\mathbf{0}}$ to $\tilde{G}$, there exists a lift

$$\tilde{x} \in X^d_{\tilde{G}}$$

of x with respect to θ such that

$$\tilde{x}_{\mathbf{0}} = w.$$

As a consequence, the square cover is maximal among covers admitting lifts of all configurations.

Corollary 5.4.1.1. [CGdMO25] Let G be a graph, let $\tilde{G}$ be a cover of G, and let $\theta : \tilde{G} \to G$ be a covering map. Assume that every configuration $x \in X^d_G$ admits a lift $\tilde{x} \in X^d_{\tilde{G}}$ such that

$$\theta \circ \tilde{x} = x.$$

Then there exists a covering map

$$\theta' : \mathcal{U}^{\square}_G \to \tilde{G}$$

satisfying

$$\theta \circ \theta' = \alpha^{\square}.$$

We also record the following uniqueness property of lifts for regular covers.

Proposition 5.4.2. [CGdMO25] Let G be a graph and $\tilde{G}$ a regular cover of G. Let

$$\theta : \tilde{G} \to G$$

be the covering map. Let $x \in X_G^d$ and suppose $\tilde{x}, \tilde{y} \in X_{\tilde{G}}^d$ are two lifts of x. Then there exists a deck transformation η such that

$$\tilde{y}_{\vec{i}} = \eta(\tilde{x}_{\vec{i}}) \quad \text{for all } \vec{i} \in \mathbb{Z}^d.$$

5.5 Graphs whose square cover is a tree

There are three graphs which satisfy the definition of a square.

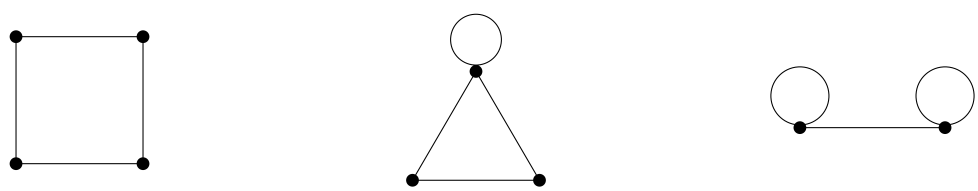

Figure 5.1: ‘Squares’

Lemma 5.5.1. The square cover of a graph G is a tree iff G does not have a square as a subgraph.

Proof. Suppose that G doesn't have a square as a subgraph. Then $\mathcal{U}_G$ is the square cover since vacuously all squares lift to squares. On the other hand if the square cover is a tree, then its fundamental group must be trivial i.e. $\Delta(G)$ is trivial. $\Delta(G)$ is trivial iff there are no squares in G. $\square$

Chapter 6

Height Functions and Slopes

6.1 Height functions.

The existence of lifts to the square cover, established in the previous chapter, allows us to quantify the ‘rigidity’ of configurations. In this section we introduce height functions and, subsequently, the notion of slope, where larger gradients correspond to increased ‘rigidity’. The general method of employing height functions were introduced by J.H. Conway and J.C. Lagarias in [CL90] and popularised by W.P. Thurston in [Thu90]. In the context of shift spaces, height functions were introduced by Klaus Schmidt in [Sch97]. We will use a version similar to that used in [Cha16].

Definition 6.1.1. *Let G be a finite connected graph and let $x \in X^d_G$. Choose a lift $\tilde{x} \in X^d_{\mathcal{U}^{\square}_G}$ of x to the square cover $\mathcal{U}^{\square}_G$. We define the associated height function*

$$h_x : \mathbb{Z}^d \times \mathbb{Z}^d \longrightarrow \mathbb{Z}$$

by

$$h_x(\vec{i}, \vec{j}) := d_{\mathcal{U}^{\square}_G}(\tilde{x}_{\vec{i}}, \tilde{x}_{\vec{j}}).$$

The existence of the lift is guaranteed by Proposition 5.4.1. By the uniqueness of lifts (to regular covers) up to deck transformations (Proposition 5.4.2), this definition is independent of the choice of lift $\tilde{x}$, and hence h_x is well defined.

Remark 6.1.2. The height function h_x is subadditive i.e. for any $x \in X_G^d$ and any $\vec{i}, \vec{j}, \vec{k} \in \mathbb{Z}^d$, we have

$$h_x(\vec{i}, \vec{j}) \ \leq \ h_x(\vec{i}, \vec{k}) + h_x(\vec{k}, \vec{j}).$$

This follows from subadditivity in the graph metric.

6.2 Slope

Definition 6.2.1 (Slope). *Let $\vec{i} = (i_1, \ldots, i_d) \in \mathbb{R}^d$ and write*

$$\lfloor \vec{i} \rfloor := (\lfloor i_1 \rfloor, \ldots, \lfloor i_d \rfloor).$$

For $n \in \mathbb{N}$, set

$$\lfloor n\vec{i} \rfloor := (\lfloor ni_1 \rfloor, \ldots, \lfloor ni_d \rfloor).$$

For $x \in X_G^d$, the slope *of x in the direction $\vec{i}$, whenever the limit exists, is defined by*

$$\mathrm{sl}_{\vec{i}}(x) := \lim_{n \to \infty} \frac{1}{n} \, h_x\big(\mathbf{0}, \ \lfloor n\vec{i} \rfloor\big).$$

Fix $\vec{i} \in \mathbb{Z}^d$ and consider the sequence of functions

$$f_n : X_G^d \longrightarrow \mathbb{N} \cup \{0\}, \qquad f_n(x) := h_x(\mathbf{0}, n\vec{i}).$$

By subadditivity of the height function, the sequence $(f_n)_{n \geq 1}$ satisfies the hypotheses of the subadditive ergodic theorem with respect to any shift-invariant probability measure on X_G^d.

Indeed, the bound

$$f_1(x) = h_x(\mathbf{0}, \vec{i}) \leq \|\vec{i}\|_1$$

holds. Moreover, if we write $T := \sigma_{\vec{i}}$, then the cocycle (subadditivity) relation becomes

$$f_{n+m}(x) = h_x(\mathbf{0}, (n+m)\vec{i}) \leq h_x(\mathbf{0}, n\vec{i}) + h_{\sigma_{n\vec{i}}x}(\mathbf{0}, m\vec{i}) = f_n(x) + f_m(T^n x).$$

Thus (f_n) forms a subadditive cocycle over the $\mathbb{Z}^d$-shift action.

The asymptotic behaviour of height functions along lattice directions therefore follows from the subadditive ergodic theorem (Theorem 3.3.4). In what follows, an ergodic measure on X_G^d will always mean a shift-invariant probability measure that is ergodic with respect to the $\mathbb{Z}^d$-shift action.

Proposition 6.2.2. [Cha16][Existence of Slopes] Let G be a finite connected graph and let μ be an ergodic measure on X_G^d. Then for every $\vec{i} \in \mathbb{Z}^d$ the limit

$$\mathrm{sl}_{\vec{i}}(x) = \lim_{n\to\infty} \frac{1}{n} h_x(\mathbf{0}, n\vec{i})$$

exists for μ-almost every x and is independent of x.

Furthermore, if $\vec{i} = (i_1, \ldots, i_d)$, then

$$\mathrm{sl}_{\vec{i}}(x) \leq \sum_{k=1}^{d} |i_k| \, \mathrm{sl}_{\vec{e}_k}(x) \qquad \text{for } \mu\text{-a.e. } x.$$

Corollary 6.2.2.1. [Cha16] Let G be a finite connected graph and let μ be an ergodic measure on X_G^d. Then for every $\vec{i} \in \mathbb{R}^d$ the limit

$$\mathrm{sl}_{\vec{i}}(x) = \lim_{n\to\infty} \frac{1}{n} h_x\big(\mathbf{0}, \lfloor n\vec{i} \rceil\big)$$

exists for μ-almost every x and is independent of x.

