Limiting techniques in measured group theory

by
László Márton Tóth

Submitted to
Central European University
Department of Mathematics and its Applications

In partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Mathematics

Supervisor: Miklós Abért
Alfréd Rényi Institute of Mathematics

Budapest, Hungary
August 2018

Doctoral thesis

László Márton Tóth

Central European University, Budapest, Hungary

Printed in Hungary

Declaration

I, the undersigned László Márton Tóth, candidate for the degree of Doctor of Philosophy at the Central European University Department of Mathematics and its Applications, declare herewith that the present thesis is based on my research and only such external information as properly credited in notes and bibliography.

I declare that no unidentified and illegitimate use was made of the work of others, and no part of this thesis infringes on any person's or institution's copyright. I also declare that no part of this thesis has been submitted to any other institution of higher education for an academic degree.

Budapest, Hungary, August 2018

László Márton Tóth

Acknowledgements

I would like to express my gratitude towards those people who have helped me during my graduate studies.

I would like to thank my advisor, Prof. Miklós Abért for introducing me to the topics, for establishing a stimulating and friendly working atmosphere at the Groups and Graphs research group at the Rényi Institute, and for giving me continuous opportunities to meet leading researchers in measured group theory from around the world. I also feel that I have learned invaluable lessons from him about navigating life as a researcher, and how a little sense of humor goes a long way.

I would like to say thank you to Ferenc Bencs, who is my coauthor on one of the papers providing the backbone of this thesis. His friendship, great ideas, enthusiasm and work ethic made our project a truly enjoyable experience.

A special thank you goes to Sam Mellick, who was my companion both literally when traveling to workshops and figuratively as a fellow student along the road towards a PhD.

I wish to thank my parents and my family for their continued assistance leading all the way up to this point in time.

Last but not least I am especially thankful to my wife, Noémi. I could always rely on her support, encouragement and unmerited belief in me. Thank you.

Contents

PREFACE

Limits of finite structures is a topic that has been rapidly developing over the last two decades. The common theme of building limiting theories has been present in probability theory, extremal graph theory and group theory, just to name the areas closest to the scope of this thesis.

The extremal graph theory side is interested in dense graphs. This line of research was initiated by Chayes, Borgs, Lovász, T. Sós and Vesztergombi [Bor+08], and through the works of the aforementioned authors together with Szegedy and Schrijver expanded rapidly during the mid to late 2000s. A thorough treatment of these results can be found in Large Networks and Graph Limits by László Lovász [Lov12].

The notion local convergence of sparse graphs was introduced by Benjamini and Schramm in [BS01]. The consequent notion of unimodularity and the Mass Transport Principle were at the heart of the hugely influential works of Aldous, Benjamini, Lyons, Peres and Schramm.

Convergence of bounded degree graphs has natural connection to the study of finitely generated groups. Sofic groups were introduced by Gromov and named by Weiss as the groups that can be approximated by finite models. It turned out that soficity is really equivalent to the Cayley graph being a Benjamini-Schramm limit of finite Schreier graphs in the labeled sense. Many long standing questions from group theory have been solved for sofic groups, for example Gottschalk's surjunctivity conjecture by Gromov [Gro99], and Kaplansky's direct finiteness conjecture by Elek and Szabó [ES04]. Lewis Bowen's work [Bow10] defining sofic entropy and using it to distinguish Bernoulli shifts over free groups opened up lines of research in dynamics that were widely thought to be inaccessible since the work of Ornstein and Weiss in the 1980s. It is a central open problem of the field whether all countable groups are sofic or not.

The shared philosophy of all these areas is that the limiting language creates a connection between the combinatorial nature of the finite structures and the analytic properties of the limiting objects. In many cases the tools available for the finite and infinite worlds reinforce each other and lead to surprising results, new notions and interesting questions.

In this thesis we exhibit examples of this approach in the field of measured group theory, where the main objects of interest are probability measure preserving (p.m.p.) actions of groups. These arise naturally as limiting objects in the sparse graph limit theory.

In Chapter 1 we investigate invariant random subgroups (IRS's) in groups acting on rooted trees. A random subgroup of a discrete group Γ is said to be an IRS, if its distribution is invariant with respect to conjugation by elements of Γ.

The limiting behavior takes place between the finite levels of our tree and its boundary consisting of infinite rays. We can connect the ergodic theory of the boundary and the combinatorics of the group acting on finite levels. We obtain an understanding of IRS's under certain branching conditions and find interesting new behavior.

To be more precise let $\mathrm{Alt}_f(T)$ denote the group of finitary even automorphisms of the d-ary rooted tree T. We prove that a nontrivial ergodic IRS of $\mathrm{Alt}_f(T)$ that acts without fixed points on the boundary of T contains a level stabilizer, in particular it is the random conjugate of a finite index subgroup. This resembles a rigidity result on IRS's of higher rank simple Lie groups [Abé+17].

When one allows fixed points on the boundary the picture becomes more complicated, but we retain a degree of control over all IRS's. Applying the technique to branch groups we prove that an ergodic IRS in a finitary regular branch group contains the derived subgroup of a generalized rigid level stabilizer. This in turn resembles a result of Vershik [Ver12] on the IRS's of FSym($\mathbb{N}$), the group of permutations of a countable set moving finitely many points.

We also prove that every weakly branch group has continuum many distinct atomless ergodic IRS's. This extends a result of Benli, Grigorchuk and Nagnibeda who exhibit a group of intermediate growth with this property.

Chapter 1 is based on the preprint [BT18] submitted to Ergodic Theory and Dynamical Systems. It is joint work with fellow CEU PhD student Ferenc Bencs.

In Chapter 2 we analyze the rank gradient of finitely generated groups with respect to sequences of subgroups of finite index that do not necessarily form a chain. The key tool here is connecting the notion of *cost* [Gab00] to the *combinatorial cost* [Ele07] through the limiting theory of local-global convergence in Theorem 2.1.1. Our combinatorial argument gives a semicontinuity result in the infinite setting.

As an application we generalize several results that were only known for chains before. In particular, we show that for a finitely generated group Γ with fixed price c, every Farber sequence has rank gradient $c - 1$. Intuitively a sequence of finite index subgroups (H_n) is Farber if it approximates the group in the sense that the finite quotient spaces Γ/H_n become better and better models of the Cayley graph of Γ locally. By adapting Lackenby's trichotomy theorem in [Lac05] to this setting, we also show that in a finitely presented amenable group, every sequence of subgroups with index tending to infinity has vanishing rank gradient. Chapter 2 is based on the preprint [AT17]. It is joint work with my advisor Miklós Abért.

Chapter 3 reports on ongoing research investigating the distortion function of p.m.p. actions. The distortion function is a secondary invariant associated to the cost, it was introduced by Abért, Gelander and Nikolov in [AGN17], who studied it for right angled groups in order to control the growth of torsion homology along Farber sequences.

This direction of research is a natural extension of Chapter 2. The main difficulty which prevents us from getting stronger results in Chapter 2 is the fact that even though we might find cheaper and cheaper ways of generating group actions, we have no control over the increasing complexity, that is the bi-Lipschitz constant. One defines the distortion function to measure this smallest necessary bi-Lipschitz constant when being ε close to the cost.

In Chapter 3 we first study how the distortion function relates to weak containment. The arguments of Abért and Weiss as well as those of Chapter 2 give a monotonicity result, although one has to take care at the countable set of discontinuity points of our functions.

We then compute the distortion function of actions of $\mathbb{Z}^d$. It turns out that these have different growth types for all $d \in \mathbb{N}$. This in particular implies that the distortion function is not an orbit equivalence invariant.

We also bound the distortion function of all lamplighter groups with finitely many states of lamps over infinite base groups by a logarithmic function.

It is an open problem to exhibit a p.m.p. action of a finitely generated group with an exponentially growing distortion function.

1

Invariant random subgroups of groups acting on rooted trees

This chapter is based on the preprint [BT18] sumbitted to Ergodic Theory and Dynamical Systems. It is joint work with Ferenc Bencs.

Abstract

We investigate invariant random subgroups in groups acting on rooted trees. Let $\mathrm{Alt}_f(T)$ be the group of finitary even automorphisms of the d-ary rooted tree T. We prove that a nontrivial ergodic IRS of $\mathrm{Alt}_f(T)$ that acts without fixed points on the boundary of T contains a level stabilizer, in particular it is the random conjugate of a finite index subgroup.

Applying the technique to branch groups we prove that an ergodic IRS in a finitary regular branch group contains the derived subgroup of a generalized rigid level stabilizer. We also prove that every weakly branch group has continuum many distinct atomless ergodic IRS's. This extends a result of Benli, Grigorchuk and Nagnibeda who exhibit a group of intermediate growth with this property.

Contents

1.1 Introduction

For a countable discrete group Γ let Sub_Γ denote the compact space of subgroups $H \leq \Gamma$, with the topology induced by the product topology on $\{0,1\}^\Gamma$. The group Γ acts on Sub_Γ by conjugation. An *invariant random subgroup* (IRS) of Γ is a Borel probability measure on Sub_Γ that is invariant with respect to the action of Γ.

Examples include Dirac measures on normal subgroups and uniform random conjugates of finite index subgroups. More generally, for any p.m.p. action $\Gamma \curvearrowright (X,\mu)$ on a Borel probability space (X,μ), the stabilizer $\mathrm{Stab}_\Gamma(x)$ of a μ-random point x defines an IRS of Γ. Abért, Glasner and Virág [AGV14] proved that all IRS's of Γ can be realized this way.

A number of recent papers have been studying the IRS's of certain countable discrete groups. Vershik [Ver12] characterized the ergodic IRS's of the group $\mathrm{FSym}(\mathbb{N})$ of finitary permutations of a countable set. In [Abé+17] the authors investigate IRS's in lattices of Lie groups. Bowen [Bow12] and Bowen-Grigorchuk-Shavchenko [BGK15] showed that there exists a large "zoo" of IRS's of non-abelian free groups and the lamplighter groups $(\mathbb{Z}/p\mathbb{Z})^n \wr \mathbb{Z}$ respectively. Thomas and Tucker-Drob [TT14; TT18] classified the ergodic IRS's of strictly diagonal limits of finite symmetric groups and inductive limits of finite alternating groups. Dudko and Medynets [DM17] extend this in certain cases to blockdiagonal limits of finite symmetric groups.

In this chapter we study the IRS's of groups of automorphisms of rooted trees. Let T be the infinite d-ary rooted tree, and let $\mathrm{Aut}(T)$ denote the group of automorphisms of T.

An *elementary* automorphism applies a permutation to the children of a given vertex, and moves the underlying subtrees accordingly. The group of *finitary* automorphisms $\mathrm{Aut}_f(T)$ is generated by the elementary automorphisms. The *finitary alternating automorphism group* $\mathrm{Alt}_f(T)$ is the one generated by even elementary automorphisms.

The group $\mathrm{Aut}(T)$ comes together with a natural measure preserving action. The *boundary* of T – denoted ∂T – is the space of infinite rays of T. It is a compact metric space with a continuous $\mathrm{Aut}(T)$ action and an ergodic invariant measure $\mu_{\partial T}$.

For some natural classes of groups IRS's tend to behave like normal subgroups. In [Abé+17] the Margulis Normal Subgroup Theorem is extended to IRS's, it is shown that every nontrivial ergodic IRS of a lattice in a higher rank simple Lie group is a random conjugate of a finite index subgroup. On the other hand, the finitary alternating permutation group $\mathrm{FAlt}(\mathbb{N})$ is simple, in particular it has no finite index subgroups, but as Vershik shows in [Ver12] it admits continuum many ergodic IRS's.

The group $\mathrm{Alt}_f(T)$ is an interesting mixture of these two worlds. It is both locally finite and residually finite, and all its nontrivial normal subgroups are level stabilizers. The Margulis Normal Subgroup Theorem does not extend to IRS's, as the stabilizer of a random boundary point gives an infinite index ergodic IRS. However, once we restrict our attention to IRS's without fixed points, the picture changes.

Theorem 1.1.1. Let H be a fixed point free ergodic IRS of $\mathrm{Alt}_f(T)$, with $d \geq 5$. Then H is the uniform random conjugate of a finite index subgroup. In other words H contains a level stabilizer.

Note, that an IRS H is *fixed point free* if it has no fixed points on ∂T almost surely. In general let $\mathrm{Fix}(H)$ denote the closed subset of fixed points of H on ∂T.

When we do not assume fixed point freeness IRS's of $\mathrm{Alt}_f(T)$ start to behave like the ones in $\mathrm{FAlt}(\mathbb{N})$. In the case of $\mathrm{FAlt}(\mathbb{N})$, any nontrivial ergodic IRS contains a specific (random) subgroup that arises by partitioning the base space in an invariant random way and then taking the direct sum of deterministic subgroups on the parts. We proceed to define a (random) subgroup of $\mathrm{Alt}_f(T)$ which highly resembles these subgroups.

