Ergod. Th. & Dynam. Sys., page 1 of 55 © The Author(s), 2023. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

doi:10.1017/etds.2023.38

Counting independent sets in amenable groups

RAIMUNDO BRICEÑO

Facultad de Matemáticas, Pontificia Universidad Católica de Chile, Santiago, Chile
(e-mail: raimundo.briceno@mat.uc.cl)

(*Received* 11 *September* 2022 *and accepted in revised form* 1 *March* 2023)

Abstract. Given a locally finite graph Γ, an amenable subgroup G of graph automorphisms acting freely and almost transitively on its vertices, and a G-invariant activity function λ, consider the free energy $f_G(\Gamma, \lambda)$ of the hardcore model defined on the set of independent sets in Γ weighted by λ. Under the assumption that G is finitely generated and its word problem can be solved in exponential time, we define suitable ensembles of hardcore models and prove the following: if $\|\lambda\|_\infty < \lambda_c(\Delta)$, there exists a randomized ϵ-additive approximation scheme for $f_G(\Gamma, \lambda)$ that runs in time $\mathrm{poly}((1+\epsilon^{-1})|\Gamma/G|)$, where $\lambda_c(\Delta)$ denotes the critical activity on the Δ-regular tree. In addition, if G has a finite index linearly ordered subgroup such that its algebraic past can be decided in exponential time, we show that the algorithm can be chosen to be deterministic. However, we observe that if $\|\lambda\|_\infty > \lambda_c(\Delta)$, there is no efficient approximation scheme, unless NP = RP. This recovers the computational phase transition for the partition function of the hardcore model on finite graphs and provides an extension to the infinite setting. As an application in symbolic dynamics, we use these results to develop efficient approximation algorithms for the topological entropy of subshifts of finite type with enough safe symbols, we obtain a representation formula of pressure in terms of random trees of self-avoiding walks, and we provide new conditions for the uniqueness of the measure of maximal entropy based on the connective constant of a particular associated graph.

Key words: amenable group, independent set, entropy, Gibbs measure, strong spatial mixing, computational phase transition
2020 Mathematics Subject Classification: 37B10, 37B40, 37D35 (Primary); 37A35, 82B20, 37A50, 37A15, 68W25 (Secondary)

Contents

1 Introduction 2
2 Preliminaries 6
3 Free energy 9
4 Gibbs measures 15
5 Trees 18
6 Orders 25

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

7 Counting 27
8 A computational phase transition in the thermodynamic limit 36
9 Reductions 44
Acknowledgements 54
References 54

1. *Introduction*

Suppose that we are given a finite simple graph $\Gamma = (V, E)$ and we are asked to count its number of *independent sets*. An independent set is a subset $I \subseteq V$ such that $(v, v') \notin E$ (that is, (v, v') is not an edge) for all $v, v' \in I$. For example, if Γ is the 4-cycle C_4 with $V = \{v_1, v_2, v_3, v_4\}$ and $E = \{(v_1, v_2), (v_2, v_3), (v_3, v_4), (v_4, v_1)\}$, it can be checked that there are exactly seven different independent sets, namely $\emptyset$, $\{v_1\}$, $\{v_2\}$, $\{v_3\}$, $\{v_4\}$, $\{v_1, v_3\}$, and $\{v_2, v_4\}$. A common generalization of this question is to ask for the 'number' of *weighted* independent sets in Γ: given a parameter $\lambda > 0$—usually called *activity* or *fugacity*—we ask for the value of the summation

$$Z_\Gamma(\lambda) := \sum_{I \in X(\Gamma)} \lambda^{|I|},$$

where $X(\Gamma)$ denotes the collection of all independent sets in Γ and $|I|$, the cardinality of a given independent set I. Notice that we recover the original problem—that is, to compute $|X(\Gamma)|$—if we set $\lambda = 1$. The sum $Z_\Gamma(\lambda)$ corresponds to the normalization factor of the probability distribution $\mathbb{P}_{\Gamma,\lambda}$ on $X(\Gamma)$ that assigns to each $I \in X(\Gamma)$ a probability proportional to $\lambda^{|I|}$, that is, the so-called *partition function* (also known as the *independence polynomial*) of the well-known *hardcore model* from statistical physics.

In general, it is not possible to compute exactly $Z_\Gamma(\lambda)$ efficiently [31], even for the case $\lambda = 1$ [56]; technically, to compute $Z_\Gamma(\lambda)$ is an NP-*hard* problem and to compute $|X(\Gamma)|$ is a #P-*complete* problem. Therefore, one may attempt to at least find ways to approximate $Z_\Gamma(\lambda)$ efficiently.

In recent years, there has been a great deal of attention given to the complexity of approximating partition functions of spin systems (e.g., see [4]). Among these systems, the hardcore model, possibly together with the Ising model [47], occupies the most important place. One of the most notable results, due to Weitz [55], and then Sly [48] and Sly and Sun [49], is the existence of a *computational phase transition* for having a *fully polynomial-time approximation scheme (FPTAS)* for the approximation of $Z_\Gamma(\lambda)$. In simple terms, Weitz developed an FPTAS, a particular kind of efficient deterministic approximation algorithm, on the family of finite graphs with bounded degree Δ provided $\lambda < \lambda_c(\Delta)$, where $\lambda_c(\Delta) := (\Delta - 1)^{(\Delta-1)}/(\Delta - 2)^\Delta$ denotes the critical activity for the hardcore model on the Δ-regular tree $\mathbb{T}_\Delta$. Conversely, a couple of years later, Sly and Sun managed to prove that the existence of even a *fully polynomial-time randomized approximation scheme (FPRAS)*—which is a probabilistic and therefore weaker version of an FPTAS—for $\lambda > \lambda_c(\Delta)$ would imply that NP = RP, the equivalence of two well-known computational complexity classes which are widely believed to be different [2].

The work of Weitz exploited a technique based on trees of self-avoiding walks and a special notion of correlation decay known as *strong spatial mixing* that, in particular,

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

holds when the graph is $\mathbb{T}_\Delta$ and $\lambda < \lambda_c(\Delta)$. Later, Sinclair *et al* [46] studied refinements of Weitz's result by considering families of finite graphs parameterized by their *connective constant* instead of their maximum degree, and established that there exists an FPTAS for $Z_\Gamma(\lambda)$ for families of graphs with connective constant bounded by μ, whenever $\lambda < \mu^\mu/(\mu-1)^{(\mu+1)}$.

Now, if Γ is an infinite graph, most of these concepts stop making sense. One way to deal with this issue is by choosing an appropriate normalization and by using the *DLR formalism*. The idea is roughly the following: suppose that we have a sequence $\{\Gamma_n\}_n$ of finite subgraphs that 'exhausts' Γ in some way. This sequence induces two other sequences: a sequence $\{Z_{\Gamma_n}(\lambda)\}_n$ of partition functions and a sequence $\{\mathbb{P}_{\Gamma_n,\lambda}\}_n$ of probability distributions. A way to extend the idea of 'number of weighted independent sets (per site)' in Γ is by considering the sequence $\{Z_{\Gamma_n}(\lambda)^{1/|\Gamma_n|}\}_n$ and hoping that it converges. Under the right assumptions on Γ and $\{\Gamma_n\}_n$, this is exactly the case and something similar happens to $\{\mathbb{P}_{\Gamma_n,\lambda}\}_n$. Moreover, there is an intimate connection between the properties of the limit measures and our ability to estimate the value of $\lim_n |\Gamma_n|^{-1} \log Z_{\Gamma_n}(\lambda)$, that is, to 'approximately count' it. We denote this limit—which, *a priori*, may depend on the sequence $\{\Gamma_n\}_n$—by $f_{\{\Gamma_n\}_n}(\Gamma, \lambda)$ and call it the *free energy* of the hardcore model (Γ, λ), one of the most crucial quantities in statistical physics [6, 17, 44].

It can be checked that if Γ is finite, to approximate the partition function $Z_\Gamma(\lambda)$ with a multiplicative error (in polynomial time) is equivalent to approximate the free energy $f_{\{\Gamma\}_n}(\Gamma, \lambda)$—where $\{\Gamma\}_n$ is the constant sequence which immediately exhausts the graph—with an additive error [30] (in polynomial time). Therefore, the problem of approximating $f_{\{\Gamma\}_n}(\Gamma, \lambda)$ recovers the problem of approximating the partition function in the finite case and, at the same time, extends the problem to the infinite setting.

The main goal of this paper is to establish a computational phase transition for the free energy on ensembles of—possibly infinite—hardcore models. In other words, we aim to prove the existence of an efficient additive approximation algorithm for the free energy when the activity is low and to establish that there is no efficient approximation algorithm for the free energy when the activity is high, unless NP = RP.

There have been many recent works related to the study of correlation decay properties and their relation to approximation algorithms for the free energy (and related quantities such as pressure, capacity, and entropy) in the infinite setting [8, 16, 34–36, 40, 53]. In this work, we put all these results in a single framework, which also encompasses the results from Weitz, Sly and Sun, and Sinclair *et al*, and at the same time generalizes them.

In 2009, Gamarnik and Katz [16] introduced what they called the *sequential cavity method*, which can be regarded as a sort of infinitary *self-reducibility* property [24]. Combining this method with Weitz's results, they managed to prove that the free energy of the hardcore model in the Cayley graph of $\mathbb{Z}^d$ with canonical generators admits a (deterministic) ϵ-additive approximation algorithm that runs in time polynomial in ϵ^{-1} whenever $\lambda < \lambda_c(2d)$, where $2d$ is the maximum degree of the graph. Related results were also proven by Pavlov in [40], who developed an approximation algorithm for the *hard square entropy*, that is, the free energy of the hardcore model in the Cayley graph of $\mathbb{Z}^2$ with canonical generators and activity $\lambda = 1$. Later, there were also some explorations due to Wang *et al* [53] in Cayley graphs of $\mathbb{Z}^2$ with respect to other generators

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

(e.g., the *non-attacking kings* system) in the context of information theory and algorithms for approximating capacities.

In this paper, we prove that all these results fit and can be generalized to hardcore models (Γ, λ) such that: (1) Γ is a locally finite graph; (2) $G \curvearrowright \Gamma$ is free and almost transitive for some countable amenable subgroup $G \leq \mathrm{Aut}(\Gamma)$; and (3) $\lambda : V \to \mathbb{Q}_{>0}$ is a—not necessarily constant—G-invariant activity function. In addition, for the algorithmic implications, we assume that G satisfies some of the recursion-theoretic assumptions described below. Given this setting, we consider a Følner sequence $\{F_n\}_n$, a fundamental domain $U_0 \subseteq V$ of $G \curvearrowright \Gamma$, and the sequence of finite subgraphs $\{\Gamma_n\}_n$ induced by $\{F_n U_0\}_n$. First, we show that $f_{\{\Gamma_n\}_n}(\Gamma, \lambda)$ is independent of $\{F_n\}_n$ and U_0, and that the limit $f_{\{\Gamma_n\}_n}(\Gamma, \lambda)$—which we denote by $f_G(\Gamma, \lambda)$ to emphasize the independence of $\{F_n\}_n$ and U_0—can be expressed as an infimum over some suitable family of finite subgraphs of Γ. Next, based on results from [9, 20], we prove in Theorem 7.1 that $f_G(\Gamma, \lambda)$ can be obtained as the pointwise limit of a Shannon–McMillan–Breiman-type ratio with regards to any Gibbs measure on $X(\Gamma)$. In Theorem 7.5, we prove that if λ is such that (Γ, λ) satisfies strong spatial mixing, then $f_G(\Gamma, \lambda)$ corresponds to the evaluation of a random information function, based on ideas about random invariant orders and the *Kieffer–Pinsker formula* for measure-theoretical entropy introduced in [1]. Then, in Theorem 7.6, using the previous representation theorem and the techniques from [55], we provide a formula for $f_G(\Gamma, \lambda)$ in terms of trees of self-avoiding walks in Γ. These first three theorems can be regarded as a preprocessing treatment of $f_G(\Gamma, \lambda)$ to obtain an *arboreal representation of free energy* to develop approximation techniques, but we believe that they are of independent interest.

Later, we consider a finitely generated amenable group G with a prescribed set of generators S such that its word problem can be solved in exponential time. This last requirement seems to be natural and many groups satisfy it (for example, any linear group, including all abelian, all nilpotent groups, and, more generally, all virtually polycyclic groups). Given a positive integer Δ and $\lambda_0 > 0$, we denote by $\mathcal{H}_G^{\Delta}(\lambda_0)$ the ensemble of hardcore models (Γ, λ) such that $G \curvearrowright \Gamma$ is free and almost transitive, the maximum degree of Γ is bounded by Δ, and the values of λ are bounded from above by λ_0. Then, in Theorem 8.5, we establish the following algorithmic implication: if $\lambda_0 < \lambda_c(\Delta)$, there exists an *additive FPRAS* on $\mathcal{H}_G^{\Delta}(\lambda_0)$ for $f_G(\Gamma, \lambda)$, where $\lambda_c(\Delta)$ denotes the critical activity on the Δ-regular tree $\mathbb{T}_\Delta$. This can be considered as a confirmation in the amenable setting of the 'algorithmic version'—as called in [55]—of [50, Conjecture 2.1]. In addition, under the extra assumption that G has a finite index linearly ordered subgroup $(H, \prec)$ such that its algebraic past $\Phi_\prec = \{g \in H : g \prec 1_G\}$ can be decided in exponential time, we prove that the algorithm can be chosen to be deterministic, that is, there exists an *additive FPTAS*. Groups that satisfy this extra condition include all finitely generated abelian groups, nilpotent groups like the Heisenberg group $H_3(\mathbb{Z})$, and solvable groups like the Baumslag–Solitar groups $BS(1, n)$. However, in Corollary 8.8, we observe that if $\lambda_0 > \lambda_c(\Delta)$, there is no additive FPRAS unless NP = RP. In particular, we obtain that the results from Weitz, Sly, and Sun correspond to the special case when G is the trivial (and orderable) group.

By an *additive FPRAS*, we mean a probabilistic algorithm that given (Γ, λ) and $\epsilon > 0$, outputs a number $\hat{f}$ such that $|f_G(\Gamma, \lambda) - \hat{f}| < \epsilon$ with probability greater than $3/4$ in time

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

polynomial in $|\Gamma/G|$ and ϵ^{-1}. Here, $|\Gamma/G|$ denotes the size of some (or any) fundamental domain of the action $G \curvearrowright \Gamma$, and therefore, all the information we need to construct Γ. However, by an *additive FPTAS*, we mean an additive FPRAS with success probability equal to 1 instead of just $3/4$. We assume throughout the paper that the standard functions and arithmetic operations of the numerical values involved can be computed exactly in one unit of time.

Finally, as an application in symbolic dynamics, we show how to use these results to establish representation formulas and efficient approximation algorithms for the topological entropy of nearest-neighbor subshifts of finite type with enough safe symbols. Also, we consider the pressure of single-site potentials with a vacuum state, which includes systems such as the *Widom–Rowlinson model* and some other weighted graph homomorphisms from Γ to any finite graph, among others. These results can also be regarded as an extension of the works by Marcus and Pavlov in $\mathbb{Z}^d$ (see [34–36]), who developed additive approximation algorithms for the entropy and free energy (or pressure) of general $\mathbb{Z}^d$-subshifts of finite type, with special emphasis in the case $d = 2$. We believe that these implications are relevant, especially in the light of results like that from Hochman and Meyerovitch. In [23], Hochman and Meyerovitch proved that the set of topological entropies that a nearest-neighbor $\mathbb{Z}^2$-subshift of finite type can achieve coincides with the set of *non-negative right-recursively enumerable real numbers*. This class of numbers includes numbers that are poorly computable or even not computable. In addition, we discuss the case of the monomer–dimer model and counting independent sets of line graphs, which is a special case that does not exhibit a phase transition. As a byproduct of our results, we also give sufficient conditions for the existence of a unique measure of maximal entropy for subshifts on arbitrary amenable groups.

We remark that our results—considering related ones, like those obtained by Gamarnik and Katz in [16]—are novel in at least three aspects.

(1) *Almost transitive framework.* The generalization to the almost transitive case provides enough flexibility so that (i) other systems (such as subshifts of finite type, matchings, etc.) can be represented through reductions in terms of independent sets in suitable graphs and (ii) the measurement of (the size of) fundamental domains as a way to measure computational complexity provides a way to obtain a computational phase transition. These aspects—to our knowledge—are new, even in the relevant case $G = \mathbb{Z}^d$, that is, the family of graphs such that $\mathbb{Z}^d$ acts almost transitively on them.

(2) *Algorithms for graphs with exponential growth.* Our approach, which provides polynomial time approximation algorithms, works for amenable groups not only of polynomial growth but also exponential growth. A relevant case that is fully explored in §8 is the family of Baumslag–Solitar groups $BS(1, n)$ for $n \geq 2$, which have exponential growth but admit even a deterministic approximation algorithm for free energy.

(3) *Lack of orderability.* If a group does not have an orderable subgroup of finite index, it is less clear how to obtain a sequential cavity method as in [16], which exploits the existence of an invariant deterministic order of the group at hand (like, for example, the lexicographic order in $\mathbb{Z}^d$). Our free energy representation formulas, in terms

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

of invariant random orders, provide a way to develop randomized approximation algorithms for groups that are not necessarily orderable.

The paper is organized as follows: in §2, we introduce the basic concepts regarding graphs, homomorphisms, independent sets, group actions, Cayley graphs, and partition functions; in §3, we rigorously define free energy based on the notion of amenability and show some robustness properties of its definition; in §4, we define Gibbs measures and relevant spatial mixing properties; in §5, we develop the formalism based on trees of self-avoiding walks and discuss some of their properties; in §6, we present the formalism of invariant (deterministic and random) orders of a group; in §7, we prove Theorems 7.1, 7.5, and 7.6, which provide a randomized sequential cavity method that allows us to obtain an arboreal representation of free energy; in §8, we prove Theorem 8.5 and establish the algorithmic implications of our results; in §9, we provide reductions that allow us to translate the problem of approximating pressure of a single-site potential and the topological entropy of a subshift into the problem of counting independent sets, and discuss other consequences that are implicit in our results.

2. *Preliminaries*

2.1. *Graphs.* A *graph* will be a pair $\Gamma = (V, E)$ such that V is a countable set—the *vertices*—and $E \subseteq V \times V$ is a symmetric relation—the *edges*. Let $\leftrightarrow$ be the equivalence relation generated by E, that is, $v \leftrightarrow v'$ if and only if there exist $n \in \mathbb{N}_0$ and $\{v_i\}_{0 \le i \le \ell}$ such that $v = v_0$, $v' = v_n$, and $(v_i, v_{i+1}) \in E$ for every $0 \le i < n$. Denote by $n(v, v')$ the smallest n with this property. This induces a notion of distance in Γ given by

$$\mathrm{dist}_\Gamma(v, v') = \begin{cases} n(v, v') & \text{if } v \leftrightarrow v', \\ +\infty & \text{otherwise.} \end{cases}$$

Given a set $U \subseteq V$, we define its *boundary* ∂U as the set $\{v \in V : \mathrm{dist}_\Gamma(v, U) = 1\}$, where $\mathrm{dist}_\Gamma(U, U') = \inf_{v \in U, v' \in U'} \mathrm{dist}_\Gamma(v, v')$. In addition, given $\ell \ge 0$ and $v \in V$, we define the *ball centered at v with radius*ℓ as $B_\Gamma(v, \ell) := \{v' \in V : \mathrm{dist}_\Gamma(v, v') \le \ell\}$.

A graph Γ is:

- *loopless*, if E is anti-reflexive (that is, there is no vertex related to itself);
- *connected*, if $v \leftrightarrow v'$ for every $v, v' \in V$; and
- *locally finite*, if every vertex is related to only finitely many vertices.

Sometimes we will write $V(\Gamma)$ and $E(\Gamma)$—instead of just V and E—to emphasize Γ.

2.2. *Homomorphisms.* Consider graphs Γ_1 and Γ_2. A *graph homomorphism* is a map $g : V(\Gamma_1) \to V(\Gamma_2)$ such that

$$(v, v') \in E(\Gamma_1) \implies (g(v), g(v')) \in E(\Gamma_2).$$

We denote by $\mathrm{Hom}(\Gamma_1, \Gamma_2)$ the set of graph homomorphisms from Γ_1 to Γ_2.

A *graph isomorphism* is a bijective map $g : V(\Gamma_1) \to V(\Gamma_2)$ such that

$$(v, v') \in E(\Gamma_1) \iff (g(v), g(v')) \in E(\Gamma_2).$$

If a map like this exists, we say that Γ_1 and Γ_2 are *isomorphic*, denoted by $\Gamma_1 \cong \Gamma_2$.

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

FIGURE 1. The graph H_0.

A *graph automorphism* is a graph isomorphism from a graph Γ to itself. We denote by $\mathrm{Aut}(\Gamma)$ the set of graph automorphisms of Γ. This set is a group when considering composition $\circ$ as the group operation and the identity map $\mathrm{id}_\Gamma : V \to V$ as the identity group element $1_{\mathrm{Aut}(\Gamma)}$. In this case, instead of writing $g_1 \circ g_2$, we will simply write $g_1 g_2$ to emphasize the group structure.

2.3. *Independent sets.* Given a subset $U \subseteq V$, the *induced subgraph* by U, denoted $\Gamma[U]$, is the graph with a set of vertices U and set of edges $E \cap (U \times U)$. A subset $I \subseteq V$ is called an *independent set* if $\Gamma[I]$ has no edges. We can also represent an independent set by its indicator function, that is, by the map $x : V \to \{0, 1\}$ so that

$$[x(v) = 1 \text{ and } (v, v') \in E] \implies x(v') = 0.$$

In addition, if we consider the finite graph $H_0 := (\{0, 1\}, \{(0, 0), (0, 1), (1, 0)\})$, then x can be also understood as a graph homomorphism from Γ to H_0 (see Figure 1).

We denote by $X(\Gamma)$ the *set of independent sets* of Γ. Notice that $X(\Gamma) \subseteq \{0, 1\}^V$ can be identified with the set $\mathrm{Hom}(\Gamma, H_0)$ and that the *empty independent set* 0^V always belongs to $X(\Gamma)$. Sometimes we will denote this independent set by 0^Γ.

2.4. *Group actions.* Let G be a subgroup of $\mathrm{Aut}(\Gamma)$. Given $g \in G$ and $v \in V$, the map $(g, v) \mapsto g \cdot v := g(v)$ is a *(left) group action*, this is to say, $1_G \cdot v = v$ and $(gg') \cdot v = g \cdot (g' \cdot v)$ for all $g' \in G$, where $1_G = 1_{\mathrm{Aut}(\Gamma)}$. In this case, we say that G *acts on* Γ and denote this fact by $G \curvearrowright \Gamma$.

The group G also acts on $\{0, 1\}^V$ by precomposition. Given $g \in G$ and $x \in \{0, 1\}^V$, consider the map $(g, x) \mapsto g \cdot x := x \circ g^{-1}$. A subset $X \subseteq \{0, 1\}^V$ is called *G-invariant* if $g \cdot X = X$ for all $g \in G$, where $g \cdot X := \{g \cdot x : x \in X\}$. Clearly, if $x \in X(\Gamma)$, then $g \cdot x$ and $g^{-1} \cdot x$ also belong to $X(\Gamma)$, since $g \in \mathrm{Aut}(\Gamma)$ and $x \in \mathrm{Hom}(\Gamma, H_0)$. Therefore, $X(\Gamma)$ is G-invariant and the action $G \curvearrowright X(\Gamma)$ is well defined.

We will usually use the letter v to denote vertices in V, the letter g to denote graph automorphisms in G, and the letter x to denote independent sets in $X(\Gamma)$. To distinguish the action of G on V from the action of G on $X(\Gamma)$, we will write gv instead of $g \cdot v$, without risk of ambiguity.

The action $G \curvearrowright \Gamma$ is always *faithful*, that is, for all $g \in G \setminus \{1_G\}$, there exists $v \in V$ such that $gv \neq v$. The *G-orbit* of a vertex $v \in V$ is the set $Gv := \{gv : g \in G\}$. The set of all G-orbits of Γ, denoted by Γ/G, is a partition of V and it is called the *quotient* of the action.

We say that a subset $\emptyset \neq U \subseteq V$ is *dynamically generating* if $GU = V$, where $GF := \{gv : g \in F, v \in U\}$ for any $F \subseteq G$, and a *fundamental domain* if it is also minimal, that

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

is, if $U' \subsetneq U$, then $GU' \subsetneq V$. The action $G \curvearrowright \Gamma$ is *almost transitive* if $|\Gamma/G| < +\infty$ and *transitive* if $|\Gamma/G| = 1$. A graph Γ is called *almost transitive* (respectively *transitive*) if $\mathrm{Aut}(\Gamma) \curvearrowright \Gamma$ is almost transitive (respectively transitive).

The *index* of a subgroup $H \leq G$, denoted by $[G : H]$, is the cardinality of the set of cosets $\{Hg : g \in G\}$. We will usually consider subgroups of finite index. In this case, we have that $|\Gamma/H| = |\Gamma/G|[G : H]$.

The *G-stabilizer* of a vertex $v \in V$ is the subgroup $\mathrm{Stab}_G(v) := \{g \in G : gv = v\}$. Notice that, since $\mathrm{Stab}_G(gv) = g\mathrm{Stab}_G(v)g^{-1}$ for every $g \in G$, we have that $|\mathrm{Stab}_G(v)| = |\mathrm{Stab}_G(v')|$ for all $v' \in Gv$. If $|\mathrm{Stab}_G(v)| < \infty$ for all v, we say that the action is *almost free*, and if $|\mathrm{Stab}_G(v)| = 1$ (that is, if $\mathrm{Stab}_G(v) = \{1_G\}$) for all v, we say that the action is *free*.

A relevant observation is that if we assume that Γ is countable and $G \curvearrowright \Gamma$ is almost transitive and almost free, then G must be a countable group. In this work, we will only consider almost free and almost transitive actions. In this case, there exists a finite fundamental domain $U_0 \subseteq V$ such that $|U_0| = |\Gamma/G|$ and, if Γ is locally finite, then Γ must have *bounded degree*, that is, there is a uniform bound on the number of vertices to which each vertex is related. In this case, we denote by $\Delta(\Gamma)$ the maximum degree among all vertices of the graph Γ.

2.5. *Transitive case: Cayley graphs.* Consider a subset $S \subseteq G$ that we assume to be *symmetric*, that is, $S = S^{-1}$, where $S^{-1} = \{s^{-1} \in S : s \in S\}$. We define the *(right) Cayley graph* as $\mathrm{Cay}(G, S) = (V, E)$, where

$$V = G \quad \text{and} \quad E = \{(g, sg) : g \in G, s \in S\}.$$

Cayley graphs are a natural construction used to represent groups in a geometric fashion. In this context, it is common to ask that $1_G \notin S$, S to be finite, and S to be *generating*, that is, $G = \langle S \rangle$, where

$$\langle S \rangle := \{s_1 \cdots s_k : s_i \in S \text{ for all } 1 \leq i \leq k \text{ and } k \in \mathbb{N}\}.$$

Groups that have a set S satisfying these conditions are called *finitely generated*. Notice that if $1_G \notin S$, then $\mathrm{Cay}(G, S)$ is loopless; if S is finite, then $\mathrm{Cay}(G, S)$ has bounded degree; and if S is generating, then $\mathrm{Cay}(G, S)$ is connected. Now, suppose that $G \curvearrowright \Gamma$ is transitive (and free). Then, there exists a symmetric set $S \subseteq G$ such that

$$\Gamma \cong \mathrm{Cay}(G, S).$$

Indeed, it suffices to take $S = \{g \in G : (v, gv) \in E\}$, where $v \in V$ is arbitrary (see [42]).