Moreover, if $\vec{i} = (i_1, i_2, \ldots, i_d)$, then

$$\mathrm{sl}_{\vec{i}}(x) \leq \sum_{k=1}^{d} |i_k| \, \mathrm{sl}_{\vec{e}_k}(x) \qquad \text{for } \mu\text{-a.e. } x.$$

Chapter 7

Entropy Minimality

We now have the machinery to deal with the question of entropy minimality of hom-shifts for graphs whose square cover is a tree.

7.1 Characterising zero entropy measures

In this section we show that existence of maximal slope in some direction is tied to the entropy, corresponding to an ergodic probability measure on X^d_G, being zero.

Lemma 7.1.1. Let G be a finite connected graph with a tree square cover and let μ be an ergodic probability measure on X^d_G. Suppose that for some $1 \leq k \leq d$ we have

$$\mathrm{sl}_{\vec{e}_k}(x) = 1 \qquad \text{for } \mu\text{-almost every } x.$$

Then μ is frozen. In particular,

$$h_\mu = 0.$$

Proof. **Step 1:** WLOG assume that $\mathrm{sl}_{\vec{e}_1}(x) = 1$ almost everywhere.

Fix $k, n \in \mathbb{N}$. By subadditivity of the height function,

$$h_x(\mathbf{0}, kn\vec{e}_1) \leq \sum_{m=0}^{n-1} h_x(km\vec{e}_1,\, k(m+1)\vec{e}_1).$$

Also,

$$h_x(km\vec{e}_1,\, k(m+1)\vec{e}_1) = h_{\sigma^{km\vec{e}_1}x}(\mathbf{0}, k\vec{e}_1).$$

Dividing by kn gives

$$\frac{1}{kn} h_x(\mathbf{0}, kn\vec{e}_1) \le \frac{1}{n}\sum_{m=0}^{n-1} \frac{1}{k} h_{\sigma^{km\vec{e}_1}x}(\mathbf{0}, k\vec{e}_1) \le 1. \tag{7.1}$$

By assumption,

$$\lim_{n\to\infty} \frac{1}{kn} h_x(\mathbf{0}, kn\vec{e}_1) = 1 \qquad \text{for } \mu\text{-a.e. } x.$$

Hence taking limits in (7.1) yields

$$\lim_{n\to\infty} \frac{1}{n}\sum_{m=0}^{n-1} \frac{1}{k} h_{\sigma^{km\vec{e}_1}x}(\mathbf{0}, k\vec{e}_1) = 1 \qquad \text{a.e.}$$

Now set

$$g(x) := \frac{1}{k} h_x(\mathbf{0}, k\vec{e}_1).$$

Then the previous term is precisely the Birkhoff average of g with respect to the measure-preserving transformation $T = \sigma^{k\vec{e}_1}$. By the ergodic theorem (Theorem 3.3.2) it follows that if $g^* = \lim_{n\to\infty} \frac{1}{n}\sum_{m=0}^{n-1} g(T^m x)$ then $\int g^* d\mu = \int g d\mu$. But g^* is 1 almost surely. Thus it follows that $\int g d\mu = 1$. In other words,

$$\int \frac{1}{k} h_x(\mathbf{0}, k\vec{e}_1)\, d\mu = 1.$$

On the other hand,

$$h_x(\mathbf{0}, k\vec{e}_1) \le k \quad \text{for all } x,$$

so $g(x) \le 1$ everywhere. Since its integral equals 1, we must have

$$h_x(\mathbf{0}, k\vec{e}_1) = k \qquad \text{for } \mu\text{-a.e. } x.$$

Finally,

$$h_x(\vec{i}, \vec{i} + k\vec{e}_1) = h_{\sigma^{\vec{i}}x}(\mathbf{0}, k\vec{e}_1) = k \tag{7.2}$$

for all $\vec{i} \in \mathbb{Z}^d$ and $k \in \mathbb{N}$, almost everywhere.

Step 2: Frozen support.

Let
$$X := \{x \in X_G^d \ : \ x \text{ satisfies } (7.2)\}.$$
By Step 1, $\mu(X) = 1$.

Claim. Every finite pattern appearing in $\mathrm{supp}(\mu)$ extends to a configuration in X.

Let a be a finite pattern appearing in $\mathrm{supp}(\mu)$ and let $[a]$ denote the corresponding cylinder set. By definition of support (Proposition 3.4.2), $\mu([a]) > 0$. Since $\mu(X) = 1$, we have $\mu(X^c) = 0$, and therefore
$$\mu([a] \cap X) = \mu([a]) - \mu([a] \cap X^c) = \mu([a]) > 0.$$
Hence $[a] \cap X \neq \varnothing$, proving the claim.

We denote $[-n, n]^d$ by B_n. Fix $n \in \mathbb{N}$ and let
$$a, b \in \mathcal{L}_{B_n \cup \partial_2 B_n}(\mathrm{supp}(\mu))$$
satisfy
$$a|_{\partial_2 B_n} = b|_{\partial_2 B_n}.$$
By the claim, there exist configurations $x, y \in X$ such that
$$x|_{B_n \cup \partial_2 B_n} = a, \qquad y|_{B_n \cup \partial_2 B_n} = b.$$
Let
$$\alpha^{\square} : \mathcal{U}_G^{\square} \to G$$
be the square covering map. Fix $\vec{j_0} \in \partial_2 B_n$. Then by Proposition 5.4.1 consider the lifts $\tilde{x}, \tilde{y}$ such that $\tilde{x}_{\vec{j_0}} = \tilde{y}_{\vec{j_0}}$. Since $\partial_2 B_n$ is connected (Note that ∂B_n is not connected which explains the choice of 2-boundary) it follows that $\tilde{x}_{\vec{j}} = \tilde{y}_{\vec{j}}$ for all $\vec{j} \in \partial_2 B_n$.

Fix $\vec{i} \in B_n$. Choose $k \in -\mathbb{N}$ such that

$$\vec{i} + k\vec{e}_1, \quad \vec{i} + (2n+2+k)\vec{e}_1 \in \partial_2 B_n.$$

By (7.2),

$$d_{\mathcal{U}_G^\square}\big(\tilde{x}_{\vec{i}+k\vec{e}_1}, \tilde{x}_{\vec{i}+(2n+2+k)\vec{e}_1}\big) = 2n+2.$$

Consider the walks

$$(\tilde{x}_{\vec{i}+k\vec{e}_1}, \ldots, \tilde{x}_{\vec{i}+(2n+2+k)\vec{e}_1})$$

and

$$(\tilde{y}_{\vec{i}+k\vec{e}_1}, \ldots, \tilde{y}_{\vec{i}+(2n+2+k)\vec{e}_1}).$$

Each walk has length $2n+2$ and joins the same endpoints, since the lifts agree on $\partial_2 B_n$. Because the graph distance between the endpoints is $2n+2$, both walks realise the distance and hence are geodesics.

Since $\mathcal{U}_G^\square$ is a tree, geodesics between two vertices are unique. Therefore the two walks coincide, and in particular

$$\tilde{x}_{\vec{i}} = \tilde{y}_{\vec{i}}.$$

As $\vec{i}$ was arbitrary, we obtain

$$\tilde{x}|_{B_n} = \tilde{y}|_{B_n}.$$

Applying $\alpha^\square$ yields

$$a|_{B_n} = b|_{B_n}.$$

This completes the proof of Step 2. Now we will complete the proof of Lemma 7.1.1.

We have that $\mathrm{supp}(\mu)$ is frozen. Consequently,

$$|\mathcal{L}_{B_n}(\mathrm{supp}(\mu))| \leq |G^{|\partial_2 B_n|}|,$$

so $h_{\mathrm{top}}(\mathrm{supp}(\mu)) = 0$, and by the variational principle, $h_\mu = 0$.