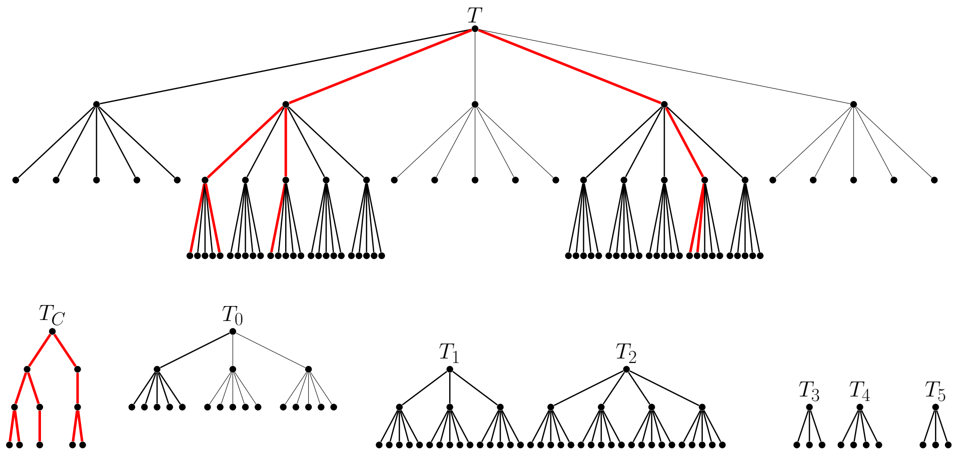

Figure 1: Decomposition of T with respect to C

Every closed subset $C \subseteq \partial T$ corresponds to a rooted subtree T_C with no leaves. The complement of T_C in T is a union of subtrees $T_0, T_1, \ldots$ as in Figure 1. Choose an integer m_i for each T_i, and let $\mathcal{L}_{m_i}(T_i)$ stand for the m_i^{th} level of the tree T_i. We define $L\big(C, (m_i)\big)$ to be the direct sum of level stabilizers in the T_i:

$$L\big(C, (m_i)\big) = \bigoplus_{i \in \mathbb{N}} \mathrm{Stab}_{\mathrm{Alt}_f(T)}\big(\mathcal{L}_{m_i}(T_i)\big).$$

It is easy to see that $\mathrm{Fix}\big(L\big(C, (m_i)\big)\big) = C$. We call such an $L\big(C, (m_i)\big)$ a *generalized congruence subgroup* with respect to the fixed point set C.

Let $\widetilde{C}$ be the translate of C with a Haar-random element from the compact group $\mathrm{Alt}(T) = \overline{\mathrm{Alt}_f(T)}$. Then $L\big(\widetilde{C}, (m_i)\big)$ becomes an ergodic IRS of $\mathrm{Alt}_f(T)$ with fixed point set $\widetilde{C}$.

Theorem 1.1.2. Let H be an ergodic IRS of $\mathrm{Alt}_f(T)$, with $d \geq 5$. Then $\mathrm{Fix}(H)$ is the Haar-random translate of a fixed closed subset C. Moreover, there exists (m_i) such that the generalized congruence subgroup $L\big(\mathrm{Fix}(H), (m_i)\big)$ is contained in H almost surely.

We can exploit our methods to prove new results on branch groups as well. We postpone the formal definition of branch groups to Section 1.2. The examples to keep in mind are the groups $\mathrm{Aut}_f(T)$, $\mathrm{Alt}_f(T)$ and groups defined by finite automata, such as the first Grigorchuk group $\mathfrak{G}$.

In [DG14] Dudko and Grigorchuk show that branch groups admit infinitely many distinct atomless (continuous) ergodic IRS's. In the ergodic case being atomless means that the measure is not supported on a finite set. In [BGN15] Benli, Grigorchuk and Nagnibeda exhibit a group of intermediate growth U_Λ with continuum many distinct atomless ergodic IRS's. We are able to find continuum many such IRS's in weakly branch groups in general.

Theorem 1.1.3. Every weakly branch group admits continuum many distinct atomless ergodic IRS's.

Note that the universal Grigorchuk group U_Λ in [BGN15] is not weakly branch, as it is not transitive on the levels. Nevertheless, by its construction it factors onto branch groups, which by Theorem 1.1.3 have continuum many distinct atomless ergodic IRS's, and those can be lifted to distinct IRS's of U_Λ. Thus Theorem 1.1.3 gives an alternate proof of the main result of [BGN15].

A key ingredient in Theorems 1.1.1 and 1.1.2 is to analyze the orbit-closures of IRS's on ∂T. For any subgroup $L \leq \mathrm{Aut}(T)$ taking the closures of orbits of L gives an equivalence relation on ∂T, that is L acts minimally on each class. It turns out that nontrivial orbit-closures of IRS's are necessarily clopen.

Theorem 1.1.4. Let H be an ergodic IRS of a countable regular branch group Γ. Then almost surely all orbit-closures of H on ∂T that are not fixed points are clopen. In particular if H is fixed point free, then H has finitely many orbit-closures on ∂T almost surely.

In a group Γ the *rigid stabilizer* of a vertex $v \in V(T)$ is the set $\mathrm{Rst}_\Gamma(v) \subseteq \Gamma$ of automorphisms that fix all vertices except the descendants of v. The rigid stabilizer of the level $\mathcal{L}_n$ is

$$\mathrm{Rst}_\Gamma(\mathcal{L}_n) = \prod_{v \in \mathcal{L}_n} \mathrm{Rst}_\Gamma(v).$$

In [Gri00, Theorem 4] Grigorchuk showed that nontrivial normal subgroups in branch groups contain the derived subgroup $\mathrm{Rst}'_\Gamma\big(\mathcal{L}_m(T)\big)$ for some $m \in \mathbb{N}$. Our next theorem can be thought of as a generalization of this statement for finitary regular branch groups.

Using the decomposition of T with respect to C above we can define a *generalized rigid level stabilizer* $L(C, m_i)$ by taking the direct sum of the rigid level stabilizers $\mathrm{Rst}_\Gamma\big(\mathcal{L}_{m_i}(T_i)\big)$ instead of the $\mathrm{Stab}_\Gamma\big(\mathcal{L}_{m_i}(T_i)\big)$ we used before. The next theorem generalizes Theorem 1.1.1 and Theorem 1.1.2 for *finitary* regular branch groups.

Theorem 1.1.5. Let Γ be a finitary regular branch group, and let H be a nontrivial ergodic IRS of Γ. Then $\mathrm{Fix}(H)$ is the Haar-random translate of a closed subset C with an element from $\overline{\Gamma}$. Also there exists (m_i) such that H almost surely contains the derived subgroup $L'\big(\mathrm{Fix}(H), (m_i)\big)$ of a generalized rigid level stabilizer. In particular if H is fixed point free, then H almost surely contains $\mathrm{Rst}'_\Gamma\big(\mathcal{L}_m(T)\big)$ for some $m \in \mathbb{N}$.

Already in the case of $\mathrm{Aut}_f(T)$ with $d = 2$ the abelianization of $\mathrm{Aut}_f(T)$ equals $(\mathbb{Z}/2\mathbb{Z})^{\mathbb{N}}$. This itself gives rise to a lot of IRS's, which makes the following consequence of Theorem 1.1.5 somewhat surprising.

Theorem 1.1.6. All ergodic fixed point free IRS's in finitary regular branch groups are supported on finitely many subgroups, and therefore are the uniform random conjugates of a subgroup with finite index normalizer.

One can think of Theorem 1.1.6 as a dual of Theorem 1.1.3. Also note that merely containing $\mathrm{Rst}'_\Gamma\big(\mathcal{L}_m(T)\big)$ does not imply finite index normalizer.

In the Grigorchuk group $\mathfrak{G}$ the elements are not finitary. In this case our methods yield a weaker result on the closures of IRS's.

Theorem 1.1.7. Let Γ be a countable regular branch group, and let H be a nontrivial ergodic *IRS* of Γ. Then there exists (m_i) such that $\overline{H}$ contains the derived subgroup

$L'(\mathrm{Fix}(H),(m_i))$ of a generalized rigid level stabilizer almost surely, where the elements of the rigid stabilizers in $L(\mathrm{Fix}(H),(m_i))$ can be chosen from $\overline{\Gamma}$ instead of Γ.

However, classifying IRS's of the discrete Grigorchuk group $\mathfrak{G}$ is still open.

Problem 1.1.8. What are the (fixed point free) ergodic IRS's of the first Grigorchuk group $\mathfrak{G}$? Is it true, that a fixed point free ergodic IRS of $\mathfrak{G}$ contains a congruence subgroup almost surely?

The structure of the chapter is as follows. We introduce the basic notions in Section 1.2 and state some lemmas leading towards Theorem 1.1.4. In Section 1.3 we investigate the actions of IRS's on the boundary and prove Theorems 1.1.3 and 1.1.4. Section 1.4 is dedicated to understanding the structure of IRS's in finitary regular branch groups and proving Theorem 1.1.5. We show how Theorems 1.1.6 and 1.1.7 follow from our earlier results in Section 1.5. In the Appendix we prove a few technical details that we postpone during the exposition.

1.2 Preliminaries

In this section we introduce the basic notions discussed in this chapter. Notation mostly follows [BGŠ03], which we recommend as an introduction to automorphisms of rooted trees and branch groups.

1.2.1 Automorphisms of rooted trees

Let T be a locally finite tree rooted at o, and let d_T denote the graph distance on T. For any vertex v the *parent* of v is the unique neighbor u of v with $d_T(u,o) = d_T(v,o) - 1$. Accordingly, the *children* of u are all the neighbors v of u with $d_T(v,o) = d_T(u,o)$. Similarly we use the phrases *ancestors* and *descendants* of a vertex v to refer to vertices that can be reached from v by taking some number of steps towards or away from the root respectively. The n^{th} level of T is the set of vertices $\mathcal{L}_n = \{v \in V(T) \mid d_T(v,o) = n\}$.

To effectively talk about automorphisms of a rooted tree T one has to distinguish the vertices. For any vertex v we fix an ordering of the children of v. In the case of the d-ary tree this corresponds to thinking of T as the set of finite length words Y^* over the alphabet Y with d letters. The empty word represents the root, and the parent of any word $w_1w_2\dots w_n$ is $w_1w_2\dots w_{n-1}$. Being an ancestor of v corresponds to being a prefix of the word corresponding to v.

An automorphism of T (which preserves the root) corresponds to a permutation of the words which preserves the prefix relation. For an element $\gamma \in \mathrm{Aut}(T)$ and a word $w \in Y^*$ we denote by w^γ the image of w under γ. For a letter $y \in Y$ we have $(wy)^\gamma = w^\gamma y'$ where y' is a uniquely determined letter in Y. The map $y \mapsto y'$ is a permutation of Y, we refer to it as the *vertex permutation* of γ at w and denote it $(w)\gamma$.

Considering all the vertex permutations $\big((w)\gamma\big)_{w\in Y^*}$ gives us the *portrait* of γ, which is a decoration of the vertices of T with elements from the symmetric group S_d. In turn any assignment of these vertex permutations – that is, every possible portrait – gives an automorphism of T. Note that one has to perform these vertex permutations "from bottom to top".

An automorphism γ is finitary, if it has finitely many nontrivial vertex permutations. It is alternating, if all are from the alternating group A_d.

Let $S_d^{\mathrm{wr}(n)}$ denote the n-times iterated permutational wreath product of the symmetric group S_d. That is, let $[d] = \{1, \ldots, d\}$ and set

$$S_d^{\mathrm{wr}(n)} = \underbrace{((S_d \wr_{[d]} \ldots) \wr_{[d]} S_d) \wr_{[d]} S_d}_{n}.$$

Then $S_d^{\mathrm{wr}(n)}$ is isomorphic to the automorphism group of the d-ary rooted tree of depth n. These groups can be embedded in $\mathrm{Aut}(T)$ as acting on the first n levels. The group Aut_f is the union of these embedded finite groups. The full automorphism group $\mathrm{Aut}(T)$ however is isomorphic to the projective limit $\varprojlim S_d^{\mathrm{wr}(n)}$ with the projections being the natural restrictions of the permutations.

The groups $\mathrm{Alt}_f(T)$ and $\mathrm{Alt}(T)$ are in a similar relationship with the finite groups $A_d^{\mathrm{wr}(n)}$.

1.2.2 The boundary of T

The *boundary* of T is the set of infinite paths starting from o, and is denoted ∂T. For two distinct paths $p_1 = (u_0, u_1, \ldots)$ and $p_2 = (v_0, v_1, \ldots)$ with $u_n, v_n \in \mathcal{L}_n$ their distance is defined to be

$$d_{\partial T}(p_1, p_2) = \frac{1}{2^k}, \text{ where } k = \max\{n \mid u_n = v_n\}.$$

Two infinite paths are close it they have a long common initial segment. This distance turns ∂T into a compact, totally disconnected metric space.

The *shadow* of v on ∂T, denoted by $\mathrm{Sh}(v)$ is the set of paths passing through v. Similarly the shadow of v on $\mathcal{L}_n$ is the set $\mathrm{Sh}_{\mathcal{L}_n}(v)$ of descendants of v in $\mathcal{L}_n$. The sets $\mathrm{Sh}(v)$ form a basis for the topology of ∂T. Define the probability measure $\mu_{\partial T}$ by setting its value on this basis:

$$\mu_{\partial T}\big(\mathrm{Sh}(v)\big) = \frac{1}{d^n} \text{ for every } v \in \mathcal{L}_n.$$

A $\mu_{\partial T}$-random point of ∂T is a random infinite word $(w_1 w_2 \ldots)$ with each letter chosen uniformly from the set Y.