We will be interested in Cayley graphs $\Gamma = \mathrm{Cay}(G, S)$ and their subgroup of automorphisms induced by group multiplication as a special and relevant case: given $g \in G$, we can define $f_g : \Gamma \to \Gamma$ as $f_g(g') = g'g$ and it is easy to check that $f_g \in \mathrm{Aut}(\Gamma)$. Then G acts (as a group, from the left) on Γ so that $g \cdot g' = f_{g^{-1}}(g') = g'g^{-1}$ for all $g' \in G$ and $G \hookrightarrow \mathrm{Aut}(\Gamma)$ by identifying g with $f_{g^{-1}}$. In addition, via this identification, G acts transitively on Γ as a subgroup of graph automorphisms. In particular, every Cayley graph is transitive.

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

2.6. *Partition functions.* Given a graph $\Gamma = (V, E)$, let us consider $\lambda : V \to \mathbb{R}_{>0}$, an *activity function*. We call the pair (Γ, λ) a *hardcore model*. We will say that a hardcore model (Γ, λ) is *finite* if Γ is finite. If $U \subseteq V$ is a finite subset, a fact that we denote by $U \Subset V$, and $x \in X(\Gamma)$ is an independent set, we define the λ-*weight* of x on U as

$$\mathrm{w}_\lambda(x, U) := \prod_{v \in U} \lambda(v)^{x(v)}$$

and the (Γ, U, λ)-*partition function* as

$$Z_\Gamma(U, \lambda) := \sum_{x \in X(\Gamma, U)} \mathrm{w}_\lambda(x, U) = \sum_{x \in X(\Gamma, U)} \prod_{v \in U} \lambda(v)^{x(v)},$$

where $X(\Gamma, U) := \{x \in X(\Gamma) : x(v) = 0 \text{ for all } v \notin U\}$ is the finite set corresponding to the subset of independent sets of Γ *supported* on U. It is easy to check that there is an identification between $X(\Gamma, U)$ and $X(\Gamma[U])$. Then, the quantity $Z_\Gamma(U, \lambda)$ corresponds to the summation of independent sets of $\Gamma[U]$ weighted by λ. In the special case $\lambda \equiv 1$, we have that $Z_\Gamma(U, 1) = |X(\Gamma, U)| = |X(\Gamma[U])|$, that is, the partition function is exactly the number of independent sets supported on U. If (Γ, λ) is finite, we will simply write $Z_\Gamma(\lambda)$ instead of $Z_\Gamma(V, \lambda)$.

Remark 2.1. Notice that if $(v, v) \in E$ or $\lambda(v) = 0$, then $Z_\Gamma(U, \lambda) = Z_\Gamma(U \setminus \{v\}, \lambda)$; due to this fact, we usually ask λ to be strictly positive and that Γ is loopless.

3. *Free energy*

Now, suppose that we have an increasing sequence $\{U_n\}_n$ of finite subsets of vertices exhausting Γ, that is, $U_n \subseteq U_{n+1}$ and $\bigcup_n U_n = V$. Tentatively, we would like to define the exponential growth rate of $Z_\Gamma(V_n, \lambda)$ as

$$\lim_n \frac{\log Z_\Gamma(U_n, \lambda)}{|U_n|}.$$

To guarantee the existence of this limit, we will provide a self-contained argument based on the particular properties of the hardcore model and *amenability*. The reader that is familiar with this kind of argument may skim over the next part and go then to §4.

3.1. *Amenability.* Let

$$\mathcal{F}(G) := \{F \subseteq G : 0 < |F| < \infty\}$$

be the set of finite non-empty subsets of G. Given $g \in G$ and $K, F \subseteq G$, we denote $Fg = \{hg : h \in F\}$, $gF = \{gh : h \in F\}$, $F^{-1} := \{g^{-1} : g \in F\}$, and $KF = \{hg : h \in K, g \in F\}$.

We say that $\{F_n\}_n \subseteq \mathcal{F}(G)$ is a *right Følner sequence* if

$$\lim_n \frac{|F_n g \triangle F_n|}{|F_n|} = 0 \quad \text{for all } g \in G,$$

where $\triangle$ denotes the symmetric difference. Similarly, $\{F_n\}_n$ is a *left Følner sequence* if

$$\lim_n \frac{|g F_n \triangle F_n|}{|F_n|} = 0 \quad \text{for all } g \in G,$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

and $\{F_n\}_n$ is a *two-sided Følner sequence* if it is both a left and a right Følner sequence. The group G is said to be *amenable* if it has a (*right* or *left*) *Følner sequence*. Notice that $\{F_n\}_n$ is left Følner if and only if $\{F_n^{-1}\}_n$ is right Følner. A Følner sequence $\{F_n\}_n$ is a *Følner exhaustion* if in addition $F_n \subseteq F_{n+1}$ and $\bigcup_n F_n = G$. Every countable amenable group has a two-sided Følner exhaustion (see [26, Theorem 4.10]).

Every virtually amenable group is amenable. Moreover, the class of amenable groups contains all finite and all abelian groups, and it is closed under the operations of taking subgroups and forming quotients, extensions, and directed unions (see [12]).

3.2. *Growth rate of independent sets.* Given $\emptyset \neq U \Subset V$, define $\varphi_U : \mathcal{F}(G) \to \mathbb{R}$ as

$$\varphi_U(F) := \log Z_\Gamma(FU, \lambda).$$

From now on, we will assume that $\lambda : V \to \mathbb{R}_{>0}$ is G-invariant, this is to say,

$$\lambda(gv) = \lambda(v) \quad \text{for all } g \in G.$$

In other words, λ is constant along the G-orbits, so it achieves at most $|\Gamma/G|$ different values. We denote by λ_+ and λ_- the maximum and minimum among these values, respectively.

Now, let W be an abstract set, M a finite subset of W, and $k \in \mathbb{N}$. We will say that a finite collection $\mathcal{K}$ of non-empty finite subsets of W, with possible repetitions, is a *k-cover* of M if $\sum_{K\in\mathcal{K}} \mathbb{1}_K \geq k\mathbb{1}_M$, where $\mathbb{1}_A : W \to \{0, 1\}$ denotes the indicator function of a set $A \subseteq W$. The following lemma is due to Downarowicz, Frej, and Romagnoli.

LEMMA 3.1. [14] *Let Y be a subset of A^n, where A is a finite set and $n \in \mathbb{N}$. Let $\mathcal{K}$ be a k-cover of the set of coordinates $M = \{1, \ldots, n\}$. For $K \in \mathcal{K}$, let $Y_K = \{y_K : y \in Y\}$, where y_K is the restriction of y to K. Then,*

$$|Y| \leq \prod_{K\in\mathcal{K}} |Y_K|^{1/k}.$$

Given $\varphi : \mathcal{F}(G) \to \mathbb{R}$, we will say that φ satisfies *Shearer's inequality* if

$$\varphi(F) \leq \frac{1}{k} \sum_{K\in\mathcal{K}} \varphi(K)$$

for all $F \in \mathcal{F}(G)$ and for all k-cover $\mathcal{K}$ of F with $K \subseteq F$ for all $K \in \mathcal{K}$. We have the following theorem.

THEOREM 3.2. [26, Theorem 4.48] *Given a countable amenable group G, suppose that $\varphi : \mathcal{F}(G) \to \mathbb{R}$ satisfies Shearer's inequality and $\varphi(Fg) = \varphi(F)$ for all $F \in \mathcal{F}(G)$ and $g \in G$. Then,*

$$\lim_n \frac{\varphi(F_n)}{|F_n|} = \inf_{F\in\mathcal{F}(G)} \frac{\varphi(F)}{|F|}$$

for any Følner sequence $\{F_n\}_n$.

Considering the two previous results, we obtain the next lemma.

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

LEMMA 3.3. *Given a fundamental domain U_0 of $G \curvearrowright \Gamma$ and $\lambda : V \to \mathbb{Q}_{>0}$ such that $\lambda(v) = p_v/q_v$ with $p_v, q_v \in \mathbb{N}$ for all $v \in V$, we have that, for any Følner sequence $\{F_n\}_n$,*

$$\lim_n \frac{\varphi(F_n)}{|F_n|} = \inf_{F \in \mathcal{F}(G)} \frac{\varphi(F)}{|F|},$$

where $\varphi : \mathcal{F}(G) \to \mathbb{R}$ is given by $\varphi(F) = \log Z_\Gamma(FU_0, \lambda) + |F| \sum_{v \in U_0} \log q_v$.

Proof. Given $F \in \mathcal{F}(G)$ and $k \in \mathbb{N}$, let $\mathcal{K}$ be a k-cover of F with $K \subseteq F$ for all $K \in \mathcal{K}$. Notice that

$$\begin{aligned} Z_\Gamma(FU_0, \lambda) &= \sum_{x \in X(\Gamma, FU_0)} \prod_{v \in FU_0} \left(\frac{p_v}{q_v}\right)^{x(v)} \\ &= \frac{1}{\prod_{v \in FU_0} q_v} \sum_{x \in X(\Gamma, FU_0)} \prod_{v \in FU_0} p_v^{x(v)} q_v^{1-x(v)}. \end{aligned}$$

Consider $q := \max_v q_v$, $p := \max_v p_v$, $A := \{-q, \ldots, -1\} \cup \{1, \ldots, p\}$, and

$$Y := \{y \in A^{FU_0} : -q_v \le y(v) \le p_v \text{ and } [y(v) \ge 1 \wedge (v, v') \in E(\Gamma)] \implies y(v') \le -1\}.$$

Notice that

$$|Y| = \sum_{x \in X(\Gamma, FU_0)} \prod_{v \in FU_0} p_v^{x(v)} q_v^{1-x(v)}.$$

Therefore, by Lemma 3.1, and noticing that $|Y_{KU_0}| = \prod_{v \in KU_0} q_v \cdot Z_\Gamma(KU_0, \lambda)$, we have that

$$\prod_{v \in FU_0} q_v \cdot Z_\Gamma(FU_0, \lambda) = |Y| \le \prod_{K \in \mathcal{K}} |Y_{KU_0}|^{1/k} \le \prod_{K \in \mathcal{K}} \left(\prod_{v \in KU_0} q_v \cdot Z_\Gamma(KU_0, \lambda) \right)^{1/k},$$

where we use that $\{KU_0 : K \in \mathcal{K}\}$ is a k-cover of FU_0. Therefore, by G-invariance of λ,

$$\begin{aligned} \varphi(F) &= \log Z_\Gamma(FU_0, \lambda) + |F| \sum_{v \in U_0} \log q_v \\ &\le \frac{1}{k} \sum_{K \in \mathcal{K}} \left(\log Z_\Gamma(KU_0, \lambda) + |K| \sum_{v \in U_0} \log q_v \right) \\ &= \frac{1}{k} \sum_{K \in \mathcal{K}} \varphi(K), \end{aligned}$$

so φ satisfies Shearer's inequality. However, by G-invariance of $X(\Gamma)$ and λ, it follows that $\varphi(Fg) = \varphi(F)$ for all $F \in \mathcal{F}(G)$ and $g \in G$. Therefore, by Theorem 3.2, we conclude. □

PROPOSITION 3.4. *Given a fundamental domain U_0 of $G \curvearrowright \Gamma$, we have that*

$$\lim_n \frac{\log Z_\Gamma(F_n U_0, \lambda)}{|F_n|} = \inf_{F \in \mathcal{F}(G)} \frac{\log Z_\Gamma(FU_0, \lambda)}{|F|}$$

for any Følner sequence $\{F_n\}_n$.

Proof. First, suppose that λ only takes rational values, that is, $\lambda : V \to \mathbb{Q}_{>0}$ so that $\lambda(v) = p_v/q_v$ for all $v \in V$. By Lemma 3.3, for $\varphi(F) = \log Z_\Gamma(FU_0, \lambda) + |F| \sum_{v \in U_0} q_v$,

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

we have that

$$\begin{aligned}\lim_n \frac{\log Z_\Gamma(F_n U_0, \lambda)}{|F_n|} + \sum_{v\in U_0} q_v &= \lim_n \frac{\varphi(F_n)}{|F_n|} \\ &= \inf_{F\in\mathcal{F}(G)} \frac{\varphi(F)}{|F|} \\ &= \inf_{F\in\mathcal{F}(G)} \frac{\log Z_\Gamma(FU_0, \lambda)}{|F|} + \sum_{v\in U_0} q_v,\end{aligned}$$

and, after canceling out $\sum_{v\in U_0} q_v$, we obtain that

$$\lim_n \frac{\log Z_\Gamma(F_n U_0, \lambda)}{|F_n|} = \inf_{F\in\mathcal{F}(G)} \frac{\log Z_\Gamma(FU_0, \lambda)}{|F|}.$$

Now, given a general λ, we can always approximate it by some G-invariant $\tilde{\lambda} : V \to \mathbb{Q}_{>0}$ arbitrarily close in the supremum norm. Given $\epsilon > 0$, pick $\tilde{\lambda}$ so that $\tilde{\lambda}(v) \le \lambda(v) \le (1+\epsilon)\tilde{\lambda}(v)$ for every v. Then,

$$\begin{aligned}\log Z_\Gamma(FU_0, \tilde{\lambda}) &\le \log Z_\Gamma(FU_0, \lambda) \\ &\le \log Z_\Gamma(FU_0, (1+\epsilon)\tilde{\lambda}) \\ &\le |FU_0| \log(1+\epsilon) + \log Z_\Gamma(FU_0, \tilde{\lambda}),\end{aligned}$$

so,

$$\frac{\log Z_\Gamma(FU_0, \tilde{\lambda})}{|F|} \le \frac{\log Z_\Gamma(FU_0, \lambda)}{|F|} \le |U_0| \log(1+\epsilon) + \frac{\log Z_\Gamma(FU_0, \tilde{\lambda})}{|F|}.$$

Therefore,

$$\begin{aligned}\liminf_n \frac{\log Z_\Gamma(F_n U_0, \lambda)}{|F_n|} &\ge \inf_{F\in\mathcal{F}(G)} \frac{\log Z_\Gamma(FU_0, \lambda)}{|F|} \\ &\ge \inf_{F\in\mathcal{F}(G)} \frac{\log Z_\Gamma(FU_0, \tilde{\lambda})}{|F|} \\ &= \lim_n \frac{\log Z_\Gamma(F_n U_0, \tilde{\lambda})}{|F_n|} \\ &\ge \limsup_n \frac{\log Z_\Gamma(F_n U_0, \lambda)}{|F_n|} - |U_0| \log(1+\epsilon),\end{aligned}$$

and since ϵ was arbitrary, we conclude. □

To fully characterize $\lim_n(\log Z_\Gamma(U_n, \lambda)/|U_n|)$, we have the following lemma.

LEMMA 3.5. *Let $\{F_n\}_n$ be a Følner sequence and U_0 a fundamental domain. Then, for any Følner sequence $\{F_n\}_n$,*

$$\lim_n \frac{|F_n U_0|}{|F_n|} = \sum_{v\in U_0} |\mathrm{Stab}_G(v)|^{-1}.$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

Proof. First, pick $v \in U_0$. Since $\mathrm{Stab}_G(v)$ is finite and $\{F_n\}_n$ is a Følner sequence, we have that $\lim_n |F_n\mathrm{Stab}_G(v)|/|F_n| = 1$. However, $F_n\mathrm{Stab}_G(v)v = F_n v$ and for each $v' \in F_n v$, there exist exactly $|\mathrm{Stab}_G(v)|$ different elements $g \in F_n\mathrm{Stab}_G(v)$ such that $gv = v'$. In other words,

$$|F_n\mathrm{Stab}_G(v)| = |F_n v||\mathrm{Stab}_G(v)|,$$

so,

$$\lim_n \frac{|F_n v|}{|F_n|} = \lim_n \frac{|F_n\mathrm{Stab}_G(v)|}{|F_n||\mathrm{Stab}_G(v)|} = |\mathrm{Stab}_G(v)|^{-1}.$$

Therefore,

$$\lim_n \frac{|F_n U_0|}{|F_n|} = \sum_{v\in U_0} \lim_n \frac{|F_n v|}{|F_n|} = \sum_{v\in U_0} |\mathrm{Stab}_G(v)|^{-1}. \qquad \square$$

Now, given a fundamental domain U_0, define

$$f_G(\Gamma, U_0, \lambda) := \inf_{F\in\mathcal{F}(G)} \frac{\log Z_\Gamma(FU_0, \lambda)}{|FU_0|},$$

which, by Proposition 3.4 and Lemma 3.5, is equal to

$$\left(\sum_{v\in U_0} |\mathrm{Stab}_G(v)|^{-1}\right)^{-1} \lim_n \frac{\log Z_\Gamma(F_n U_0, \lambda)}{|F_n|}$$

for any Følner sequence $\{F_n\}_n$ and, in particular, for any Følner exhaustion. Notice that, since $GU_0 = V$, the sequence $\{U_n\}_n$ defined as $U_n = F_n U_0$ is an exhaustion of V in the sense for which we were looking. Now we will see that $f_G(\Gamma, U_0, \lambda)$ is independent of U_0.

PROPOSITION 3.6. *Given two fundamental domains U_0 and U_0' of $G \curvearrowright \Gamma$, we have that*

$$f_G(\Gamma, U_0, \lambda) = f_G(\Gamma, U_0', \lambda).$$

Proof. Since $V = GU_0 = GU_0'$, there must exist $K, K' \in \mathcal{F}(G)$ such that $U_0' \subseteq KU_0$ and $U_0 \subseteq K'U_0'$. Then, for every $F \in \mathcal{F}(G)$,

$$\begin{aligned} FU_0 \triangle FU_0' &= (FU_0 \setminus FU_0') \cup (FU_0' \setminus FU_0) \\ &\subseteq (FK'U_0' \setminus FU_0') \cup (FKU_0 \setminus FU_0) \\ &= (FK' \setminus F)U_0' \cup (FK \setminus F)U_0. \end{aligned}$$

Therefore, $|FU_0 \triangle FU_0| \le |FK' \setminus F||U_0'| + |FK \setminus F||U_0|$. Now, notice that for $U, U' \Subset V$, we always have that:

(1) $Z_\Gamma(U \cup U', \lambda) \le Z_\Gamma(U, \lambda) \cdot Z_\Gamma(U', \lambda)$, provided $U \cap U' = \emptyset$;
(2) $Z_\Gamma(U, \lambda) \le Z_\Gamma(U', \lambda)$, provided $U \subseteq U'$; and
(3) $Z_\Gamma(U, \lambda) \le (2\max\{1, \lambda_+\})^{|U|}$,

so

$$\begin{aligned} \log Z_\Gamma(FU_0, \lambda) &\le \log Z_\Gamma(FU_0 \cap FU_0', \lambda) + \log Z_\Gamma(FU_0 \setminus FU_0', \lambda) \\ &\le \log Z_\Gamma(FU_0', \lambda) + \log Z_\Gamma(FU_0 \triangle FU_0', \lambda) \end{aligned}$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

$$\leq \log Z_\Gamma(FU_0', \lambda) + |FU_0 \triangle FU_0'| \log(2 \max\{1, \lambda_+\})$$
$$\leq \log Z_\Gamma(FU_0', \lambda) + (|FK' \setminus F||U_0'| + |FK \setminus F||U_0|) \log(2 \max\{1, \lambda_+\}).$$

Finally, since $|U_0| = |U_0'|$ and $|FU_0| = |F||U_0|$, it follows by amenability that

$$\begin{aligned} f_G(\Gamma, U_0, \lambda) &= \lim_n \frac{\log Z_\Gamma(F_n U_0, \lambda)}{|F_n U_0|} \\ &\leq \lim_n \frac{\log Z_\Gamma(F_n U_0', \lambda)}{|F_n U_0'|} \\ &\quad + \lim_n \left(\frac{|F_n K' \setminus F_n|}{|F_n|} + \frac{|F_n K \setminus F_n|}{|F_n|} \right) \log(2 \max\{1, \lambda_+\}) \\ &= f_G(\Gamma, U_0', \lambda), \end{aligned}$$

and by symmetry of the argument, we conclude. □

Then, we can consistently define the *Gibbs* (Γ, λ)*-free energy according to* G as

$$f_G(\Gamma, \lambda) := f_G(\Gamma, U_0, \lambda),$$

where U_0 is an arbitrary fundamental domain of $G \curvearrowright \Gamma$. In addition, it is easy to see that if G_1 and G_2 act almost transitively on Γ, and the G_1-orbits and G_2-orbits coincide, that is, $G_1 v = G_2 v$ for all $v \in V$, then

$$f_{G_1}(\Gamma, \lambda) = f_{G_2}(\Gamma, \lambda).$$

In particular, we have that $f_G(\Gamma, \lambda)$ is equal for all G acting transitively on Γ. Then, we can define the *Gibbs* (Γ, λ)*-free energy* as

$$f(\Gamma, \lambda) := \inf_{\emptyset \neq U \Subset V} \frac{\log Z_\Gamma(U, \lambda)}{|U|},$$

which is a quantity that only depends on the graph Γ and the activity function λ, and satisfies that $f(\Gamma, \lambda) = f_G(\Gamma, \lambda)$ for any $G \leq \mathrm{Aut}(\Gamma)$ acting transitively on Γ.

Remark 3.7. In the almost transitive case, $f_G(\Gamma, \lambda)$ does not necessarily coincide with $f(\Gamma, \lambda)$ for G acting almost transitively: consider the graph Γ obtained by taking the disjoint union of $\Gamma_1 = \mathrm{Cay}(\mathbb{Z}, \emptyset)$ and $\Gamma_2 = \mathrm{Cay}(\mathbb{Z}, \{1, -1\})$, and the constant activity function $\lambda \equiv 1$. Then, $f_{\mathbb{Z}}(\Gamma_1, 1) = \log 2$ and $f_{\mathbb{Z}}(\Gamma_2, 1) = \log((1 + \sqrt{5})/2)$, so

$$f_{\mathbb{Z}}(\Gamma, 1) = \frac{1}{2}(f_{\mathbb{Z}}(\Gamma_1, 1) + f_{\mathbb{Z}}(\Gamma_2, 1)) > f_{\mathbb{Z}}(\Gamma_2, 1) \geq \inf_{\emptyset \neq U \Subset V} \frac{\log Z_\Gamma(U, 1)}{|U|}.$$

The value of $f_{\mathbb{Z}}(\Gamma_2, 1)$ corresponds to the *topological entropy* of the *golden mean shift* (see [27, Example 4.1.4] and §9).

The main theme of this paper will be to explore our ability to approximate $f_G(\Gamma, \lambda)$. From now, and without much loss of generality, we will assume that $G \curvearrowright \Gamma$ is free (see §9.7 for a reduction of the almost free case to the free case).

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

4. *Gibbs measures*

Given a graph $\Gamma = (V, E)$, consider the set $\{0, 1\}^V$ endowed with the product topology and the set $X(\Gamma)$, with the subspace topology. The set of independent sets $X(\Gamma)$ is a compact and metrizable space. A base for the topology is given by the *cylinder sets*

$$[x_U] := \{x' \in X(\Gamma) : x'_U = x_U\}$$

for $U \Subset V$ and $x \in X(\Gamma)$, where x_U denotes the *restriction* of x from V to U. If U is a singleton $\{v\}$, we will omit the brackets and simply write x_v and the same convention will hold in analogous instances. Given $W \subseteq V$, we denote by $\mathcal{B}_W$ the smallest σ-algebra generated by

$$\{[x_U] : U \Subset W, x \in X(\Gamma)\},$$

and by $\mathcal{B}_\Gamma$ the *Borel σ-algebra*, which corresponds to $\mathcal{B}_V$.

Let $\mathcal{M}(X(\Gamma))$ be the set of Borel probability measures $\mathbb{P}$ on $X(\Gamma)$. We say that $\mathbb{P}$ is *G-invariant* if $\mathbb{P}(A) = \mathbb{P}(g \cdot A)$ for all $A \in \mathcal{B}_\Gamma$ and $g \in G$, and *G-ergodic* if $g \cdot A = A$ for all $g \in G$ implies that $\mathbb{P}(A) \in \{0, 1\}$. We will denote by $\mathcal{M}_G(X(\Gamma))$ and $\mathcal{M}_G^{\text{erg}}(X(\Gamma))$ the set of G-invariant and the subset of G-invariant measures that are G-ergodic, respectively.

For $\mathbb{P} \in \mathcal{M}(X(\Gamma))$, define the *support* of $\mathbb{P}$ as

$$\text{supp}(\mathbb{P}) := \{x \in X(\Gamma) : \mathbb{P}([x_U]) > 0 \text{ for all } U \Subset V\}.$$

Given $\emptyset \neq U \Subset V$ and $y \in X(\Gamma)$, we define π_U^y to be the probability distribution on $X(\Gamma, U)$ given by

$$\pi_U^y(x) := \mathrm{w}_\lambda^y(x, U) Z_\Gamma^y(U, \lambda)^{-1},$$

where $\mathrm{w}_\lambda^y(x, U) = \mathrm{w}_\lambda(x, U)\mathbb{1}_{[y_{U^c}]}(x)$ and $Z_\Gamma^y(U, \lambda) = \sum_x \mathrm{w}_\lambda^y(x, U)$. In other words, to each independent set x supported on U, we associate a probability proportional to its λ-weight over U, $\prod_{v \in U} \lambda(v)^{x(v)}$, provided x_U is *compatible* with y_{U^c}, in the sense that the element $z \in \{0, 1\}^V$ such that $z_U = x_U$ and $z_{U^c} = y_{U^c}$ is an independent set.

Now, given an activity function $\lambda : V \to \mathbb{R}_{>0}$, consider the hardcore model (Γ, λ) and the collection $\pi_{\Gamma,\lambda} = \{\pi_U^y : U \Subset V, y \in X(\Gamma)\}$. We call $\pi_{\Gamma,\lambda}$ the *Gibbs (Γ, λ)-specification*. A measure $\mathbb{P} \in \mathcal{M}(X(\Gamma))$ is called a *Gibbs measure (for (Γ, λ))* if for all $U \Subset V$, $U' \subseteq U$, and $x \in X(\Gamma)$,

$$\mathbb{P}([x_{U'}]|\mathcal{B}_{U^c})(y) = \pi_U^y([x_{U'}]) \quad \mathbb{P}\text{-almost surely in } y,$$

where $\pi_U^y([x'_U])$ denotes the marginalization

$$\pi_U^y([x'_U]) = \sum_{x' \in X(\Gamma[U]) : x'_{U'} = x_{U'}} \pi_U^y(x')$$

and $\mathbb{P}(A|\mathcal{B}_U) = \mathbb{E}_\mathbb{P}(\mathbb{1}_A|\mathcal{B}_U)$ for $A \in \mathcal{B}_\Gamma$. We denote by $\mathcal{M}_{\text{Gibbs}}(\Gamma, \lambda)$ the set of Gibbs measures for (Γ, λ).

An important question in statistical physics is whether the set of Gibbs measures is empty or not, and if not, whether there is a unique or multiple Gibbs measures [17].

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

4.1. *The locally finite case.* The model described in [17, Example 4.16] can be understood as an attempt to formalize the idea of a system where there is a single particle 1 (uniformly distributed) or none, that is, 0 everywhere. There, it is proven that this model cannot be represented as a Gibbs measure. This example can be also viewed as a hardcore model in a countable graph that is *complete* (that is, there is an edge between any pair of different vertices) and, in particular, in a non-locally finite graph. In other words, there exist examples of non-locally finite graphs Γ such that the (Γ, λ)-specification $\pi_{\Gamma,\lambda}$ has no Gibbs measure.

From now on, we will always assume that Γ is locally finite. In this case, the existence of Gibbs measures is guaranteed (see [10, 13]) and, moreover, every Gibbs measure must be a *Markov random field* that is *fully supported*.