□

7.2 Full Supports

In this section, we will show that for a finite connected graph G (we do not consider the trivial cases of a single vertex with a self-loop or a single edge since they have been considered in Chapter 2) with a tree square cover, every ergodic shift-invariant probability measure adapted to X^d_G with positive entropy is fully supported. This will constitute a proof for Theorem 1.0.1 using the following argument.
By Theorem 3.6.5, every measure of maximal entropy on X^d_G is a uniform shift-invariant Markov random field adapted to X^d_G. The topological entropy of X^d_G is positive (Proposition 2.3.5 and Proposition 2.3.6), barring the trivial cases which had been dealt with in Chapter 2. By the Variational Principle (Theorem 3.5.5), these measures have positive entropy. If we check that ergodic measures adapted to X^d_G are fully supported, then by Theorem 3.6.6, all the measures of maximal entropy are fully supported and then Proposition 3.5.6 implies Theorem 1.0.1.

Some of the ideas underlying the following proof are inspired by [Cha16] and by Lemma 6.4 of [CM16].

Step 1: Non maximal slopes.

Let μ be an ergodic shift-invariant probability measure adapted to X^d_G, such that $h_\mu > 0$. If for some $1 \leq k \leq d$ we had

$$\mathrm{sl}_{\vec{e}_k}(\mu) = 1,$$

then by Lemma 7.1.1 the measure μ would be frozen and hence have zero entropy, contradicting $h_\mu > 0$. Therefore

$$\mathrm{sl}_{\vec{e}_k}(\mu) < 1 \qquad \text{for all } 1 \leq k \leq d.$$

Set

$$\theta := \max_{1\leq k\leq d} \mathrm{sl}_{\vec{e}_k}(\mu),$$

so that $\theta < 1$.

By Corollary 6.2.2.1, for every $\vec{v} \in \mathbb{R}^d$ the slope $\mathrm{sl}_{\vec{v}}(\mu)$ exists almost everywhere and satisfies

$$\mathrm{sl}_{\vec{v}}(\mu) \leq \sum_{k=1}^{d} |v_k| \, \mathrm{sl}_{\vec{e}_k}(\mu).$$

Let $\mathcal{S}^{d-1}$ denote the unit sphere in $\mathbb{R}^d$ with respect to the ℓ^1 norm. Then in particular, if $\vec{v} \in \mathcal{S}^{d-1}$,

$$\mathrm{sl}_{\vec{v}}(\mu) \leq \theta < 1.$$

Thus slopes in all directions are uniformly bounded away from 1.

Step 2: Boundary heights.

Let $\theta < 1$ be as obtained in Step 1 and fix β such that

$$\theta < \beta < 1.$$

Choose $\varepsilon > 0$ small enough so that

$$\theta + 3\varepsilon < \beta.$$

Since $\mathcal{S}^{d-1}$ is compact in the ℓ^1-metric, there exist finitely many vectors $\vec{v}_1, \ldots, \vec{v}_t \in \mathcal{S}^{d-1}$ such that for every $\vec{v} \in \mathcal{S}^{d-1}$ there exists $1 \leq i \leq t$ satisfying

$$\|\vec{v} - \vec{v}_i\|_1 < \varepsilon.$$

Uniform convergence via Egoroff's theorem.

For each $1 \leq i \leq t$ we have

$$\frac{1}{n} h_x\big(\mathbf{0}, \lfloor n\vec{v}_i \rfloor\big) \longrightarrow \mathrm{sl}_{\vec{v}_i}(\mu) \leq \theta \quad \text{for } \mu\text{-almost every } x.$$

Fix $\varepsilon > 0$. For each i, apply Egoroff's theorem with parameter ε/t. Since μ is a probability measure, for each $1 \leq i \leq t$ there exists a measurable set $A_i \subset X_G^d$ such that

$$\mu(A_i) > 1 - \frac{\varepsilon}{t},$$

and an integer N_i for which convergence is uniform on A_i. In particular, for all

$x \in A_i$ and all $n \geq N_i$,

$$\frac{1}{n} h_x\big(\mathbf{0}, \lfloor n\vec{v}_i \rfloor\big) \leq \theta + \varepsilon.$$

Let

$$A := \bigcap_{i=1}^{t} A_i, \qquad N := \max_{1 \leq i \leq t} N_i.$$

Then

$$\mu(A) \geq 1 - \sum_{i=1}^{t} \frac{\varepsilon}{t} = 1 - \varepsilon.$$

Moreover, for every $x \in A$, every $n \geq N$, and every $1 \leq i \leq t$,

$$\frac{1}{n} h_x\big(\mathbf{0}, \lfloor n\vec{v}_i \rfloor\big) \leq \theta + \varepsilon.$$

Extension to all boundary points.

Fix $x \in A$ and $n \geq N$. Let $\vec{i} \in \partial D_n$. Then,

$$\vec{v} := \frac{\vec{i}}{n} \in \mathcal{S}^{d-1}.$$

Choose $1 \leq i_0 \leq t$ such that $\|\vec{v} - \vec{v}_{i_0}\|_1 < \varepsilon$. Then,

$$\|\vec{i} - n\vec{v}_{i_0}\|_1 = n\|\vec{v} - \vec{v}_{i_0}\|_1 < \varepsilon n.$$

Using subadditivity of the height function,

$$h_x(\mathbf{0}, \vec{i}) \leq h_x\big(\mathbf{0}, \lfloor n\vec{v}_{i_0} \rfloor\big) + h_x\big(\lfloor n\vec{v}_{i_0} \rfloor, \vec{i}\big).$$

The first term satisfies $h_x\big(\mathbf{0}, \lfloor n\vec{v}_{i_0} \rfloor\big) \leq (\theta + \varepsilon)n$ by the uniform estimate on A. For the second term, since the height function is 1-Lipschitz with respect to the ℓ^1 metric,

$$h_x\big(\lfloor n\vec{v}_{i_0} \rfloor, \vec{i}\big) \leq \|\vec{i} - \lfloor n\vec{v}_{i_0} \rfloor\|_1 \leq \|\vec{i} - n\vec{v}_{i_0}\|_1 + d < \varepsilon n + d,$$

where d accounts for the rounding error (It is the d in $\mathbb{Z}^d$).

Therefore,

$$h_x(\mathbf{0}, \vec{i}) \leq (\theta + 2\varepsilon)n + d.$$

In order to deduce the desired bound $h_x(\mathbf{0}, \vec{i}) \leq (\theta + 3\varepsilon)n$, it suffices to ensure that

$$(\theta + 2\varepsilon)n + d \leq (\theta + 3\varepsilon)n.$$

i.e.

$$d \leq \varepsilon n.$$

Since d is independent of n, choose $N_1 := \left\lceil \frac{d}{\varepsilon} \right\rceil$. Then for all $n \geq N_1$ we have $d \leq \varepsilon n$, and therefore

$$(\theta + 2\varepsilon)n + d \leq (\theta + 3\varepsilon)n.$$

Replacing the previous threshold N by $N' := \max(N, N_1)$, we have that for all $n \geq N'$,

$$h_x(\mathbf{0}, \vec{i}) \leq (\theta + 3\varepsilon)n.$$

Since $\theta + 3\varepsilon < \beta$, we conclude that for every $x \in A$, every $n \geq N'$, and every $\vec{i} \in \partial D_n$,

$$h_x(\mathbf{0}, \vec{i}) \leq \beta n.$$

In particular, we can choose $x \in A \cap \operatorname{supp}(\mu)$.

Step 3: Arbitrary finite patterns are in the support.

Let $F \subset \mathbb{Z}^d$ be finite and let $P \in \mathcal{L}_F(X_G^d)$ be any globally-allowed pattern in X_G^d. We show that P appears in $\operatorname{supp}(\mu)$.