As $\gamma \in \mathrm{Aut}(T)$ permutes the vertices, it induces a bijection on ∂T, so we have an action of $\mathrm{Aut}(T)$ on ∂T. This action is by isometries and preserves the measure $\mu_{\partial T}$.

The objects in relation of the tree considered in here include vertices $v \in V(T)$, points $x \in \partial T$, closed subsets $C \subseteq \partial T$ and later 3-colorings of the vertices $\varphi : V(T) \to \{r, g, b\}$. For any such object z let z^γ denote its translate by γ.

1.2.3 Topology on $\mathrm{Aut}(T)$

We equip $\mathrm{Aut}(T)$ with the topology of pointwise convergence. This can be metrized by the following distance:

$$d_{\mathrm{Aut}(T)}(\gamma_1, \gamma_2) = \frac{1}{2^k}, \text{ where } k = \max\{n \mid \gamma_1|_{\mathcal{L}_n} = \gamma_2|_{\mathcal{L}_n}\}.$$

Two automorphisms are close if they act the same way on a deep level of T. This metric turns $\mathrm{Aut}(T)$ into a compact, totally disconnected group.

The action $\mathrm{Aut}(T) \curvearrowright \partial T$ is continuous in the first coordinate as well. For any subgroup $H \leq \mathrm{Aut}(T)$ the set $\mathrm{Fix}(H)$ is closed in ∂T, and similarly for any set $C \subseteq \partial T$ its pointwise stabilizer $\mathrm{Stab}_\Gamma(C)$ is closed in $\mathrm{Aut}(T)$.

For a subgroup $\Gamma \leq \mathrm{Aut}(T)$ its closure $\overline{\Gamma}$ is a closed subgroup of $\mathrm{Aut}(T)$, and therefore it is compact. We note that $\overline{\mathrm{Aut}_f(T)} = \mathrm{Aut}(T)$ and $\overline{\mathrm{Alt}_f(T)} = \mathrm{Alt}(T)$. Even though the groups Γ we are considering are discrete, their closures in $\mathrm{Aut}(T)$ always carry a unique Haar probability measure.

For any object z in relation to the tree we write $\widetilde{z}$ for its Haar random translate, that is z^γ where $\gamma \in \overline{\Gamma}$ is chosen randomly according to the Haar measure.

1.2.4 Fixed points and orbit-closures in ∂T

We aim to understand the IRS's of Γ through their actions on ∂T. The first step is to look at the set of fixed points. The boundary $(\partial T, d_{\partial T})$ is a compact metric space, so let $(\mathcal{C}, d_H)$ denote the compact space of closed subsets of ∂T with the Hausdorff metric.

Lemma 1.2.1. The map $H \mapsto \mathrm{Fix}(H)$ is a measurable and Γ-equivariant map from Sub_Γ to $(\mathcal{C}, d_H)$.

Equivariance is trivial, while the proof of the measurability is a standard argument. We postpone it to the Appendix.

Lemma 1.2.1 implies that the fixed points of the IRS constitute a Γ-invariant random closed subset of ∂T. We will also consider the orbit-closures of the subgroup on ∂T. For a subgroup $H \leq \mathrm{Aut}(T)$ let $\mathcal{O}_H$ denote the set of orbit-closures of the action $H \curvearrowright \partial T$. It is easy to see that $\mathcal{O}_H$ is a partition of ∂T into closed subsets. Note that all fixed points are orbit-closures. Denote by $\mathcal{O}$ the space of all possible orbit-closure partitions on ∂T, i.e. $\mathcal{O} = \{\mathcal{O}_H \mid H \leq \mathrm{Aut}(T)\}$. This $\mathcal{O}$ is a subset of all the possible partitions of ∂T.

As earlier, we would like to argue that the map $H \mapsto \mathcal{O}_H$ is a measurable map, with respect to the appropriate measurable structure on $\mathcal{O}$. This allows us to associate to our IRS a Γ-invariant random partition (into closed subsets) of ∂T. We will then analyze these invariant random objects on the boundary.

To this end we introduce a metric on the space $\mathcal{O}$. Denote by $\mathcal{O}_{H,\mathcal{L}_n}$ the partition of $\mathcal{L}_n$ into H-orbits. As $\mathcal{L}_n$ is finite, there is no need to take closure here.

Definition. Let $P = \mathcal{O}_H \in \mathcal{O}$ be the orbit-closure partition of H and $n \in \mathbb{N}$. Then let P_n be the orbit-structure of H on $\mathcal{L}_n$, i.e.

$$P_n = \mathcal{O}_{H,\mathcal{L}_n}.$$

For $P \neq Q \in \mathcal{O}$ let

$$d_{\mathcal{O}}(P,Q) = \min_{n \in \mathbb{N}} \left\{ \frac{1}{2^n} \mid P_n = Q_n \right\}.$$

Observe that if $P_n = Q_n$, then $P_{n-1} = Q_{n-1}$, so the above distance measures how deep one has to go in the tree to see that two partitions are distinct. This definition turns $(\mathcal{O}, d_{\mathcal{O}})$ into a metric space. To check that distinct points cannot have zero distance we argue that if $x = (v_0, v_1, \ldots)$ and $y = (u_0, u_1, \ldots)$ are two rays such that v_n and u_n are in the same orbit in $\mathcal{L}_n$ for all n, then y is indeed in the closure of the orbit of x.

The group $\mathrm{Aut}(T)$ acts on $\mathcal{O}$ in a natural way by shifting the sets of the partition. The resulting partition is again in $\mathcal{O}$ because $(\mathcal{O}_H)^\gamma = \mathcal{O}_{H^\gamma}$ for $\gamma \in \mathrm{Aut}(T)$.

Lemma 1.2.2. The map $H \mapsto \mathcal{O}_H$ is measurable and Γ-equivariant.

Again, equivariance is obvious, and measurability is proved in the Appendix.

1.2.5 Invariant random objects on ∂T

Now we study invariant random closed subsets and partitions on the boundary. We show that the invariance can be extended to $\overline{\Gamma}$, which carries a Haar measure. Ergodic objects turn out to be random translates according to this Haar measure.

Lemma 1.2.3. Every Γ-invariant random closed subset of ∂T is in fact $\overline{\Gamma}$-invariant. Similarly a Γ-invariant random $P \subseteq \mathcal{O}$ is $\overline{\Gamma}$-invariant.

Proof. Let $P(\mathcal{C})$ denote the set of probability measures on $\mathcal{C}$. The action of $\overline{\Gamma}$ on ∂T gives rise to a translation action on $(\mathcal{C}, d_H)$, which in turn gives rise to an action on $P(\mathcal{C})$.

We claim that this action $\overline{\Gamma} \times P(\mathcal{C}) \to P(\mathcal{C})$ is continuous in both coordinates with respect to the pointwise convergence topology on $\overline{\Gamma}$ and the weak star topology on $P(\mathcal{C})$.

The weak topology on $P(\mathcal{C})$ is metrizable by the Lévy - Prokhorov metric, which is defined as follows:

$$\pi(\mu, \nu) = \inf\{\varepsilon > 0 \mid \mu(A) \leq \nu(A^\varepsilon) + \varepsilon, \text{ and } \nu(A) \leq \mu(A^\varepsilon) + \varepsilon \text{ for all } A \subseteq \mathcal{C} \text{ Borel}\}.$$

Here A^ε denotes the set elements of $\mathcal{C}$ with d_H distance at most ε from A.

If γ_1 and γ_2 agree on the first n levels of T, then for every $x \in \partial T$ we have $d(\gamma_1 x, \gamma_2 x) \leq 1/2^n$. This implies, that for a compact set $C \in \mathcal{C}$ we have

$$d_H(\gamma_1 C, \gamma_2 C) \leq 1/2^n.$$

This in turn implies that for all $A \subseteq \mathcal{C}$ Borel we have $\gamma_1 A \subseteq \gamma_2 A^{1/2^n}$ and vice versa.

This means, that $\big((\gamma_1)_*\mu\big)(A) = \mu(\gamma_1^{-1}A) \leq \mu(\gamma_2^{-1}A^{1/2^n}) = \big((\gamma_2)_*\mu\big)(A^{1/2^n})$, so as a consequence $\pi\big((\gamma_1)_*\mu, (\gamma_2)_*\mu\big) \leq 1/2^n$. That is, the action is continuous in the first coordinate.

Continuity in the second coordinate is an easy exercise, as it turns out that the elements of $\overline{\Gamma}$ act by isometries on $(\partial T, d_{\partial T})$, $(\mathcal{C}, d_H)$ and $\big(P(\mathcal{C}), \pi\big)$ respectively.

As Γ is a dense subset of $\overline{\Gamma}$, by continuity we can say that if some $\mu \in P(\mathcal{C})$ is Γ-invariant then it is also $\overline{\Gamma}$-invariant, thus proving the statement for invariant random closed subsets.

The proof for invariant random partitions follows the exact same steps after substituting $(\mathcal{C}, d_H)$ with $(\mathcal{O}, d_{\mathcal{O}})$ everywhere. ■

Remark. The fact that the same lemma holds with the same proof for closed subsets and partitions is not a coincidence. A closed subset C can be thought of as a partition into the two sets C and C^c (the complement might not be closed). While it is not generally true that this partition is in $\mathcal{O}$ – it might not arise as an orbit-closure partition of some $H \subseteq \mathrm{Aut}(T)$ – but it still can be approximated on the finite levels. Indeed define C_n to be the set of vertices v on $\mathcal{L}_n$ with $\mathrm{Sh}(v) \cap C \neq \emptyset$. The C_n correspond to the $(1/2^n)$-neighborhoods of C in ∂T, and so the Hausdorff distance of closed subsets coincides with the distance $d_{\mathcal{O}}$ we could define using these C_n.

Lemma 1.2.4. Any ergodic $\overline{\Gamma}$-invariant random closed subset of ∂T is the γ translate of a fixed closed subset C, where $\gamma \in \overline{\Gamma}$ is a uniform random element chosen according to the Haar measure. Similarly an ergodic $\overline{\Gamma}$-invariant partition from $\mathcal{O}$ is the Haar-random translate of some fixed $P \in \mathcal{O}$.

Proof. We introduce an equivalence relation on closed subsets of ∂T: we say that $C_1 \sim C_2$ if and only if there is an automorphism $\gamma \in \overline{\Gamma}$ such that $C_1^\gamma = C_2$. Let $[C]$ denote the equivalence class of C.

Define the following metric on equivalence classes that measures how well one can overlap two arbitrary sets from the classes:

$$d([C_1],[C_2]) = \min_{\gamma \in \overline{\Gamma}} \{d_H(C_1^\gamma, C_2)\}.$$

The minimum exists by compactness of $\overline{\Gamma}$, and standard arguments using the fact that $\overline{\Gamma}$ acts by isometries on $(\mathcal{C}, d_H)$ show that this is well defined and indeed a metric.

The function $C \to [C]$ is measurable (in fact continuous) and $\overline{\Gamma}$-invariant, hence it is almost surely a constant by the ergodicity of the measure.

In other words the measure is concentrated on one equivalence class, say $[C]$. However $[C]$ is a homogeneous space of $\overline{\Gamma}$, i.e. the action of $\overline{\Gamma}$ on $[C]$ is the same as $\overline{\Gamma} \curvearrowright \overline{\Gamma}/\mathrm{Stab}_{\overline{\Gamma}}(C)$. $\mathrm{Stab}_{\overline{\Gamma}}(C)$ is a closed and therefore compact subgroup of $\overline{\Gamma}$, and as such $\overline{\Gamma}/\mathrm{Stab}_{\overline{\Gamma}}(C)$ carries a unique invariant measure. Of course picking a random translate of C is an invariant measure, so the two must coincide.

The result for partitions again follows word for word, by writing P for C and $(\mathcal{O}, d_{\mathcal{O}})$ for $(\mathcal{C}, d_H)$ everywhere. ∎

Remark. A way to put the previous lemmas into a general framework is the following: let G be a metrizable compact group acting continuously on a compact metric space (X, d) and let $\Gamma \leq G$ be a dense subgroup. Then any Γ-invariant measure on X is also G-invariant. Moreover if the metric d is G-compatible, then any ergodic Γ-invariant measure on X has the distribution of a random G-translate of a fix element in X.

If the IRS $H \leq \Gamma$ is ergodic, then so is the associated invariant random closed subset. This means that $\mathrm{Fix}(H)$ is the random translate of a fixed closed subset C. Similarly $\mathcal{O}_H$ is the random translate of some partition P.

1.2.6 Branch Groups

For a vertex v of T let T_v denote the induced subtree of T on v and its descendants. We denote by $\mathrm{Stab}_\Gamma(v)$ the stabilizer of v in Γ. Every $\gamma \in \mathrm{Stab}_\Gamma(v)$ acts on T_v by an automorphism, which we denote γ_v. Then $U_v = \{\gamma_v \mid \gamma \in \mathrm{Stab}_\Gamma(v)\}$ is a subgroup of $\mathrm{Aut}(T_v)$. U_v is the group of automorphisms of T_v that are realized by some element of Γ.