Indeed, it can be checked that $\pi_{\Gamma,\lambda}$ is an example of a *Markovian specification* (see [17, Example 8.24]). In this case, any Gibbs measure $\mathbb{P} \in \mathcal{M}_{\text{Gibbs}}(\Gamma, \lambda)$ satisfies the following *local Markov property*:

$$\mathbb{P}([x_U]|\mathcal{B}_{U^c})(y) = \mathbb{P}([x_U]|\mathcal{B}_{\partial U})(y) \quad \mathbb{P}\text{-almost surely in } y$$

for any $U \Subset V$ and $x \in X(\Gamma)$. In other words, $\mathbb{P}$ is a *Markov random field*, so any event supported on a finite set conditioned to a specific value on its boundary is independent of events supported on the complement.

In addition, it can be checked that any Gibbs measure $\mathbb{P}$ must be *fully supported*, that is, $\text{supp}(\mathbb{P}) = X(\Gamma)$. Indeed, it suffices to check that $X(\Gamma) \subseteq \text{supp}(\mathbb{P})$; the other direction follows directly from the definition of $\pi_{\Gamma,\lambda}$ and Gibbs measures. Now, given $x \in X(\Gamma)$ and $U \Subset V$, we would like to check that $\mathbb{P}([x_U]) > 0$. To prove this, observe that given $x \in X(\Gamma)$, we have that $z \in \{0, 1\}^V$ defined as $z_U = x_U$, $z_{\partial U} \equiv 0$, and $z_{W^c} = y_{W^c}$ always belongs to $X(\Gamma)$ for any $y \in X(\Gamma)$, where $W = U \cup \partial U$. In particular, $\pi_W^y(z) > 0$ for any $y \in X(\Gamma)$. Then, considering that $\partial(W^c)$ is finite,

$$\begin{aligned}
\mathbb{P}([x_U]) &\geq \mathbb{P}([z_W]) \\
&= \sum_{y \in X(\Gamma, \partial W): \mathbb{P}([y_{\partial W}]) > 0} \mathbb{P}([z_W]|[y_{\partial W}])\mathbb{P}([y_{\partial W}]) \\
&= \sum_{y \in X(\Gamma, \partial W): \mathbb{P}([y_{\partial W}]) > 0} \pi_W^y(z)\mathbb{P}([y_{\partial W}]) \\
&\geq \pi_W^{y^*}(z)\mathbb{P}([y^*_{\partial W}]) > 0,
\end{aligned}$$

since $\mathbb{P}$ is a probability measure and there must exist $y^* \in X(\Gamma)$ such that $\mathbb{P}([y^*_{\partial W}]) > 0$. In other words, $X(\Gamma)$ satisfies the property (D*) introduced in [41, 1.14 Remark], which guarantees full support.

4.2. *Spatial mixing and uniqueness.* Given a Gibbs (Γ, λ)-specification $\pi_{\Gamma,\lambda}$, we define two *spatial mixing* properties fundamental to this work.

Definition 4.1. We say that a hardcore model (Γ, λ) exhibits *strong spatial mixing (SSM)* if there exists a *decay rate* function $\delta : \mathbb{N} \to \mathbb{R}_{\geq 0}$ such that $\lim_{\ell \to \infty} \delta(\ell) = 0$ and for all

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

$U \Subset V$, $v \in U$, and $y, z \in X(\Gamma)$,

$$|\pi_U^y([0^v]) - \pi_U^z([0^v])| \leq \delta(\mathrm{dist}_\Gamma(v, D_U(y, z))),$$

where $[0^v]$ denotes the event that the vertex v takes the value 0 and

$$D_U(y, z) := \{v' \in U^c : y(v') \neq z(v')\}.$$

This definition is equivalent (see [35, Lemma 2.3]) to the—*a priori*—stronger following property: for all $U' \subseteq U \Subset V$ and $x, y, z \in X(\Gamma)$,

$$|\pi_U^y([x_{U'}]) - \pi_U^z([x_{U'}])| \leq |U'|\delta(\mathrm{dist}_\Gamma(U', D_U(y, z))).$$

Similarly, we say that (Γ, λ) exhibits *weak spatial mixing (WSM)* if for all $U' \subseteq U \Subset V$ and $x, y, z \in X(\Gamma)$,

$$|\pi_U^y([x_{U'}]) - \pi_U^z([x_{U'}])| \leq |U'| \cdot \delta(\mathrm{dist}_\Gamma(U', U^c)).$$

Clearly, SSM implies WSM. Moreover, it is well known that, in this context, WSM (and therefore, SSM) implies uniqueness of Gibbs measures [54]. In other words, $\mathcal{M}_{\mathrm{Gibbs}}(\Gamma, \lambda) = \{\mathbb{P}_{\Gamma,\lambda}\}$, where $\mathbb{P}_{\Gamma,\lambda}$ denotes the unique Gibbs measure for (Γ, λ). In this case, $\mathbb{P}_{\Gamma,\lambda}$ is always $\mathrm{Aut}(\Gamma)$-invariant.

We say that (Γ, λ) exhibits *exponential SSM* (respectively *exponential WSM*) if there exist constants $C, \alpha > 0$ such that $\pi_{\Gamma,\lambda}$ exhibits SSM (respectively WSM) with decay rate function $\delta(n) = C \cdot \exp(-\alpha \cdot n)$.

Given $U \subseteq V$, we denote by $\Gamma \setminus U$ the subgraph induced by $V \setminus U$, that is, $\Gamma[V \setminus U]$. We have the following result due to Gamarnik and Katz.

PROPOSITION 4.1. [16, Proposition 1] *If a hardcore model (Γ, λ) satisfies SSM, then so does the hardcore model (Γ', λ) for any subgraph Γ' of Γ. The same assertion applies to exponential SSM. Moreover, for every $U \subseteq V$ and $v \in V \setminus U$, the following identity holds:*

$$\mathbb{P}_{\Gamma,\lambda}([0^v]|[0^U]) = \mathbb{P}_{\Gamma\setminus U,\lambda}([0^v]),$$

where $\mathbb{P}_{\Gamma,\lambda}$ and $\mathbb{P}_{\Gamma\setminus U,\lambda}$ are the unique Gibbs measures for (Γ, λ) and $(\Gamma \setminus U, \lambda)$, respectively, and $[0^U]$ denotes the event that all the vertices in U take the value 0. In particular, $\mathbb{P}_{\Gamma,\lambda}([0^v]|[0^U])$ is always well defined, even if U is infinite.

Remark 4.2. Notice that any event of the form $[x_U]$ can be translated into an event of the form $[0^{U'}]$ for a suitable set U': it suffices to define $U' = U \cup \partial\{v \in U : x_U(v) = 1\}$ since, deterministically, every neighbor of a vertex colored 1 must be 0, so Proposition 4.1 still holds for more general events. We also remark that in [16], it is assumed that λ is a constant function. Here we drop this assumption, but it is direct to check that the same proof of [16, Proposition 1] also applies to the more general non-constant case.

4.3. *Families of hardcore models.* We will denote by $\mathcal{H}$ the family of hardcore models (Γ, λ) such that Γ is a countable locally finite graph and λ is any activity function $\lambda : V(\Gamma) \to \mathbb{R}_{>0}$.

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

Given a countable group G, we will denote by $\mathcal{H}_G$ the set of hardcore models (Γ, λ) in $\mathcal{H}$ for which G is isomorphic to some subgroup of $\mathrm{Aut}(\Gamma)$ such that $G \curvearrowright \Gamma$ is free and almost transitive and $\lambda : V(\Gamma) \to \mathbb{R}_{>0}$ is a G-invariant activity function.

Given a positive integer Δ, we will denote by $\mathcal{H}^{\Delta}$ the set of hardcore models (Γ, λ) in $\mathcal{H}$ such that $\Delta(\Gamma) \leq \Delta$. Notice that any hardcore model defined on the Δ-regular (infinite) tree $\mathbb{T}_{\Delta}$ belongs to $\mathcal{H}^{\Delta}$.

Given $\lambda_0 > 0$, we will denote by $\mathcal{H}(\lambda_0)$ the family of hardcore models (Γ, λ) in $\mathcal{H}$ such that $\lambda_+ \leq \lambda_0$.

We will also combine the notation for these families in the natural way; for example, $\mathcal{H}_G^{\Delta}(\lambda_0)$ will denote the set of hardcore models (Γ, λ) in $\mathcal{H}$ such that $G \curvearrowright \Gamma$ is free and almost transitive, λ is G-invariant, $\Delta(\Gamma) \leq \Delta$, and $\lambda_+ \leq \lambda_0$.

5. *Trees*

Given a graph Γ, a *trail* w in Γ is a finite sequence $w = (v_1, \ldots, v_n)$ of vertices such that consecutive vertices are adjacent in Γ and the edges (v_i, v_{i+1}) involved are not repeated. For a fixed vertex $v \in V(\Gamma)$, the *tree of self-avoiding walks starting from* v, denoted by $T_{\mathrm{SAW}}(\Gamma, v)$, is defined as follows.

(1) Consider the set W_0 of trails starting from v that repeat no vertex and the set W_1 of trails that repeat a single vertex exactly once and then stop (that is, the set of non-backtracking walks that end immediately after performing a cycle). We define $T_{\mathrm{SAW}}(\Gamma, v)$ to be a rooted tree with root $\rho = (v)$ such that the set of vertices $V(T_{\mathrm{SAW}}(\Gamma, v))$ is $W_0 \cup W_1$ and the set of (undirected) edges $E(T_{\mathrm{SAW}}(\Gamma, v))$ corresponds to all the pairs (w, w') such that w' is a one-vertex extension of w or *vice versa*. In simple words, $T_{\mathrm{SAW}}(\Gamma, v)$ is a rooted tree that represents all self-avoiding walks in Γ that start from v. It is easy to check that the set of leaves of $T_{\mathrm{SAW}}(\Gamma, v)$ contains W_1, but they are not necessarily equal (e.g., see vertex b in Figure 2).

(2) For $u \in V(\Gamma)$, consider an arbitrary ordering $\partial\{u\} = \{u_1, \ldots, u_d\}$ of its neighbors. Given $w \in W_1$, we can represent this walk as a sequence

$$w = (v, \ldots, u, u_i, \ldots, u_j, u),$$

with $u_i, u_j \in \partial\{u\}$. Notice that $i \neq j$, since we are not repeating edges. Considering this, we condition the 'terminal' trail w to be 1 (*occupied*) if $i < j$ and to be 0 (*unoccupied*) if $i > j$, inducing the corresponding effect of this conditioning in the graph (that is, removing the vertex and its neighbors or just removing the vertex, respectively).

Given a hardcore model (Γ, λ), a vertex $v \in V(\Gamma)$, a subset $U \subseteq V(\Gamma)$, and an independent set $x \in X(\Gamma)$, we are interested in computing the marginal probability that v is unoccupied in Γ given the partial configuration x_U, that is, $\mathbb{P}_{\Gamma,\lambda}([0^v]|[x_U])$. Notice that if (Γ, λ) satisfies SSM (which includes the particular but relevant case of Γ being finite), then this probability is always well defined due to Proposition 4.1, even if U is infinite.

To understand better $\mathbb{P}_{\Gamma,\lambda}([0^v]|[x_U])$, we consider $(T_{\mathrm{SAW}}(\Gamma, v), \overline{\lambda})$ to be the hardcore model where $\overline{\lambda}(w) = \lambda(u)$ for every trail w ending in u. In this context, a condition x_U in (Γ, λ) is translated into the condition $\overline{x_U}$ in $T_{\mathrm{SAW}}(\Gamma, v)$, whose support is the set $W(U)$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

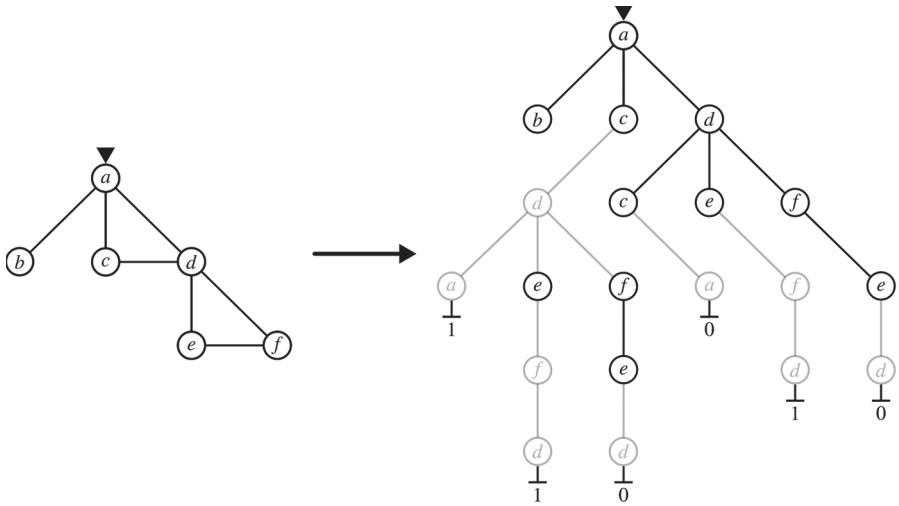

FIGURE 2. A representation of a graph Γ and its corresponding tree of self-avoiding walks $T_{\mathrm{SAW}}(\Gamma, v)$ including the conditioning of terminal trails ($\perp$). Here, the order of each neighborhood is alphabetical and every trail/vertex is represented by the final vertex of the trail in Γ starting from $v = a$. See also [55] for an explanation of the same picture.

of trails w that end in u for some $u \in U$, and $\overline{x}(w) = x(u)$ for all these w. We have the following result from [55], that we adapt to the more general non-constant λ case and we include its proof for completeness.

THEOREM 5.1. [55, Theorem 3.1] *For every finite hardcore model* (Γ, λ)*, every* $v \in V(\Gamma)$*, and* $U \subseteq V(\Gamma)$*,*

$$\mathbb{P}_{\Gamma,\lambda}([0^v]|[x_U]) = \mathbb{P}_{T_{\mathrm{SAW}}(\Gamma,v),\overline{\lambda}}([0^\rho]|[\overline{x_U}]).$$

Proof. Instead of probabilities, we work with the ratios

$$R_{\Gamma,\lambda}(v, x_U) := \frac{\mathbb{P}_{\Gamma,\lambda}([1^v]|[x_U])}{\mathbb{P}_{\Gamma,\lambda}([0^v]|[x_U])},$$

where if $v \in U$ and $x_U(v)$ is equal to 1 or 0, we let $R_{\Gamma,\lambda}(v, x_U)$ be ∞ or 0, respectively. Notice that

$$\mathbb{P}_{\Gamma,\lambda}([0^v]|[x_U]) = \frac{1}{1 + R_{\Gamma,\lambda}(v, x_U)} \quad \text{and} \quad \mathbb{P}_{\Gamma,\lambda}([1^v]|[x_U]) = \frac{R_{\Gamma,\lambda}(v, x_U)}{1 + R_{\Gamma,\lambda}(v, x_U)}.$$

Given a finite tree T rooted at ρ, let us denote by $\{\rho_1, \ldots, \rho_d\}$ the set of neighbors $\partial\{\rho\}$ of ρ and by T_i, for $i = 1, \ldots, d$, the corresponding subtrees starting from ρ_i, that is, $V(T) = \{\rho\} \cup V(T_1) \cup \cdots \cup V(T_d)$. If we have a condition x_U on U, we define $U_i = U \cap V(T_i)$ and $x_{U_i} = (x_U)|_{U_i}$. Considering this, we have that

$$\begin{aligned} R_{T,\lambda}(\rho, x_U) &= \frac{\mathbb{P}_{T,\lambda}([1^\rho]|[x_U])}{\mathbb{P}_{T,\lambda}([0^\rho]|[x_U])} \\ &= \frac{\lambda(\rho) \cdot Z^{x_U}_{T\setminus\{\rho\cup\partial\{\rho\}\}}(\lambda)}{Z^{x_U}_T(\lambda)} \cdot \frac{Z^{x_U}_T(\lambda)}{Z^{x_U}_{T\setminus\{\rho\}}(\lambda)} \end{aligned}$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

$$= \lambda(\rho) \cdot \prod_{i=1}^{d} \frac{Z_{T_i \setminus \{\rho_i\}}^{x_{U_i}}(\lambda_i)}{Z_{T_i}^{x_{U_i}}(\lambda)}$$
$$= \lambda(\rho) \cdot \prod_{i=1}^{d} \mathbb{P}_{T_i,\lambda}([0^{\rho_i}]|[x_{U_i}])$$
$$= \lambda(\rho) \cdot \prod_{i=1}^{d} \frac{1}{1 + R_{T_i,\lambda}(\rho_i, x_{U_i})},$$

where

$$Z_{\Gamma}^{x_U}(\lambda) := \sum_{y \in X(\Gamma): y_U = x_U} \prod_{v \in V(\Gamma)} \lambda(v)^{y(v)}.$$

Notice that this gives us a linear recursive procedure for computing $R_{T,\lambda}(\rho, x_U)$, and therefore $\mathbb{P}_{T,\lambda}([0^{\rho}]|[x_U])$, with base cases: $R_{T,\lambda}(\rho, x_U) = 0$ or $+\infty$ if ρ is fixed, and $R_{T,\lambda}(\rho, x_U) = \lambda(\rho)$ if ρ is free and isolated.

Now, consider an arbitrary hardcore model (Γ, λ) and $v \in V(\Gamma)$ with neighbors $\partial\{v\} = \{u_1, \ldots, u_d\}$. We consider the auxiliary hardcore model (Γ', λ'), where:

- $V(\Gamma') = V(\Gamma) \setminus \{v\} \cup \{v_1, \ldots, v_d\}$;
- $E(\Gamma') = E(\Gamma) \setminus \{(v, u_i)\}_{i=1,\ldots,d} \cup \{(v_i, u_i)\}_{i=1,\ldots,d}$;
- $\lambda'(v_i) = \lambda(v)^{1/d}$ for $i = 1, \ldots, d$, and $\lambda'(u) = \lambda(u)$ otherwise.

Notice that

$$R_{\Gamma,\lambda}(v, x_U) = \frac{\mathbb{P}_{\Gamma,\lambda}([1^v]|[x_U])}{\mathbb{P}_{\Gamma,\lambda}([0^v]|[x_U])}$$
$$= \frac{\mathbb{P}_{\Gamma',\lambda'}([1^{\{v_1,\ldots,v_d\}}]|[x_U])}{\mathbb{P}_{\Gamma',\lambda'}([0^{\{v_1,\ldots,v_d\}}]|[x_U])}$$
$$= \prod_{i=1}^{d} \frac{\mathbb{P}_{\Gamma',\lambda'}([0^{\{v_1,\ldots,v_{i-1}\}}1^{\{v_i,\ldots,v_d\}}]|[x_U])}{\mathbb{P}_{\Gamma',\lambda'}([0^{\{v_1,\ldots,v_i\}}1^{\{v_{i+1},\ldots,v_d\}}]|[x_U])}$$
$$= \prod_{i=1}^{d} \frac{\mathbb{P}_{\Gamma',\lambda'}([1^{v_i}]|[x_U z_i])}{\mathbb{P}_{\Gamma',\lambda'}([0^{v_i}]|[x_U z_i])}$$
$$= \prod_{i=1}^{d} R_{\Gamma',\lambda'}(v_i, x_U z_i),$$

where $z_i = 0^{\{v_1,\ldots,v_{i-1}\}}1^{\{v_{i+1},\ldots,v_d\}}$ and $x_U z_i$ is the concatenation of x_U and z_i. Now, since v_i is connected only to u_i, notice that

$$R_{\Gamma',\lambda'}(v_i, x_U z_i) = \frac{\lambda'(v_i) \cdot Z_{\Gamma' \setminus \{v_i, u_i\}}^{x_U z_i}(\lambda')}{Z_{\Gamma' \setminus \{v_i\}}^{x_U z_i}(\lambda')} = \frac{\lambda^{1/d}(v)}{1 + R_{\Gamma' \setminus \{v_i\},\lambda'}(u_i, x_U z_i)}.$$

Therefore,

$$R_{\Gamma,\lambda}(v, x_U) = \prod_{i=1}^{d} \frac{\lambda^{1/d}(v)}{1 + R_{\Gamma' \setminus \{v_i\},\lambda'}(u_i, x_U z_i)} = \lambda(v) \cdot \prod_{i=1}^{d} \frac{1}{1 + R_{\Gamma' \setminus \{v_i\},\lambda'}(u_i, x_U z_i)}.$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

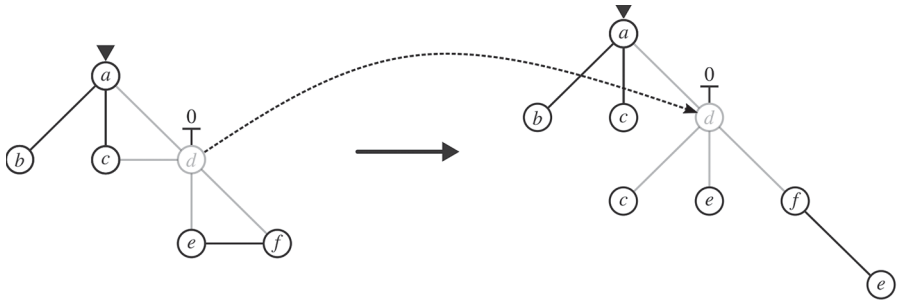

FIGURE 3. A condition ($\top$) on Γ and its representation on the tree of self-avoiding walks $T_{\mathrm{SAW}}(\Gamma, v)$ for $v = a$. The only relevant portion of the tree for computing the marginal probability associated to the root is its connected component.

Notice that the previous recursion can increase the original number of vertices, but the number of free vertices always decreases, so the recursion ends. Then, we have that:

(1) $R_{T,\lambda}(v, x_U) = \lambda(\rho) \cdot f(R_{T_1,\lambda}(\rho_1, x_{U_1}), \ldots, R_{T_d,\lambda}(\rho_d, x_{U_d}))$; and

(2) $R_{\Gamma,\lambda}(v, x_U) = \lambda(v) \cdot f(R_{\Gamma'\setminus\{v_1\},\lambda'}(u_1, x_U\tau_1), \ldots, R_{\Gamma'\setminus\{v_d\},\lambda'}(u_d, x_U\tau_d))$,

where $f(r_1, \ldots, r_d) = \prod_{i=1}^{d}(1/(1 + r_i))$. Now we proceed by induction in the number of free vertices. We can consider the base case where there are no free vertices (besides v) and the theorem is trivial. Then, if we know that the theorem is true when we have n free vertices, we prove it for $n + 1$. Notice that if $R_{\Gamma',\lambda'}(v, x_U)$ involves $n + 1$ free vertices, then $R_{\Gamma'\setminus\{v\},\lambda'}(v_i, x_U z_i)$ involves n free vertices, so by the induction hypothesis,

$$R_{\Gamma'\setminus\{v_i\},\lambda}(u_i, x_U z_i) = R_{T_{\mathrm{SAW}}(\Gamma,v),\bar{\lambda}}(\rho_i, \overline{x_U z_i}).$$

Then, noticing that the rooted subtree (T_i, ρ_i) and the condition $\overline{x_U z_i}$ give exactly the tree of self-avoiding walks of $\Gamma'\setminus\{v_i\}$ starting from u_i under the condition $x_U z_i$, we are done. □

Remark 5.2. The recursions presented in the proof of Theorem 5.1 give us a recursive procedure to compute the marginal probability of the root ρ of a tree T being occupied which requires linear time with respect to the size of the tree. However, if Γ is such that $\Delta(\Gamma) \leq \Delta$, then $T_{\mathrm{SAW}}(\Gamma, v)$ is a subtree of $\mathbb{T}_\Delta$ and its size of $T_{\mathrm{SAW}}(\Gamma, v)$ can be (at most) exponential in the size of Γ. Since hardcore models are Markov random fields and we are interested in the sensitivity of the root ρ associated to v, we only need to consider the graph obtained after pruning all the subtrees below $W(U)$ (see Figure 3).

Before stating the main results concerning hardcore models and strong spatial mixing, we will establish the following bounds.

LEMMA 5.3. *Given a finite hardcore model* $(\Gamma, V) \in \mathcal{H}^\Delta$ *and* $v \in V$, *we have that*

$$0 < \frac{1}{1 + \lambda_+} \leq \mathbb{P}_{\Gamma,\lambda}([0^v]) \leq \frac{(1 + \lambda_+)^\Delta}{\lambda_- + (1 + \lambda_+)^\Delta} < 1$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

and

$$0 < \frac{\lambda_-}{\lambda_- + (1+\lambda_+)^\Delta} \le \mathbb{P}_{\Gamma,\lambda}([1^v]) \le \frac{\lambda_+}{1+\lambda_+} < 1.$$

Proof. Notice that, since a 1 at v forces 0s in $\partial\{v\}$,

$$\mathbb{P}_{\Gamma,\lambda}([1^v]) = \mathbb{P}_{\Gamma,\lambda}([1^v]|[0^{\partial\{v\}}])\mathbb{P}_{\Gamma,\lambda}([0^{\partial\{v\}}]) \le \mathbb{P}_{\Gamma,\lambda}([1^v]|[0^{\partial\{v\}}]),$$

so, considering that $\mathbb{P}_{\Gamma,\lambda}$ is a Markov random field and $\lambda/(1+\lambda)$ is increasing in $\lambda > 0$, we obtain that

$$\mathbb{P}_{\Gamma,\lambda}([1^v]) \le \mathbb{P}_{\Gamma,\lambda}([1^v]|[0^{\partial\{v\}}]) = \frac{\lambda(v)}{1+\lambda(v)} \le \frac{\lambda_+}{1+\lambda_+} < 1$$

and

$$\mathbb{P}_{\Gamma,\lambda}([0^v]) = 1 - \mathbb{P}_{\Gamma,\lambda}([1^v]) \ge 1 - \frac{\lambda_+}{1+\lambda_+} = \frac{1}{1+\lambda_+} > 0.$$

However, by Theorem 5.1, without loss of generality, we can suppose that Γ is a tree rooted at v. Then, if Γ_i denotes the ith subtree of Γ rooted at $v_i \in \partial\{v\}$,

$$\begin{aligned}\frac{\mathbb{P}_{\Gamma,\lambda}([1^v])}{\mathbb{P}_{\Gamma,\lambda}([0^v])} &= \lambda(v)\cdot\prod_{i=1}^{d}\frac{1}{1+\mathbb{P}_{(\Gamma_i,\lambda)}([1^{v_i}])/\mathbb{P}_{(\Gamma_i,\lambda)}([0^{v_i}])}\\ &\ge \lambda_-\cdot\prod_{i=1}^{d}\frac{1}{1+\lambda_+/(1+\lambda_+)/1/(1+\lambda_+)} \ge \frac{\lambda_-}{(1+\lambda_+)^\Delta}.\end{aligned}$$

Therefore, since $\mathbb{P}_{\Gamma,\lambda}([0^v]) = 1 - \mathbb{P}_{\Gamma,\lambda}([1^v])$, we have that

$$\mathbb{P}_{\Gamma,\lambda}([1^v]) \ge \frac{\lambda_-}{\lambda_- + (1+\lambda_+)^\Delta} > 0 \quad\text{and}\quad \mathbb{P}_{\Gamma,\lambda}([0^v]) \le \frac{(1+\lambda_+)^\Delta}{\lambda_- + (1+\lambda_+)^\Delta} < 1. \qquad \square$$

We define the *critical activity function* $\lambda_c : [2, +\infty) \to (0, +\infty]$ as

$$\lambda_c(t) := \frac{(t-1)^{(t-1)}}{(t-2)^t}.$$

We have the following result.