Choose $\delta > 0$ such that

$$\beta < 1 - \delta < 1.$$

Fix n sufficiently large so that $F \subset D_{\delta n}$ and $n > N'$, and $x \in \operatorname{supp}(\mu)$ be a configuration as chosen in Step 2. Let $u \in F$ and $v \in \partial D_n$.

$$\|u - v\|_1 \geq |\|u\|_1 - \|v\|_1| \geq n - \delta n = (1 - \delta)n.$$

Let $\tilde{x}$ be a lift of x. Lift the pattern P to a map $\tilde{P} : F \to \mathcal{U}_G^\square$ satisfying $\alpha^\square \circ \tilde{P} = P$. Since P is globally-allowed, the lift exists. Also, $\tilde{P}$ is 1-Lipschitz.

We now define a partial map

$$f : \partial D_n \cup F \to \mathcal{U}_G^{\square}$$

by

$$f|_{\partial D_n} = \tilde{x}|_{\partial D_n}, \qquad f|_F = \tilde{P}.$$

On the other hand, by the triangle inequality in the tree,

$$d_{\mathcal{U}_G^{\square}}\big(\tilde{P}(u), \tilde{x}(v)\big) \le d_{\mathcal{U}_G^{\square}}\big(\tilde{P}(u), \tilde{x}_0\big) + d_{\mathcal{U}_G^{\square}}\big(\tilde{x}_0, \tilde{x}(v)\big).$$

The second term is at most βn due to our choice of x.

Since F is finite, the set

$$\{d_{\mathcal{U}_G^{\square}}(\tilde{P}(u), \tilde{x}_{\mathbf{0}}) : u \in F\}$$

is finite. Define

$$M := \max_{u \in F} d_{\mathcal{U}_G^{\square}}(\tilde{P}(u), \tilde{x}_{\mathbf{0}}).$$

Since $\beta < 1 - \delta$, the inequality

$$M + \beta n \le (1 - \delta) n$$

holds for all sufficiently large n. Thus the map f is 1-Lipschitz on its domain.

Since $\mathcal{U}_G^{\square}$ is a tree, it is bipartite $\mathbb{Z}^d$-Kirszbraun. If the parities of the patterns on F and ∂D_n do not match, we can reconcile them by shifting the pattern P by a unit vector. Since $\mathrm{supp}(\mu)$ is shift-invariant, this does not cause any problems. Hence f extends to a graph homomorphism

$$\tilde{f} : \mathbb{Z}^d \to \mathcal{U}_G^{\square}.$$

Projecting back to G gives a configuration

$$y := \alpha^{\square} \circ \tilde{f} \in X_G^d.$$

By construction,

$$y|_F = P, \qquad y|_{\partial D_n} = x|_{\partial D_n}.$$

Thus y differs from x only on the finite set D_n. Since μ is adapted, its support is closed under finite modifications, and hence

$$y \in \operatorname{supp}(\mu).$$

Therefore P appears in $\operatorname{supp}(\mu)$.

Chapter 8

Further Directions

8.1 Tilings and height functions

In this chapter we discuss a possible extension of the ideas developed in this thesis to tilings of $\mathbb{Z}^2$ by translates of the two tiles with dimensions $1 \times m$ and $n \times 1$. This is joint work with Aditya Thorat and Nishant Chandgotia.

Our eventual goal is to understand whether methods analogous to those used for hom-shifts can be applied in this setting, with the hope of proving entropy minimality for these tiling spaces as well. A first step in this direction is to understand when prescribed height values on a finite subset of $\mathbb{Z}^2$ can be extended to a global tiling. One may view such an extension criterion as playing a role analogous to that of a Kirszbraun-type theorem in the hom-shift setting.

This chapter is inspired by the work of Martin Tassy in his PhD dissertation, in particular Theorem 6 of [Tas12]. The original work contains some gaps and typos, and we provide an alternative proof here.

Consider tilings of $\mathbb{Z}^2$ by translates of the $1 \times m$ and $n \times 1$ rectangles $(m, n > 1)$, where the vertices of the tiles lie in $\mathbb{Z}^2$. These tilings admit a natural height-function description, introduced by Conway, Lagarias, and Thurston (see [CL90, Thu90]). The height function is defined on the lattice vertices $\mathbb{Z}^2$ and takes values in a group generated by horizontal and vertical increments, with relations expressing the fact that traversing the boundary of an $m \times 1$ tile or a $1 \times n$ tile produces no net change

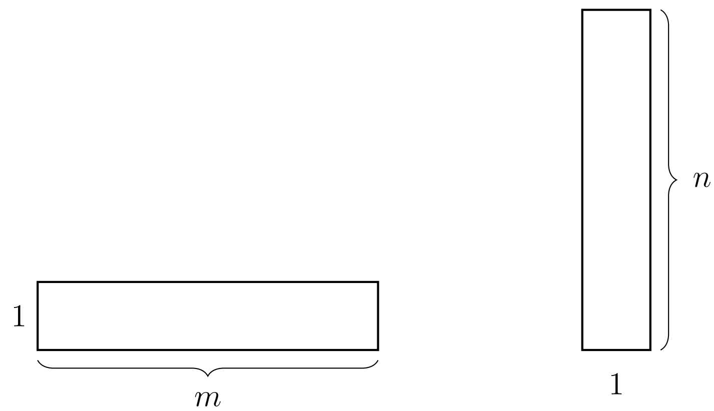

Figure 8.1: The $1 \times m$ and $n \times 1$ tiles.

in height. Thus one may begin with the group

$$\langle h, v \mid h^m v h^{-m} v^{-1},\ v^n h v^{-n} h^{-1} \rangle.$$

We will describe this height function more formally below. However, for the extension problem considered here, it is more convenient to work with the quotient

$$G = \langle h, v \mid h^m = v^n = 1 \rangle \cong \mathbb{Z}_m * \mathbb{Z}_n.$$

This quotient retains the algebraic information relevant for the compatibility conditions on height differences, while being much easier to handle. Claire and Richard Kenyon had used these height functions to come up with a linear-time algorithm to decide if a polygon can be tiled by these tiles in [KK92].

Given a tiling τ of $\mathbb{Z}^2$ by $1 \times m$ and $n \times 1$ rectangles, one obtains a G-valued height function in the usual way. Choose a basepoint $x_0 \in \mathbb{Z}^2$, an element $g \in G$, and set $h(\tau, x_0) = g$. For any other point $x \in \mathbb{Z}^2$, choose a lattice path from x_0 to x along the boundaries of the tiles and read off a word in the generators h and v by recording horizontal and vertical steps along the path. Passing one unit to the right contributes a factor of h, passing one unit to the left contributes h^{-1}, passing one unit upward contributes v, and passing one unit downward contributes v^{-1}. The relations imposed by the tiles ensure that the resulting group element depends only on the endpoint x and not on the choice of path. We multiply the obtained word by

g on the left since that is the height value we chose for the basepoint. In this way one obtains a well-defined function

$$h(\tau, \cdot) : \mathbb{Z}^2 \to G.$$

Equivalently, the height difference between two points $x, y \in \mathbb{Z}^2$ is obtained by evaluating the group element associated to any lattice path following tile boundaries from x to y. The relations $h^m = 1$ and $v^n = 1$ reflect the fact that traversing across the length of a horizontal $1 \times m$ tile or a vertical $n \times 1$ tile produces no net change in height in the quotient group G.