The trees we are considering are regular, so T_v is canonically isomorphic to T. (The isomorphism preserves the ordering of the vertices on each level. If we think of the vertices as finite words over a fixed alphabet, then this isomorphism just deletes the initial segment of each word in T_v.) This identification of the trees allows us to compare the action of G on T to the action of U_v on T_v. In particular we say that Γ is a *self-similar* group, if U_v is equal to G for all $v \in V(T)$ (under the above identification of the trees they act on).

For a vertex $v \in V(T)$ let $\mathrm{Rst}_\Gamma(v)$ denote the *rigid stabilizer* of v, that is the subgroup of elements of Γ that fix every vertex except the descendants of v. Clearly $\mathrm{Rst}_\Gamma(v) \leq \mathrm{Stab}_\Gamma(v)$. For a subset of vertices $V \subseteq \mathcal{L}_n$ the rigid stabilizer of the set is $\mathrm{Rst}_\Gamma(V) = \prod_{v \in V} \mathrm{Rst}_\Gamma(v)$.

Throughout the chapter we will be able to prove statements in varying levels of generality, so we introduce several notions of branching. In all cases we assume Γ to be transitive on all levels. We say that Γ is *weakly branch*, if all rigid vertex stabilizers $\mathrm{Rst}_\Gamma(v)$ are nontrivial. We say that the group Γ is *branch*, if for all n the rigid level stabilizer $\mathrm{Rst}_\Gamma(\mathcal{L}_n)$ is a finite index subgroup of Γ. Finally we define *regular branch groups*.

Definition. Suppose the self-similar group Γ has a finite index subgroup K. The group K^d is a subgroup of $\mathrm{Aut}(T)$, each component acting independently on T_{v_i} where $\{v_1, \dots, v_d\}$ are the vertices on $\mathcal{L}_1$. We say that Γ is a *regular branch group* over K, if K contains K^d as a finite index subgroup.

In self-similar groups the action on any subtree T_v is the same as on T, however for some $v_1, v_2 \in \mathcal{L}_n$ we might not be able to move T_{v_1} and T_{v_2} independently. This independence (up to finite index) is what we required in the definition above. The following lemma is straightforward and we leave the proof to the reader.

Lemma 1.2.5. Having finite index and being a direct product remains to be true after taking closures:

1. Let $K \subseteq \Gamma$ be a subgroup of finite index. Then $\overline{K}$ is a finite index subgroup of $\overline{\Gamma}$;
2. $\underbrace{\overline{K \times \cdots \times K}}_{d} = \underbrace{\overline{K} \times \cdots \times \overline{K}}_{d}$.

1.3 Fixed points and orbit-closures of IRS's

In this section we prove Theorems 1.1.3 and 1.1.4. To a closed subset C on the boundary one can associate two natural subgroup of Γ, the pointwise stabilizer of C and the setwise stabilizer of C. The pointwise stabilizer gives us a big "zoo" of IRS's when we choose C as a $\overline{\Gamma}$-invariant closed subset, proving Theorem 1.1.3. The setwise stabilizer will play a key role in the proof of Theorem 1.1.4.

In order to investigate these stabilizers we introduce a coloring to encode C on the tree T. The coloring will help analyzing the Haar random translate $\widetilde{C}$.

1.3.1 Closed subsets of the boundary

To every closed subset of the boundary $C \subseteq \partial T$ we associate a vertex coloring $\varphi : V(T) \to \{r, g, b\}$ with 3 colors: red, green and blue. If a vertex has its shadow completely in C, then color it red. If it has its entire shadow in the complement of C, then color it blue. Otherwise color it green.

$$\varphi(v) = \begin{cases} r, & \text{if } \mathrm{Sh}(v) \subseteq C; \\ b, & \text{if } \mathrm{Sh}(v) \cap C = \emptyset; \\ g, & \text{otherwise.} \end{cases}$$

All descendants of a red vertex are red, and similarly all descendants of blue vertices are blue. On the other hand all ancestors of a green vertex are green.

C being clopen is equivalent to saying that after some level all vertices are either red or blue. So if C is not clopen, then there are green vertices on all the levels. Using König's lemma we see that there is an infinite ray with vertices colored green. This ray corresponds to a boundary point of C. As the complement of C is open, we get that every vertex on this infinite ray has a blue descendant.

Lemma 1.3.1. Let $\Gamma \subseteq \mathrm{Aut}(T)$ be a group of automorphisms that is transitive on every level. Let $\varphi : V(T) \to \{r, g, b\}$ be a vertex coloring with the colors red, green and blue, and suppose it satisfies the above properties, namely:

1. descendants of red and blue vertices are red and blue respectively;
2. ancestors of green vertices are green; (This formally follows from 1.)
3. there is an infinite ray $(u_0, u_1, \ldots)$ of green vertices such that for each u_i there exists some descendant of u_i which is blue.

Then φ has infinitely many Γ-translates.

Proof. The root u_0 is colored green. It has a blue descendant, say on the level n_1. We denote this blue descendant w_{n_1}. By the transitivity assumption there is some $\gamma_1 \in \Gamma$, such that $\gamma_1(w_{n_1}) = u_{n_1}$. Furthermore, u_{n_1} has a blue descendant, on some level n_2, we denote it w_{n_2}. We choose $\gamma_2 \in \Gamma$ such that $\gamma_1(w_{n_1}) = u_{n_1}$, and so on. One can easily check, that moving φ with the different γ_i yields different colorings. Indeed $\varphi^{\gamma_i}(u_{n_j}) = g$ for all $j < i$, and $\varphi^{\gamma_i}(u_{n_i}) = b$, and this shows that the φ^{γ_i} are pairwise distinct. See Figure 2. ■

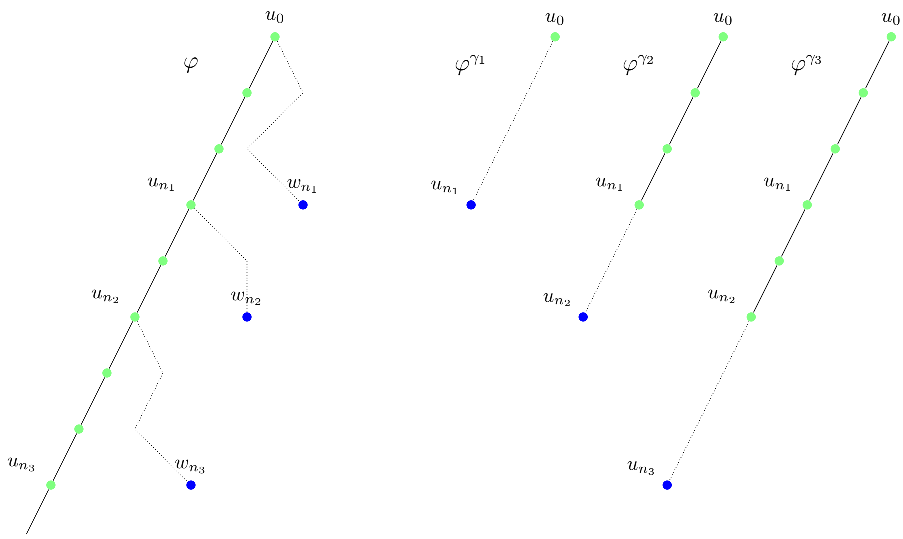

Figure 2: Distinct colorings

Corollary 1.3.2. Let Γ and φ be as in Lemma 1.3.1. Then the uniform (Haar) random $\overline{\Gamma}$-translate of φ is an atomless measure on the space of all 3-vertex-colorings.

Proof. If there was some translate φ^g, $g \in \overline{\Gamma}$ which occurred with positive probability, then all its Γ-translates would occur with the same positive probability. Furthermore φ^g would also satisfy the assumptions of Lemma 1.3.1, which then implies that it has infinitely many Γ-translates, and they would have infinite total measure, which is a contradiction. ■

Corollary 1.3.3. If C is not clopen, then its random $\overline{\Gamma}$-translate $\widetilde{C}$ is an atomless measure on $(\mathcal{C}, d_H)$.

1.3.2 Continuum many distinct atomless ergodic IRS's in weakly branch groups

Proof of Theorem 1.1.3. We argue that for any closed subset $C \subseteq \partial T$ the random subgroup $\mathrm{Stab}_\Gamma(\widetilde{C})$ is an ergodic IRS. This follows from $\widetilde{C}$ being an ergodic invariant random closed subset.

We also claim that if $[C_1] \neq [C_2]$, then the corresponding IRS's are distinct. To prove this we first observe that in weakly branch groups taking the stabilizer $\mathrm{Stab}_\Gamma(C)$ of a closed subset C, and then looking at the fixed points of that subset we get back C.

Lemma 1.3.4. For any $C \subseteq \partial T$ closed we have $\mathrm{Fix}\big(\mathrm{Stab}_\Gamma(C)\big) = C$.

Proof. The key idea – present in [BGN15, Proposition 8] and earlier works credited there – is to show, that for any $x \notin C$, with $x = (u_0, u_1, \dots)$ we can find some n large enough such that $\mathrm{Sh}(u_n) \cap C = \emptyset$, and some $\gamma \in \mathrm{Rst}_\Gamma(u_n)$ with $x^\gamma \neq x$.

Indeed such an n exists as the complement of C is open. By weak branching there exists some $\gamma_0 \in \mathrm{Rst}_\Gamma(u_n)$ moving some descendant of u denoted v to $v' \neq v$ on $\mathcal{L}_m$, $m \geq n$. By transitivity we can find some $\eta \in \mathrm{Stab}_\Gamma(u_n)$ with $v = u_m^\eta$. Now $u_m^{\eta\gamma_0\eta^{-1}} = (v')^{\eta^{-1}} \neq u_m$, so $\gamma = \eta\gamma_0\eta^{-1} \in \mathrm{Rst}_\Gamma(u_n)$, and $x^\gamma \neq x$ as witnessed on $\mathcal{L}_m$.

As $\mathrm{Sh}(u_n) \cap C = \emptyset$ we have $\mathrm{Rst}_\Gamma(u_n) \subseteq \mathrm{Stab}_\Gamma(C)$. The existence of γ shows that $x \notin \mathrm{Fix}\big(\mathrm{Stab}_\Gamma(C)\big)$, which implies $\mathrm{Fix}\big(\mathrm{Stab}_\Gamma(C)\big) \subseteq C$, which is the nontrivial inclusion. ■

To show that $[C_1] \neq [C_2]$ implies that $\mathrm{Stab}_\Gamma(\widetilde{C_1})$ and $\mathrm{Stab}_\Gamma(\widetilde{C_2})$ distinct simply consider the function $H \mapsto [\mathrm{Fix}(H)]$ on ergodic IRS's. Using Lemma 1.3.4 we have

$$\big[\mathrm{Fix}\big(\mathrm{Stab}_\Gamma(\widetilde{C_1})\big)\big] = [\widetilde{C_1}] = [C_1].$$

This implies that the constructed IRS are distinct. By Corollary 1.3.3 we know that if C is not clopen then $\widetilde{C}$ is atomless. Then Lemma 1.3.4 implies that $\mathrm{Stab}_\Gamma\widetilde{C}$ is an atomless IRS.

There are continuum many non-$\overline{\Gamma}$-equivalent closed (but not clopen) subsets of ∂T, as one can construct a closed subset C_r with $\mu_{\partial T}(C_r) = r$ for any $r \in [0, 1]$, and if r is irrational then C is not clopen. ■

1.3.3 Random colorings in regular branch groups

We proceed to prove a stronger versions of Corollary 1.3.2 for the case when the group is regular branch.

Let Γ be a regular branch group over K. Consider the finite index subgroup $\overline{K^{d^n}} \leq \mathrm{Stab}_{\overline{\Gamma}}(\mathcal{L}_n) = \cap_{v \in \mathcal{L}_n} \mathrm{Stab}_{\overline{\Gamma}}(v)$, and let $\{t_1, \dots, t_l\}$ be a transversal to $\overline{K^{d^n}}$ in $\overline{\Gamma}$. We can think of a random element γ of $\overline{\Gamma}$ as $\gamma = \gamma_0 \cdot k$, where γ_0 is chosen uniformly from the transversal and k is chosen according to the Haar measure on $\overline{K^{d^n}}$.

Take φ to be a 3-vertex-coloring as in Lemma 1.3.1. Let $\widetilde{\varphi} = \varphi^\gamma$ denote the translate of φ by the Haar random group element γ. Conditioning on $\gamma_0 = t_i$ we get a conditional distribution $(\widetilde{\varphi}|\gamma_0 = t_i)$. Note that this random coloring is always the same up to the n^{th} level, and t_i already determines where the random translate of the infinite green ray $(u_0, u_1, \dots)$ intersects $\mathcal{L}_n$, namely at $v = u_n^{t_i}$.

Lemma 1.3.5. The restriction of the random coloring $(\widetilde{\varphi}|\gamma_0 = t_i)$ to T_v is atomless.