PROPOSITION 5.4. [25] *For every $\Delta \in \mathbb{N}$, the hardcore model $(\mathbb{T}_\Delta, \lambda_0)$ exhibits WSM if and only if $\lambda_0 \le \lambda_c(\Delta)$. If the inequality is strict, then $(\mathbb{T}_\Delta, \lambda_0)$ exhibits exponential WSM with a decay rate δ involving constants that depend on Δ and λ_0.*

We summarize, in the following theorem, the main results from [55] that relate the correlation decay in $(\mathbb{T}_\Delta, \lambda_0)$ with the correlation decay in $\mathcal{H}^\Delta(\lambda_0)$. Here again, as in Theorem 5.1, the results in [55] are focused on the constant activity case. However, we can also adapt the results to the non-constant case by considering that the main tool used in [55] to prove them is [55, Theorem 4.1], which is based on hardcore models with non-constant activity functions.

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

THEOREM 5.5. [55, Theorems 2.3 and 2.4] *Fix* $\Delta \in \mathbb{N}$ *and* $\lambda_0 > 0$. *Then:*

(1) *if* $(\mathbb{T}_\Delta, \lambda_0)$ *exhibits WSM with decay rate* δ, *then* $(\mathbb{T}_\Delta, \lambda_0)$ *exhibits SSM with rate*

$$\frac{(1+\lambda_0)(\lambda_0 + (1+\lambda_0)^\Delta)}{\lambda_0}\delta;$$

(2) *if* $(\mathbb{T}_\Delta, \lambda_0)$ *exhibits SSM with decay rate* δ, *then* (Γ, λ) *exhibits SSM with rate* δ *for every* $(\Gamma, \lambda) \in \mathcal{H}^\Delta(\lambda_0)$.

Then, combining Proposition 5.4 and Theorem 5.5, we have that if $\lambda_0 \leq \lambda_c(\Delta)$, then every hardcore model $(\Gamma, \lambda) \in \mathcal{H}^\Delta(\lambda_0)$ exhibits SSM with the same decay rate δ that would be exponential if the inequality is strict. In addition, observe that if (Γ, λ) is a hardcore model such that $(T_{\mathrm{SAW}}(\Gamma, v), \overline{\lambda})$ exhibits SSM with decay rate δ for every $v \in V(\Gamma)$, then (Γ, λ) exhibits SSM with decay rate δ as well. This follows from Theorem 5.1, since SSM is a property that depends on finitely supported events and the probabilities involved can be translated into probabilities defined on finite hardcore models which at the same time can be translated into events on finite subtrees of $T_{\mathrm{SAW}}(\Gamma, v)$. Considering this, we have the following theorem, which can be understood as a generalization of Theorem 5.1 to the infinite setting.

THEOREM 5.6. *Given a hardcore model* (Γ, λ) *and* $v \in V(\Gamma)$ *such that* $(T_{\mathrm{SAW}}(\Gamma, v), \overline{\lambda})$ *exhibits SSM, then for every* $x \in X(\Gamma)$ *and* $U \subseteq V(\Gamma)$,

$$\mathbb{P}_{\Gamma,\lambda}([0^v]|[x_U]) = \mathbb{P}_{T_{\mathrm{SAW}}(\Gamma,v),\overline{\lambda}}([0^\rho]|[\overline{x_U}]).$$

Proof. Assume that $(T_{\mathrm{SAW}}(\Gamma, v), \overline{\lambda})$ exhibits SSM with decay rate δ. Then, for every $\ell \in \mathbb{N}$,

$$\begin{aligned}
\mathbb{P}_{\Gamma,\lambda}([0^v]|[x_U]) &= \sum_{w \in \{0,1\}^{\partial B_\Gamma(v,\ell)\setminus U}} \mathbb{P}_{\Gamma,\lambda}([0^v]|[x_U w])\mathbb{P}_{\Gamma,\lambda}([w]|[x_U]) \\
&\leq \sum_{w \in \{0,1\}^{\partial B_\Gamma(v,\ell)\setminus U}} (\mathbb{P}_{\Gamma,\lambda}([0^v]|[x_U 0^{\partial B_\Gamma(v,\ell)\setminus U}]) + \delta(\ell))\mathbb{P}_{\Gamma,\lambda}([w]|[x_U]) \\
&= \mathbb{P}_{\Gamma,\lambda}([0^v]|[x_U 0^{\partial B(v,\ell)\setminus U}]) + \delta(\ell)
\end{aligned}$$

and, similarly,

$$\mathbb{P}_{\Gamma,\lambda}([0^v]|[x_U]) \geq \mathbb{P}_{\Gamma,\lambda}([0^v]|[x_U 0^{\partial B_\Gamma(v,\ell)\setminus U}]) - \delta(\ell).$$

Therefore, since $\lim_{\ell\to\infty} \delta(\ell) = 0$,

$$\mathbb{P}_{\Gamma,\lambda}([0^v]|[x_U]) = \lim_{\ell\to\infty} \mathbb{P}_{\Gamma,\lambda}([0^v]|[x_U 0^{\partial B_\Gamma(v,\ell)\setminus U}]),$$

and, by the same argument,

$$\mathbb{P}_{(T_{\mathrm{SAW}}(\Gamma,v),\overline{\lambda})}([0^\rho]|[\overline{x_U}]) = \lim_{\ell\to\infty} \mathbb{P}_{(T_{\mathrm{SAW}}(\Gamma,v),\overline{\lambda})}([0^\rho]|[\overline{x_U} 0^{\partial B_{T_{\mathrm{SAW}}(\Gamma,v)}(\rho,\ell)\setminus W(U)}]).$$

Considering this, the Markov random field property, and Proposition 4.1, we have that

$$\begin{aligned}
\mathbb{P}_{\Gamma,\lambda}([0^v]|[x_U]) &= \lim_{\ell\to\infty} \mathbb{P}_{\Gamma,\lambda}([0^v]|[x_U 0^{\partial B_\Gamma(v,\ell)\setminus U}]) \\
&= \lim_{\ell\to\infty} \mathbb{P}_{\Gamma\cap B(v,\ell),\lambda}([0^v]|[x_{U\cap B_\Gamma(v,\ell)}])
\end{aligned}$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

$$= \lim_{\ell\to\infty} \mathbb{P}_{T_{\mathrm{SAW}}(\Gamma\cap B_{T_{\mathrm{SAW}}(\Gamma,v)}(v,\ell),v),\bar{\lambda}}([0^{\rho}]|[\overline{x_{U\cap B_{\Gamma}(v,\ell)}}])$$
$$= \lim_{\ell\to\infty} \mathbb{P}_{T_{\mathrm{SAW}}(\Gamma,v),\bar{\lambda}}([0^{\rho}]|[\overline{x_U}0^{\partial B_{T_{\mathrm{SAW}}(\Gamma,v)}(\rho,\ell)\setminus W(U)}])$$
$$= \mathbb{P}_{T_{\mathrm{SAW}}(\Gamma,v),\bar{\lambda}}([0^{\rho}]|[\overline{x_U}]). \qquad \square$$

Notice that Theorem 5.6 requires that $(T_{\mathrm{SAW}}(\Gamma, v), \bar{\lambda})$ exhibits SSM rather than the graph (Γ, λ), since SSM on $T_{\mathrm{SAW}}(\Gamma, v)$ may be a stronger condition than SSM on Γ. A key fact is that if $(\mathbb{T}_{\Delta}, \lambda_0)$ exhibits SSM, then (T, λ_0) exhibits SSM for every subtree T of $\mathbb{T}_{\Delta}$. Then, since for every Γ with $\Delta(\Gamma) \leq \Delta$ we have that $T_{\mathrm{SAW}}(\Gamma, v)$ is a subtree of $\mathbb{T}_{\Delta}$, it follows that $(T_{\mathrm{SAW}}(\Gamma, v), \bar{\lambda})$, and therefore (Γ, λ_0), exhibit SSM. Considering this, we have the following corollary.

COROLLARY 5.7. *Fix $\Delta \in \mathbb{N}$. Then, every $(\Gamma, \lambda) \in \mathcal{H}^{\Delta}(\lambda_c(\Delta))$ exhibits SSM and for every $v \in V(\Gamma)$, $x \in X(\Gamma)$, and $U \subseteq V(\Gamma)$,*

$$\mathbb{P}_{\Gamma,\lambda}([0^v]|[x_U]) = \mathbb{P}_{T_{\mathrm{SAW}}(\Gamma,v),\bar{\lambda}}([0^{\rho}]|[\overline{x_U}]).$$

Since we are ultimately interested in studying the interplay between the SSM property on $T_{\mathrm{SAW}}(\Gamma, v)$ and Γ, we may wonder whether it is really necessary to have control over the full Δ-regular tree $\mathbb{T}_{\Delta}$. In [46], a refinement of this fact was proved by considering the *connective constant* of the graphs involved.

5.1. *Connective constant.* Given a graph Γ, a vertex v, and an integer k, let $N_{\Gamma}(v, k)$ denote the *number of self-avoiding walks* in Γ of length k starting from v. Following [46], we consider the *connective constant* $\mu(\mathcal{G})$ of a family of finite graphs $\mathcal{G}$. Here, $\mu(\mathcal{G})$ is defined as the infimum over all $\mu > 0$ for which there exist $a, c > 0$ such that for any $\Gamma \in \mathcal{G}$ and any $v \in V(\Gamma)$, it holds that $\sum_{k=1}^{\ell} N_{\Gamma}(v, k) \leq c\mu^{\ell}$ for all $\ell \geq a \log|V(\Gamma)|$. This definition extends the more usual definition of connective constant for a single infinite almost transitive graph Γ, which is given by

$$\mu(\Gamma) := \max_{v\in V(\Gamma)} \lim_{\ell\to\infty} N_{\Gamma}(v, \ell)^{1/\ell}.$$

Indeed, if Γ is almost transitive, then $\mu(\Gamma) = \mu(\mathcal{G}(\Gamma))$, where $\mathcal{G}(\Gamma)$ denotes the family of finite subgraphs of Γ. Notice that $\mu(\Gamma)$ exists due to Fekete's lemma and that, if Γ is connected, then $\mu(\Gamma) = \lim_{\ell\to\infty} N_{\Gamma}(v, \ell)^{1/\ell}$ for arbitrary v. Roughly, the connective constant measures the growth rate of the number of self-avoiding walks according to their length or, equivalently, the *branching* of $T_{\mathrm{SAW}}(\Gamma, v)$. In general, it is not an easy task to compute $\mu(\Gamma)$ (e.g., see [15]).

Considering this, we extend the definition of strong spatial mixing to families of graphs as follows: given a family of graphs $\mathcal{G}$ and a family of activity functions $\Lambda = \{\lambda^{\Gamma}\}_{\Gamma\in\mathcal{G}}$ with $\lambda^{\Gamma} : V(\Gamma) \to \mathbb{R}_{>0}$, we say that $(\mathcal{G}, \Lambda)$ satisfies strong spatial mixing if there exists a decay rate function $\delta : \mathbb{N} \to \mathbb{R}_{\geq 0}$ such that $\lim_{\ell\to\infty} \delta(\ell) = 0$ and for all $\Gamma \in \mathcal{G}$, for all $U \Subset V(\Gamma)$, $v \in U$, and $y, z \in X(\Gamma)$,

$$|\pi^{y}_{\Gamma,U}([0^v]) - \pi^{z}_{\Gamma,U}([0^v])| \leq \delta(\mathrm{dist}_{\Gamma}(v, D_U(y, z))),$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

where $\pi^{y}_{\Gamma,U}$ denotes the specification element corresponding to the hardcore model $(\Gamma, \lambda^{\Gamma})$. We translate into this language the following result from [46].

THEOREM 5.8. [46] *Let $\mathcal{G}$ be a family of almost transitive locally finite graphs and $\Lambda = \{\lambda^{\Gamma}\}_{\Gamma\in\mathcal{G}}$ a set of activity functions such that*

$$\sup_{\Gamma\in\mathcal{G}} \lambda^{\Gamma}_{+} < \lambda_c(\mu(\mathcal{G}) + 1).$$

Then, $(\mathcal{G}, \Lambda)$ exhibits exponential SSM.

Notice that if a graph has maximum degree Δ, then $\mu(\Gamma) \leq \Delta - 1$. In addition, observe that $N_{\Gamma}(v, \ell) = N_{T_{\mathrm{SAW}}(\Gamma,v)}(\rho, \ell)$. We have the following corollary.

COROLLARY 5.9. *If (Γ, λ) is a hardcore model such that*

$$\lambda_{+} < \lambda_c(\mu(\Gamma) + 1),$$

then $(T_{\mathrm{SAW}}(\Gamma, v), \overline{\lambda})$ exhibits (exponential) SSM for every $v \in V(\Gamma)$. In particular, (Γ, λ) exhibits (exponential) SSM and for every $v \in V(\Gamma)$, $x \in X(\Gamma)$, and $U \subseteq V(\Gamma)$,

$$\mathbb{P}_{\Gamma,\lambda}([0^{v}]|[x_U]) = \mathbb{P}_{T_{\mathrm{SAW}}(\Gamma,v),\overline{\lambda}}([0^{\rho}]|[\overline{x_U}]).$$

6. *Orders*

We have already explored the main combinatorial and measure-theoretical tools that we require to establish the main results. In this section, we present some concepts of a more group-theoretical nature, namely, our ability to order a given group.

6.1. *Orderable groups.*

Let $\prec$ be a strict total order on G. We say that $\prec$ is an *invariant (right) order* if, for all $h_1, h_2, g \in G$,

$$h_1 \prec h_2 \Longrightarrow h_1 g \prec h_2 g.$$

We call the pair $(G, \prec)$ a *(linearly) ordered group*. The associated *algebraic past* of an ordered group $(G, \prec)$ is the set $\Phi_{\prec} := \{g \in G : g \prec 1_G\}$ and it is a semigroup which satisfies that

$$G = \Phi_{\prec} \sqcup \{1_G\} \sqcup \Phi_{\prec}^{-1}.$$

Notice that $h \prec g \iff hg^{-1} \in \Phi_{\prec}$, so $\Phi_{\prec}$ fully determines $\prec$ and *vice versa*.

The class of orderable groups contains all torsion-free abelian groups and it is closed under the operations of taking subgroups, and forming extensions, arbitrary direct products, and directed unions.

A group G is called *virtually orderable (respectively ordered)* if there exists an orderable (respectively ordered) subgroup $H \leq G$ of finite index $[G : H]$. Notice that if $G \curvearrowright \Gamma$ is almost transitive and free with fundamental domain U_0, then $H \curvearrowright \Gamma$ is also almost transitive and free with fundamental domain KU_0, where $K \in \mathcal{F}(G)$ is any finite set of representatives. In particular, $|\Gamma/H| = |\Gamma/G|[G : H]$. For this reason, since we are interested in almost transitive actions, there is no loss of generality if, given a virtually orderable group, we assume that it is just orderable: the free energies $f_G(\Gamma, \lambda)$ and

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

$f_H(\Gamma, \lambda)$ will be equal and the only effect of passing to a finite-index subgroup will be that the size of the fundamental domain of the action is multiplied by a constant factor (that is, the index of the subgroup $[G : H]$). This last point is relevant, since the size of the fundamental domain will play a role later when measuring the computational complexity of some problems.

Given a finitely generated group G, a generating set S, and its corresponding Cayley graph $\Gamma = \mathrm{Cay}(G, S)$, we define the volume growth function as $g_\Gamma(n) = |B_\Gamma(1_G, n)|$. We say that G has *polynomial growth* if $g_\Gamma(n) \leq p(n)$ for some polynomial p. It is well known that groups with polynomial growth are amenable and a classic result due to Gromov asserts that they are virtually nilpotent [19]. Without further detail, from Schreier's lemma, it is also well known that finite index subgroups of finitely generated groups are also finitely generated [32, Proposition 4.2] and finitely generated nilpotent groups have a torsion-free nilpotent subgroup with finite index [43, Proposition 2]. From this, and since torsion-free nilpotent groups are orderable [39, p. 37], it follows that any finitely generated group G with polynomial growth is amenable and virtually orderable. In particular, all our results that apply to amenable and virtually orderable groups will hold for groups of polynomial growth, but they will also hold in groups of super-polynomial—namely, exponential—growth. This includes solvable groups that are not virtually nilpotent [38] and, more concretely, cases like the Baumslag–Solitar groups $\mathrm{BS}(1, n)$ that can also be ordered. However, not every amenable group is virtually orderable; for example, the direct sum of countably many non-trivial finite groups always results in a countable group that is amenable but not virtually orderable. To address these cases, we introduce a randomized generalization of invariant orders.

6.2. *Random orders.* Consider now the set of relations $\{0, 1\}^{G\times G}$ endowed with the product topology and the closed subset $\mathrm{Ord}(G)$ of strict total orders $\prec$ on G. We will consider the action $G \curvearrowright \mathrm{Ord}(G)$ given by

$$h_1(g \cdot \prec)h_2 \iff (h_1 g) \prec (h_2 g)$$

for $h_1, h_2, g \in G$ and $\prec \in \mathrm{Ord}(G)$. An *invariant random order* on G is a G-invariant Borel probability measure on $\mathrm{Ord}(G)$. Notice that a fixed point for the action $G \curvearrowright \mathrm{Ord}(G)$ corresponds to a (deterministic) invariant order on G. The space of invariant random orders will be denoted by $\mathcal{M}_G(\mathrm{Ord}(G))$.

Invariant random orders were introduced in [1] to answer problems about predictability in topological dynamics through what they called the *Kieffer–Pinsker formula* for the Kolmogorov–Sinai entropy of a group action.

Now, as in the deterministic case, we can also define a notion of past for the group. An *invariant random past* on G is a random function $\tilde{\Phi} : G \to \{0, 1\}^G$ or, equivalently, a Borel probability measure on $(\{0, 1\}^G)^G$ that satisfies, for almost every instance of $\tilde{\Phi}$, the following properties:

(1) for all $g \in G$, the condition $g \notin \tilde{\Phi}(g)$ holds;
(2) for all $g, h \in G$, if $g \in \tilde{\Phi}(h)$, then $\tilde{\Phi}(g) \subseteq \tilde{\Phi}(h)$;
(3) if $g \neq h$, then either $g \in \tilde{\Phi}(h)$ or $h \in \tilde{\Phi}(g)$; and

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

(4) for all $g \in G$, the random subsets $\tilde{\Phi}(g)$ and $\tilde{\Phi}(1_G)g$ have the same distribution.

Notice that if $\prec$ is an invariant random order, then the random function $g \mapsto \{h \in G : h \prec g\}$ defines an invariant random past.

In contrast to deterministic invariant orders, every countable group G admits at least one invariant random total order. Namely, consider the random process $(\chi_g)_{g\in G}$ of independent random variables such that each χ_g has uniform distribution on $[0, 1]$. This process is invariant and each realization of it induces an order on G almost surely. We call such order the *uniform random order*.

7. *Counting*

From now on, given $(\Gamma, \lambda) \in \mathcal{H}_G$, we always assume that there is some (or any) fixed fundamental domain U_0 for $G \curvearrowright \Gamma$ and we introduce the auxiliary function $\phi_\lambda : X(\Gamma) \to \mathbb{R}$ given by

$$\phi_\lambda(x) = \frac{1}{|\Gamma/G|} \sum_{v\in U_0} x(v) \log \lambda(v).$$

7.1. *A pointwise Shannon–McMillan–Breiman-type theorem.* The next theorem establishes a pointwise Shannon–McMillan–Breiman-type theorem for Gibbs measures (related results can be found in [9, 20]). To prove it, we use the pointwise ergodic theorem [28], which requires tyhe Følner sequence $\{F_n\}_n$ to be *tempered*, a technical condition that is satisfied by every Følner sequence up to a subsequence and that we will assume without further detail.

THEOREM 7.1. *Let G be a countable amenable group. For every $(\Gamma, \lambda) \in \mathcal{H}_G$ and every $\mathbb{P} \in \mathcal{M}_{\mathrm{Gibbs}}(\Gamma, \lambda)$,*

$$\lim_n \left[-\frac{1}{|F_n U_0|} \log \mathbb{P}([x_{F_n U_0}]) \right] = -\int \phi_\lambda \, d\mathbb{Q} + f_G(\Gamma, \lambda) \quad \mathbb{Q}(x)\text{-almost surely in } x$$

for any tempered Følner sequence $\{F_n\}_n$ and any $\mathbb{Q} \in \mathcal{M}_G^{\mathrm{erg}}(X(\Gamma))$.

Proof. Consider the sets $U_n = F_n U_0$ and $M_n = U_n \cup \partial U_n$. Notice that, by amenability, $\lim_{n\to\infty} |M_n|/|U_n| = 1$. Indeed, define $K = \{g \in G : \mathrm{dist}_\Gamma(U_0, gU_0) \le 1\}$. Then, $1_G \in K$ and $U_0 \cup \partial U_0 \subseteq KU_0$. Since Γ is locally finite and the action is free, K is finite. In addition, $U_n \cup \partial U_n \subseteq F_n K U_0$. Therefore, by amenability,

$$1 \le \lim_n \frac{|U_n \cup \partial U_n|}{|U_n|} \le \lim_n \frac{|F_n K U_0|}{|F_n U_0|} = 1,$$

so $\lim_n |\partial U_n|/|U_n| = 0$. Fix independent sets $x \in X(\Gamma[U_n])$, $z_1, z_2 \in X(\Gamma[M_n \setminus U_n])$, and $y \in X(\Gamma)$ such that $xz_i y \in X(\Gamma)$ for $i = 1, 2$. Then,

$$\begin{aligned}\frac{\pi^y_{M_n}(xz_1)}{\pi^y_{M_n}(xz_2)} &= \frac{\mathrm{w}^y_\lambda(xz_1, M_n) Z^y_\Gamma(M_n, \lambda)^{-1}}{\mathrm{w}^y_\lambda(xz_2, M_n) Z^y_\Gamma(M_n, \lambda)^{-1}} \\ &= \frac{\mathrm{w}_\lambda(xz_1, M_n) 1_{[y_{M_n^c}]}(xz_1)}{\mathrm{w}_\lambda(xz_2, M_n) 1_{[y_{M_n^c}]}(xz_2)}\end{aligned}$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

$$= \frac{\prod_{v\in M_n\setminus U_n} \lambda(v)^{z_1(v)}}{\prod_{v\in M_n\setminus U_n} \lambda(v)^{z_2(v)}}$$
$$\leq \prod_{v\in M_n\setminus U_n} \lambda_+^{z_1(v)}\lambda_-^{-z_2(v)}.$$

Taking $z_2 = 0^\Gamma$ and adding over all possible z_1, we obtain that

$$1 \leq \frac{\pi^y_{M_n}([x_{U_n}])}{\pi^y_{M_n}(x_{U_n}0^{M_n\setminus U_n})} = \sum_{\substack{z\in\{0,1\}^{M_n\setminus U_n}:\\ x_{U_n}z\in X(\Gamma[M_n])}} \frac{\pi^y_{M_n}(x_{U_n}z)}{\pi^y_{M_n}(x_{U_n}0^{M_n\setminus U_n})} \leq (2\max\{1,\lambda_+\})^{|M_n\setminus U_n|}.$$

However, we have that

$$\frac{\pi^y_{M_n}([x_{U_n}0^{M_n\setminus U_n}])}{\pi^{0^\Gamma}_{M_n}([x_{U_n}0^{M_n\setminus U_n}])} = \frac{\mathrm{w}^y_\lambda(x_{U_n}0^{M_n\setminus U_n}, M_n)Z^y_\Gamma(M_n,\lambda)^{-1}}{\mathrm{w}^{0^\Gamma}_\lambda(x_{U_n}0^{M_n\setminus U_n}, M_n)Z^{0^\Gamma}_\Gamma(M_n,\lambda)^{-1}} = \frac{Z^y_\Gamma(M_n,\lambda)}{Z_\Gamma(M_n,\lambda)},$$

since

$$\mathrm{w}^y_\lambda(x_{U_n}0^{M_n\setminus U_n}, M_n) = \mathrm{w}^{0^\Gamma}_\lambda(x_{U_n}0^{M_n\setminus U_n}, M_n) = \mathrm{w}_\lambda(x_{U_n}, U_n)$$

and $Z^{0^\Gamma}_\Gamma(M_n,\lambda) = Z_\Gamma(M_n,\lambda)$. In addition,

$$1 \leq \frac{Z^y_\Gamma(M_n,\lambda)}{Z_\Gamma(M_n,\lambda)} \leq (2\max\{1,\lambda_+\})^{|M_n\setminus U_n|},$$

since

$$\begin{aligned}
Z^y_\Gamma(M_n,\lambda) &\leq Z_\Gamma(M_n,\lambda)\\
&= \sum_{x\in X(\Gamma[U_n])} \sum_{\substack{z\in X(\Gamma[M_n\setminus U_n]):\\ xz\in X(\Gamma[M_n])}} \prod_{v\in U_n}\lambda(v)^{x(v)} \prod_{v\in M_n\setminus U_n}\lambda(v)^{z(v)}\\
&\leq \sum_{x\in X(\Gamma[U_n])} \sum_{\substack{z\in X(\Gamma[M_n\setminus U_n]):\\ xz\in X(\Gamma[M_n])}} \prod_{v\in U_n}\lambda(v)^{x(v)} \max\{1,\lambda_+\}^{|M_n\setminus U_n|}\\
&\leq 2^{|M_n\setminus U_n|}\max\{1,\lambda_+\}^{|M_n\setminus U_n|} \sum_{x\in X(\Gamma[U_n])} \prod_{v\in U_n}\lambda(v)^{x(v)}\\
&= (2\max\{1,\lambda_+\})^{|M_n\setminus U_n|} Z_\Gamma(U_n,\lambda)\\
&\leq (2\max\{1,\lambda_+\})^{|M_n\setminus U_n|} Z^y_\Gamma(M_n,\lambda).
\end{aligned}$$

Therefore,

$$\begin{aligned}
1 &\leq \frac{\pi^y_{M_n}([x_{U_n}])}{\pi^{0^\Gamma}_{M_n}([x_{U_n}0^{M_n\setminus U_n}])}\\
&= \frac{\pi^y_{M_n}([x_{U_n}])}{\pi^y_{M_n}([x_{U_n}0^{M_n\setminus U_n}])}\frac{\pi^y_{M_n}([x_{U_n}0^{M_n\setminus U_n}])}{\pi^{0^\Gamma}_{M_n}([x_{U_n}0^{M_n\setminus U_n}])}
\end{aligned}$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

$$= \frac{\pi^y_{M_n}([x_{U_n}])}{\pi^y_{M_n}([x_{U_n}0^{M_n\setminus U_n}])} \frac{Z^y_\Gamma(M_n,\lambda)}{Z_\Gamma(M_n,\lambda)}$$
$$\leq (2\max\{1,\lambda_+\})^{2|M_n\setminus U_n|}.$$

In particular, since $\mathbb{P}([x_{F_nU_0}]) = \mathbb{E}_{\mathbb{P}}(\mathbb{P}([x_{U_n}]|\mathcal{B}_{M_n^c})(y)) = \mathbb{E}_{\mathbb{P}}(\pi^y_{M_n}([x_{U_n}]))$, we have that

$$1 \leq \frac{\mathbb{P}([x_{F_nU_0}])}{\pi^{0^\Gamma}_{M_n}([x_{U_n}0^{M_n\setminus U_n}])} \leq (2\max\{1,\lambda_+\})^{2|M_n\setminus U_n|},$$

so

$$|\log \mathbb{P}([x_{F_nU_0}]) - \log \pi^{0^\Gamma}_{M_n}(x_{U_n}0^{M_n\setminus U_n})| \leq 2|M_n\setminus U_n|\log(2\max\{1,\lambda_+\}).$$

Now, since $\mathrm{w}^{0^\Gamma}_\lambda(x_{M_n}, M_n) = \mathrm{w}_\lambda(x_{M_n}, M_n)$ for every x, we have that

$$\pi^{0^\Gamma}_{M_n}(x_{U_n}0^{M_n\setminus U_n}) = \mathrm{w}_\lambda(x_{U_n}0^{M_n\setminus U_n}, M_n)Z_\Gamma(M_n,\lambda)^{-1} = \mathrm{w}_\lambda(x_{U_n}, U_n)Z_\Gamma(M_n,\lambda)^{-1}.$$

Therefore,

$$|\log \mathbb{P}([x_{F_nU_0}]) - (\log \mathrm{w}_\lambda(x_{U_n}, U_n) - \log Z_\Gamma(M_n,\lambda))| \leq 2|M_n\setminus U_n|\log(2\max\{1,\lambda_+\}),$$

so

$$\lim_{n\to\infty} \left|-\frac{\log \mathbb{P}([x_{F_nU_0}])}{|U_n|} + \left(\frac{\log \mathrm{w}_\lambda(x_{U_n}, U_n)}{|U_n|} - \frac{\log Z_\Gamma(M_n,\lambda)}{|U_n|}\right)\right|$$
$$\leq \lim_n 2\frac{|M_n\setminus U_n|}{|U_n|}\log(2\max\{1,\lambda_+\}) = 0,$$

and we conclude that

$$-\lim_n \frac{\log \mathbb{P}([x_{F_nU_0}])}{|F_nU_0|} = \lim_n \frac{-\log \mathrm{w}_\lambda(x_{U_n}, U_n)}{|U_n|} + \frac{\log Z_\Gamma(M_n,\lambda)}{|U_n|}$$
$$= -\lim_n \left(\frac{1}{|F_n|}\sum_{g\in F_n}\frac{1}{|U_0|}\sum_{v\in gU_0} x(v)\log\lambda(v)\right) + f_G(\Gamma,\lambda),$$

where we have used that $Z_\Gamma(U_n,\lambda) \leq Z_\Gamma(M_n,\lambda) \leq (2\max\{1,\lambda_+\})^{|M_n\setminus U_n|}Z_\Gamma(U_n,\lambda)$. Finally, notice that

$$\frac{1}{|F_n|}\sum_{g\in F_n}\frac{1}{|U_0|}\sum_{v\in gU_0} x(v)\log\lambda(v) = \frac{1}{|F_n|}\sum_{g\in F_n}\phi_\lambda(g\cdot x)$$

and by the pointwise ergodic theorem, we obtain that

$$\lim_n \frac{1}{|F_n|}\sum_{g\in F_n}\phi_\lambda(g\cdot x) = \int \phi_\lambda\, d\mathbb{Q} \quad \mathbb{Q}\text{-almost surely},$$

so

$$-\lim_n \frac{\log \mathbb{P}([x_{F_nU_0}])}{|F_nU_0|} = -\int \phi_\lambda\, d\mathbb{Q} + f_G(\Gamma,\lambda),$$

and we conclude the proof. □

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

We have the following lemma.