8.2 A necessary and sufficient condition for tileability

Let

$$A = \{x_1, x_2, \ldots, x_N\} \subset \mathbb{Z}^2,$$

and suppose that for each $x_i \in A$ we are given a height value $g_i \in G$. For $i \neq j$, the element $g_i^{-1} g_j \in G$ admits a unique reduced expression in free product normal form, which we write as

$$g_i^{-1} g_j = h^{\alpha_1} v^{\beta_1} h^{\alpha_2} v^{\beta_2} \cdots h^{\alpha_\ell} v^{\beta_\ell},$$

where

$$\begin{aligned}
\alpha_t &= \alpha_t(x_i, x_j) \in [1, m-1] \quad && \text{for } 2 \leq t \leq \ell = \ell(x_i, x_j), \\
\alpha_1 &= \alpha_1(x_i, x_j) \in [0, m-1], && \\
\beta_t &= \beta_t(x_i, x_j) \in [1, n-1] \quad && \text{for } 1 \leq t \leq \ell - 1, \\
\beta_\ell &= \beta_\ell(x_i, x_j) \in [0, n-1]. &&
\end{aligned}$$

Thus all exponents are taken to be non-negative, with the first exponent of h and the last exponent of v allowed to be zero so as to include all possible reduced words.

We now state a necessary and sufficient condition, proved below, for the prescribed values $\{g_1, \ldots, g_N\}$ to extend to the height function of a tiling of $\mathbb{Z}^2$ by

$1 \times m$ and $n \times 1$ rectangles. In particular, when $S \subset \mathbb{Z}^2$ is a finite simply connected region, this criterion can be used to determine whether S admits such a tiling. From the point of view of the present thesis, the significance of this result is that it suggests a possible analogue of the Kirszbraun property in this tiling setting, and hence a possible route toward proving entropy minimality for these tiling spaces.

Theorem 8.2.1. For $i \neq j$, write

$$x_j - x_i = (a_{ij}, b_{ij}).$$

Then there exists a tiling τ of $\mathbb{Z}^2$ by $1 \times m$ and $n \times 1$ rectangles such that

$$h(\tau, x_i) = g_i \qquad \text{for all } 1 \leq i \leq N$$

if and only if for every pair $i \neq j$, there exist integers $k = k(i,j)$ and $k' = k'(i,j)$, with at least one of k and k' non-positive, such that the following hold:

1. If $a_{ij} \geq 0$ and $b_{ij} \geq 0$, then
$$\sum_t \alpha_t(x_i, x_j) - a_{ij} = km, \qquad \sum_t \beta_t(x_i, x_j) - b_{ij} = k'n.$$

2. If $a_{ij} \leq 0$ and $b_{ij} \geq 0$, then
$$\sum_t \big(m - \alpha_t(x_i, x_j)\big) + a_{ij} = km, \qquad \sum_t \beta_t(x_i, x_j) - b_{ij} = k'n.$$

3. If $a_{ij} \leq 0$ and $b_{ij} \leq 0$, then
$$\sum_t \big(m - \alpha_t(x_i, x_j)\big) + a_{ij} = km, \qquad \sum_t \big(n - \beta_t(x_i, x_j)\big) + b_{ij} = k'n.$$

4. If $a_{ij} \geq 0$ and $b_{ij} \leq 0$, then
$$\sum_t \alpha_t(x_i, x_j) - a_{ij} = km, \qquad \sum_t \big(n - \beta_t(x_i, x_j)\big) + b_{ij} = k'n.$$

Proof.

Necessity:

We first prove necessity. Suppose that there exists a tiling τ of $\mathbb{Z}^2$ such that

$$h(\tau, x_i) = g_i \qquad \text{for } 1 \leq i \leq N.$$

Case 1a: $x_j - x_i = (a_{ij}, 0)$, where $a_{ij} \geq 0$. Since τ is a tiling of $\mathbb{Z}^2$, there exists a non-backtracking path along tile boundaries joining x_i to x_j for which every horizontal step is in the positive x-direction, as shown in Figure 8.2. To show the existence of such a path, note that it is possible when going to a point, which is one unit to the right of the current point.

This path records the change in the height function between x_i and x_j. Moreover, the sum of the exponents of h and v in the reduced word for $g_i^{-1} g_j$ can be no greater than the corresponding sums in the word traced by this path, when the latter is viewed in the free group $\langle h, v \rangle$. The difference between the sum of the exponents of h and a_{ij} is a multiple of m, equal to m times the number of disjoint straight horizontal segments of length m traversed by the path. For the vertical direction, note that the net vertical displacement between x_i and x_j is zero. Thus every upward segment must be paired with a downward segment. Such pairs combine to give contributions in multiples of n. Moreover, cancellations in the reduced word also occur in blocks of length n, and hence only change the total exponent by multiples of n. It follows that

$$\sum_t \alpha_t(x_i, x_j) - a_{ij} = km, \qquad \sum_t \beta_t(x_i, x_j) = k'n,$$

where $k \leq 0$ and $k' \geq 0$.

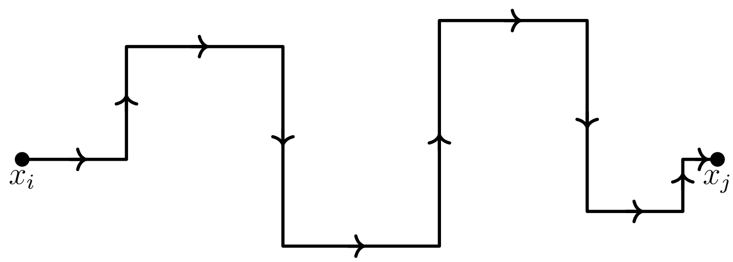

Figure 8.2: A non-backtracking path from x_i to x_j along tile boundaries in which every horizontal step is to the right.

Case 1b: $x_j - x_i = (0, b_{ij})$, where $b_{ij} \geq 0$.

By symmetry with Case 1a, we obtain

$$\sum_t \alpha_t(x_i, x_j) = km, \qquad \sum_t \beta_t(x_i, x_j) - b_{ij} = k'n,$$

where $k \geq 0$ and $k' \leq 0$.

Case 1c: $x_j - x_i = (a_{ij}, b_{ij})$, where $a_{ij} \geq 0$ and $b_{ij} \geq 0$.

Consider an up-right path in the tiling τ starting at x_i, where at each step one moves either upward or rightward. Choose one direction and continue along the tiling as far as possible; then switch direction and continue in the same way. Such a path must either reach a point of the form $x_i + (a_{ij}, b)$ with $0 \leq b \leq b_{ij}$, or a point of the form $x_i + (a, b_{ij})$ with $0 \leq a \leq a_{ij}$. WLOG, assume the former occurs, and write

$$w_{ij} = x_i + (a_{ij}, b).$$

Then

$$g_i^{-1} g_j = h(\tau, x_i)^{-1} h(\tau, x_j) = \big(h(\tau, x_i)^{-1} h(\tau, w_{ij})\big)\big(h(\tau, w_{ij})^{-1} h(\tau, x_j)\big).$$

Since the path from x_i to w_{ij} is an up-right path, if

$$h(\tau, x_i)^{-1} h(\tau, w_{ij}) = \prod_t h^{\alpha'_t} v^{\beta'_t},$$

then

$$\sum_t \alpha'_t(x_i, w_{ij}) - a_{ij} = pm, \qquad \sum_t \beta'_t(x_i, w_{ij}) - b = p'n,$$

where both $p \leq 0$ and $p' \leq 0$.

Similarly, by Case 1b, if

$$h(\tau, w_{ij})^{-1} h(\tau, x_j) = \prod_t h^{\alpha''_t} v^{\beta''_t},$$

then

$$\sum_t \alpha''_t(w_{ij}, x_j) = qm, \qquad \sum_t \beta''_t(w_{ij}, x_j) - (b_{ij} - b) = q'n,$$

where $q \geq 0$ and $q' \leq 0$.

Now consider the product

$$\left(\prod_t h^{\alpha'_t} v^{\beta'_t}\right) \left(\prod_t h^{\alpha''_t} v^{\beta''_t}\right)$$

as a word in the free group $\langle h, v \rangle$. The reduced word for $g_i^{-1} g_j$ in G can only have smaller total exponent sums in h and v. Hence, if

$$g_i^{-1} g_j = \prod_t h^{\alpha_t} v^{\beta_t},$$

then

$$\sum_t \beta_t(x_i, x_j) - b_{ij} = k'n$$

for some $k' \leq 0$. Together with Case 1a, this gives the required condition in the case $a_{ij} \geq 0$ and $b_{ij} \geq 0$.