Proof. The coloring φ^{t_i} restricted to T_v satisfies the assumptions of Lemma 1.3.1. Hence its Γ-orbit is infinite. As K is finite index in Γ, its K orbit is also infinite. Hence when it is randomly translated with an element from $\overline{K}$ (corresponding to v in $\overline{K}^{d^n}$) the resulting random coloring is atomless. ∎

Lemma 1.3.6. Fix any isomorphism $f : V(T_v) \to V(T_{v'})$ between T_v and $T_{v'}$ for some $v' \in \mathcal{L}_n$. Then the probability that f respects the colorings we get by restricting $(\widetilde{\varphi}|\gamma_0 = t_i)$ to T_v and $T_{v'}$ respectively is 0.

$$\mathbb{P}\Big[\big((\widetilde{\varphi}|\gamma_0 = t_i)|_{T_v}\big)^f = (\widetilde{\varphi}|\gamma_0 = t_i)|_{T_{v'}}\Big] = 0.$$

Proof. The restricted colorings $\varphi^{t_i}|_{T_v}$ and $\varphi^{t_i}|_{T_{v'}}$ are translated by the random elements $k_1, k_2 \in \overline{K}$ respectively. These k_1 and k_2 are independent since they are two coordinates of a Haar random element from $\overline{K}^{d^n}$. Furthermore we know from Lemma 1.3.5 that $(\varphi^{t_i}|_{T_v})^{k_1}$ is atomless, and hence $\big((\varphi^{t_i}|_{T_v})^{k_1}\big)^f$ is also atomless. This together with the independence of k_1 and k_2 implies that

$$\mathbb{P}\Big[\big((\varphi^{t_i}|_{T_v})^{k_1}\big)^f = (\varphi^{t_i}|_{T_{v'}})^{k_2}\Big] = 0.$$

∎

1.3.4 Proof of Theorem 1.1.4

The idea of the proof is to show that taking the *setwise* stabilizer of a Haar random translate $\widetilde{C}$ of a closed but not clopen subset C has a fixed point in $\widetilde{C}$. With some considerations one can apply this to the orbit-closures of H, which are setwise stabilized by H.

Proposition 1.3.7. Let Γ be a countable regular branch group over K. Suppose C is a closed subset of ∂T. Consider the IRS $L \leq \Gamma$ obtained by taking the setwise stabilizer of $\widetilde{C}$, which is the uniform $\overline{\Gamma}$-translate of C. If C is not clopen, then L has a fixed point in $\widetilde{C}$ almost surely.

Proof. Associate the coloring $\varphi : V(T) \to \{r, g, b\}$ to C as before: vertices with shadows contained in C are colored red, vertices with shadows in the complement are colored blue, everything else is colored green. As automorphisms move the set C the coloring moves with it.

Choose a point $x_0 \in \partial T$ which is on the boundary of C, that is $x_0 \in C \setminus \mathrm{int}(C)$. Being a boundary point means that every vertex on the path $(u_0, u_1, \ldots)$ corresponding to x_0 is green, and we can find a blue vertex among the descendants of u_i for all i.

Let $\widetilde{C}$, $\widetilde{\varphi}$ and $\widetilde{x}_0$ denote the uniform random translates of C, φ and x_0 respectively.

Fix an element $\eta \in \Gamma$. We will study the probability that η stabilizes $\widetilde{C}$ and does not fix $\widetilde{x}_0$, and conclude that it is 0. If η stabilizes $\widetilde{C}$ then it preserves $\widetilde{\varphi}$.

First assume η is finitary, that is we can find a level n with vertices $\mathcal{L}_n = \{v_1, \ldots, v_{d^n}\}$ such that η moves the subtrees T_{v_i} hanging off the n^{th} level rigidly. The condition that $\widetilde{x}_0$ is moved has to be witnessed on $\mathcal{L}_n$. Assume $v_1, \ldots, v_l$ are moved by η and $v_{l+1}, \ldots, v_{d^n}$ are fixed.

Let us assume that $\widetilde{\varphi}$ is preserved by η, and the ray corresponding to $\widetilde{x}_0$ is moved by η. Then $\widetilde{x}_0 \cap \mathcal{L}_n = v_i$ for some $i \leq l$ with $v_j = \eta v_i \neq v_i$. Conditioning on this v_i we are looking for the probability that $(\widetilde{\varphi}|_{T_{v_i}})^\eta = \widetilde{\varphi}|_{T_{v_j}}$. However, as the colorings are uniform random translates on T_{v_i} and T_{v_j} respectively, the probability of the coloring appearing in the exact same way under two points is 0. In the case of $\Gamma = \mathrm{Alt}_f(T)$ this is an easy consequence of Corollary 1.3.2. There are finitely many choices of v_i, so the probability of moving $\widetilde{x}_0$ while stabilizing $\widetilde{C}$ is 0.

As Γ is countable this means that with probability 1 the whole setwise stabilizer of $\widetilde{C}$ fixes $\widetilde{x}_0$.

In the general case when Γ is a regular branch group we condition on $\gamma_0 = t_i$ as in Lemma 1.3.6, and with f the canonical isomorphism between T_{v_i} and T_{v_j} we conclude that the conditional probability of η preserving $(\widetilde{\varphi}|\gamma_0 = t_i)$ is 0. There are finitely many choices for t_i, so again we conclude that the probability of moving $\widetilde{x}_0$ while stabilizing $\widetilde{C}$ is 0.

When η is not finitary there are two points where the above argument fails:

i) the trees T_{v_i} are not moved rigidly;

ii) the n^{th} level might not witness that $\widetilde{x}_0$ is moved by η.

Notice however that i) is not a real problem as we have fixed η and this fixes an isomorphism between T_{v_i} and T_{v_j}. The full generality of Lemma 1.3.6 (with $f = \eta|_{T_{v_i} \to T_{v_j}}$) ensures that the probability of randomizing η-compatible colorings for T_{v_i} and T_{v_j} is 0 even if γ is not finitary.

To work our way around ii) we notice that the probability of $\widetilde{x}_0$ being moved by η but this not being witnessed on $\mathcal{L}_n$ tends to 0 as $n \to \infty$. The set of fixed points of η is the decreasing intersection of the shadows of its fixed points on the finite levels. So the probability of $\widetilde{x}_0$ being in the shadow of the fixed points of $\mathcal{L}_n$ but outside $\mathrm{Fix}(\eta)$ converges to 0. This means that repeating the argument for all $n \in \mathbb{N}$ we get $\mathbb{P}[\eta \text{ moves } \widetilde{x}_0, \text{ but preserves } \widetilde{\varphi}] = 0$. ■

Proof of Theorem 1.1.4. By ergodicity and Lemma 1.2.4 we know that there exists a $P \in \mathcal{O}$, such that $\widetilde{P}$ has the same distribution as $\mathcal{O}_H$. Let us choose a closed set C which is not a single point from the partition P. We aim to use Proposition 1.3.7 to conclude that C is clopen. For that we will couple H and $\widetilde{C}$ such that $H \leq L$ holds almost surely, where L is the setwise stabilizer IRS of $\widetilde{C}$. Then H moving all points of C implies the same for L, which then through Proposition 1.3.7 implies that C is clopen.

Let $X = (P, C) \in \mathcal{O} \times \mathcal{C}$. Consider the diagonal action of $\overline{\Gamma}$ on $\mathcal{O} \times \mathcal{C}$. Let $\widetilde{X}$ be the Haar random translate of X. This way we obtained that the first coordinate of $\widetilde{X}$ has the same distribution as $\mathcal{O}_H$, the second coordinate has the same distribution as $\widetilde{C}$, and the second coordinate is always a closed subset in the partition given by the first coordinate.

Now we use the transfer theorem (see Theorem 6.10. of [Kal02]) to obtain a random element C_H of $\mathcal{C}$, such that $(\mathcal{O}_H, C_H) \stackrel{d}{=} \widetilde{X}$. The first coordinate of $\widetilde{X}$ always contains the second, therefore $C_H \in \mathcal{O}_H$ and clearly $C_H \stackrel{d}{=} \widetilde{C}$. Choosing L to be the setwise stabilizer of C_H concludes the proof. ■

1.4 IRS's in regular branch groups

Our goal is to understand all IRS's H of Γ. Let $\widetilde{C} = \mathrm{Fix}(H)$. Lemmas 1.2.3 and 1.2.4 tell us that $\widetilde{C}$ is the γ translate of a fixed closed subset $C \subseteq \partial T$, where $\gamma \in \overline{\Gamma}$ is Haar random. First we exhibit some concrete examples which are worth to keep in mind and to motivate the decomposition of the tree in Subsection 1.4.2. We study the action of H on the parts in Subsection 1.4.3. The last two subsections contain the proof of the main theorem of this chapter.

1.4.1 Examples

We show a few examples to keep in mind. For simplicity let $d = 5$, and $\Gamma = \mathrm{Alt}_f(T)$. Recall that in this group the normal subgroups are the level stabilizers $\mathrm{Stab}_\Gamma(\mathcal{L}_n)$, and the quotients are the finite groups $A_d^{\mathrm{wr}(n)}$.

Example 1.4.1. Pick $n \in \mathbb{N}$, and a finite subgroup $L \leq A_d^{\mathrm{wr}(n)}$. Let $\widetilde{L}$ be the uniform random conjugate of L in $A_d^{\mathrm{wr}(n)}$, and H be the preimage of $\widetilde{L}$ under the quotient map, that is $H = \widetilde{L} \cdot \mathrm{Stab}_\Gamma(\mathcal{L}_n)$. Then H is an ergodic fixed point free IRS of Γ. Note that this construction also works if G is only eventually d-ary, i.e. vertices on the first few levels might have different number of children.

Theorem 1.1.1 states that all ergodic fixed point free IRS of $\mathrm{Alt}_f(T)$ are listed in Example 1.4.1. We give a very broad outline of the proof for this case in the hope that it makes the subsequent proof of the stronger Theorem 1.1.5 more transparent and motivates Proposition 1.4.6 that we state beforehand.

Outline of proof of Theorem 1.1.1. By Theorem 1.1.4 we know that an ergodic fixed point free IRS H has finitely many clopen orbit-closures on the boundary. A deep enough level $\mathcal{L}_{k_0}$ witnesses this partition into clopen sets, and H acts transitively on the different parts on each $\mathcal{L}_n$ with $n \geq k_0$.

This means that we can find fixed elements supported above some level $\mathcal{L}_k$ ($k \geq k_0$) generating the orbits on $\mathcal{L}_{k_0}$ that are in H with positive probability. In the finite groups $A_d^{\mathrm{wr}(n)}$ (with $n \geq k$) this property translates to having a fixed subgroup L containing many conjugates of fixed elements. One can show that if n is sufficiently large this implies L containing a whole level stabilizer $\mathrm{Stab}_{A_d^{\mathrm{wr}(n)}}(\mathcal{L}_m)$ for some $m \geq k$ which does not depend on the choice of n.

Using that $\mathrm{Alt}_f(T)$ is the union of the $A_d^{\mathrm{wr}(n)}$ with some additional analysis of ergodic components one can show that actually H contains $\mathrm{Stab}_{\mathrm{Alt}_f(T)}(\mathcal{L}_m)$ almost surely. ■

Example 1.4.2. Pick a random point $x \in \partial T$, this will be the single fixed point of the IRS H. Deleting the edges of the ray $(u_0, u_1, \ldots)$ corresponding to x from T we get infinitely many disjoint trees, where the roots u_n have degree 4, while the rest of the vertices have 5 children. Pick any fixed point free IRS for each of these trees as in example 1.4.1, randomize them independently and take their direct sum to be H. This construction works with other random fixed point sets instead of a single point as well.

Example 1.4.3. A modification of the previous example is the following. Let $x \in \partial T$ be random as before, and do the exact same thing for all the trees hanging of the ray $(u_0, u_1, \ldots)$ except for the first two, T_1 and T_2 rooted at u_0 and u_1 respectively. The finitary alternating automorphism groups of these trees are $\mathrm{Alt}_f(T) \wr A_4$. Now pick an (ergodic) fixed point free IRS of the finitary alternating and bi-root-preserving automorphism group of $T_1 \cup T_2$, which is $(\mathrm{Alt}_f(T) \wr A_4) \times (\mathrm{Alt}_f(T) \wr A_4)$, and use this to randomize H on $T_1 \cup T_2$. We will show that this is different from the previous examples. When we pick an IRS of $(\mathrm{Alt}_f(T) \wr A_4) \times (\mathrm{Alt}_f(T) \wr A_4)$ we pick some $n \in \mathbb{N}$, assume that the stabilizers of the n^{th} levels in T_1 and T_2 are in the IRS, and pick a random conjugate of some $L \leq (\mathrm{Alt}_f(T) \wr A_4) \times (\mathrm{Alt}_f(T) \wr A_4)$ to extend the stabilizer. If we pick for example $L = \{(\gamma, \gamma) \mid \gamma \in (\mathrm{Alt}_f(T) \wr A_4)\}$, then the IRS we construct will not be the direct product of IRS's on the two components, because the "top" parts of the subgroups are coupled together. Taking a random conjugate of L makes the coupling random as well, but nonetheless in every realization of H there is some nontrivial dependence between the actions of H on T_1 and T_2.

1.4.2 Decomposition of T

The set of fixed points $\widetilde{C}$ corresponds to a subtree $T_{\widetilde{C}}$, which is the union of all the rays corresponding to the points of $\widetilde{C}$. All elements of H fix all vertices of the tree $T_{\widetilde{C}}$, so understanding H requires us to focus on the rest of T.

We will decompose T according to the subtree $T_{\widetilde{C}}$. Note that the following decomposition is slightly different to the one in the introduction as it is easier to work with.