LEMMA 7.2. *Given* $\Delta \in \mathbb{N}$ *and* $(\Gamma, \lambda) \in \mathcal{H}^{\Delta}$, *there exists a constant* $C = C(\Delta, \lambda_-, \lambda_+) > 0$ *such that for every* $U \Subset V(\Gamma)$, $x \in X(\Gamma)$, *and* $v \in U$,

$$\pi_U^x([x_v]) \geq C.$$

Proof. Fix $(\Gamma, \lambda) \in \mathcal{H}_G^{\Delta}$, $U \Subset V(\Gamma)$, $x \in X(\Gamma)$, and $v \in U$. Notice that if $U^c \cap \partial\{v\} \neq \emptyset$ and $u \in U^c \cap \partial\{v\}$ is such that $x(u) = 1$, then necessarily $x_v = 0^v$ and $\pi_U^x([x_v]) = 1$. However, if $U \cap \partial\{v\} = \emptyset$, then $\pi_U^x([x_v]) = \mathbb{P}_{\Gamma[U'],\lambda}([x_v])$ for $U' = U \setminus \{u \in U : x(u') = 1 \text{ for some } u' \in U^c \cap \partial\{u\}\}$, so, by Lemma 5.3,

$$\pi_U^x([x_v]) = \mathbb{P}_{\Gamma[U'],\lambda}([x_v]) \geq \min\left\{\frac{1}{1+\lambda_+}, \frac{\lambda_-}{\lambda_- + (1+\lambda_+)^{\Delta}}\right\} \geq \frac{\min\{1, \lambda_-\}}{\lambda_- + (1+\lambda_+)^{\Delta}} > 0.$$

Therefore, by taking $C = \min\{1, \min\{1, \lambda_-\}/(\lambda_- + (1+\lambda_+)^{\Delta})\} = \min\{1, \lambda_-\}/(\lambda_- + (1+\lambda_+)^{\Delta})$, we conclude. □

7.2. *A randomized sequential cavity method.* Suppose now that $(\Gamma, \lambda) \in \mathcal{H}_G$ is such that the Gibbs (Γ, λ)-specification satisfies SSM and let $\mathbb{P}$ be the unique Gibbs measure. Considering this, we define the function $I_{\mathbb{P}} : X(\Gamma) \times (\{0,1\}^G)^G \to \mathbb{R}$ given by

$$I_{\mathbb{P}}(x, \Phi) := \limsup_{n\to\infty} I_{\mathbb{P},n}(x, \Phi),$$

where

$$I_{\mathbb{P},n}(x, \Phi) := -\frac{1}{|\Gamma/G|} \log \mathbb{P}([x_{U_0}]|[x_{(\Phi(1_G)\cap F_n)U_0}])$$

and $\{F_n\}_n$ is any exhaustion of G (not necessarily Følner). Here, $\Phi \in (\{0,1\}^G)^G$ should be understood as an instance of a random invariant past and $I_{\mathbb{P},n}$ as a rudimentary information function. In this sense, notice that $I_{\mathbb{P},n}(x, \Phi)$ only depends on the values of x restricted to $\Phi(1_G) \in \{0,1\}^G$, so Φ carries redundant information. This redundancy will play a role in the next results when taking averages.

LEMMA 7.3. *If* $(\Gamma, \lambda) \in \mathcal{H}_G$ *is such that the Gibbs* (Γ, λ)*-specification satisfies SSM and* $\mathbb{P}$ *is the unique Gibbs measure, then the function* $I_{\mathbb{P}}$ *is measurable, non-negative, defined everywhere, and bounded. Moreover,*

$$I_{\mathbb{P}}(x, \Phi) = -\frac{1}{|\Gamma/G|} \log \mathbb{P}([x_{U_0}]|[x_{\Phi(1_G)U_0}]).$$

Proof. Since $\mathbb{P}([x_{U_0}]|[x_{(\Phi(1_G)\cap F_n)U_0}])$ depends on finitely many coordinates in both $X(\Gamma)$ and $(\{0,1\}^G)^G$, $0 < \mathbb{P}([x_{U_0}]|[x_{(\Phi(1_G)\cap F_n)U_0}]) < 1$, and $-\log(\cdot)$ is a continuous function, $I_{\mathbb{P},n}$ is measurable and since $I_{\mathbb{P}}$ is a limit superior, it is measurable as well.

By SSM and Proposition 4.1,

$$\lim_{n\to\infty} \mathbb{P}([x_{U_0}]|[x_{(\Phi(1_G)\cap F_n)U_0}]) = \mathbb{P}([x_{U_0}]|[x_{\Phi(1_G)U_0}])$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

is always a well-defined limit. By Lemma 7.2 and since $(\Gamma, \lambda) \in \mathcal{H}_G^{\Delta}$ for some Δ, there exists a constant $C > 0$ such that for every $v \in V$, $U \Subset V \setminus \{v\}$, and $x \in X(\Gamma)$,

$$\pi_U^x([x_v]) \geq C.$$

Now, combined with the SSM property, this implies that for every $v \in V$, $U \subseteq V$, and $x \in X(\Gamma)$, we have that

$$\mathbb{P}([x_v]|[x_U]) \geq C.$$

Indeed, if $v \in U$, this is direct, since $\mathbb{P}([x_v]|[x_U]) = 1$. However, if $v \notin U$, by SSM,

$$\mathbb{P}([x_v]|[x_U]) = \lim_{\ell\to\infty} \mathbb{P}([x_v]|[x_{U\cup B_\Gamma(v,\ell)^c}]) = \lim_{\ell\to\infty} \pi^x_{U\cup B_\Gamma(v,\ell)^c}([x_v]) \geq \lim_{\ell\to\infty} C = C.$$

Therefore, by conditioning and iterating, we obtain that

$$1 \geq \mathbb{P}([x_{U_0}]|[x_{\Phi(1_G)U_0}]) = \prod_{i=1}^{|U_0|} \mathbb{P}([x_{v_i}]|[x_{\Phi U_0\cup\{v_1,\dots,v_{i-1}\}}]) \geq \prod_{i=1}^{|U_0|} C = C^{|\Gamma/G|},$$

so

$$0 \leq I_{\mathbb{P}}(x, \Phi) = -\log \mathbb{P}([x_{U_0}]|[x_{\Phi(1_G)U_0}]) \leq -|\Gamma/G| \log C < +\infty,$$

that is, $I_{\mathbb{P}}(x, \Phi)$ is bounded. □

Following [1], given an invariant random past $\tilde{\Phi} : G \to \{0, 1\}^G$ on G with law $\tilde{\nu} \in \mathcal{M}_G((\{0, 1\}^G)^G)$, we denote

$$\mathbb{E}_{\tilde{\Phi}} f(\tilde{\Phi}) = \int f(\tilde{\Phi}) d\tilde{\nu}(\tilde{\Phi})$$

for $f \in L^1(\tilde{\nu})$. Now, since $I_{\mathbb{P}}$ is measurable, non-negative, and bounded, we have that for every $\mathbb{Q} \in \mathcal{M}_G(X(\Gamma))$, the function $I_{\mathbb{P}}$ is integrable with respect to $\mathbb{Q} \times \tilde{\nu}$ and by Tonelli's theorem, the function $\mathbb{E}_{\tilde{\Phi}} I_{\mathbb{P}} : X(\Gamma) \to \mathbb{R}$ is integrable, defined $\mathbb{Q}$-almost everywhere, and satisfies that

$$\int \mathbb{E}_{\tilde{\Phi}} I_{\mathbb{P}}(x, \tilde{\Phi})\, d\mathbb{Q}(x) = \mathbb{E}_{\tilde{\Phi}} \int I_{\mathbb{P}}(x, \tilde{\Phi})\, d\mathbb{Q}(x).$$

We call $\mathbb{E}_{\tilde{\Phi}} I_{\mathbb{P}}$ the *random* $\mathbb{P}$*-information function* (with respect to $\tilde{\Phi}$).

LEMMA 7.4. *For any tempered Følner sequence $\{F_n\}_n$, any invariant random past $\tilde{\Phi}$, and any $\mathbb{Q} \in \mathcal{M}_G(X(\Gamma))$,*

$$\lim_n \left[-\frac{1}{|F_n U_0|} \log \mathbb{P}([x_{F_n U_0}]) - \frac{1}{|F_n|} \sum_{g\in F_n} \mathbb{E}_{\tilde{\Phi}} I_{\mathbb{P}}(g \cdot x, \tilde{\Phi}) \right] = 0$$

for $\mathbb{Q}$-almost every x.

Proof. Fix a (tempered) Følner sequence $\{F_n\}_n$. By the properties of $\tilde{\Phi}$, for $\tilde{\nu}$-almost every instance Φ, every F_n can be ordered as $g_1, \dots, g_{|F_n|}$ so that $\Phi(g_i) \cap F_n = \{g_1, \dots, g_{i-1}\}$. Then,

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

$$\mathbb{P}([x_{F_nU_0}]) = \prod_{g\in F_n} \mathbb{P}([x_{gU_0}]|[x_{(\Phi(g)\cap F_n)U_0}]).$$

Given $\ell > 0$, let $K_\ell = \{g \in G : \mathrm{dist}_\Gamma(U_0, gU_0) \le \ell\}$. If $g \in \mathrm{Int}_{K_\ell}(F_n) = \{g \in G : K_\ell g \subseteq F_n\}$, then

$$|\mathbb{P}([x_{gU_0}]|[x_{(\Phi(g)\cap F_n)U_0}]) - \mathbb{P}([x_{gU_0}]|[x_{(\Phi(g)\cap K_\ell g)U_0}])| \le |U_0|\delta(\ell)$$

and

$$|\mathbb{P}([x_{gU_0}]|[x_{(\Phi(g)\cap K_\ell g)U_0}]) - \mathbb{P}([x_{gU_0}]|[x_{\Phi(g)U_0}])| \le |U_0|\delta(\ell),$$

so

$$|\mathbb{P}([x_{gU_0}]|[x_{(\Phi(g)\cap F_n)U_0}]) - \mathbb{P}([x_{gU_0}]|[x_{\Phi(g)U_0}])| \le 2|U_0|\delta(\ell).$$

However, by Lemma 7.2 and the discussion after it, for every $g \in G$ and $U \subseteq V$,

$$\mathbb{P}([x_{gU_0}]|[x_U]) \ge C^{|\Gamma/G|} > 0.$$

Therefore, by the mean value theorem,

$$|\log \mathbb{P}([x_{gU_0}]|[x_{(\Phi(g)\cap F_n)U_0}]) - \log \mathbb{P}([x_{gU_0}]|[x_{\Phi(g)U_0}])| \le \frac{2|U_0|}{C^{|\Gamma/G|}}\delta(\ell).$$

Notice that

$$\begin{aligned}\log \mathbb{P}([x_{F_nU_0}]) = &\sum_{g\in \mathrm{Int}_{K_\ell}(F_n)} \log \mathbb{P}([x_{gU_0}]|[x_{(\Phi(g)\cap F_n)U_0}])\\ &+ \sum_{g\in F_n\setminus \mathrm{Int}_{K_\ell}(F_n)} \log \mathbb{P}([x_{gU_0}]|[x_{(\Phi(g)\cap F_n)U_0}]),\end{aligned}$$

so

$$\begin{aligned}|\log \mathbb{P}([x_{FU_0}]) - \sum_{g\in F_n} \log \mathbb{P}([x_{gU_0}]|[x_{\Phi(g)U_0}])| \le |\mathrm{Int}_{K_\ell}(F_n)|\frac{2|U_0|}{C^{|\Gamma/G|}}\delta(\ell)\\ + 2|F_n \setminus \mathrm{Int}_{K_\ell}(F_n)| \log(C^{-|\Gamma/G|}).\end{aligned}$$

Now, given $\epsilon > 0$, there exists ℓ and n_0 such that for every $n \ge n_0$,

$$\left|\frac{\log \mathbb{P}([x_{F_nU_0}])}{|F_n|} - \frac{1}{|F_n|}\sum_{g\in F_n} \log \mathbb{P}([x_{gU_0}]|[x_{\Phi(g)U_0}])\right| \le \epsilon.$$

By G-invariance of $\mathbb{P}$,

$$\begin{aligned}\mathbb{P}([x_{gU_0}]|[x_{\Phi(g)U_0}]) &= \mathbb{P}(g^{-1}\cdot[(g\cdot x)_{U_0}]|[x_{\Phi(g)U_0}])\\ &= \mathbb{P}([(g\cdot x)_{U_0}]|g\cdot[x_{\Phi(g)U_0}])\\ &= \mathbb{P}([(g\cdot x)_{U_0}]|[(g\cdot x)_{\Phi(g)g^{-1}U_0}]),\end{aligned}$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

and combining this fact with the previous estimate, we obtain that

$$\left|\frac{\log \mathbb{P}([x_{F_n U_0}])}{|F_n|} - \frac{1}{|F_n|}\sum_{g\in F_n} \log \mathbb{P}([(g\cdot x)_{U_0}]|[(g\cdot x)_{\Phi(g)g^{-1}U_0}])\right| \le \epsilon.$$

Integrating against $\tilde{\nu}$, we obtain that, for $\mathbb{Q}$-almost every x,

$$\left|\frac{\log \mathbb{P}([x_{F_n U_0}])}{|F_n|} - \frac{1}{|F_n|}\sum_{g\in F_n} \mathbb{E}_{\tilde{\Phi}} \log \mathbb{P}([(g\cdot x)_{U_0}]|[(g\cdot x)_{\tilde{\Phi}(g)g^{-1}U_0}])\right| \le \epsilon$$

and since $\tilde{\Phi}(g)$ has the same distribution as $\tilde{\Phi}(1_G)g$, we get that

$$\begin{aligned}\mathbb{E}_{\tilde{\Phi}} \log \mathbb{P}([(g\cdot x)_{U_0}]|[(g\cdot x)_{\tilde{\Phi}(g)g^{-1}U_0}]) &= \mathbb{E}_{\tilde{\Phi}} \log \mathbb{P}([(g\cdot x)_{U_0}]|[(g\cdot x)_{\tilde{\Phi}(1_G)gg^{-1}U_0}])\\ &= \mathbb{E}_{\tilde{\Phi}} \log \mathbb{P}([(g\cdot x)_{U_0}]|[(g\cdot x)_{\tilde{\Phi}(1_G)U_0}])\\ &= -|\Gamma/G|\mathbb{E}_{\tilde{\Phi}} I_{\mathbb{P}}(g\cdot x, \tilde{\Phi}),\end{aligned}$$

so

$$\left|-\frac{\log \mathbb{P}([x_{F_n U_0}])}{|F_n U_0|} - \frac{1}{|F_n|}\sum_{g\in F_n} \mathbb{E}_{\tilde{\Phi}} I_{\mathbb{P}}(g\cdot x, \tilde{\Phi})\right| \le \epsilon,$$

and since ϵ is arbitrary and the limit exists $\mathbb{Q}$-almost surely, we conclude. □

We have the following representation theorem for free energy, which can be regarded as a randomized generalization of the results in [9, 16, 36] tailored for the specific case of the hardcore model. Notice that, in contrast with [9, 16, 36], the representation holds in every amenable group and not just virtually orderable groups.

THEOREM 7.5. *Let* $(\Gamma, \lambda) \in \mathcal{H}_G$ *such that the Gibbs* (Γ, λ)*-specification satisfies SSM and* $\mathbb{P}$ *is the unique Gibbs measure. Then,*

$$f_G(\Gamma, \lambda) = \int (\mathbb{E}_{\tilde{\Phi}} I_{\mathbb{P}} + \phi_\lambda)\, d\mathbb{Q}$$

for any $\mathbb{Q} \in \mathcal{M}_G(X(\Gamma))$ *and for any invariant random past* $\tilde{\Phi}$ *of* G*. In particular,*

$$f_G(\Gamma, \lambda) = \mathbb{E}_{\tilde{\Phi}} I_{\mathbb{P}}(0^\Gamma, \tilde{\Phi}).$$

Proof. First, notice that if the statement holds for every $\mathbb{Q} \in \mathcal{M}_G^{\mathrm{erg}}(X(\Gamma))$, then it holds for every $\mathbb{Q} \in \mathcal{M}_G(X(\Gamma))$ by the ergodic decomposition theorem. Then, without loss of generality, we can assume that $\mathbb{Q}$ is G-ergodic. Considering this, by Theorem 7.1, for $\mathbb{Q}$-almost every x,

$$\lim_n \left[-\frac{1}{|F_n U_0|}\log \mathbb{P}([x_{F_n}])\right] = -\int \phi_\lambda\, d\mathbb{Q} + f_G(\Gamma, \lambda).$$

By Lemma 7.4, for $\mathbb{Q}$-almost every x,

$$\lim_n \left[-\frac{1}{|F_n U_0|}\log \mathbb{P}([x_{F_n}])\right] = \lim_n \frac{1}{|F_n|}\sum_{g\in F_n} \mathbb{E}_{\tilde{\Phi}} I_{\mathbb{P}}(g\cdot x, \tilde{\Phi}).$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

Therefore, for $\mathbb{Q}$-almost every x,

$$f_G(\Gamma, \lambda) - \int \phi_\lambda \, d\mathbb{Q} = \lim_n \frac{1}{|F_n|} \sum_{g \in F_n} \mathbb{E}_{\tilde{\Phi}} I_{\mathbb{P}}(g \cdot x, \tilde{\Phi}).$$

Integrating against $\mathbb{Q}$, we obtain that

$$\begin{aligned}
\int \lim_n \frac{1}{|F_n|} \sum_{g \in F_n} \mathbb{E}_{\tilde{\Phi}} I_{\mathbb{P}}(g \cdot x, \tilde{\Phi}) \, d\mathbb{Q} &= \lim_n \int \frac{1}{|F_n|} \sum_{g \in F_n} \mathbb{E}_{\tilde{\Phi}} I_{\mathbb{P}}(g \cdot x, \tilde{\Phi}) \, d\mathbb{Q} \\
&= \lim_n \frac{1}{|F_n|} \sum_{g \in F_n} \mathbb{E}_{\tilde{\Phi}} \int I_{\mathbb{P}}(g \cdot x, \tilde{\Phi}) \, d\mathbb{Q} \\
&= \lim_n \frac{1}{|F_n|} \sum_{g \in F_n} \mathbb{E}_{\tilde{\Phi}} \int I_{\mathbb{P}}(x, \tilde{\Phi}) \, d\mathbb{Q} \\
&= \mathbb{E}_{\tilde{\Phi}} \int I_{\mathbb{P}}(x, \tilde{\Phi}) \, d\mathbb{Q} \\
&= \int \mathbb{E}_{\tilde{\Phi}} I_{\mathbb{P}}(x, \tilde{\Phi}) \, d\mathbb{Q},
\end{aligned}$$

where the first equality is due to the dominated convergence theorem, the second and last equalities are due to Tonelli's theorem, and the third equality is due to the G-invariance of $\mathbb{Q}$. We conclude that

$$f_G(\Gamma, \lambda) = \int (\mathbb{E}_{\tilde{\Phi}} I_{\mathbb{P}} + \phi_\lambda) \, d\mathbb{Q}.$$

In particular, if $\mathbb{Q} = \delta_{0^\Gamma} \in \mathcal{M}_G(X(\Gamma))$, the Dirac measure supported on 0^Γ, then

$$f_G(\Gamma, \lambda) = \mathbb{E}_{\tilde{\Phi}} I_{\mathbb{P}}(0^\Gamma, \tilde{\Phi}) + \phi_\lambda(0^\Gamma) = \mathbb{E}_{\tilde{\Phi}} I_{\mathbb{P}}(0^\Gamma, \tilde{\Phi}). \qquad \square$$

7.3. *An arboreal representation of free energy.* The following theorem tell us that, under some special conditions, $f_G(\Gamma, \lambda)$ can be expressed using $|\Gamma/G|$ terms depending on the probability that the roots of some particular trees are unoccupied. Roughly, the trees involved are the trees of self-avoiding walks that are rooted at the vertices of a given fundamental domain and explore the graph Γ without entering to the 'past graph' induced by an invariant random past.

THEOREM 7.6. *Let $(\Gamma, \lambda) \in \mathcal{H}_G$ such that the Gibbs $(T_{\mathrm{SAW}}(\Gamma, v), \overline{\lambda})$-specification satisfies SSM for every $v \in V(\Gamma)$ and let $v_1, \ldots, v_{|\Gamma/G|}$ be an arbitrary ordering of a fundamental domain U_0. Given an invariant random past $\tilde{\Phi}$ of G, denote by $\Gamma_i(\tilde{\Phi})$ the random graph given by $\Gamma \setminus (\tilde{\Phi}(1_G)U_0 \cup \{v_1, \ldots, v_{i-1}\})$. Then,*

$$f_G(\Gamma, \lambda) = -\frac{1}{|\Gamma/G|} \sum_{i=1}^{|\Gamma/G|} \mathbb{E}_{\tilde{\Phi}} \log \mathbb{P}_{T_{\mathrm{SAW}}(\Gamma_i(\tilde{\Phi}), v_i), \overline{\lambda}}([0^{\rho_i}]),$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

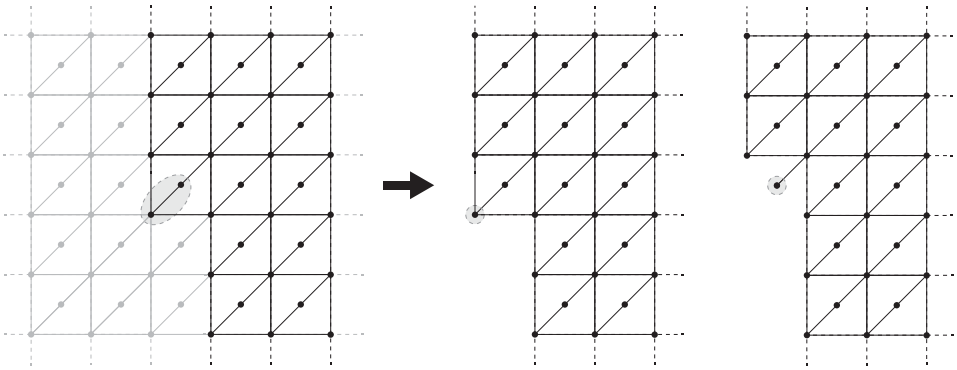

FIGURE 4. The graph $\Gamma \setminus \Phi_{\prec} U_0$ and the corresponding graphs $\Gamma_i(\Phi_{\prec})$ for $G = \mathbb{Z}^2$ and the lexicographic order $\prec$.

where ρ_i denotes the root of $T_{\mathrm{SAW}(\Gamma_i(\Phi),v_i)}$. In particular, if $\prec$ is a deterministic invariant order of G,

$$f_G(\Gamma, \lambda) = -\frac{1}{|\Gamma/G|} \sum_{i=1}^{|\Gamma/G|} \log \mathbb{P}_{T_{\mathrm{SAW}(\Gamma_i(\Phi_{\prec}),v_i),\bar{\lambda}}}([0^{\rho_i}]).$$

Proof. By Theorem 7.5, we know that

$$f_G(\Gamma, \lambda) = \mathbb{E}_{\tilde{\Phi}} I_{\mathbb{P}_{\Gamma,\lambda}}(0^{\Gamma}, \tilde{\Phi}) = -\frac{1}{|\Gamma/G|} \mathbb{E}_{\tilde{\Phi}} \log \mathbb{P}_{\Gamma,\lambda}([0^{U_0}] | [0^{\tilde{\Phi}(1_G)U_0}]).$$

By iterating conditional probabilities, linearity of expectation, and Proposition 4.1 (see Figure 4),

$$\begin{aligned}
\mathbb{E}_{\tilde{\Phi}} \log \mathbb{P}_{\Gamma,\lambda}([0^{U_0}] | [0^{\tilde{\Phi}(1_G)U_0}]) &= \sum_{i=1}^{|\Gamma/G|} \mathbb{E}_{\tilde{\Phi}} \log \mathbb{P}_{\Gamma,\lambda}([0^{v_i}] | [0^{\tilde{\Phi}(1_G)U_0 \cup \{v_1,\ldots,v_{i-1}\}}]) \\
&= \sum_{i=1}^{|\Gamma/G|} \mathbb{E}_{\tilde{\Phi}} \log \mathbb{P}_{\Gamma \setminus (\tilde{\Phi}(1_G)U_0 \cup \{v_1,\ldots,v_{i-1}\}),\lambda}([0^{v_i}]) \\
&= \sum_{i=1}^{|\Gamma/G|} \mathbb{E}_{\tilde{\Phi}} \log \mathbb{P}_{\Gamma_i(\tilde{\Phi}),\lambda}([0^{v_i}]).
\end{aligned}$$

Finally, by Theorem 5.6, we obtain

$$\sum_{i=1}^{|\Gamma/G|} \mathbb{E}_{\tilde{\Phi}} \log \mathbb{P}_{\Gamma_i(\tilde{\Phi}),\lambda}([0^{v_i}]) = \sum_{i=1}^{|\Gamma/G|} \mathbb{E}_{\tilde{\Phi}} \log \mathbb{P}_{T_{\mathrm{SAW}(\Gamma_i(\tilde{\Phi}),v_i),\bar{\lambda}}}([0^{\rho_i}]).$$

In particular, if $\prec$ is a deterministic invariant order on G (see Figure 5), we have that

$$f_G(\Gamma, \lambda) = -\frac{1}{|\Gamma/G|} \sum_{i=1}^{|\Gamma/G|} \log \mathbb{P}_{T_{\mathrm{SAW}(\Gamma_i(\Phi_{\prec}),v_i),\bar{\lambda}}}([0^{\rho_i}]). \qquad \square$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

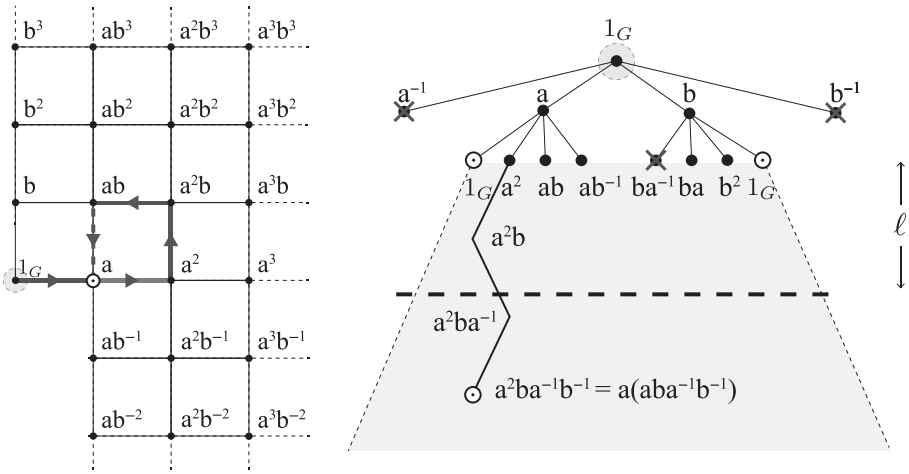

FIGURE 5. A representation of $T_{\mathrm{SAW}(\Gamma_i(\Phi_\prec),v_i)}$ and a logarithmic depth truncation for $G = \mathbb{Z}^2$ and $\Gamma = \mathrm{Cay}(\mathbb{Z}^2, \{\pm(1,0), \pm(0,1)\})$.