For the rest of the cases, if the difference is non-negative in some coordinate, we proceed as in Case 1. For coordinates where the difference is negative, consider instead the path from x_j to x_i, where

$$g_j^{-1} g_i = \prod_{t'} v^{n-\beta_{t'}} h^{m-\alpha_{t'}}, \qquad \text{where } t' = \ell + 1 - t.$$

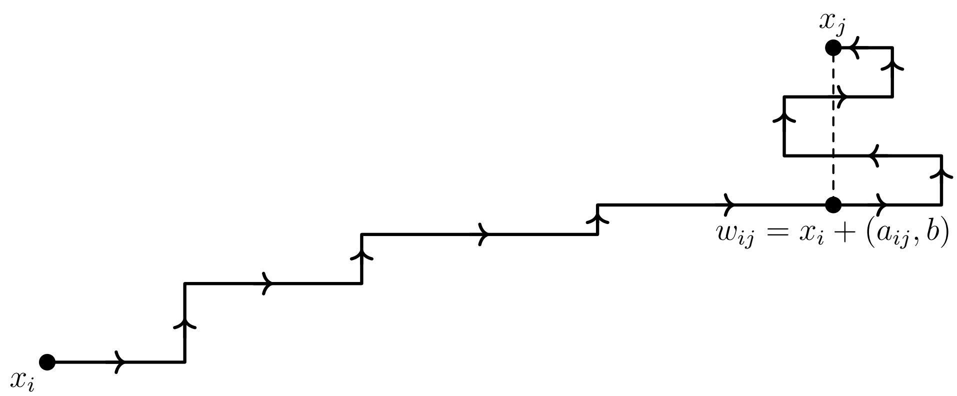

Figure 8.3: An up-right path from x_i to $w_{ij} = x_i + (a_{ij}, b)$, followed by a non-backtracking path from w_{ij} to x_j whose vertical steps are all upward.

This establishes the necessity of the stated condition. We now prove that it is also sufficient.

Sufficiency:

For an element

$$g = h^{\alpha_1} v^{\beta_1} \cdots h^{\alpha_\ell} v^{\beta_\ell}$$

written in reduced form with non-negative exponents as above, define

$$H(g) := \sum_{t=1}^{\ell} \alpha_t, \qquad V(g) := \sum_{t=1}^{\ell} \beta_t.$$

Assume now that the prescribed values satisfy the conditions in the statement of the theorem. We will show that they extend to a height function on all of $\mathbb{Z}^2$. It is enough to show that any admissible height assignment on a finite set can be extended to one additional point.

Let $x_0 \in \mathbb{Z}^2 \setminus A$. For each $1 \le i \le N$, define

$$C_i := \{\, g \in G \mid g \text{ satisfies the condition corresponding to } x_0 - x_i \,\}.$$

Then a value $g \in G$ is an admissible choice for the height at x_0 if and only if $g_i^{-1} g \in C_i$ for every $1 \le i \le N$. Equivalently, the set of all admissible values at x_0 is $\bigcap_{i=1}^{k} g_i C_i$. Thus it suffices to show that this intersection is non-empty.

We first prove that

$$g_i C_i \cap g_j C_j \neq \varnothing \qquad \text{for all } i \neq j.$$

It is enough to consider the case $x_j - x_i = (a, b)$ with $a > 0$, $b > 0$, since the remaining cases are analogous. Write

$$g_i^{-1} g_j = h^{\alpha_1} v^{\beta_1} \cdots h^{\alpha_\ell} v^{\beta_\ell}.$$

The hypothesis implies that

$$\sum_t \alpha_t - a = km, \qquad \sum_t \beta_t - b = k'n,$$

where at least one of k and k' is non-positive. Starting at x_i, trace the path in the lattice determined by the reduced word for $g_i^{-1} g_j$, and let w be its endpoint. The assumptions imply that the horizontal and vertical displacements from x_j to w are multiples of m and n respectively, and also that w is either to the left or below that of x_j. Hence w and x_j are diagonally opposite corners of a rectangle whose side lengths are multiples of m and n, and which only intersects the path at its boundary. This rectangle admits a tiling by $1 \times m$ and $n \times 1$ rectangles compatible with the traced path. Extending this tiling outside the rectangle then yields a tiling of all of $\mathbb{Z}^2$. This yields a height value at x_0 which is compatible with both g_i and g_j. Hence

$$g_i C_i \cap g_j C_j \neq \varnothing.$$

To pass from pairwise intersections to a common intersection, we use the structure of the commutator subgroup of G. By Proposition 4 of Chapter I in [Ser03], the kernel of the canonical map

$$\mathbb{Z}_m * \mathbb{Z}_n \to \mathbb{Z}_m \times \mathbb{Z}_n$$

which is the commutator subgroup $K = [G, G]$, is a free group on the generators

$$\{[h^i, v^j] \mid 1 \leq i \leq m-1,\ 1 \leq j \leq n-1\}.$$

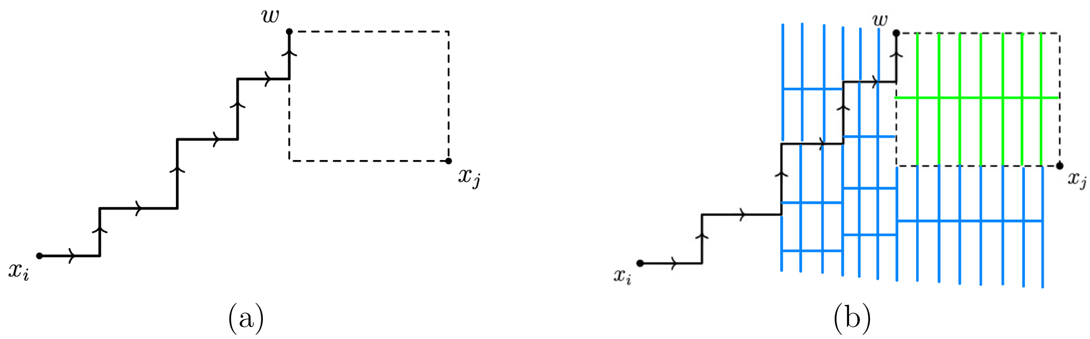

Figure 8.4: (a) The path determined by the reduced word for $g_i^{-1}g_j$ starting at x_i. The endpoint w and the point x_j are diagonally opposite corners of a rectangle whose side lengths are multiples of m and n. (b) A compatible tiling.

Hence the Cayley graph of K with respect to this generating set and its inverses is a tree.

Now observe that if $g_1, g_2 \in G$ satisfy

$$H(g_1) \equiv H(g_2) \pmod{m}, \qquad V(g_1) \equiv V(g_2) \pmod{n},$$

then

$$H(g_1^{-1}g_2) \equiv 0 \pmod{m}, \qquad V(g_1^{-1}g_2) \equiv 0 \pmod{n},$$

and hence $g_1^{-1}g_2 \in K$. Thus two elements of G lie in the same coset of K if and only if they have the same parity class modulo (m, n). Since $G/K \cong \mathbb{Z}_m \times \mathbb{Z}_n$, we have that $[G : K] = mn$ i.e. there are exactly mn cosets.

Choose a transversal

$$T = \{t_1, \dots, t_{mn}\}$$

for the left cosets of K, so that

$$G = \bigsqcup_{r=1}^{mn} t_r K.$$

Each set C_i is contained in a single parity class, hence in a single coset $t_r K$. After identifying that coset with K by left multiplication by t_r^{-1}, we may regard C_i as a subset of the Cayley graph of K. Likewise, each translate g_iC_i lies in a single coset and may be viewed as a subset of the same tree.