On $\mathcal{L}_n$ denote the set of fixed vertices $F_n = V(T_{\widetilde{C}}) \cap \mathcal{L}_n$. Remove all edges $E(T_{\widetilde{C}})$ from T, the remaining graph T' is a union of trees.

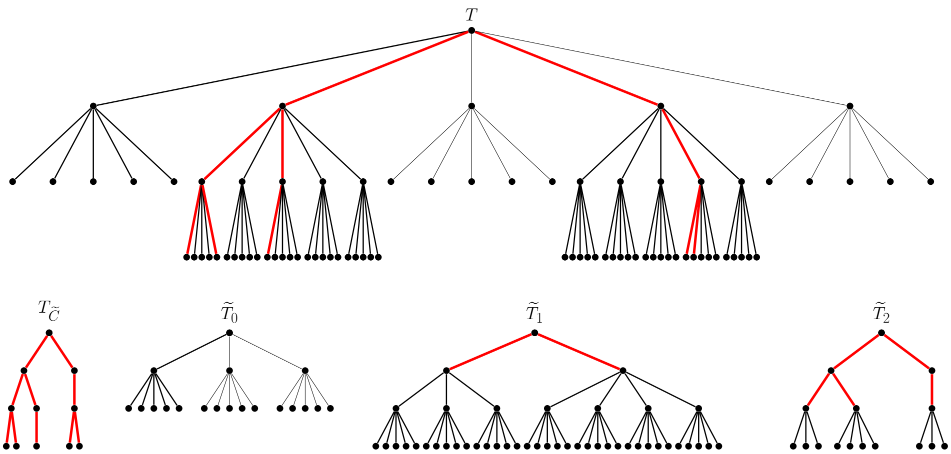

Figure 3: Decomposition of T with respect to $\widetilde{C}$

Let $\widetilde{T}_0$ be the connected component of T' containing the root of T. In other words it is the tree starting at the root in T'. In general let $\widetilde{T}_n$ be constructed as follows. The first n levels on $\widetilde{T}_n$ will be the same as the first n levels of $T_{\widetilde{C}}$, and beyond that select the connected components of T' containing the vertices of F_n. The vertices of $\widetilde{T}_n$ are exactly the vertices of T that can be reached from the root by taking n steps in $T_{\widetilde{C}}$ and then some number of steps in T'. See Figure 3.

The boundary ∂T decomposes as well. Clearly $\partial T_{\widetilde{C}} = \widetilde{C}$, and

$$\partial T = \widetilde{C} \cup \partial \widetilde{T}_0 \cup \partial \widetilde{T}_1 \cup \dots$$

Each $\partial \widetilde{T}_i$ is H-invariant, and a clopen and therefore compact subset of ∂T. It is the union of clopen orbit-closures from $\mathcal{O}_H$ because of Theorem 1.1.4, so it is the union of fintely many.

In the remaining part of this section we will prove that for any $C \in \mathcal{O}_H$ there exists some number $m^* \in \mathbb{N}$ and a subset $C_{m^*} \subseteq \mathcal{L}_{m^*}$ with $\mathrm{Sh}(C_{m^*}) = C$ such that $\mathrm{Rst}'_\Gamma(C_{m^*}) \leq H$. This m^* does not depend on the realization of $\mathcal{O}_H$, only on the equivalence class $[C]$.

Using this for the finitely many orbit-closures that constitute $\partial\widetilde{T}_i$ and taking a maximum yields that for some $m_i \geq i$ we have $\mathrm{Rst}'_\Gamma(\mathcal{L}_{m_i}(\widetilde{T}_i)) \subseteq H$. Knowing this for all i yields

$$\bigoplus_{i\in\mathbb{N}} \mathrm{Rst}'_\Gamma(\mathcal{L}_{m_i}(\widetilde{T}_i)) \subseteq H,$$

which is equivalent to the statement of Theorem 1.1.5.

1.4.3 The action of H on the $\widetilde{T}_i$

Before we turn to proving Theorem 1.1.5 we argue that all IRS's resemble the previous examples in the sense that their projections on the $\widetilde{T}_n$ are fixed point free IRS's in $\mathrm{Stab}_\Gamma(\widetilde{T}_n)$.

While the $\widetilde{T}_n$ are random, the isomorphism type of each $\widetilde{T}_n$ is always the same because of ergodicity, and $\widetilde{T}_n$ can appear in finitely many Γ-equivalent ways in T. Let $T_n^1, T_n^2, \dots T_n^{l(n)}$ denote the possible realizations of $\widetilde{T}_n$, and note that $\mathbb{P}[\widetilde{T}_n = T_n^i]$ is the same for all $i \in \{1, \dots, l(n)\}$.

Let $\varphi_n : H \to \mathrm{Stab}_\Gamma(\widetilde{T}_n)$ denote the restriction function:

$$\varphi_n(h) = h|_{\widetilde{T}_n}.$$

The function φ_n is also random, but it only depends on $\widetilde{T}_n$, so once we condition H on $\widetilde{T}_n$ the function φ_n is well defined.

Proposition 1.4.4. The random subgroup $\varphi_n\big((H \mid \widetilde{T}_n = T_n^i)\big)$ is a fixed point free IRS in $\mathrm{Stab}_\Gamma(T_n^i)$.

Proof. For a fixed subgroup $L \leq \Gamma$ let $T_n(L)$ denote the deterministic subtree defined the same way as $\widetilde{T}_n$ was for H. The set $\{L \leq \Gamma \mid T_n(L) = T_n^i\}$ is invariant under the conjugation action of $\mathrm{Stab}_\Gamma(T_n^i) \leq \Gamma$, so the invariance of the random subgroup H implies the invariance of the conditioned subgroup $(H \mid \widetilde{T}_n = T_n^i)$. This IRS is fixed point free because all fixed points of H are in $T_{\widetilde{C}}$. ■

Remark. One might be tempted to prove the more general Theorem 1.1.5 by first proving the more transparent fixed point free case and then using Proposition 1.4.4 on the individual subtrees, where H acts fixed point freely. However, we do not see this approach to work. Instead with some mild additional technical difficulties we present the proof for the more general case.

1.4.4 IRS's in finite subgroups of Γ

Let Γ_n stand for the elements of Γ that only have nontrivial vertex permutations above $\mathcal{L}_n$.

Lemma 1.4.5. For n large enough we have $[\Gamma_n : (K \cap \Gamma_n)] \leq [\Gamma : K]$.

Proof. Fix a transversal for K. All elements in the transversal are finitary, so choose n such that all are supported above $\mathcal{L}_n$. Then the translates of $(K \cap \Gamma_n)$ with this transversal cover Γ_n. ■

Let $\gamma \in \Gamma$, and $v \in \mathcal{L}_k$. The *section* of γ at v is the automorphism $[\gamma]_v$ we get by restricting the portrait of γ to the rooted subtree T_v consisting of v and its descendants. That is, the vertex permutations of $[\gamma]_v$ are $(u)\gamma$ for every $u \in T_v$ and the identity permutation otherwise. We think of $[\gamma]_v$ as the automorphism on T_v carried out by γ before all the vertex permutations above the level $\mathcal{L}_k$ take place.

Suppose $s \in \Gamma_k$, and let $L \subseteq \Gamma_n$ where $k < n$. Let $\widetilde{L}$ denote the uniform random Γ_n-conjugate of L, which is an IRS of Γ_n. Furthermore, assume that $\mathbb{P}[s \in \widetilde{L}] \geq c > 0$, which is equivalent to

$$\frac{\left|\{\gamma \in \Gamma_n \mid s^\gamma \in L\}\right|}{|\Gamma_n|} \geq c.$$

Let $R \subseteq \Gamma_n$ be a transversal for the subgroup $\mathrm{Rst}_{\Gamma_n}(\mathcal{L}_k)$. By choosing the optimal one, we can find $\bar{\gamma} \in R$ such that

$$\frac{\left|\left\{(\sigma_{v_1}, \ldots, \sigma_{v_{dk}}) \in \mathrm{Rst}_{\Gamma_n}(\mathcal{L}_k) \;\middle|\; s^{\bar{\gamma}(\sigma_{v_1}, \ldots, \sigma_{v_{dk}})} \in L\right\}\right|}{|\mathrm{Rst}_{\Gamma_n}(\mathcal{L}_k)|} \geq c. \tag{1.1}$$

Here $(\sigma_{v_1}, \ldots, \sigma_{v_{dk}})$ stands for the element of $\mathrm{Rst}_{\Gamma_n}(\mathcal{L}_k)$ that pointwise fixes $\mathcal{L}_k$, and has sections $\sigma_{v_i} \in \mathrm{Rst}_{\Gamma_n}(v_i)$ at the vertices $v_i \in \mathcal{L}_k$.

Let $\bar{s} = s^{\bar{\gamma}}$, and let the cycles of $\bar{s}$ on $\mathcal{L}_k$ be $C_1, \ldots, C_r$ and let $C_i = (u_1^i u_2^i \ldots u_{l(i)}^i)$, $l(i)$ denotes the length of the cycle C_i, and $\bar{s}(u_j^i) = u_{j+1}^i$. We use the convention that $u_{l(i)+1}^i = u_1^i$. Assume that $l(1) \geq l(2) \geq \ldots \geq l(r)$ and let t be the largest index for which $l(t) \geq 3$. Then $C = C_1 \cup \ldots \cup C_t \subseteq \mathcal{L}_k$ is the union of $\bar{s}$-orbits of length at least 3 on $\mathcal{L}_k$.

The next proposition shows that if n is large enough, then L has to contain the double commutator of some rigid level stabilizer under C, where the depth of this level does not depend on n.

Proposition 1.4.6. Let k, s and c be fixed. Then there exists some $m > k$ and $n_0 > m$ such that for any $n \geq n_0$, L and corresponding $\bar{\gamma}$ satisfying (1.1) above we have $\mathrm{Rst}''_{\Gamma_n}\big(\mathrm{Sh}_{\mathcal{L}_m}(C)\big) \subseteq L$.

Proof. Let $\sigma = (\sigma_{v_1}, \ldots, \sigma_{v_{dk}})$. Fix σ_{v_i} for all $v_i \notin C$, and let the rest of the coordinates $\sigma_{u_j^i}$ vary over $\mathrm{Rst}_{\Gamma_n}(u_j^i)$. Choosing a maximum over all choices of the fixed σ_{v_i} we can assume that

$$\frac{\left|\left\{(\sigma_{u_j^i})_{i,j=1}^{r,l(i)} \in \mathrm{Rst}_{\Gamma_n}(C) \;\middle|\; \bar{s}^\sigma \in L\right\}\right|}{|\mathrm{Rst}_{\Gamma_n}(C)|} \geq c.$$

Consider the conjugates $\bar{s}^\sigma$, more precisely what their sections are at the vertices u_j^i:

$$\left[\bar{s}^{(\sigma_{v_1}, \ldots, \sigma_{v_{dk}})}\right]_{u_j^i} = \sigma_{u_j^i} \cdot (\sigma_{u_{j+1}^i})^{-1}. \tag{1.2}$$

Fix one $\eta = (\eta_{v_1}, \ldots, \eta_{v_{dn-1}}) \in \mathrm{Rst}_{\Gamma_n}(\mathcal{L}_k)$ with $\eta_{v_i} = \sigma_{v_i}$ for all $v_i \notin C$ and $\bar{s}^\eta \in L$. Let $\sigma_{u_j^i}$ run through $\mathrm{Rst}_{\Gamma_n}(u_j^i)$, and consider $\bar{s}^\sigma \cdot (\bar{s}^\eta)^{-1}$. All these elements fix $\mathcal{L}_k$ pointwise, and their sections are

$$\left[\bar{s}^\sigma \cdot (\bar{s}^\eta)^{-1}\right]_{u_j^i} = \sigma_{u_j^i} \cdot (\sigma_{u_{j+1}^i})^{-1} \cdot \left(\eta_{u_j^i} \cdot (\eta_{u_{j+1}^i})^{-1}\right)^{-1}.$$

Observe that the sections are trivial over $v_i \notin C$.

We will discard one vertex from each C_i, and focus on the sections we see on the rest. Let $D_i = C_i \setminus \{u_1^i\}$.

Consider the sections of $\bar{s}^\sigma$ at the vertices in D_i as the sections $(\sigma_{u_1^i}, \dots, \sigma_{u_{l(i)}^i})$ run through $\mathrm{Rst}_{\Gamma_n}(C_i)$. We claim that the sections $\left([\bar{s}^\sigma]_{u_2^i}, \dots, [\bar{s}^\sigma]_{u_{l(i)}^i}\right)$ run through $\mathrm{Rst}_{\Gamma_n}(D_i)$.

Indeed, given any sections $\left([\bar{s}^\sigma]_{u_j^i}\right)_{j=2}^{l(i)}$, and any choice of $\sigma_{u_2^i}$ we can sequentially choose the $\sigma_{u_{j+1}^i}$ according to (1.2) to get the given sections at $j = 2, 3, \dots, l(i)$. The last choice is $\sigma_{u_1^i}$, which ensures $[\bar{s}^\sigma]_{u_{l(i)}^i}$ is correct. The last remaining section $[\bar{s}^\sigma]_{u_1^i}$ is already determined at this point, so we cannot hope to surject onto the whole $\mathrm{Rst}_{\Gamma_n}(C_i)$.