8. *A computational phase transition in the thermodynamic limit*

Given an amenable countable group G, we are interested in having an algorithm to efficiently approximate $f_G(\Gamma, \lambda)$ in some uniform way over $\mathcal{H}_G$.

Let $\mathbb{M} \subseteq \mathcal{H}_G$ be a family of hardcore models. We will say that $\mathbb{M}$ admits an *additive fully polynomial-time randomized approximation scheme (additive FPRAS)* for $f_G(\Gamma, \lambda)$ if there is a probabilistic algorithm such that given an input $(\Gamma, \lambda) \in \mathbb{M}$ and $\epsilon > 0$, outputs $\hat{f}$ with

$$\mathbb{P}[|f_G(\Gamma, \lambda) - \hat{f}| \leq \epsilon] \geq \tfrac{3}{4}$$

in polynomial time in $|\langle \Gamma, \lambda \rangle|$ and ϵ^{-1}, where $|\langle \Gamma, \lambda \rangle|$ denotes the length of any reasonable representation $\langle \Gamma, \lambda \rangle$ of (Γ, λ). Similarly, we will say that $\mathbb{M}$ admits an *additive fully polynomial-time approximation scheme (additive FPTAS)* for $f_G(\Gamma, \lambda)$ if there is a deterministic additive FPRAS with null failure probability.

An additive FPRAS and an additive FPTAS will be what we will regard as an efficient and uniform approximation algorithm for $f_G(\Gamma, \lambda)$, random and deterministic, respectively.

Remark 8.1. The constant $\frac{3}{4}$ in the definition of additive FPRAS is the standard choice for minimum success probability but it can be replaced by any constant bounded away from $\frac{1}{2}$ without any sensible change in the definition. To not have to deal with numerical details about the representation of λ, we will always implicitly assume that the values taken by λ have a bounded number of digits uniformly on $\mathbb{M}$.

8.1. *Weitz's algorithm and a computational phase transition.*

Observe that if G is the trivial group $\{1\}$, then $\mathcal{H}_{\{1\}}$ is exactly the family of finite hardcore models. In this case, we have that

$$f_G(\Gamma, \lambda) = \frac{\log Z_\Gamma(\lambda)}{|V(\Gamma)|},$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

and we can translate an approximation of $Z_\Gamma(\lambda)$ into an approximation of $f_G(\Gamma, \lambda)$ and *vice versa*.

In this finite context, it is common to consider a *fully polynomial-time approximation scheme*. Given a family $\mathbb{M} \subseteq \mathcal{H}_{\{1\}}$ of finite hardcore models, we will say that $\mathbb{M}$ admits an *FPTAS* for $Z_\Gamma(\lambda)$ if there is an algorithm such that, given an input $(\Gamma, \lambda) \in \mathbb{M}$ and $\epsilon > 0$, outputs $\hat{Z}$ with

$$Z_\Gamma(\lambda)e^{-\epsilon} \leq \hat{Z} \leq Z_\Gamma(\lambda)e^{\epsilon}$$

in polynomial time in $|V(\Gamma)|$ and ϵ^{-1}. An FPTAS is regarded as an efficient and uniform approximation algorithm for $Z_\Gamma(\lambda)$.

If we take logarithms and divide by $|V(\Gamma)|$ in the previous equation, we obtain that

$$|f_G(\Gamma, \lambda) - \hat{f}| = \left| \frac{\log Z_\Gamma}{|V(\Gamma)|} - \frac{\log \hat{Z}}{|V(\Gamma)|} \right| \leq \frac{\epsilon}{|V(\Gamma)|},$$

where $\hat{f} = \log \hat{Z}/|V(\Gamma)|$, so an FTPAS for $Z_\Gamma(\lambda)$ is equivalent to an additive FPTAS for $f_G(\Gamma, \lambda)$, since a polynomial in $|V(\Gamma)|$ and ϵ^{-1} is also a polynomial in $|\langle \Gamma, \lambda \rangle|$ and $|V(\Gamma)|\epsilon^{-1}$ and *vice versa*. The same correspondence holds between the natural randomized counterparts (FPRAS and additive FPRAS, respectively).

We will fix a positive integer Δ and $\lambda_0 > 0$. Given these parameters, we aim to develop a fully polynomial-time additive approximation on $\mathbb{M} = \mathcal{H}_G^\Delta(\lambda_0)$.

The main theorem in [55] was the development of an FPTAS for $Z_\Gamma(\lambda_0)$ on $\mathcal{H}_{\{1\}}^\Delta(\lambda_0)$ for $\lambda_0 < \lambda_c(\Delta)$. It is not difficult to see that the theorem extends to non-constant activity functions λ. Then, and also translated into the language of free energy, we have the following result.

THEOREM 8.2. [55] *For every $\Delta \in \mathbb{N}$ and $0 < \lambda_0 < \lambda_c(\Delta)$, there exists an FPTAS (respectively additive FPTAS) on $\mathcal{H}_{\{1\}}^\Delta(\lambda_0)$ for $Z_\Gamma(\lambda)$ (respectively for $f_G(\Gamma, \lambda)$).*

This theorem was subsequently refined in [46] by considering connective constants instead of maximum degree Δ. A very interesting fact is that when classifying graphs according to their maximum degree, then Theorem 8.2 is in some sense optimal due to the following theorem.

THEOREM 8.3. [48, 49] *For every $\Delta \geq 3$ and $\lambda_0 > \lambda_c(\Delta)$, there does not exist an FPRAS (respectively an additive FPRAS) on $\mathcal{H}_{\{1\}}^\Delta(\lambda_0)$ for $Z_\Gamma(\lambda_0)$ (respectively for $f_G(\Gamma, \lambda_0)$), unless* NP = RP.

Remark 8.4. Notice that the lack of existence of an FPRAS (respectively additive FPRAS) directly implies the lack of existence of an FPTAS (respectively additive FPTAS).

The combination of Theorems 8.2 and 8.3 is what is regarded as a *computational phase transition*. We aim to extend these theorems to the infinite setting. A theoretical advantage about considering $f_G(\Gamma, \lambda)$ instead of $Z_\Gamma(\lambda)$ is that the free energy still makes sense in infinite graphs and, at the same time, recovers the theory for $Z_\Gamma(\lambda)$ in the finite case.

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

8.2. *An extension of Weitz's algorithm to the infinite setting.* For algorithmic purposes, in this section, we only consider finitely generated groups G with a fixed symmetric set of generators S. In this case, if $(\Gamma, \lambda) \in \mathcal{H}_G^\Delta$, then it suffices to know $\Gamma[U_0]$ for some fundamental domain U_0 and $L = \{(v_1, s, v_2) \in U_0 \times S \times U_0 : (v_1, sv_2) \in E(\Gamma)\}$ to fully reconstruct the graph Γ. In particular, the size of the necessary information to reconstruct Γ is bounded by a polynomial in Δ and $|\Gamma/G|$. In addition, given a G-invariant activity function $\lambda : V(\Gamma) \to \mathbb{Q}_{>0}$ (that is, that only takes positive rational values), we only need to know $\lambda|_{U_0}$ to recover λ, that is, just $|\Gamma/G|$ many rational numbers. Therefore, in this context, the length $|\langle\Gamma, \lambda\rangle|$ of the representation $\langle\Gamma, \lambda\rangle$ of a hardcore model (Γ, λ) will be polynomial in Δ and $|\Gamma/G|$.

First, we are interested in being able to generate in an effective way balls of arbitrary radius in $\mathrm{Cay}(G, S)$. Given an input word $w \in S^*$, we will assume that we can decide whether $e(w) = 1_G$ or not in time $\exp(O(|w|))$, where $|w|$ denotes the length of w, $e : S^* \to G$ is the usual evaluation map, and the O-notation regards $|S|$ and Δ as constants. In other words, we will assume that *the word problem of G can be solved in exponential time*. By problems that can be solved in exponential time, we mean the set of decision problems that can be solved by a deterministic Turing machine in time $\exp(O(n))$. This complexity class is sometimes known as E; it contains P and it is strictly contained in EXP.

Now, if the word problem can be solved in exponential time, then $\mathrm{Cay}(G, S)$ is constructible in time $\exp(O(\ell))$ as well (see [37, Theorem 5.10]); this is to say, given $\ell > 0$, we can generate $B_{\mathrm{Cay}(G,S)}(1_G, \ell)$ in time $\exp(O(\ell))$. Having that, it is possible to construct $\Gamma[B_{\mathrm{Cay}(G,S)}(1_G, \ell)U_0]$ in time $O(\mathrm{poly}(|\Gamma/G|)\exp(O(\ell)))$ by identifying each $g \in B_{\mathrm{Cay}(G,S)}(1_G, \ell)$ with a copy $\Gamma[gU_0]$ of $\Gamma[U_0]$ and by connecting it to other adjacent copies according to L.

In the ordered case, we will also consider the situation where *the algebraic past of* $(G, \prec)$ *can be decided in exponential time*, that is, given an input word $w \in S^*$, we will consider that we can decide whether $e(w) \in \Phi_\prec$ or not in time $\exp(O(|w|))$. Notice that this implies that the word problem can be solved in time $\exp(O(|w|))$, since $e(w) = 1_G$ if and only if $e(w) \notin \Phi_\prec$ and $e(w^{-1}) \notin \Phi_\prec$. In particular, and since $|B_{\mathrm{Cay}(G,S)}(1_G, \ell)| \leq |S|^\ell$, by identifying and removing all the copies $\Gamma[gU_0]$ with $g \in \Phi_\prec$ in $\Gamma[B_{\mathrm{Cay}(G,S)}(1_G, \ell)U_0]$, we can construct $\Gamma[(B_{\mathrm{Cay}(G,S)}(1_G, \ell) \setminus \Phi_\prec)U_0]$ in time $O(\mathrm{poly}(|\Gamma/G|)\exp(O(\ell)))$. Recall that we are not losing generality if we assume G to be ordered instead of just virtually ordered.

Considering all this, we have the following key theorem.

THEOREM 8.5. *Let G be a finitely generated amenable group such that its word problem can be solved in exponential time. Then, for every* $\Delta \in \mathbb{N}$ *and* $0 < \lambda_0 < \lambda_c(\Delta)$, *there exists an additive FPRAS on* $\mathcal{H}_G^\Delta(\lambda_0)$ *for* $f_G(\Gamma, \lambda)$. *If, in addition, G is orderable and has an algebraic past that can be decided in exponential time, then the algorithm can be chosen to be deterministic, that is, an additive FPTAS.*

Proof. Pick (Γ, λ) as in the statement and enumerate as $U_0 = \{v_1, \ldots, v_n\}$ the fundamental domain of $G \curvearrowright \Gamma$. Denote $n = |\Gamma/G|$. Then, by Theorem 7.5,

$$f_G(\Gamma, \lambda) = \sum_{i=1}^{n} -\mathbb{E}_{\tilde{\Phi}} \log \mathbb{P}_{\Gamma_i(\tilde{\Phi}),\lambda}([0^{v_i}]),$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

where $\Gamma_i(\tilde{\Phi}) = \Gamma \setminus (\tilde{\Phi}(1_G)U_0 \cup \{v_1, \ldots, v_{i-1}\})$ and $\tilde{\Phi}$ is any invariant random past of G. Let $f_i = -\mathbb{E}_{\tilde{\Phi}} \log \mathbb{P}_{\Gamma_i(\tilde{\Phi}),\lambda}([0^{v_i}])$. Given $\epsilon > 0$, our goal is to generate numbers $\hat{f}_i$ such that

$$\hat{f}_i - \frac{\epsilon}{n} \leq f_i \leq \hat{f}_i + \frac{\epsilon}{n}$$

for every $i = 1, \ldots, n$. If we manage to obtain these approximations, we have that

$$\sum_{i=1}^{n} \hat{f}_i - \epsilon \leq \sum_{i=1}^{n} f_i = f_G(\Gamma, \lambda) \leq \sum_{i=1}^{n} \hat{f}_i + \epsilon,$$

so $\hat{f} := \sum_{i=1}^{n} \hat{f}_i$ will be the required approximation. By SSM in Γ, we have that, for every $\ell > 0$,

$$|\log \mathbb{P}_{\Gamma_i(\tilde{\Phi}),\lambda}([0^{v_i}]) - \log \mathbb{P}_{\Gamma_i(\tilde{\Phi})\cap B_\Gamma(v_i,\ell),\lambda}([0^{v_i}])| \leq C \exp(-\alpha\ell),$$

and, again by SSM but now on the tree of self-avoiding walks,

$$|\log \mathbb{P}_{\Gamma_i(\tilde{\Phi})\cap B_\Gamma(v_i,\ell),\lambda}([0^{v_i}]) - \log \mathbb{P}_{T_{\mathrm{SAW}}(\Gamma_i(\tilde{\Phi})\cap B_\Gamma(v_i,\ell),v_i)\cap B(\rho_i,\ell),\bar{\lambda}}([0^{\rho_i}])| \leq C \exp(-\alpha\ell),$$

where C is a constant that depends on Δ and λ_0. Notice that $T_{\mathrm{SAW}}(\Gamma_i(\tilde{\Phi}) \cap B_\Gamma(v_i, \ell), v_i) \cap B(\rho_i, \ell) = T_{\mathrm{SAW}}(\Gamma_i(\tilde{\Phi}), v_i) \cap B(\rho_i, \ell)$, so

$$|f_i - \mathbb{E}_{\tilde{\Phi}} \log \mathbb{P}_{T_{\mathrm{SAW}}(\Gamma_i(\tilde{\Phi}),v_i)\cap B(\rho_i,\ell),\bar{\lambda}}([0^{\rho_i}])| \leq 2C \exp(-\alpha\ell).$$

To conclude, it suffices to define $\hat{f}_i$ as $\mathbb{E}_{\tilde{\Phi}} \log \mathbb{P}_{T_{\mathrm{SAW}}(\Gamma_i(\tilde{\Phi}),v_i)\cap B(\rho_i,\ell),\bar{\lambda}}([0^{\rho_i}])$ with ℓ so that $2C \exp(-\alpha\ell) \leq \frac{\epsilon}{n}$ and show that each approximation $\hat{f}_i$ can be efficiently computed. Notice that we can pick ℓ to be $\ell^* = \lceil 1/\alpha \log(2Cn\epsilon^{-1})\rceil$.

Let us first assume that $\tilde{\Phi}$ is a (deterministic) algebraic past $\Phi_\prec$ of G that can be decided in exponential time. The general probabilistic case will be a slight variation of this case. To compute $\hat{f}_i$, first generate the ball $\Gamma[B_{\mathrm{Cay}(G,S)}(1_G, \ell^*)U_0]$, which takes time $\mathrm{poly}(n) \exp(O(\ell^*)) = \mathrm{poly}((1 + \epsilon^{-1})n)$. Notice that $B_\Gamma(v_i, \ell) \subseteq B_{\mathrm{Cay}(G,S)}(1_G, \ell^*)U_0$. Next, remove the vertices that are at a distance greater than ℓ^* to v_i and those which belong to $\Phi_\prec U_0 \cup \{v_1, \ldots, v_{i-1}\}$. This procedure also takes time $\mathrm{poly}(n) \exp(O(\ell^*))$, since $\Phi_\prec$ can be decided in exponential time. Having this, construct the tree of self-avoiding walks $T_i := T_{\mathrm{SAW}}(\Gamma_i(\Phi_\prec), v_i) \cap B(\rho_i, \ell^*)$ (which is a subtree of the tree of self-avoiding walks of $\Gamma_i(\Phi_\prec)$). Using the recursive procedure, compute $\mathbb{P}_{T_i,\bar{\lambda}}([0^{\rho_i}])$ and then compute its logarithm. For every $i = 1, \ldots, n$,

$$|T_i| \leq \Delta^{\ell^*} \leq (C(1 + \lambda_0)(1 + \epsilon^{-1})n)^{\log \Delta/\alpha},$$

which is also a bound for the order of time required for computing $\hat{f}_i$, because T_i is a tree (see Figure 5). Finally, since we require to do this procedure n times for each $i = 1, \ldots, n$, we have that the total order of the algorithm is still $\mathrm{poly}((1 + \epsilon^{-1})n)$, that is, a polynomial in n and ϵ^{-1}, where the constants involved depend only on Δ, λ_0, and $|S|$. This gives the desired additive FPTAS in the ordered case.

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

Now, in the general not necessarily ordered case, we consider the following variation of the previous algorithm. Let $\tilde{\Phi}_{\text{unif}}$ be the uniform random order in G. Then,

$$\begin{aligned} &-\mathbb{E}_{\tilde{\Phi}_{\text{unif}}} \log \mathbb{P}_{\Gamma_i(\tilde{\Phi}_{\text{unif}})\cap B_\Gamma(v_i,\ell),\lambda}([0^{v_i}]) \\ &= -\frac{1}{2^m} \sum_{\Phi \subseteq B_{\text{Cay}(G,S)}(1_G,\ell)\setminus\{1_G\}} \log \mathbb{P}_{\Gamma_i(\Phi)\cap B_\Gamma(v_i,\ell),\lambda}([0^{v_i}]) \end{aligned}$$

with $m = |B_{\text{Cay}(G,S)}(v_i, \ell)| - 1$. Consider the random variable X with probability distribution $\mathbb{P}(X = -\log \mathbb{P}_{\Gamma_i(\Phi)\cap B_\Gamma(v_i,\ell),\lambda}([0^{v_i}])) = 1/2^m$ for $\Phi \subseteq B_{\text{Cay}(G,S)}(1_G, \ell) \setminus \{1_G\}$. Then, $\mathbb{E}[X] = -\mathbb{E}_{\tilde{\Phi}_{\text{unif}}} \log \mathbb{P}_{\Gamma_i(\tilde{\Phi}_{\text{unif}})\cap B_\Gamma(v_i,\ell),\lambda}([0^{v_i}])$. Due to Lemma 5.3, we have that

$$0 < \log\left(1 + \frac{\lambda_-}{(1+\lambda_+)^\Delta}\right) \le X \le \log(1+\lambda_+) < \infty.$$

In particular,

$$\text{Var}(X) = \mathbb{E}[X^2] - \mathbb{E}[X]^2 \le (\log(1+\lambda_+))^2 - \left(\log\left(1 + \frac{\lambda_-}{(1+\lambda_+)^\Delta}\right)\right)^2 =: V(\lambda, \Delta) < \infty.$$

Now, let $\{X_1, \ldots, X_N\}$ be a random sample of size N of the variable X and define the sample average $\bar{X}_N \equiv 1/N \sum_{j=1}^N X_j$. Notice that $\mathbb{E}[\bar{X}_N] = \mathbb{E}[X]$ and $\text{Var}(\bar{X}_N) = 1/N\text{Var}(X) \le V(\lambda, \Delta)/N$. Therefore, by Chebyshev's inequality, for any $\delta > 0$,

$$\mathbb{P}\left(|\bar{X}_N - \mathbb{E}[X]| \le \left(\frac{V(\lambda, \Delta)}{N\delta}\right)^{1/2}\right) \ge 1 - \delta.$$

We are interested in having $|\bar{X}_N - \mathbb{E}[X]| \le \epsilon/n$ with probability greater than $\frac{3}{4}$. To guarantee this, we need to take a number of samples N so that

$$\left(\frac{V(\lambda, \Delta)}{N\delta}\right)^{1/2} \le \frac{\epsilon}{n},$$

and δ such that $1 - \delta \ge \frac{3}{4}$, that is, it suffices to take $N \ge 4V(\lambda, \Delta)(n\epsilon^{-1})^2$. Notice that we need to take a number of samples polynomial in n and ϵ^{-1}, and that each sample can be computed in polynomial time, exactly as in the ordered case. This gives the desired additive FPRAS in the general case. □

Remark 8.6. Notice that Theorem 8.5 holds for groups of exponential growth, despite it involving a polynomial time algorithm. In addition, in virtue of Theorem 5.8, the families of graphs in Theorem 8.5 could be parameterized according to their connective constant instead of the maximum degree.

Next, we reduce the problem of approximating the partition function of a finite hardcore model to the problem of approximating the free energy of a hardcore model in $\mathcal{H}_G$.

PROPOSITION 8.7. *Let G be an amenable group. Then, for every $\Delta \ge 3$ and $\lambda_0 > \lambda_c(\Delta)$, there does not exist an additive FPRAS on $\mathcal{H}_G^\Delta(\lambda_0)$ for $f_G(\Gamma, \lambda_0)$, unless* NP = RP.

Proof. Suppose that we have an additive FPRAS on $\mathcal{H}_G^\Delta(\lambda_0)$ for $f_G(\Gamma, \lambda)$ for some amenable group G. Then, we claim that we would have an FPRAS on $\mathcal{H}_{\{1\}}^\Delta(\lambda_0)$ for

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

$Z_\Gamma(\lambda_0)$, contradicting Theorem 8.3. Indeed, given the input $(\Gamma, \lambda) \in \mathcal{H}^{\Delta}_{\{1\}}(\lambda_0)$, it suffices to consider the graph Γ^G made out of copies $\{\Gamma_g\}_{g\in G}$ of Γ indexed by $g \in G$, where v_g denotes the copy in Γ_g of the vertex v in Γ. Then, there is a natural action $G \curvearrowright \Gamma^G$ consisting of just translating copies of vertices, that is, $gv_h = v_{hg}$, and a fundamental domain of the action is $U_0 = V(\Gamma_{1_G})$. Therefore, since $|\Gamma^G/G| = |V(\Gamma)|$, if we could ϵ-approximate in an additive way $f_G(\Gamma, \lambda_0)$ in polynomial time in $|\Gamma^G/G|$ and ϵ^{-1}, then we would be able to ϵ-approximate in a multiplicative way $Z_\Gamma(\lambda_0)$ in polynomial time in $|V(\Gamma)|$ and ϵ^{-1}, because

$$\begin{aligned} f_G(\Gamma, \lambda_0) &= \inf_{F\in\mathcal{F}(G)} \frac{\log Z_{\Gamma^G}(FU_0, \lambda_0)}{|FU_0|} \\ &= \inf_{F\in\mathcal{F}(G)} \frac{\log Z_{\Gamma^G}(U_0, \lambda_0)^{|F|}}{|F||U_0|} \\ &= \frac{\log Z_{\Gamma^G}(U_0, \lambda_0)}{|U_0|} \\ &= \frac{\log Z_\Gamma(\lambda_0)}{|V(\Gamma)|}, \end{aligned}$$

but this contradicts Theorem 8.3. □

Considering Theorem 8.5 and Proposition 8.7, we have the following corollary.

COROLLARY 8.8. *Let G be a finitely generated amenable group such that its word problem can be solved in exponential time. Then, for every $\Delta \geq 3$ and $\lambda_0 > 0$, if $\lambda_0 < \lambda_c(\Delta)$, there exists an additive FPRAS on $\mathcal{H}^{\Delta}_G(\lambda_0)$ for $f_G(\Gamma, \lambda)$. If, in addition, G has a finite index orderable subgroup such that its algebraic past can be decided in exponential time, then the algorithm can be chosen to be deterministic, that is, there exists an additive FPTAS on $\mathcal{H}^{\Delta}_G(\lambda_0)$ for $f_G(\Gamma, \lambda)$. However, if $\lambda_0 > \lambda_c(\Delta)$, there is no additive FPRAS on $\mathcal{H}^{\Delta}_G(\lambda_0)$ for $f_G(\Gamma, \lambda)$, unless* NP = RP.

Remark 8.9. Notice that Corollary 8.8 still holds for $\Delta = 1$ and $\Delta = 2$. The first case is trivial and in the second case, there is no phase transition and the conditions for the existence of an additive FPRAS (respectively additive FPTAS) hold for every λ_0.

To ask that the word problem can be solved in exponential time seems to be a natural requirement for having an efficient algorithm for approximating $f_G(\Gamma, \lambda)$ and, fortunately, there are several classes of finitely generated groups which satisfy this condition.

Example 8.10. Lipton and Zalcstein [29] proved that every linear group over a field of characteristic zero has a word problem that can be solved in logarithmic space. This result was extended by Simon [45] to linear groups over a field of prime characteristic. In particular, the word problem of all finitely generated amenable linear groups—which by the Tits alternative [52] must be virtually solvable—can be solved in logarithmic space, and therefore polynomial time.

Due to a result from Mal'tsev [33], all solvable subgroups of the integer general linear group $\mathrm{GL}_d(\mathbb{Z})$ are polycyclic (that is, solvable groups in which every subgroup is finitely

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

generated) and virtually polycyclic groups coincide with the class of polycyclic-by-finite groups, which are always finitely presented, residually finite, and have many other desirable algorithmic properties (see [5]). However, Auslander [3] and Swan [51] proved that any polycyclic group is a subgroup of the integer general linear group. This shows that the class of polycyclic groups is a general and natural setup for the application of our results, since they are amenable, finitely generated, and their word problem can be solved in polynomial time. Examples of polycyclic groups include all finitely generated abelian groups and all finitely generated nilpotent groups.

To understand how to guarantee the existence of an algebraic past that can be decided in exponential time, we start by observing two basic facts: (1) the group of integers $\mathbb{Z}$ is orderable with the natural order and its algebraic past can be decided in linear time and (2) the following lemma.

LEMMA 8.11. *Let $(H_1, \prec_1)$ and $(H_2, \prec_2)$ be two ordered groups and let G be a finitely generated group which is an extension of H_2 by H_1, that is, there is a short exact sequence*

$$1 \longrightarrow H_1 \xrightarrow{\iota} G \xrightarrow{\pi} H_2 \longrightarrow 1.$$

Then, G can be ordered by considering the algebraic past $\Phi := \iota(\Phi_1) \cup \pi^{-1}(\Phi_2)$, where Φ_i denotes the algebraic past of H_i for $i = 1, 2$. In particular, if $\iota(\Phi_1)$ and $\pi^{-1}(\Phi_2)$ can be decided in exponential time, then Φ can be decided in exponential time as well.