We now show that each such subset is convex. Consider first the case

$$x_0 - x_i = (a, b) \qquad \text{with } a > 0,\ b > 0,$$

and assume for simplicity that

$$a \equiv 0 \pmod{m}, \qquad b \equiv 0 \pmod{n},$$

so that $C_i \subset K$. Then

$$C_i = \{\, g \in K \mid H(g) = a + km,\ V(g) = b + k'n,\ \text{with at least one of } k, k' \leq 0 \,\}.$$

Equivalently,

$$C_i = K \cap \big(\{g \in G : H(g) \leq a\} \cup \{g \in G : V(g) \leq b\}\big).$$

We claim that C_i is convex in the Cayley graph of K.

Indeed, begin at the identity vertex and follow any non-backtracking path in the Cayley graph of K. Suppose that at some stage one reaches, for the first time, a vertex for which both

$$H(g) > a \qquad \text{and} \qquad V(g) > b.$$

It is enough to show that along any continuation of this path which does not immediately retrace the last edge, neither H nor V can decrease. This will imply that once the path leaves C_i, it can never re-enter, and hence C_i is convex.

The generators of K are

$$[h^i, v^j] = h^i v^j h^{m-i} v^{n-j}, \qquad 1 \leq i \leq m-1, \quad 1 \leq j \leq n-1,$$

together with their inverses

$$[h^i, v^j]^{-1} = v^j h^i v^{n-j} h^{m-i}.$$

Let

$$q = [h^i, v^j] = h^i v^j h^{m-i} v^{n-j},$$

and let $r \neq q^{-1}$ be another generator or inverse generator. If the reduced word for r begins with h, then there is no cancellation between the end of q and the beginning of r, and therefore both H and V increase.

Now suppose that

$$r = v^{j'} h^{i'} v^{n-j'} h^{m-i'}.$$

Then

$$qr = h^i v^j h^{m-i} v^{j'-j \bmod n} h^{i'} v^{n-j'} h^{m-i'}.$$

If $j' \neq j$, then $H(qr) \geq H(q)$, while

$$V(qr) = n + (j - j') + [(j' - j) \bmod n],$$

which is equal to either n or $2n$, and in particular is not smaller than $V(q)$. If $j' = j$, then

$$qr = h^i v^j h^{i'-i \bmod m} v^{n-j} h^{m-i'},$$

and since $r \neq q^{-1}$ we have $i' \neq i$. Hence again H does not decrease and V remains unchanged. The remaining cases are entirely similar. Now let

$$g_0, g_1, \ldots, g_r$$

be a non-backtracking path in the Cayley graph of K. Writing

$$g_{i+1} = g_i s_i,$$

the reduced word for g_i ends with some generator q when written as a concatenation of generators as dictated by the non-backtracking path to reach that vertex on the graph. Since the path is non-backtracking we have $s_i \neq q^{-1}$. Reduction in the product $g_i s_i$ can therefore occur only at the interface between the final generator of g_i and the initial generator of s_i. No cancellation propagates further into the reduced word for g_i. Hence the preceding calculations for products of two generators apply at every step along the path. It follows that along any non-backtracking path in the Cayley graph of K, H and V are non-decreasing. Therefore C_i is convex.

If a and b are not both divisible by m and n, then C_i lies not in K itself but in some coset $t_r K$. After identifying that coset with K by left translation, the same argument applies verbatim. Thus in every case the set C_i, and therefore also each translate $g_i C_i$, is convex in the relevant copy of the Cayley tree of K.

We have already shown that the family $\{g_i C_i\}_{i=1}^k$ is pairwise intersecting. Since these sets are convex subsets of a tree, the Helly property for trees implies that

$$\bigcap_{i=1}^{k} g_i C_i \neq \varnothing.$$

(Note that although we originally defined the Helly property for graphs in terms of balls, the proof of Proposition 4.3.1 was done for subtrees, which are precisely the convex subsets of a tree.) Hence there exists a value $g \in G$ which may be assigned to x_0 so that the enlarged height assignment on $A \cup \{x_0\}$ still satisfies the condition of the theorem.

By an argument similar to that in Chapter 4, we obtain height values on all of $\mathbb{Z}^2$ satisfying the same pairwise conditions.

Reconstructing the tiling:

It remains to show that the resulting global height function corresponds to a tiling of $\mathbb{Z}^2$ by $1 \times m$ and $n \times 1$ rectangles.

Step 1: We begin by analysing the height differences between neighbouring vertices. Let $e_1 = (1, 0), e_2 = (0, 1)$ and let $x, y \in \mathbb{Z}^2$ be neighbouring vertices with

$$y = x + e_1.$$

Then the conditions satisfied are,

$$\sum_t \alpha_t(x,y) - 1 = km, \qquad \sum_t \beta_t(x,y) = k'n,$$

where at least one of k and k' is non-positive.

Since all exponents are non-negative, we have

$$\sum_t \alpha_t(x,y) \geq 0, \qquad \sum_t \beta_t(x,y) \geq 0.$$

Moreover, $\sum_t \beta_t(x,y) \in n\mathbb{Z}_{\geq 0}$.

Case 1 ($\sum_t \beta_t(x,y) \geq 2n$): In this case $k' \geq 2$, so we must have that $k \leq 0$ i.e.

$$\sum_t \alpha_t(x,y) = 1 + km \leq 1.$$

To satisfy the condition for this case the word for the height difference between x and y must also have at least three blocks of v's with blocks of h's between every consecutive pair (since every block of v's can have length at most $n-1$). But then we would have $\sum_t \alpha_t(x,y) \geq 2$, a contradiction.

Case 2 ($\sum_t \beta_t(x,y) \in \{0,n\}$): Here, suppose that $\sum_t \alpha_t(x,y) \neq 1$. Then

$$\sum_t \alpha_t(x,y) \geq 1 + m,$$

i.e. $k \geq 1$. We must then have that $k' = 0 \implies \sum_t \beta_t(x,y) = 0$. The reduced word would be of the form $h^{\sum_t \alpha_t(x,y)}$. But then we can only have

$$0 \leq \sum_t \alpha_t(x,y) \leq m - 1,$$

which is a contradiction. Thus we must have that $\sum_t \alpha_t(x,y) = 1$.

Therefore, for every horizontal edge from x to $x + e_1$,

$$h(x)^{-1}h(x + e_1) = v^a hv^{n-a}$$

for some

$$0 \leq a \leq n - 1.$$

That is, height differences across horizontal unit edges are precisely conjugates of h. Similarly, height differences across vertical unit edges are precisely conjugates of v.

Step 2: Take neighbouring vertices x and $y = x + e_1$, and suppose

$$h(x)^{-1}h(y) = v^a hv^{n-a}$$

for some $1 \leq a \leq n - 1$.

Set

$$x' = x + e_2, \qquad y' = y + e_2.$$

By the previous part, there exists $a' \in \{0, \dots, n - 1\}$ such that

$$h(x')^{-1}h(y') = v^{a'} hv^{n-a'}.$$

Similarly, there exist $b, b' \in \{0, \dots, m - 1\}$ such that

$$h(x)^{-1}h(x') = h^b vh^{m-b},$$

and

$$h(y)^{-1}h(y') = h^{b'} vh^{m-b'}.$$

Comparing the two paths from x to y', we obtain

$$(v^a hv^{n-a})(h^{b'} vh^{m-b'}) = (h^b vh^{m-b})(v^{a'} hv^{n-a'}).$$

Rearranging gives

$$(h^b v h^{m-b})^{-1}(v^a h v^{n-a})(h^{b'} v h^{m-b'}) = v^{a'} h v^{n-a'}.$$

The right-hand side is a reduced word with exactly three alternating blocks. Hence the left-hand side must reduce to such a word as well. This is possible only if the outer h-blocks cancel completely, forcing

$$b = b' = 0.$$

In that case,

$$v^{n-1}(v^a h v^{n-a})v = v^{a-1} h v^{n-(a-1)},$$

where the indices are interpreted modulo n. Thus

$$a' \equiv a - 1 \pmod{n}.$$

Therefore, if two horizontal edges are vertically adjacent and one of them has height difference with three alternating blocks, then their conjugacy indices differ by -1 modulo n. By symmetry, if two vertical edges are horizontally adjacent and one of them has height difference with three alternating blocks, then their conjugacy indices differ by -1 modulo m.