We can do this independently for each D_i. Let

$$D = \bigcup_i D_i.$$

The sections of $\bar{s}^\sigma$ over the index set D give $\mathrm{Rst}_{\Gamma_n}(D)$ as the sections $(\sigma_{u_j^i})$ run through $\mathrm{Rst}_{\Gamma_n}(C)$.

The fact that a fixed positive proportion of these conjugates are in L ensures that when we consider $\bar{s}^\sigma \cdot (\bar{s}^\eta)^{-1}$ we get that a fixed proportion of the elements of $\mathrm{Rst}_{\Gamma_n}(D)$ are seen in L_0, where $L_0 \subseteq \mathrm{Stab}_L(\mathcal{L}_k)$ is the set of elements with trivial sections outside C. Let $\pi_D : \mathrm{Stab}_L(\mathcal{L}_k) \to \Gamma_{n-k}^D$ denote the projection to the coordinates in D. Formally we get

$$\left|\pi_D(L_0)\right| \geq c \cdot |\mathrm{Rst}_{\Gamma_n}(D)|.$$

We have $\pi_D(L_0) \leq (\Gamma_{n-k})^{|D|}$. Since $(K \cap \Gamma_{n-k})^{|D|} \leq \mathrm{Rst}_{\Gamma_n}(D)$, using Lemma 1.4.5 we get that the index of $\pi_D(L_0)$ in $\Gamma_{n-k}^{|D|}$ is bounded:

$$\left[(\Gamma_{n-k})^{|D|} : \pi_D(L_0)\right] \leq \left\lceil \frac{1}{c} \right\rceil \cdot [\Gamma : K]^{|D|}.$$

This means we can find some $N \lhd (\Gamma_{n-k})^{|D|}$ such that $N \leq \pi_D(L_0)$ and

$$\left[(\Gamma_{n-k})^{|D|} : N\right] \leq \left(\left\lceil \frac{1}{c} \right\rceil \cdot [\Gamma : K]^{|D|}\right)!$$

The bound on the index of N does not depend on n, only on k, s and c. The bounded index ensures, that we can find some m_0 such that for each index $u \in D$ we can find an element $\varphi \in N$ such that $\pi_u(\varphi) \notin \mathrm{Stab}_{\Gamma_{n_k}}\left(\mathcal{L}_{m_0}(T_u)\right)$. Let $m = k + m_0$. Choose $n_0 > m$ such that Γ_{n_0-k} acts transitively on $\mathcal{L}_{m_0}(T_u)$.

Using Grigorchuk's standard argument from [BGŠ03, Lemma 5.3] and [Gri00, Theorem 4] we pick some $w \in \mathcal{L}_{m_0}(T_u)$ not fixed by φ, elements f and g from $\mathrm{Rst}_{\Gamma_n}(uw)$ and argue that the commutator $[[\varphi, f], g] = [f, g]$ is in N. This shows $\mathrm{Rst}'_{\Gamma_n}(uw) \subseteq N$. If $n \geq n_0$ then G_{n-k} is transitive on $\mathcal{L}_{m_0}(T_u)$, so we get

$$\mathrm{Rst}'_{\Gamma_n}\left(\mathcal{L}_{m_0}(T_u)\right) \subseteq N.$$

Repeating the argument of the previous paragraph for all $u \in D$ we get

$$\mathrm{Rst}'_{\Gamma_n}\left(\mathrm{Sh}_{\mathcal{L}_m}(D)\right) \subseteq N \subseteq \pi_D(L_0).$$

We now repeat this discussion, but we discard different points from the orbits: let $E_i = (C_i) \setminus \{u_2^i\}$ and $E = \bigcup_i E_i$. We have

$$\mathrm{Rst}'_{\Gamma_n}\big(\mathrm{Sh}_{\mathcal{L}_m}(E)\big) \subseteq \pi_E(L_0).$$

We claim that $\mathrm{Rst}''_{\Gamma_n}\big(\mathrm{Sh}_{\mathcal{L}_m}(D \cap E)\big) \subseteq L$. Indeed, let $u^i_j \in C_i$, $j \neq 1, 2$. By the above we see that for any $\varphi \in \mathrm{Rst}'_{\Gamma_n}\big(\mathrm{Sh}_{\mathcal{L}_m}(u^i_j)\big)$ we have $h_1 \in L_0$ such that $\pi_D(h_1)_{u^i_j} = \varphi$ and all other coordinates of $\pi_D(h_1)$ are the identity. Similarly we have $h_2 \in L_0$ such that $\pi_E(h_2)_{u^i_j} = \psi$ and all other coordinates of $\pi_E(h_1)$ are the identity. Since $\mathcal{L}_k \setminus D$ and $\mathcal{L}_k \setminus E$ are disjoint the commutator $[h_1, h_2] \in L_0$ has all identity coordinates except for the one corresponding to u^i_j which is $[\varphi, \psi]$.

We have managed to take care of the points u^i_j where $j \neq 1, 2$. To cover the remaining points as well we need one more way to discard points from the orbits. Namely F, where we discard the third vertex u^i_3 from every C_i. Using the fact that $(D \cap E) \cup (E \cap F) \cup (D \cap F) = C$ we get that $\mathrm{Rst}''_{\Gamma_n}\big(\mathrm{Sh}_{\mathcal{L}_m}(C)\big) \subseteq L$, which finishes the proof. ∎

1.4.5 Proof of the main result

Proof of Theorem 1.1.5. During the proof we will have to choose deeper and deeper levels in T. For the convenience of the reader we summarized these choices in Figure 4.

Let $\Gamma_n \subseteq \Gamma$ denote the elements of Γ that are supported on the first n levels. Suppose that H is an ergodic IRS of Γ.

By Theorem 1.1.4 we know that the all nontrivial orbit-closures of H on ∂T are clopen. For every clopen set C there exists a smallest integer k_C such that C is the union of shadows of points on $\mathcal{L}_{k_C}$. Clearly k_C does not change when C is translated by some automorphism. For the random subgroup H and a fixed $k_0 \in \mathbb{N}$ we can collect the clopen sets C from $\mathcal{O}_H$ with $k_C < k_0$, let C_{H,k_0} be the union of these. This set moves together with H when conjugating by some $\gamma \in \Gamma$:

$$C_{H^\gamma,k_0} = (C_{H,k_0})^\gamma.$$

For $n \geq k_0$ let $V_n \subset \mathcal{L}_n$ be the set of points whose shadow make up C_{H,k_0}. As C_{H,k_0} moves with H, so does V_n. V_{k_0} is a union of orbits of H, let those orbits be denoted $V^i_{k_0}$, where $i \in \{1, \ldots, j\}$ and

$$V_{k_0} = \bigcup_{i=1}^{j} V^i_{k_0}.$$

Let $V^i_n = \mathrm{Sh}_{\mathcal{L}_n}(V^i_{k_0})$. The fact that H acts minimally on the components of C_{H,k_0} translates to saying that H acts transitively on each V^i_n. Notice that since we collected clopen sets C with k_C strictly less then k_0 we ensured that $V^i_{k_0}$ contains at least d points for all i.

For every realization of H we can choose finitely many elements of H that already show that H acts transitively on the $V^i_{k_0}$. These finitely many elements are all finitary, so there is some n_H, which might depend on the realization of H, such that all those finitely many elements are in Γ_{n_H}.

This function n_H is not necessarily conjugation-invariant, so it need not be constant merely by ergodicity. However one can find some $k \geq k_0$ such that the $V^i_{k_0}$ are distinct orbits of $H_k = H \cap \Gamma_k$ on $\mathcal{L}_{k_0}$ with probability $1 - \varepsilon$. This k is a deterministic number, it does not depend on the realization of H.

Enlist all the possible subsets $S_1, \dots, S_N$ of Γ_k that generate a realization of the $V^i_{k_0}$ as orbits on $\mathcal{L}_{k_0}$. Clearly there are finitely many. The probability that $S_i \subseteq H$ cannot always be 0, otherwise we would contradict the previous paragraph. So we can find some finite set S of elements of Γ_k and some sets $U^i_{k_0} \subseteq \mathcal{L}_{k_0}$ such that the $U^i_{k_0}$ are a realization of the $V^i_{k_0}$, S is in H with probability $p > 0$ and the $U^i_{k_0}$ are orbits of S.

By replacing S with $\langle S \rangle$ we may assume that S is a subgroup of Γ_k, as $S \subseteq H$ and $\langle S \rangle \subseteq H$ are the same events.

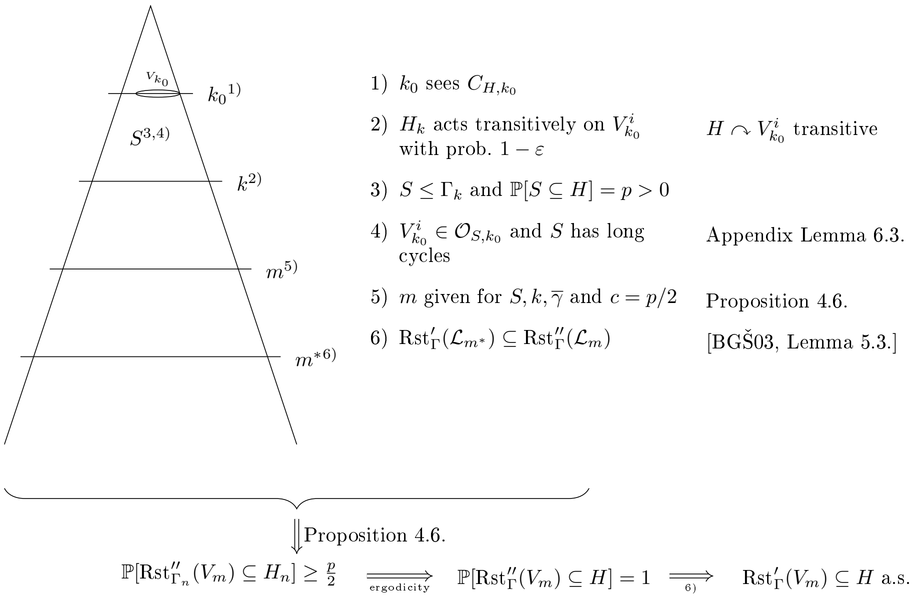

Figure 4: Choice of levels

As $|U^i_{k_0}| \ge d$, we know that all vertices of U_{k_0} are moved by some $s \in S$. As a consequence the same holds for U_k: for every vertex $v \in U_k$ there is some $s \in S$ such that $v \ne v^s$. However, we will need a stronger technical assumption on S to make our argument work. We will assume that for every $v \in U_k$ we can find some $s \in S$ such that v, v^s and v^{s^2} are distinct, that is v is part of a cycle of length at least 3 in the cycle decomposition of s. In Lemma 1.6.3 in the Appendix we show that one can indeed find such a k and S.

Let $H_n = H \cap \Gamma_n$, for $n \ge k$. The random subgroup H_n is clearly an IRS of Γ_n, however it need not be ergodic, i.e. the uniform random conjugate of a fixed subgroup in Γ_n. As $S \le \Gamma_k \le \Gamma_n$ we have $\mathbb{P}[S \le H_n] = p$.

Lemma 1.4.7. In the ergodic decomposition of H_n the measure of components that contain S with probability at least $p/2$ is at least $p/2$.

Proof. Denote the ergodic components of H_n by $H_n^1, \dots, H_n^r$. Assume H_n^i has weight q_i in the decomposition, and contains S with probability p_i. By ordering appropriately we can also assume $p_1, \dots, p_l \le p/2$ and $p_{l+1}, \dots p_r < p/2$.

$$p = \sum_{i=1}^r q_i p_i = \left(\sum_{i=1}^l q_i\right) \cdot 1 + \left(\sum_{i=l+1}^r q_i\right) \cdot \frac{p}{2} \le \left(\sum_{i=1}^l q_i\right) + \frac{p}{2},$$

$$\frac{p}{2} \le \sum_{i=1}^l q_i.$$

So the weight of components containing S with probability at least $p/2$ is at least $p/2$. ■

Choose an ergodic component of H_n which contains S with probability at least $p/2$. This ergodic component is the uniform random conjugate of a fixed subgroup $L \leq \Gamma_n$.