Proof. Consider the set $\Phi := \iota(\Phi_1) \cup \pi^{-1}(\Phi_2)$. It suffices to check that Φ is a semigroup (that is, $\Phi^2 \subseteq \Phi$) and $G = \Phi \sqcup \{1_G\} \sqcup \Phi^{-1}$. Indeed, since $\iota(H_1) = \pi^{-1}(1_{H_2})$, we have that

$$\begin{aligned}
G &= \pi^{-1}(H_2) \\
&= \pi^{-1}(\Phi_2 \sqcup \{1_{H_2}\} \sqcup \Phi_2^{-1}) \\
&= \pi^{-1}(\Phi_2) \sqcup \pi^{-1}(1_{H_2}) \sqcup \pi^{-1}(\Phi_2^{-1}) \\
&= \pi^{-1}(\Phi_2) \sqcup \iota(H_1) \sqcup \pi^{-1}(\Phi_2)^{-1} \\
&= \pi^{-1}(\Phi_2) \sqcup \iota(\Phi_1 \sqcup \{1_{H_1}\} \sqcup \Phi_1^{-1}) \sqcup \pi^{-1}(\Phi_2)^{-1} \\
&= [\iota(\Phi_1) \sqcup \pi^{-1}(\Phi_2)] \sqcup \iota(1_{H_1}) \sqcup [\iota(\Phi_1)^{-1} \sqcup \pi^{-1}(\Phi_2)^{-1}] \\
&= [\iota(\Phi_1) \sqcup \pi^{-1}(\Phi_2)] \sqcup \{1_G\} \sqcup [\iota(\Phi_1) \sqcup \pi^{-1}(\Phi_2)]^{-1} \\
&= \Phi \sqcup \{1_G\} \sqcup \Phi^{-1}.
\end{aligned}$$

However, since $\iota(\Phi_1) \subseteq \iota(H_1) = \pi^{-1}(1_{H_2})$ and $\Phi_i^2 \subseteq \Phi_i$ for $i = 1, 2$,

$$\begin{aligned}
\Phi^2 &= [\iota(\Phi_1) \sqcup \pi^{-1}(\Phi_2)]^2 \\
&= \iota(\Phi_1^2) \cup \iota(\Phi_1)\pi^{-1}(\Phi_2) \cup \pi^{-1}(\Phi_2)\iota(\Phi_1) \cup \pi^{-1}(\Phi_2^2) \\
&\subseteq \iota(\Phi_1) \cup \pi^{-1}(1_{H_2})\pi^{-1}(\Phi_2) \cup \pi^{-1}(\Phi_2)\pi^{-1}(1_{H_2}) \cup \pi^{-1}(\Phi_2) \\
&= \iota(\Phi_1) \cup \pi^{-1}(\Phi_2) \cup \pi^{-1}(\Phi_2) \cup \pi^{-1}(\Phi_2) \\
&= \Phi.
\end{aligned}$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

Therefore, Φ is an algebraic past for G and it induces the invariant order $\prec$, where $g \prec h$ for $h, g \in G$ if and only if $\pi(gh^{-1}) \prec_2 1_{H_2}$ or $[\pi(gh^{-1}) = 1_{H_2}$ and $gh^{-1} \in \iota(\Phi_1)]$. In particular, if $\iota(\Phi_1)$ and $\pi^{-1}(\Phi_2)$ can be decided in exponential time, then it is direct that Φ can also be decided in exponential time. □

The previous lemma can be used as a tool for constructing algebraic pasts that can be decided in exponential time in new groups out of simpler ones. We have the following example.

Example 8.12. By the fundamental theorem of finitely generated abelian groups, every finitely generated abelian group G is isomorphic to a group of the form $\mathbb{Z}^d \oplus \mathbb{Z}_{q_1} \oplus \cdots \oplus \mathbb{Z}_{q_t}$, where $d \geq 0$ is the rank and $q_1, \ldots, q_t$ are powers of prime numbers. In particular, $[G : \mathbb{Z}^d]$ is finite, so $\mathbb{Z}^d$ is a finite index subgroup of G. However, $\mathbb{Z}^d$ is an orderable group. A canonical presentation of $\mathbb{Z}^d$ is given by

$$\langle a_1, \ldots, a_d \mid \{[a_i, a_j] : 1 \leq i, j \leq d\}\rangle,$$

where $[g, h] = ghg^{-1}h^{-1}$ is the commutator of g and h. A normal form for $\mathbb{Z}^d$ is given by $\{a_1^{i_1} \cdots a_d^{i_d} : i_1, \ldots, i_d \in \mathbb{Z}\}$ and, given any word $w \in \{a_1^{\pm 1}, \ldots, a_d^{\pm 1}\}^*$, it takes linear time to obtain its normal form. A canonical order of $\mathbb{Z}^d$ is the lexicographic order $\prec$, where we declare $a_1^{i_1} \cdots a_d^{i_d} \prec a_1^{j_1} \cdots a_d^{j_d}$ if $i_k < j_k$ for some $1 \leq k \leq d$ and $i_m = j_m$ for $m < k$, where $<$ is the usual order in $\mathbb{Z}$. It is easy to see that it can be decided in polynomial time whether $a_1^{i_1} \cdots a_d^{i_d} \prec a_1^0 \cdots a_d^0$ or not. An alternative way to see this is through Lemma 8.11, by observing that $\mathbb{Z}^d$ is an extension of $\mathbb{Z}^{d-1}$ by $\mathbb{Z}$ and proceed inductively until reaching the base case $\mathbb{Z}^1 = \mathbb{Z}$.

Another illustrative example is the discrete Heisenberg group:

$$H_3(\mathbb{Z}) := \left\{ \begin{pmatrix} 1 & i & k \\ 0 & 1 & j \\ 0 & 0 & 1 \end{pmatrix} : i, j, k \in \mathbb{Z} \right\} \subseteq \mathrm{SL}_3(\mathbb{Z}).$$

The group $H_3(\mathbb{Z})$ is a non-abelian nilpotent (and therefore amenable with polynomial growth) finitely generated group. A presentation of $H_3(\mathbb{Z})$ is given by

$$\langle a, b, c \mid [a, c], [b, c], [a, b]c^{-1}\rangle,$$

where we identify a, b, and c with

$$\begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}, \quad \text{and} \quad \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix},$$

respectively. A normal form for $H_3(\mathbb{Z})$ is given by $\{b^j c^k a^i : i, j, k \in \mathbb{Z}\}$, where

$$b^j c^k a^i = \begin{pmatrix} 1 & i & k \\ 0 & 1 & j \\ 0 & 0 & 1 \end{pmatrix}.$$

It is not difficult to check that given a word $w \in \{a^{\pm 1}, b^{\pm 1}, c^{\pm 1}\}^*$ in its normal form and a generator $s \in \{a^{\pm 1}, b^{\pm 1}, c^{\pm 1}\}$, it takes linear time to write sw in its normal form.

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

Observe that it is enough to measure how much time it takes this particular operation and then proceed inductively. Now, it is known that $H_3(\mathbb{Z})$ is an extension of $\mathbb{Z}^2$ by $\mathbb{Z}$, that is,

$$1 \longrightarrow \mathbb{Z} \xrightarrow{\iota} H_3(\mathbb{Z}) \xrightarrow{\pi} \mathbb{Z}^2 \longrightarrow 1,$$

with $\iota : \mathbb{Z} \longrightarrow H_3(\mathbb{Z})$ and $\pi : H_3(\mathbb{Z}) \longrightarrow \mathbb{Z}^2$ given by $\iota(k) = c^k$ and $\pi(b^j c^z a^i) = (i, j)$, respectively. Considering Lemma 8.11 and that $\mathbb{Z}$ and $\mathbb{Z}^2$ have algebraic pasts $\Phi_{\mathbb{Z}}$ and $\Phi_{\mathbb{Z}^2}$, respectively, such that $\iota(\Phi_{\mathbb{Z}})$ and $\pi^{-1}(\Phi_{\mathbb{Z}^2})$ can be decided in exponential time, we conclude that $H_3(\mathbb{Z})$ also has an algebraic past that can be decided in exponential time. More concretely, this algebraic past $\Phi_{H_3(\mathbb{Z})}$ is defined by declaring that $b^j c^k a^i \in \Phi_{H_3(\mathbb{Z})}$ if and only if (1) $i < 0$ or (2) $i = 0$ and $j < 0$ or (3) $i = j = 0$ and $k < 0$, which takes linear time to decide.

Finally, one other example is the case of the Baumslag–Solitar group $BS(1, 2)$. A presentation of $BS(1, 2)$ is given by $\langle a, b \mid bab^{-1}a^{-2}\rangle$, where we identify a and b with the linear functions $x \mapsto 2x$ and $x \mapsto x + 1$ in Homeo($\mathbb{R}$), respectively. The group $BS(1, 2)$ is a non-nilpotent solvable (and therefore amenable with exponential growth) finitely generated group. A normal form for $BS(1, 2)$ is given by $\{a^{-j}b^{2k+1}a^{j+i} : i, j, k \in \mathbb{Z}\} \cup \{a^i : i \in \mathbb{Z}\}$ and it can also be checked that given a word $w \in \{a^{\pm 1}, b^{\pm 1}\}^*$ in its normal form and a generator $s \in \{a^{\pm 1}, b^{\pm 1}\}$, it takes polynomial time to write sw in its normal form. It is known that $BS(1, 2)$ is an extension of $\mathbb{Z}$ by $\mathbb{Z}[\frac{1}{2}]$, the group of dyadic rationals, that is,

$$1 \longrightarrow \mathbb{Z}\left[\tfrac{1}{2}\right] \xrightarrow{\iota} BS(1, 2) \xrightarrow{\pi} \mathbb{Z} \longrightarrow 1,$$

with $\iota : \mathbb{Z}[\frac{1}{2}] \longrightarrow BS(1, 2)$ and $\pi : BS(1, 2) \longrightarrow \mathbb{Z}$ given by $\iota((2k + 1)/2^j) = a^{-j}b^{2k+1}a^j$, and $\pi(a^{-j}b^{2k+1}a^{j+i}) = i$ and $\pi(a^i) = i$, respectively. Then, due to Lemma 8.11 and the fact that $\mathbb{Z}[\frac{1}{2}]$ and $\mathbb{Z}$ have algebraic pasts $\Phi_{\mathbb{Z}[1/2]}$ and $\Phi_{\mathbb{Z}}$, respectively, such that $\iota(\Phi_{\mathbb{Z}[1/2]})$ and $\pi^{-1}(\Phi_{\mathbb{Z}})$ can be decided in exponential time, we conclude that $BS(1, 2)$ also has an algebraic past $\Phi_{BS(1,2)}$ that can be decided in exponential time. More concretely, this algebraic past is defined by declaring that $a^{-j}b^{2k+1}a^{j+i} \in \Phi_{BS(1,2)}$ if and only if (1) $i < 0$ or (2) $i = 0$ and $k < 0$, which takes linear time to decide. This construction can be easily generalized to the group $BS(1, n)$ given by the presentation $\langle a, b \mid bab^{-1}a^{-n}\rangle$.

The previous facts about word problems and algebraic pasts give us general conditions for efficiently generating $\Gamma[B_{\mathrm{Cay}(G,S)}(1_G, \ell)U_0]$ and $\Gamma[(B_{\mathrm{Cay}(G,S)}(1_G, \ell) \setminus \Phi_{\prec})U_0]$, respectively.

9. *Reductions*

In this section, we provide a set of reductions which exploit the combinatorial properties of independent sets and relate the results already obtained for hardcore models with other systems.

9.1. *G-subshifts and conjugacies.* Given a countable group G and a finite set Σ endowed with the discrete topology, the *full shift* is the set Σ^G of maps $\omega : G \to \Sigma$ endowed with the product topology. We define the *G-shift* as the group action $G \times \Sigma^G \to \Sigma^G$ given by

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

$(g, \omega) \mapsto g \cdot \omega$, where $(g \cdot \omega)(h) = \omega(hg)$ for all $h \in G$. A *G-subshift* Ω is a G-invariant closed subset of Σ^G.

Given two G-subshifts Ω_1 and Ω_2, we say that a map $\varphi : \Omega_1 \to \Omega_2$ is a *conjugacy* if it is bijective, continuous, and G-equivariant, that is, $g \cdot \varphi(x) = \varphi(g \cdot x)$ for every $\omega \in \Omega_1$ and $g \in G$. In this context, these maps are characterized as *sliding block codes* (e.g., see [12, 27]) and provide a notion of isomorphism between G-subshifts.

Any G-subshift Ω is characterized by the existence of a family of forbidden patterns $\mathfrak{F} \subseteq \bigcup_{F \in \mathcal{F}(G)} \Sigma^F$ such that $\Omega = X_{\mathfrak{F}}$, where

$$X_{\mathfrak{F}} = \{\omega \in \Sigma^G : (g \cdot x)_F \notin \mathfrak{F} \text{ for all } g \in G\}.$$

If the family $\mathfrak{F}$ can be chosen to be finite, we say that Ω is a *G-subshift of finite type* (*G-SFT*). Given a finite set $S \subseteq G$, we can consider a family of $|\Sigma| \times |\Sigma|$ binary matrices $\mathbf{M} = \{M_s\}_{s \in S}$ with rows and columns indexed by the elements of Σ, and define the set

$$\Omega_{\mathbf{M}} = \{\omega \in \Sigma^G : M_s(\omega(g), \omega(sg)) = 1 \text{ for all } g \in G, s \in S\}.$$

The set $\Omega_{\mathbf{M}}$ is a special kind of G-SFT known as *nearest neighbor (n.n.) G-SFT*. It is known that for every G-SFT there exists a conjugacy to an n.n. G-SFT, so we are not losing much generality by considering n.n. G-SFTs instead of general G-SFTs.

We say that an n.n. G-SFT $\Omega_{\mathbf{M}}$ has a *safe symbol* if there exists $a \in \Sigma$ such that a can be adjacent to any other symbol $b \in \Sigma$. Formally, this means that for all $s \in S$ and $b \in \Sigma$, $M_s(a, b) = M_s(b, a) = 1$.

9.2. *Entropy and potentials.* Given a G-subshift Ω, we define its *topological entropy* as

$$h_G(\Omega) := \lim_n \frac{\log|\Omega_{F_n}|}{|F_n|},$$

where $\{F_n\}_n$ is a Følner sequence and $\Omega_F = \{\omega_F : \omega \in \Omega\}$ is the set of restrictions of points in Ω to the set $F \subseteq G$. It is known that the definition of $h_G(\Omega)$ is independent of the choice of Følner sequence and is also a *conjugacy invariant*, that is, if $\varphi : \Omega_1 \to \Omega_2$ is a conjugacy, then $h_G(\Omega_1) = h_G(\Omega_2)$.

A *potential* is any continuous function $\phi : \Omega \to \mathbb{R}$. Given a potential, we define the *pressure* as

$$p_G(\phi) := \lim_n \frac{\log|Z_{F_n}(\phi)|}{|F_n|},$$

where $Z_{F_n}(\phi) = \sum_{w \in \Omega_F} \sup_{\omega \in [w]} \exp(\sum_{g \in F} \phi(g \cdot \omega))$. Notice that $p_G(0) = h_G(\Omega)$.

A *single-site potential* is any potential that only depends on the value of ω at 1_G, that is, ω_{1_G}. In other words, and without risk of ambiguity, we can think that a single-site potential is just a function $\phi : \Sigma \to \mathbb{R}$. In this case, $Z_{F_n}(\phi)$ has the following simpler expression:

$$Z_{F_n}(\phi) = \sum_{w \in \Omega_F} \prod_{g \in F} \exp(\phi(w(g))).$$

In this context, we will say that a symbol $a \in \Sigma$ is a *vacuum state* if a is a safe symbol and $\phi(a) = 0$.

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

9.3. *From a hardcore model to an n.n. G-SFT with a vacuum state.* Let (Γ, λ) be a hardcore model in $\mathcal{H}_G$. If $G \curvearrowright \Gamma$ is transitive, then $\Gamma = \mathrm{Cay}(G, S)$ for some finite symmetric set $S \subseteq G$. Then, it is easy to see that if $\Sigma = \{0, 1\}$ and, for all $s \in S$,

$$M_s = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix},$$

then $\Omega_{\mathbf{M}}$ coincides with the set $X(\Gamma)$ and 0 is a safe symbol. In addition, there is a natural relationship between the activity function λ and the single-site potential given by $\phi(0) = 0$ and $\phi(1) = \log \lambda(v)$, where v is some (or any) vertex v. In other words, if $G \curvearrowright \Gamma$ is transitive, then (Γ, λ) corresponds to an n.n. G-SFT with a vacuum state.

More generally, if $G \curvearrowright \Gamma$ is almost transitive, then (Γ, λ) can also be interpreted as an n.n. G-SFT with a vacuum state. Indeed, consider the set $\Sigma_\Gamma = X(\Gamma[U_0])$, that is, the set of independent sets of the subgraph $\Gamma[U_0]$ induced by some fundamental domain U_0. Since Γ is locally finite and $G \curvearrowright \Gamma$ is free, there must exist a finite set $S \subseteq G \setminus \{1_G\}$ such that SU_0 contains all the vertices adjacent to U_0. Considering this, we define a collection of matrices $\mathbf{M}_\Gamma = \{M_s\}_{s \in S}$, where

$$M_s(x, x') = \begin{cases} 1 & \text{if } xx' \in X(\Gamma[U_0 \cup sU_0]), \\ 0 & \text{otherwise,} \end{cases}$$

and xx' denotes the concatenation of the independent set x of $\Gamma[U_0]$ and the independent set x' of $\Gamma[sU_0]$. In other words, $M_s(x, x') = 1$ if and only if the union of the independent set x and the independent set x' is also an independent set of $\Gamma[U_0 \cup sU_0]$.

Then, there is a natural identification between $\Omega_{\mathbf{M}_\Gamma} \subseteq \Sigma_\Gamma^G$ and $X(\Gamma)$. In particular, the symbol $0^{U_0} \in X(\Gamma[U_0])$ plays the role of a safe symbol in $\Omega_{\mathbf{M}_\Gamma}$. Moreover, we can define the single-site potential $\phi_\lambda : \Omega_{\mathbf{M}_\Gamma} \to \mathbb{R}$ given by $\phi_\lambda(\omega) = \sum_{v \in U_0} \omega_{1_G}(v) \log \lambda(v)$. Then, for every $F \in \mathcal{F}(G)$,

$$\begin{aligned} Z_F(\phi_\lambda) &= \sum_{w \in \Omega_F} \prod_{g \in F} \exp(\phi_\lambda(w(g))) \\ &= \sum_{w \in X(\Gamma[FU_0])} \prod_{g \in F} \exp\left(\sum_{v \in U_0} w(g)(v) \log \lambda(v) \right) \\ &= \sum_{w \in X(\Gamma[FU_0])} \prod_{g \in F} \prod_{v \in U_0} \exp(w(g)(v) \log \lambda(v)) \\ &= \sum_{w \in X(\Gamma[FU_0])} \prod_{v \in FU_0} \lambda(v)^{w(v)} \\ &= Z_\Gamma(FU_0, \lambda). \end{aligned}$$

Therefore, $p_G(\Omega_\Gamma, \phi_\lambda) = f_G(\Gamma, \lambda)$. In the language of dynamics, for every almost transitive and locally finite graph Γ, there exists an n.n. G-SFT with a safe symbol Ω_Γ such that $G \curvearrowright X(\Gamma)$ and $G \curvearrowright \Omega_\Gamma$ are conjugated. Moreover, this gives us a way to identify any hardcore model $(\Gamma, \lambda) \in \mathcal{H}_G$ with the corresponding G-SFT Ω_Γ and the single-site potential ϕ_λ.

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

9.4. *From an n.n. G-SFT with a vacuum state to a hardcore model.* Conversely, given an n.n. G-SFT Ω and a potential with a vacuum state, we can translate this scenario into a hardcore model. Indeed, consider the graph Γ_Ω defined as follows:

- for every $g \in G$, consider a finite graph Γ_g isomorphic to $K_{|\Sigma|}$, the complete graph with $|\Sigma|$ vertices. In other words, for each $g \in G$ and for each $a \in \Sigma$, there will be a vertex $v_{g,a} \in V(\Gamma_g)$ and for every $a \neq b$, the edge $(v_{g,a}, v_{g,b})$ will belong to $E(\Gamma_g)$;
- the graph Γ_Ω will be the union of all the finite graphs Γ_g plus some extra edges;
- for every $s \in S$ and $a, b \in \Sigma$, we add the edge $(v_{1_G,a}, v_{s,b})$ if and only if $M_s(a, b) = 0$;
- we define $\lambda_\phi : \Gamma_\Omega \to \mathbb{R}_{>0}$ as $\lambda_\phi(v_{g,a}) = \exp(\phi(a))$ for every $g \in G$ and $a \in \Sigma$.

Then, G acts on Γ_Ω in the natural way and $V(\Gamma_{1_G})$ corresponds to a fundamental domain of the action $G \curvearrowright \Gamma_\Omega$. In the language of dynamics, for every n.n. G-SFT with a safe symbol Ω, there exists an almost transitive and locally finite graph Γ_Ω such that $G \curvearrowright \Omega$ and $G \curvearrowright X(\Gamma_\Omega)$ are conjugated. Moreover, it is clear that

$$f_G(\Gamma_\Omega, \lambda_\phi) = p_G(\Omega, \phi),$$

so in particular, all the representation and approximation theorems for free energy of hardcore models can be used to represent and approximate the pressure of n.n. G-SFTs Ω and potentials ϕ with a vacuum state, provided $(\Gamma_\Omega, \lambda_\phi)$ satisfies the corresponding hypotheses. Relevant cases like the *Widom–Rowlinson model* [18] and graph homomorphisms from Γ to any finite graph with some vertex (which plays the role of a safe symbol) connected to every other vertex fall into this category.

9.5. *Topological entropy and constraintedness of n.n. G-SFTs with safe symbols.* Let $\Omega \subseteq \Sigma^G$ be an n.n. G-SFT with $|\Sigma| = n_s + n_u$, where n_s denotes the number of safe symbols in Σ and n_u denotes the number of symbols that are not safe symbols (*unsafe*). Consider the n.n. G-SFT $\Omega_{n_u} \subseteq \Sigma_{n_u}^G$ obtained after collapsing all the safe symbols in Σ into a single one, so that the $|\Sigma_{n_u}| = 1 + n_u$, and construct the graph $\Gamma_{\Omega_{n_u}}$. Then, given $F \in \mathcal{F}(G)$, we have that

$$\begin{aligned}
|\Omega_F| &= \sum_{x \in X(\Gamma_{\Omega_{n_u}}, FU_0)} \prod_{v \in FU_0} 1^{x(v)} n_s^{1-x(v)} \\
&= \sum_{x \in X(\Gamma_{\Omega_{n_u}}, FU_0)} \prod_{v \in FU_0} \left(\frac{1}{n_s}\right)^{x(v)} n_s \\
&= n_s^{|FU_0|} \sum_{x \in X(\Gamma_{\Omega_{n_u}}, FU_0)} \prod_{v \in FU_0} \left(\frac{1}{n_s}\right)^{x(v)} \\
&= n_s^{|FU_0|} Z_{\Gamma_{\Omega_{n_u}}}\left(FU_0, \frac{1}{n_s}\right),
\end{aligned}$$

so, considering that $n_u = |U_0| = |\Gamma_{\Omega_{n_u}}/G|$,

$$\begin{aligned}
h_G(\Omega) &= \lim_n \frac{\log|\Omega_{F_n}|}{|F_n|} \\
&= |U_0| \log n_s + \lim_n \frac{Z_{\Gamma_{\Omega_{n_u}}}(FU_0, 1/n_s)}{|F_n|}
\end{aligned}$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

$$= |U_0| \log n_s + |U_0| f_G(\Gamma_{\Omega_{n_u}}, 1/n_s)$$
$$= n_u(\log n_s + f_G(\Gamma_{\Omega_{n_u}}, 1/n_s)).$$

Therefore, to understand and approximate $h_G(\Omega)$ reduces to study the hardcore model on Γ_{Ω,n_s} with constant activity $1/n_s$. In particular, if

$$\frac{1}{n_s} < \lambda_c(\mu(\Gamma_{\Omega_{n_u}})),$$

the hardcore model $(\Gamma_{\Omega,n_s}, 1/n_s)$ satisfies exponential SSM and the theory developed in the previous sections applies. This motivates the definition of the *constraintedness* of an n.n. G-SFT Ω as the connective constant of $\Gamma_{\Omega_{n_u}}$, that is,

$$\mu(\Omega) := \mu(\Gamma_{\Omega_{n_u}}),$$

which can be regarded as a measure of how constrained is Ω (the higher $\mu(\Omega)$, the more constrained it is). Notice that if

$$\frac{1}{n_s} < \lambda_c(\mu(\Omega) + 1),$$

then $(\Gamma_{\Omega_{n_u}}, 1/n_s)$ satisfies exponential SSM. In particular, Ω_{n_u} has a unique *measure of maximal entropy* and therefore, Ω also has unique measure of maximal entropy, namely, the pushforward measure (see [11, 21]). Moreover, the topological entropy of Ω_{n_u} has an arboreal representation and can be approximated efficiently. Since $\mu(\Omega) \leq \Delta(\Gamma_{\Omega_{n_u}}) - 1$, we have that it suffices that

$$\frac{1}{n_s} < \lambda_c(\Delta(\Gamma_{\Omega_{n_u}})),$$

For example, the n.n. G-SFT Ω represented in Figure 6 satisfies that $\Delta(\Gamma_{\Omega_{n_u}}) = 6$ and $\lambda_c(6) = 5^5/4^6 = 3125/4096$; then, if $n_s > 4096/3125 = 1.31072$, we see that it suffices to have two copies of the safe symbol 0 to have exponential SSM.

In general, since each vertex of the fundamental domain is connected to $n_u - 1$ vertices in the clique and to at most n_u vertices for each element s in the generating set S, we see that each vertex in $\Gamma_{\Omega_{n_u}}$ is connected to at most $(n_u - 1) + |S|n_u$ other vertices. Then, we can estimate that

$$\Delta(\Gamma_{\Omega_{n_u}}) \leq (|S| + 1)n_u - 1,$$

so, in particular, if

$$\frac{1}{n_s} < \lambda_c((|S| + 1)n_u - 1),$$

exponential SSM holds (and therefore, again, uniqueness of measure of maximal entropy). This last equation and its relationship with the constraintedness of Ω has a similar flavor to the relationship between the percolation threshold $p_c(\mathbb{Z}^d)$ of the $\mathbb{Z}^d$ lattice and the concept of *generosity* for $\mathbb{Z}^d$-SFTs introduced in [21] by Häggström.

Remark 9.1. It may be the case that an n.n. G-SFT $\Omega \subseteq \Sigma^G$ with a safe symbol could be represented by a graph Γ which is better in terms of connectedness or maximum degree

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

0	1	2	1	3	2	1
0	0	1	0	1	0	2
2	3	0	1	2	1	3
3	0	0	2	3	2	1
1	0	1	3	0	0	3
2	1	3	2	0	0	0
3	2	0	3	2	1	3

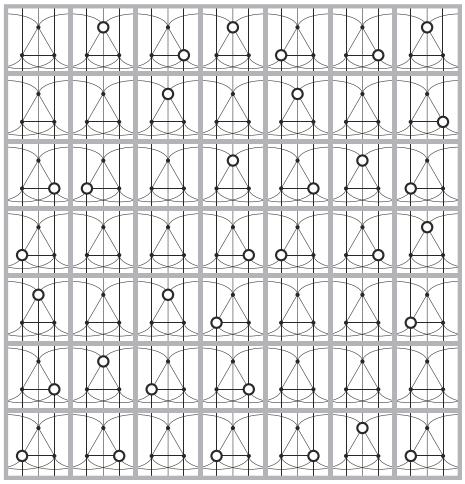

FIGURE 6. On the left, a sample of a configuration in the n.n. SFT Ω corresponding to proper 3-colorings of $\mathrm{Cay}(\mathbb{Z}^2, \{\pm(1,0), \pm(0,1)\})$ plus a safe symbol 0, where each square corresponds to an element of $\mathbb{Z}^2$. On the right, the independent set in the graph Γ_Ω representing the configuration in Ω.