We now define a tiling associated to the height function h. A unit cell is said to be covered by a vertical tile if one of its horizontal edges has height difference

$$v^a h v^{n-a}$$

with

$$1 \le a \le n - 1.$$

Similarly, a unit cell is said to be covered by a horizontal tile if one of its vertical edges has height difference

$$h^b v h^{m-b}$$

with

$$1 \leq b \leq m - 1.$$

It is easy to see from the previous calculation that there is no overlap between distinct tiles.

It remains to show that every unit cell is covered by a tile. Let x_1, x_2, x_3, x_4 be the vertices of a unit cell listed clockwise, with x_1 the lower-left corner. Suppose that every edge difference is of the form of a single block. Then traversing the boundary of the cell gives

$$h(x_1) = h(x_1)hvh^{-1}v^{-1},$$

which is impossible in $\mathbb{Z}_m * \mathbb{Z}_n$.

Hence at least one edge of the cell must have height difference with three alternating blocks. Therefore every cell is covered by either a horizontal tile or a vertical tile. Since overlaps are impossible, these tiles form a tiling of $\mathbb{Z}^2$. By construction, this tiling induces the original height function h.

$\square$

Bibliography

[BC08] Hans-Jürgen Bandelt and Victor Chepoi. Metric graph theory and geometry: a survey. In *Surveys on Discrete and Computational Geometry*, pages 49–86. American Mathematical Society, Providence, RI, 2008.

[BPS10] Mike Boyle, Ronnie Pavlov, and Michael Schraudner. Multidimensional sofic shifts without separation and their factors. *Transactions of the American Mathematical Society*, 362(9):4617–4653, 2010.

[CGdMO25] Nishant Chandgotia, Silvère Gangloff, Benjamin Hellouin de Menibus, and Piotr Oprocha. Undecidability of the block gluing classes of homshifts, 2025.

[Cha16] Nishant Chandgotia. Four-cycle free graphs, height functions, the pivot property and entropy minimality. *Ergodic Theory and Dynamical Systems*, 37(4):1102–1132, March 2016.

[CL90] J.H Conway and J.C Lagarias. Tiling with polyominoes and combinatorial group theory. *Journal of Combinatorial Theory, Series A*, 53(2):183–208, 1990.

[CM16] Nishant Chandgotia and Tom Meyerovitch. Markov random fields, markov cocycles and the 3-colored chessboard. *Ergodic Theory and Dynamical Systems*, 36(8):2435–2462, 2016.

[CPT18] Nishant Chandgotia, Igor Pak, and Martin Tassy. Kirszbraun-type theorems for graphs, 2018.

[Din70] E. I. Dinaburg. The relation between topological entropy and metric entropy. *Soviet Mathematics Doklady*, 11:13–16, 1970.

[EW11] Manfred Einsiedler and Thomas Ward. *Ergodic Theory with a View Towards Number Theory*, volume 259 of *Graduate Texts in Mathematics*. Springer, 2011.

[Geo88] Hans-Otto Georgii. *Gibbs Measures and Phase Transitions*, volume 9 of *de Gruyter Studies in Mathematics*. Walter de Gruyter & Co., Berlin, 1988.

[GJKS25] Su Gao, Steve Jackson, Edward Krohne, and Brandon Seward. *Continuous Combinatorics of Abelian Group Actions*. Number 1573 in Memoirs of the American Mathematical Society. American Mathematical Society, Providence, RI, 2025.

[Goo71] Tim N. T. Goodman. Relating topological entropy and measure entropy. *Bulletin of the London Mathematical Society*, 3(2):176–180, 1971.

[Hat02] A. Hatcher. *Algebraic Topology*. Algebraic Topology. Cambridge University Press, 2002.

[Hel23] Eduard Helly. Über mengen konvexer körper mit gemeinschaftlichen punkten. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 32:175–176, 1923. in German.

[Kel98] Gerhard Keller. *Equilibrium States in Ergodic Theory*, volume 42 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1998.

[Kir34] M. D. Kirszbraun. Über die zusammenziehende und lipschitzsche transformationen. *Fundamenta Mathematicae*, 22:77–108, 1934.

[KK92] Claire Kenyon and Richard Kenyon. Tiling a polygon with rectangles. In *Proceedings of the 33rd Annual Symposium on Foundations of Computer Science (FOCS 1992)*, pages 610–619, 1992.

[KM25] Krzysztof Kapulkin and Udit Mavinkurve. The fundamental group in discrete homotopy theory. *Advances in Applied Mathematics*, 164:102838, 2025.

[LM21] Douglas Lind and Brian Marcus. *An Introduction to Symbolic Dynamics and Coding*. Cambridge University Press, Cambridge, United Kingdom, 2 edition, 2021.

[LR69] O. E. Lanford and D. Ruelle. Observables at infinity and states with short range correlations in statistical mechanics. *Communications in Mathematical Physics*, 13(3):194–215, 1969.

[McS34] Edward J. McShane. Extension of range of functions. *Bulletin of the American Mathematical Society*, 40:837–842, 1934.

[PS97] Karl Petersen and Klaus Schmidt. Symmetric gibbs measures. *Transactions of the American Mathematical Society*, 349(7):2775–2811, 1997.

[QT00] A. N. Quas and P. B. Trow. Subshifts of multi-dimensional shifts of finite type. *Ergodic Theory and Dynamical Systems*, 20(3):859–874, 2000.

[Rue04] David Ruelle. *Thermodynamic Formalism: The Mathematical Structure of Equilibrium Statistical Mechanics*. Cambridge Mathematical Library. Cambridge University Press, 2 edition, 2004.

[Sar19] Omri Sarig. Lecture notes on ergodic theory. https://www.wisdom.weizmann.ac.il/~sarigo/ergodic_notes.pdf, 2019. Accessed: 2025-09-03.

[Sch97] Klaus Schmidt. Invariant cocycles, random tilings and the super-k and strong markov properties. *Transactions of the American Mathematical Society*, 349(7):2813–2825, 1997.

[Sch10] Michael Schraudner. Projectional entropy and the electrical wire shift. *Discrete and Continuous Dynamical Systems*, 26(1):333–346, 2010.

[Ser03] Jean-Pierre Serre. *Trees*. Springer Monographs in Mathematics. Springer, Berlin, Heidelberg, 2 edition, 2003.

[Tas12] Martin Tassy. *Tiling by Bars*. PhD thesis, Brown University, Providence, RI, 2012.

[Thu90] William P. Thurston. Conway's tiling groups. *Am. Math. Monthly*, 97(8):757–756, September 1990.

[Val45] F. A. Valentine. A lipschitz condition preserving extension for a vector function. *American Journal of Mathematics*, 67(1):83–93, 1945.

[Wal82] Peter Walters. *An Introduction to Ergodic Theory*, volume 79 of *Graduate Texts in Mathematics*. Springer, New York, 1982.

[Whi34] Hassler Whitney. Analytic extensions of differentiable functions defined in closed sets. *Transactions of the American Mathematical Society*, 36(1):63–89, 1934.

[Wro17] Marcin Wrochna. Square-free graphs are multiplicative. *Journal of Combinatorial Theory, Series B*, 122:479–507, 2017.