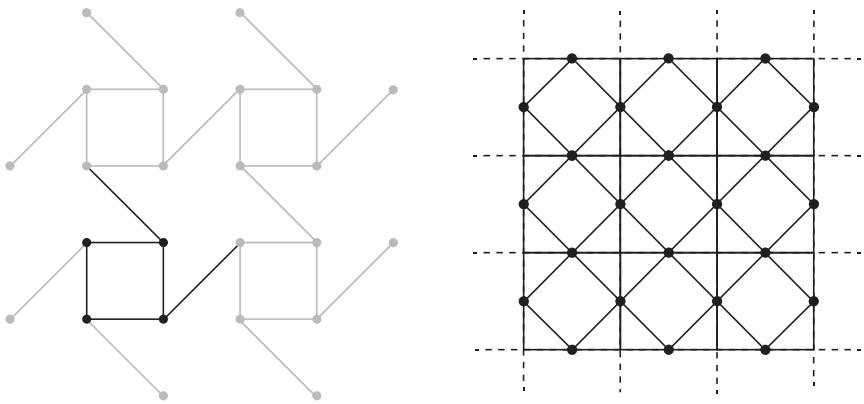

FIGURE 7. On the left, an almost transitive graph Γ with $\Gamma[U_0] \cong C_4$, the 4-cycle. On the right, a portion of the line graph of $\mathrm{Cay}(\mathbb{Z}^2, \{\pm(1,0), \pm(0,1)\})$.

compared with the canonical representation Γ_Ω, since we could encode Σ using other fundamental domains, with a lower connectivity than the complete graph. For example, the n.n. $\mathbb{Z}^2$-SFT Ω_Γ corresponding to the graph Γ on the left in Figure 7 has seven symbols (the seven independent sets of the 4-cycle), including a safe one. However, the canonical graph representation of Ω_Γ, that is, the graph Γ_{Ω_Γ}, has a fundamental domain consisting of a clique with six vertices, without considering extra connections. In particular, we see that both Γ and Γ_{Ω_Γ} represent Ω, but $\Delta(\Gamma) = 3 < 6 \leq \Delta(\Gamma_{\Omega_\Gamma})$. This motivates a finer notion of constraintedness, namely,

$$\tilde{\mu}(\Omega) = \inf\{\mu(\Gamma) : \Gamma \text{ represents } \Omega_{n_u}\},$$

and the aforementioned results would still hold if we replace $\mu(\Omega)$ by $\tilde{\mu}(\Omega)$. Notice that a fundamental domain U_0 has at least $|U_0| + 1$ independent sets (the empty one and all

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

the singletons). In particular, this implies that $\tilde{\mu}(\Omega)$ is a minimum, since we only need to optimize over graphs Γ with a fundamental domain U_0 such that $X(\Gamma[U_0]) = |\Sigma_{n_u}|$.

9.6. *The monomer–dimer model and line graphs.* Given a graph $\Gamma = (V, E)$, we say that two different edges $e_1, e_2 \in E$ are *incident* if they have one vertex in common. A *matching* in Γ is a subset M of E without incident edges. In a total parallel with the hardcore model case, we can represent a matching with an indicator function $m : E \to \{0, 1\}$, denote the set of matchings of Γ by $X^e(\Gamma)$, and define the associated partition function for some activity function $\lambda : E \to \mathbb{R}_{>0}$ as

$$Z^e_\Gamma(\lambda) = \sum_{m \in X^e(\Gamma)} \prod_{e \in E} \lambda(e)^{m(e)}.$$

The pair (Γ, λ) is called the *monomer–dimer model* and, as for the case of the hardcore model, we can define its associated free energy and Gibbs measures for a Gibbs specification adapted to this case.

An important feature of the monomer–dimer model is that, despite all its similarities with the hardcore model, it exhibits the SSM property for all values of λ [7] and, in particular, there is no phase transition [22].

Considering this, most of the results presented in this paper, in particular those related to representation and approximation, can be adapted to counting matchings (see [16] for a particular case), and there will not be a phase transition. One way to see this is through the *line graph* $L(\Gamma)$ of the given graph Γ. Indeed, if we define $L(\Gamma)$ as the graph with the set of vertices E and the set of edges containing all the adjacent edges in E, it is direct to see that there is a correspondence between matchings in Γ and independent sets in $L(\Gamma)$, that is,

$$Z^e_\Gamma(\lambda) = Z_{L(\Gamma)}(\lambda).$$

In particular, this tell us that all the results in our paper that involve some restriction on λ apply to every graph that can be obtained as a line graph of another one without restriction on λ. For example, the graph Γ represented on the right in Figure 7 corresponds to the line graph of the Cayley graph of $\mathbb{Z}^2$ with canonical generators, that is,

$$\Gamma = L(\mathrm{Cay}(\mathbb{Z}^2, \{\pm(1, 0), \pm(0, 1)\})).$$

Then, this observation implies that we can represent and approximate $f_{\mathbb{Z}^2}(\Gamma, \lambda)$ for every $\mathbb{Z}^2$-invariant activity function λ on Γ.

9.7. *From almost free to free.* Suppose that $G \curvearrowright \Gamma_*$ is almost free, that is, $|\mathrm{Stab}_G(v)| < \infty$ for all v. We proceed to show how to reduce the computation of $f_G(\Gamma, \lambda)$ to the computation of $f_G(\Gamma_*, \lambda_*)$ for some free action $G \curvearrowright \Gamma_*$, where Γ_* and λ_* are some suitable auxiliary graph and activity function, respectively. We consider this as another example of the versatility of independent sets for representing many phenomena.

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

Given a graph Γ and an almost free action $G \curvearrowright \Gamma$, let Γ_* be the new graph obtained by setting

$$V(\Gamma_*) = \{(v, s) : v \in V(\Gamma), s \in \mathrm{Stab}_G(v)\},$$

and

$$E(\Gamma_*) = \bigcup_{\substack{v \in V(\Gamma) \\ s, s' \in \mathrm{Stab}_G(v) \\ s \neq s'}} \{((v, s), (v, s'))\} \cup \bigcup_{\substack{(v,v') \in E(\Gamma) \\ s \in \mathrm{Stab}_G(v) \\ s' \in \mathrm{Stab}_G(v')}} \{((v, s), (v', s'))\}.$$

In simple words, for each v in $V(\Gamma)$, there are $|\mathrm{Stab}_G(v)|$ copies of v in $V(\Gamma_*)$ such that:

(1) the $|\mathrm{Stab}_G(v)|$ copies of v form a clique in Γ_*;

(2) for $(v, v') \in E(\Gamma)$, each copy of v is connected to all copies of v' in Γ_*.

Next, consider the activity function $\lambda_* : V(\Gamma_*) \to \mathbb{R}_{>0}$ given by

$$\lambda_*((v, s)) = \lambda(v)/|\mathrm{Stab}_G(v)|.$$

Notice that for every $U \Subset V(\Gamma)$, we have that

$$Z_\Gamma(U, \lambda) = Z_{\Gamma_*}(U_*, \lambda_*),$$

where $U_* \Subset V(\Gamma_*)$ denotes the set of all copies of vertices in U. Indeed, notice that each independent set in $\Gamma_*[U_*]$ can be naturally identified with a unique independent set in $\Gamma[U]$: for $x' \in X(\Gamma_*)$, we can define $x \in X(\Gamma)$ as $x(v) = 1$ if and only if there exists $s \in \mathrm{Stab}_G(v)$ so that $x'((v, s)) = 1$. Conversely, if Γ is finite, each independent set $x \in X(\Gamma)$ can be identified with $\prod_{v \in \Gamma} |\mathrm{Stab}_G(v)|^{x(v)}$ copies in $X(\Gamma_*)$. Therefore,

$$\begin{aligned} Z_{\Gamma_*}(U_*, \lambda_*) &= \sum_{x' \in X(\Gamma_*[U_*])} \prod_{(v,s) \in U_*} \lambda_*((v, s))^{x'((v,s))} \\ &= \sum_{x \in X(\Gamma[U])} \prod_{v \in U} |\mathrm{Stab}_G(v)|^{x(v)} \prod_{v \in U} \left(\frac{\lambda(v)}{|\mathrm{Stab}_G(v)|} \right)^{x(v)} \\ &= \sum_{x \in X(\Gamma[U])} \prod_{v \in U} \lambda(v)^{x(v)} \\ &= Z_\Gamma(U, \lambda). \end{aligned}$$

Now, pick a fundamental domain $U_0 \subseteq V$ of $G \curvearrowright \Gamma$ and, for each $v \in U_0$, consider the set of left cosets $\{g\mathrm{Stab}_G(v) : g \in G\}$ with a fixed set of representatives $R(v)$, that is, $G = \bigsqcup_{g \in R(v)} g\mathrm{Stab}_G(v)$ for each $v \in U_0$. Without loss of generality, $1_G \in R(v)$ for every v. Next, given $v \in U_0$ and $h \in G$, define $t_{v,h}$ as the unique element in $\mathrm{Stab}_G(v)$ such that $ht_{v,h}^{-1} \in R(v)$. In addition, given another element $r \in G$, define $\psi^v_{h \to r} := ht_{v,h}^{-1} t_{v,r} h^{-1}$. Notice that if $t \in \mathrm{Stab}_G(v)$, then $t_{v,ht} = t_{v,h}t$ and $\psi^v_{ht \to rt} = \psi^v_{h \to r}$. Next, consider the action of G on Γ_* that, given a vertex $(u, s) \in V(\Gamma_*)$ and $h \in G$, it is defined as

$$h \cdot (u, s) = (hu, hs\psi^v_{g \to hg} h^{-1}),$$

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

where $u \in Gv$ for $v \in U_0$ and g is some (or any) element in G such that $u = gv$. First, let us check that this is indeed an action. If $v \in U_0$, then

$$\begin{aligned} r \cdot (h \cdot (v, 1_G)) &= r \cdot (hv, h\psi^v_{1_G \to h} h^{-1}) \\ &= r \cdot (hv, ht_{v,h}h^{-1}) \\ &= (rhv, rht_{v,h}h^{-1}\psi^v_{h \to rh} r^{-1}) \\ &= (rhv, rht_{v,h}h^{-1}(ht^{-1}_{v,h}t_{v,rh}h^{-1})r^{-1}) \\ &= (rhv, rht_{v,rh}(rh)^{-1}) \\ &= (rh) \cdot (v, 1_G). \end{aligned}$$

Therefore, if (u, s) is arbitrary, with $u = gv$ and $s = gtg^{-1}$ for $g \in R(v)$ and $t \in \mathrm{Stab}_G(v)$, we have that

$$\begin{aligned} r \cdot (h \cdot (u, s)) &= r \cdot (h \cdot ((gt) \cdot (v, 1_G)) \\ &= r \cdot ((hgt) \cdot (v, 1_G)) \\ &= (rhgt) \cdot (v, 1_G) \\ &= (rh) \cdot ((gt) \cdot (v, 1_G)) \\ &= (rh) \cdot (u, s). \end{aligned}$$

Now, let us check that the action is free. If $h \cdot (v, 1_G) = (v, 1_G)$ for $v \in U_0$, then $(hv, ht_{v,h}h^{-1}) = (v, 1_G)$, so $h \in \mathrm{Stab}_G(v)$ and $t_{v,h} = 1_G$. Therefore, $h = 1_G$. In the general case, if $h \cdot (gv, s) = (gv, s)$ with $u = gv$ and $s = gtg^{-1}$ for $g \in R(v)$ and $t \in \mathrm{Stab}_G(v)$, we have that $(hgt) \cdot (v, 1_G) = h \cdot ((gt) \cdot (v, 1_G)) = (gt) \cdot (v, 1_G)$, so $(t^{-1}g^{-1}hgt) \cdot (v, 1_G) = (v, 1_G)$ and, by the previous step, $t^{-1}g^{-1}hgt = 1_G$ or, equivalently, $h = 1_G$.

Now, if $G \curvearrowright \Gamma$ is almost transitive, then $G \curvearrowright \Gamma_*$ is almost transitive, too, with $|\Gamma_*/G| = |\Gamma/G|$. Indeed, if U_0 is a fundamental domain for $G \curvearrowright \Gamma$, we have that $U_0 \times \{1_G\}$ is a fundamental domain for $G \curvearrowright \Gamma_*$, since for every $(u, s) \in V(\Gamma_*)$, there exists a unique $v \in U_0$, $g \in G$, and $t \in \mathrm{Stab}_G(v)$ such that $u = gv$ and $s = gtg^{-1}$, so $(gt) \cdot (v, 1_G) = (u, s)$.

Finally, if $K = \bigcup_{v \in U_0} \mathrm{Stab}_G(v)$, observe that $U_0 \times \{1_G\} \subseteq (U_0)_* \subseteq K(U_0 \times \{1_G\})$, so

$$Z_{\Gamma_*}(F_n(U_0 \times \{1_G\}), \lambda_*) \leq Z_{\Gamma_*}(F_n(U_0)_*, \lambda_*) \leq Z_{\Gamma_*}(F_n K(U_0 \times \{1_G\}), \lambda_*)$$

and, since $Z_\Gamma(F_n U_0, \lambda) = Z_{\Gamma_*}((F_n U_0)_*, \lambda_*) = Z_{\Gamma_*}(F_n(U_0)_*, \lambda_*)$, applying logarithms, dividing by $|F_n U_0|$, and taking the limit in n, we obtain that

$$f_G(\Gamma, U_0, \lambda) = f_G(\Gamma_*, U_0 \times \{1_G\}, \lambda_*),$$

where we have used that $\{F_n K\}_n$ is a Følner sequence, $\lim_n |F_n K|/|F_n| = 1$, and $|F_n(U_0 \times \{1_G\})| = |F_n U_0| = |F_n||U_0|$.

9.8. *Spectral radius of matrices and occupation probabilities on trees.* A curious consequence of the hardcore model representation of an n.n. G-SFT with a safe symbol

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

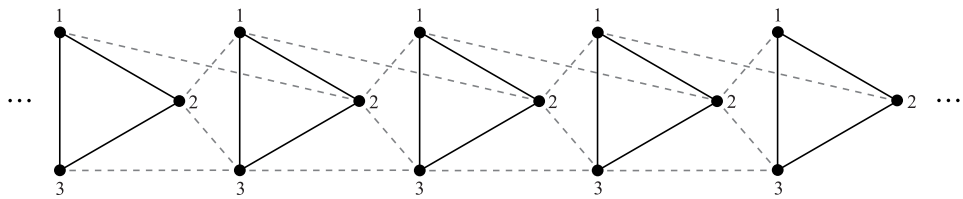

FIGURE 8. A graph representation of the 0–1 matrix M.

is that when $G = \mathbb{Z}$ and $S = \{1\}$, then $h_{\mathbb{Z}}(\Omega)$ has a well-known characterization in terms of the transition matrix $M = M_1$ [27].

If M is irreducible and aperiodic, there is always a unique stationary Markov chain $\mathbb{P}_M$ associated to M such that $\log \lambda_M = h_{\mathbb{Z}}(\Omega_M)$, where λ_M denotes the Perron eigenvalue of M and we consider the natural invariant order in $\mathbb{Z}$.

Now, if M is a matrix such that the ith row and the ith column have no zeros, then M is irreducible and aperiodic. In fact, the ith symbol, let us call it a, is a safe symbol. In this case, we have that

$$\log \lambda_M = h_{\mathbb{Z}}(\Omega_M) = -\log \mathbb{P}_M([a^0]|[a^{-\mathbb{N}}]) = -\log \mathbb{P}_M([a^0]|[a^{-1}]).$$

Therefore, $\lambda_M = 1/\mathbb{P}_M(a^0|a^{-1})$ and to compute the spectral radius of M reduces to compute $\mathbb{P}_M(a^0|a^{-1})$. For example, consider the following matrix:

$$M = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 1 & 1 & 1 & 0 \end{pmatrix},$$

where a is the symbol associated to the first row. Given this matrix M, we can always construct a graph representation Γ_{Ω_M} of Ω_M as in Figure 8.

Now, it is known that $|\mathbb{P}_M(0^0|0^{-1}) - \mathbb{P}_M(0^0|0^{\{n,-1\}})|$ goes to zero exponentially fast as n goes to infinity and, since every Markov chain is a Markov random field, we have that

$$\begin{aligned}
\mathbb{P}_M(a^0|a^{\{n,-1\}}) &= \mathbb{P}_{\Gamma_{\Omega_M}[\{0,\dots,n-1\}U_0],1}(0^{U_0}) \\
&= \sum_{i=1}^{U_0} \mathbb{P}_{\Gamma_{\Omega_M}[\{0,\dots,n-1\}U_0],1}(0^{v_i}|0^{\{v_1,\dots,v_{i-1}\}}) \\
&= \sum_{i=1}^{U_0} \mathbb{P}_{\Gamma_{\Omega_M}[\{0,\dots,n-1\}U_0\setminus\{v_1,\dots,v_{i-1}\}],1}(0^{v_i}) \\
&= \sum_{i=1}^{U_0} \mathbb{P}_{T_{\mathrm{SAW}}\Gamma_{\Omega_M}[\{0,\dots,n-1\}U_0\setminus\{v_1,\dots,v_{i-1}\}],1}(0^{v_i}).
\end{aligned}$$

This gives us an arboreal representation and a method to compute the spectral radius of any matrix M satisfying the conditions described above that we believe could be of independent interest.

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

Acknowledgements. I would like to thank Brian Marcus for his valuable suggestions after reading a first draft of this work, Tom Meyerovitch for a helpful discussion about invariant random orders, and the anonymous referee for a careful reading of the manuscript and many useful comments and suggestions. I acknowledge the support of ANID/FONDECYT de Iniciación en Investigación 11200892.

References

[1] A. Alpeev, T. Meyerovitch and S. Ryu. Predictability, topological entropy and invariant random orders. *Proc. Amer. Math. Soc.* **149** (2021), 1443–1457.

[2] S. Arora and B. Barak. *Computational Complexity – A Modern Approach*. Cambridge University Press, New York, NY, 2009.

[3] L. Auslander. On a problem of Philip Hall. *Ann. of Math. (2)* **86** (1967), 112–116.

[4] A. Barvinok. *Combinatorics and Complexity of Partition Functions (Algorithms and Combinatorics, 30)*. Springer, Cham, 2016.

[5] G. Baumslag, F. B. Cannonito, D. J. S. Robinson and D. Segal. The algorithmic theory of polycyclic-by-finite groups. *J. Algebra* **142**(1) (1991), 118–149.

[6] R. J. Baxter. Hard hexagons: exact solution. *J. Phys. A* **13**(3) (1980), L61.

[7] M. Bayati, D. Gamarnik, D. Katz, C. Nair and P. Tetali. Simple deterministic approximation algorithms for counting matchings. *Proc. Thirty-Ninth Annual ACM Symp. on Theory of Computing*. ACM, New York, 2007, pp. 122–127.

[8] R. Briceño. The topological strong spatial mixing property and new conditions for pressure approximation. *Ergod. Th. & Dynam. Sys.* **38**(5) (2018), 1658–1696.

[9] R. Briceño. An SMB approach for pressure representation in amenable virtually orderable groups. *J. Anal. Math.* **142** (2020), 421–451.

[10] G. R. Brightwell and P. Winkler. Graph homomorphisms and phase transitions. *J. Combin. Theory Ser. B* **77**(2) (1999), 221–262.

[11] R. Burton and J. E. Steif. New results on measures of maximal entropy. *Israel J. Math.* **89** (1995), 275–300.

[12] T. Ceccherini-Silberstein and M. Coornaert. *Cellular Automata and Groups*. Springer, Berlin, 2010.

[13] P. Dobrushin. The description of a random field by means of conditional probabilities and conditions of its regularity. *Theory Probab. Appl.* **13**(2) (1968), 197–224.

[14] T. Downarowicz, B. Frej and P.-P. Romagnoli. Shearer's inequality and infimum rule for Shannon entropy and topological entropy. *Dynamics and Numbers*. Eds. S. Kolyada, M. Möller, P. Moree and T. Ward. American Mathematical Society, Providence, RI, 2016, pp. 63–75.

[15] H. Duminil-Copin and S. Smirnov. The connective constant of the honeycomb lattice equals $\sqrt{2+\sqrt{2}}$. *Ann. of Math. (2)* **175** (2012), 1653–1665.

[16] D. Gamarnik and D. Katz. Sequential cavity method for computing free energy and surface pressure. *J. Stat. Phys.* **137**(2) (2009), 205–232.

[17] H.-O. Georgii. *Gibbs Measures and Phase Transitions*, 2nd edn. de Gruyter, Berlin, 2011.

[18] H.-O. Georgii, O. Häggström and C. Maes. *The Random Geometry of Equilibrium Phases (Phase Transitions and Critical Phenomena, 18)*. Academic Press, London, 2001.

[19] M. Gromov. Groups of polynomial growth and expanding maps. *Publ. Math. Inst. Hautes Études Sci.* **53**(1) (1981), 53–78.

[20] B. M. Gurevich and A. A. Tempelman. A Breiman type theorem for Gibbs measures. *J. Dyn. Control Syst.* **13**(3) (2007), 363–371.

[21] O. Häggström. On phase transitions for subshifts of finite type. *Israel J. Math.* **94** (1996), 319–352.

[22] O. J. Heilmann and E. H. Lieb. Theory of monomer-dimer systems. *Comm. Math. Phys.* **25**(3) (1972), 190–232.

[23] M. Hochman and T. Meyerovitch. A characterization of the entropies of multidimensional shifts of finite type. *Ann. of Math. (2)* **171**(3) (2010), 201–2038.

[24] M. R. Jerrum, L. G. Valiant and V. V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theor. Comput. Sci.* **43** (1986), 169–188.

[25] F. P. Kelly. Stochastic models of computer communication systems. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **47**(3) (1985), 379–395.

[26] D. Kerr and H. Li. *Ergodic Theory. Independence and Dichotomies*. Springer, Cham, 2016.

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press

[27] D. Lind and B. Marcus. *An Introduction to Symbolic Dynamics and Coding*. Cambridge University Press, Cambridge, 1995.
[28] E. Lindenstrauss. Pointwise theorems for amenable groups. *Invent. Math.* **146**(2) (2001), 259–295.
[29] R. J. Lipton and Y. Zalcstein. Word problems solvable in logspace. *J. Assoc. Comput. Mach.* **24**(3) (1977), 522–526.
[30] J. Liu, A. Sinclair and P. Srivastava. The Ising partition function: zeros and deterministic approximation. *J. Stat. Phys.* **174**(2) (2019), 287–315.
[31] M. Luby and E. Vigoda. Approximately counting up to four. *Proc. Twenty-Ninth Annual ACM Symp. on Theory of Computing*. ACM, New York, 1997, pp. 682–687.
[32] R. C. Lyndon and P. E. Schupp. *Combinatorial Group Theory*. Springer-Verlag, Berlin, 2001.
[33] A. I. Mal'tsev. On some classes of infinite soluble groups. *Mat. Sb.* **70**(3) (1951), 567–588.
[34] B. Marcus and R. Pavlov. Approximating entropy for a class of $\mathbb{Z}^2$ Markov random fields and pressure for a class of functions on $\mathbb{Z}^2$ shifts of finite type. *Ergod. Th. & Dynam. Sys.* **33** (2013), 186–220.
[35] B. Marcus and R. Pavlov. Computing bounds for entropy of stationary $\mathbb{Z}^d$ Markov random fields. *SIAM J. Discrete Math.* **27**(3) (2013), 1544–1558.
[36] B. Marcus and R. Pavlov. An integral representation for topological pressure in terms of conditional probabilities. *Israel J. Math.* **207**(1) (2015), 395–433.
[37] J. Meier. *Groups, Graphs and Trees: An Introduction to the Geometry of Infinite Groups (London Mathematical Society Student Texts, 73)*. Cambridge University Press, Cambridge, 2008.
[38] J. Milnor. Growth of finitely generated solvable groups. *J. Differential Geom.* **2**(4) (1968), 447–449.
[39] R. Mura and A. Rhemtulla. *Orderable Groups (Lecture Notes in Pure and Applied Mathematics, 27)*. Marcel Dekker, Inc., New York, 1977.
[40] R. Pavlov. Approximating the hard square entropy constant with probabilistic methods. *Ann. Probab.* **40**(6) (2012), 2362–2399.
[41] D. Ruelle. *Thermodynamic Formalism. The Mathematical Structure of Equilibrium Statistical Mechanics (Cambridge Mathematical Library, Encyclopedia of Mathematics and its Applications, 5)*, 2nd edn. Cambridge University Press, Cambridge, 2004.
[42] G. Sabidussi. On a class of fixed-point-free graphs. *Proc. Amer. Math. Soc.* **9**(5) (1958), 800–804.
[43] D. Segal. *Polycyclic Groups (Cambridge Tracts in Mathematics, 82)*. Cambridge University Press, Cambridge, 1983.
[44] B. Simon. *The Statistical Mechanics of Lattice Gases*. Vol. I. Princeton University Press, Princeton, 1993.
[45] H. U. Simon. Word problems for groups and context free recognition. *Fundamentals of Computation Theory (Proc. Conf. Algebraic, Arithmetic and Categorical Methods in Computational Theory, BerlinWendisch-Rietz, 1979) (Mathematical Research, 2)*. Akademie-Verlag, Berlin, 1979, pp. 417–422.
[46] A. Sinclair, P. Srivastava, D. Štefankovič and Y. Yin. Spatial mixing and the connective constant: Optimal bounds. *Probab. Theory Related Fields* **168**(1–2) (2017), 153–197.
[47] A. Sinclair, P. Srivastava and M. Thurley. Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs. *J. Stat. Phys.* **155** (2014), 666–686.
[48] A. Sly. Computational transition at the uniqueness threshold. *51th Annual IEEE Symp. on Foundations of Computer Science*. IEEE Computer Society, Los Alamitos, CA, 2010, pp. 287–296.
[49] A. Sly and N. Sun. The computational hardness of counting in two-spin models on d-regular graphs. *53rd Annual IEEE Symp. on Foundations of Computer Science*. IEEE Computer Society, Los Alamitos, CA, 2012, pp. 361–369.
[50] A. D. Sokal. A personal list of unsolved problems concerning lattice gases and antiferromagnetic Potts models. *Markov Process. Related Fields* **7** (2001), 21–38.
[51] R. G. Swan. Representations of polycyclic groups. *Proc. Amer. Math. Soc.* **18**(3) (1967), 573–574.
[52] J. Tits. Free subgroups in linear groups. *J. Algebra* **20**(2) (1972), 250–270.
[53] Y.-K. Wang, Y. Yin and S. Zhong. Approximate capacities of two-dimensional codes by spatial mixing. *2014 IEEE Int. Symp. on Information Theory*. IEEE, New York, 2014, pp. 1061–1065.
[54] D. Weitz. Combinatorial criteria for uniqueness of Gibbs measures. *Random Structures Algorithms* **27**(4) (2005), 445–475.
[55] D. Weitz. Counting independent sets up to the tree threshold. *Proc. Thirty-Eighth Annual ACM Symp. on Theory of Computing*. ACM, New York, 2006, pp. 140–149.
[56] M. Xia and W. Zhao. #3-Regular bipartite planar vertex cover is #P-complete. *Int. Conf. on Theory and Applications of Models of Computation*. Eds. J.-Y. Cai, S. B. Cooper and A. Li. Springer, Berlin, 2006, pp. 356–364.

https://doi.org/10.1017/etds.2023.38 Published online by Cambridge University Press