We have $\mathbb{P}[S \in \widetilde{L}] \geq p/2 > 0$. In other words L contains at least a $p/2$ proportion of the Γ_n-conjugates of S. By a "maximum is at least as large as the average" argument we can find some $\bar{\gamma}$ from the transversal of $\mathrm{Rst}_{\Gamma_n}(\mathcal{L}_k)$ such that

$$\frac{\left|\left\{(\sigma_{v_1}, \ldots, \sigma_{v_{d^k}}) \in \mathrm{Rst}_{\Gamma_n}(\mathcal{L}_k) \;\middle|\; S^{\bar{\gamma}(\sigma_{v_1}, \ldots, \sigma_{v_{d^k}})} \in L\right\}\right|}{|\mathrm{Rst}_{\Gamma_n}(\mathcal{L}_k)|} \geq \frac{p}{2}.$$

We now use Proposition 1.4.6 for all $s \in S$ with k, $\bar{\gamma}$ defined above and $c = p/2$. As the cycles of length at least 3 of elements of $S^{\bar{\gamma}}$ cover $(U_k)^{\bar{\gamma}}$ we get that for some fixed m and large enough n we have

$$\mathrm{Rst}''_{\Gamma_n}\big((U_m)^{\bar{\gamma}}\big) \subseteq L.$$

It is clear that $(U_{k_0})^{\gamma}$ is the realization of V_{k_0} corresponding to the realization L of H_n, so we (almost surely) have $\mathrm{Rst}''_{\Gamma_n}(V_m) \subseteq \widetilde{L}$. By Lemma 1.4.7 this means that

$$\mathbb{P}\big[\mathrm{Rst}''_{\Gamma_n}(V_m) \subseteq H_n\big] \geq \frac{p}{2}.$$

As $\big(\mathrm{Rst}''_{\Gamma_n}(V_m) \subseteq H_n\big) \Leftrightarrow \big(\mathrm{Rst}''_{\Gamma_n}(V_m) \subseteq H\big)$ we have

$$\mathbb{P}\big[\mathrm{Rst}''_{\Gamma_n}(V_m) \subseteq H\big] \geq \frac{p}{2}.$$

We get this for all n large enough. Since $\mathrm{Rst}''_{\Gamma_n}(V_m) \subseteq \mathrm{Rst}''_{\Gamma_{n+1}}(V_m)$ the events in question form a decreasing chain, and for the intersection we get

$$\mathbb{P}\big[\mathrm{Rst}''_{\Gamma}(V_m) \subseteq H\big] \geq \frac{p}{2}.$$

As H is ergodic the above implies

$$\mathbb{P}\big[\mathrm{Rst}''_{\Gamma}(V_m) \subseteq H\big] = 1.$$

Clearly $\mathrm{Rst}''_{\Gamma}(\mathcal{L}_m) \lhd \Gamma$, so using [BGŠ03, Lemma 5.3] we can find some $m^* \geq m$ such that $\mathrm{Rst}'_{\Gamma}(\mathcal{L}_{m^*}) \subseteq \mathrm{Rst}''_{\Gamma}(\mathcal{L}_m)$. This also means that $\mathrm{Rst}'_{\Gamma}(V_{m^*}) \subseteq \mathrm{Rst}''_{\Gamma}(V_m)$, so

$$\mathbb{P}\big[\mathrm{Rst}'_{\Gamma}(V_{m^*}) \subseteq H\big] = 1.$$

The number m^* only depended on the IRS H and the choice of k_0. Repeating this argument for all $k_0 \in \mathbb{N}$ covers all clopen sets from $\mathcal{O}_H$, which as discussed in part 1.4.2 proves Theorem 1.1.5. ■

1.5 Corollaries of Theorem 1.1.5

In this section we prove Theorem 1.1.6 and sketch the proof of Theorem 1.1.7.

1.5.1 Fixed point free IRS's

To motivate the following result let us recall Theorem 1.1.1, which states that any ergodic IRS of $\mathrm{Alt}_f(T)$ with $d \geq 5$ contains a whole level stabilizer, in particular H is a random conjugate of a finite indexed subgroup. In other words the measure defining the IRS is atomic. As it turns out the fixed point free case of Theorem 1.1.5 implies this for fixed point free ergodic IRS's of countable, finitary regular branch groups as well.

Proof of Theorem 1.1.6. By Theorem 1.1.5 we know that an ergodic almost surely fixed point free IRS H contains $\mathrm{Rst}'_\Gamma(\mathcal{L}_m)$ for some $m \in \mathbb{N}$.

IRS's of Γ containing the normal subgroup $\mathrm{Rst}'_\Gamma(\mathcal{L}_m)$ are in one-to-one correspondence with IRS's of the quotient $G = \Gamma/\mathrm{Rst}'_\Gamma(\mathcal{L}_m)$, which in this case is of the form $A \rtimes F$ where A is the abelian group $\mathrm{Rst}_\Gamma(\mathcal{L}_m)/\mathrm{Rst}'_\Gamma(\mathcal{L}_m)$, and F is the finite group $\Gamma/\mathrm{Rst}_\Gamma(\mathcal{L}_m)$. As Γ is assumed to finitary both Γ and G are countable.

Let $\widehat{H} = H/\mathrm{Rst}'_\Gamma(\mathcal{L}_m) \leq G$ be the image of H in G. It is an ergodic IRS of G. Let $\widehat{H}_0 = \widehat{H} \cap A$, which is also an ergodic IRS of G. We see that $\widehat{H}_0 \subseteq A$ is an ergodic random subgroup with distribution invariant under conjugation by elements of G. As A is abelian and F finite, it is clearly the uniform random F-conjugate of some subgroup $L_0 \leq A$. This shows that $\widehat{H}_0$ can only obtain finitely many possible values.

We claim that once $\widehat{H}_0$ is fixed, there are only countably many possible choices for $\widehat{H}$. Indeed we have to choose a coset of $\widehat{H}_0$ in G for all $f \in F$, which can do in only countably many different ways.

This shows that the support of $\widehat{H}$ is countable, but there is no ergodic invariant measure on a countably infinite set, so the support is finite. ■

1.5.2 IRS's in non-finitary branch groups

In this subsection we will sketch the proof of Theorem 1.1.7. This theorem is not a direct consequence (as far as we see) of Theorem 1.1.5, but one can alter the proof to obtain the desired theorem. First of all let us fix $\pi_n : \Gamma \to S_d^{\mathrm{wr}(n)}$ to be the projection from Γ to the automorphism group of the d-ary tree of depth n, which is the restriction of elements to the first n levels. The main conceptional difference is that we are trying to understand the group Γ through the groups $\pi_n(\Gamma)$ instead of Γ_n. The statement that we conclude in this case is weaker.

Our aim is to present only the spine of the proof, as the reasoning is very similar to the proof of Theorem 1.1.5 and we leave the details to the reader. In fact some technical details such as the ergodicity of H_n and the fact that H_n already acts transitively on the V_n^i makes this proof easier.

Proof of Theorem 1.1.7. Let $G_n = \pi_n(\Gamma)$, fix $k \in \mathbb{N}$ and let $C_{H,k}$ be the union of clopen orbit-closures C from $\mathcal{O}_H$ in ∂T with $k_C < k$. For any $n \geq k$ let $V_n \subseteq \mathcal{L}_n$ be the set of points whose shadow make up $C_{H,k}$. We can decompose V_k into H-orbits, denoted by

$$V_k = \cup_{i=1}^{j} V_k^i.$$

Observe that for any realization of H one can find at most $|V_k|$ many elements in H that already show that H acts transitively on each V_k^i. This means that we can find an $S \subset \Gamma$ of size at most $|V_k|$, such that S generates a realization of V_k on $\mathcal{L}_k$ and

$$\mathbb{P}[S \subseteq H] = p > 0.$$

Denote by U_k^i the realization of V_k^i generated by S. As before we can ensure that for any $v \in U_k$ there is an $s \in S$, such that v, v^s and v^{s^2} are distinct by replacing k and S if necessary. (See the Remark after the proof of Lemma 1.6.3 in the Appendix.)

For every n let $H_n = \pi_n(H) \leq G_n$. The random subgroup H_n is an ergodic IRS of G_n, therefore there exists an $L_n \leq G_n$ such that H_n is an uniform random conjugate of L_n. Since

$$\mathbb{P}[\pi_n(S) \subseteq \widetilde{L_n}] = \mathbb{P}[\pi_n(S) \subseteq H_n] \geq p,$$

we have an element $\overline{\gamma}$ from the transversal of $\mathrm{Rst}_{G_n}(\mathcal{L}_k)$ in G_n such that

$$\frac{\left|\left\{(\sigma_{v_1}, \ldots, \sigma_{v_{d^k}}) \in \mathrm{Rst}_{G_n}(\mathcal{L}_k) \;\middle|\; \pi_n(S)^{\bar{\gamma}(\sigma_{v_1}, \ldots, \sigma_{v_{d^k}})} \in L_n\right\}\right|}{|\mathrm{Rst}_{G_n}(\mathcal{L}_k)|} \geq p.$$

By following the argument in Proposition 1.4.6 but replacing Γ_n by G_n one can prove that there exists some m such that for any n large enough

$$\mathrm{Rst}''_{G_n}((U_m)^{\bar{\gamma}}) \subseteq L_n.$$

Therefore

$$\mathbb{P}[\mathrm{Rst}''_{G_n}(V_m) \subseteq H_n] \geq p > 0,$$

which by ergodicity implies

$$\mathbb{P}[\mathrm{Rst}''_{G_n}(V_m) \subseteq H_n] = 1.$$

Again we can find an $m^* \geq m$, such that $\mathrm{Rst}'_{\Gamma}(V_{m^*}) \subseteq \mathrm{Rst}''_{\Gamma}(V_m)$, therefore

$$\mathbb{P}[\pi_n(\mathrm{Rst}'_{\Gamma}(V_{m^*})) \subseteq \pi_n(H)] = 1.$$

This means that for any $g \in \mathrm{Rst}'_{\Gamma}(V_{m^*})$ there exists a sequence $h_n \in H$, such that $\pi_n(h_n) = \pi_n(g)$, which implies that $\overline{\mathrm{Rst}'_{\Gamma}(V_{m^*})} \subseteq \overline{H}$ with probability 1.

On the other hand $\overline{\mathrm{Rst}'_{\Gamma}(V_{m^*})} \supseteq \overline{\mathrm{Rst}_{\Gamma}(V_{m^*})}'$. We claim that $\mathrm{Rst}_{\overline{\Gamma}}(V_{m^*}) = \overline{\mathrm{Rst}_{\Gamma}(V_{m^*})}$. Indeed, $\mathrm{Rst}_{\overline{\Gamma}}(\mathcal{L}_{m^*})$ is finite index in $\overline{\Gamma}$ which implies that it is open. Using this one can show that $\mathrm{Rst}_{\overline{\Gamma}}(\mathcal{L}_{m^*}) = \overline{\mathrm{Rst}_{\Gamma}(\mathcal{L}_{m^*})}$, which implies the same for $V_{m^*} \subseteq \mathcal{L}_{m^*}$.

Putting this together we get

$$\mathrm{Rst}'_{\overline{\Gamma}}(V_{m^*}) \subseteq \overline{H}$$

with probability 1. ■

Note that this result on closures is possibly weaker than our earlier results. It is not clear even in the fixed point free case in $\mathrm{Alt}_f(T)$ if for some $L \leq \mathrm{Alt}_f(T)$ the closure $\overline{L}$ containing a level stabilizer implies the same for L.

Problem 1.5.1. Let $L \leq \mathrm{Alt}_f(T)$ be a subgroup such that $\pi_n(L) = A_d^{\mathrm{wr}(n)}$ for all n. Does it follow that $L = \mathrm{Alt}_f(T)$?

In other words: is there a subgroup $L \neq \mathrm{Alt}_f(T)$ which is dense in $\mathrm{Alt}(T)$? We saw that this cannot happen with positive probability when L is invariant random.

The answer to Problem 1.5.1 is negative in the case of $\mathrm{Aut}(T)$. In the case of the binary tree let L be the subgroup of elements with an even number of nontrivial vertex permutations. Generally for arbitrary d let L be the subgroup of elements whose vertex permutations multiply up to an alternating element. This L is not the whole group, yet dense in $\mathrm{Aut}(T)$. Of course the really relevant question in this case would involve the containment of derived subgroups of level stabilizers.

1.6 Appendix

In this section we prove the technical statements that we postponed during the rest of this chapter.

1.6.1 Measurability of maps

Proof of Lemma 1.2.1.

A closed subset C can be approximated on the finite levels. Define C_n to be the set of vertices v on $\mathcal{L}_n$ with $\mathrm{Sh}(v) \cap C \neq \emptyset$. The C_n correspond to the $(1/2^n)$-neighborhoods of C in ∂T.

We show that any preimage of a ball in $(\mathcal{C}, d_H)$ is measurable in Sub_Γ. Let $C \in \mathcal{C}$ and $n \in \mathbb{N}$ be fixed. Then the ball

$$B_{1/2^n}(C) = \{C' \in \mathcal{C} \mid C_n = C'_n\},$$

therefore its preimage is

$$X = \{H \in \mathrm{Sub}_\Gamma \mid \mathrm{Fix}(H)_n = C_n\}.$$

We say that a finite subset $S \subseteq \Gamma$ witnesses C_n, if the subgroup they generate has no fixed points in $\mathcal{L}_n \setminus C_n$. Clearly every C_n has a witness of cardinality at most $|\mathcal{L}_n \setminus C_n|$. Let W_{C_n} be the set of possible witnesses of C_n of size at most $|\mathcal{L}_n \setminus C_n|$:

$$W_{C_n} = \big\{S \subseteq \Gamma \;\big|\; |S| \leq |\mathcal{L}_n \setminus C_n| \text{ and } S \text{ has no fixed points in } \mathcal{L}_n \setminus C_n\big\}.$$

Let us define $F_{C_n} \subseteq \Gamma$ to be the set of forbidden group elements, which do not fix C_n. These are the elements that cannot be in any $H \in X$.

Observe that both W_{C_n} and F_{C_n} are countable, since Γ is countable, and X can be obtained as

$$X = \bigcup_{S \in W_{C_n}} \bigcap_{g \in F_{C_n}} \{H \in \mathrm{Sub}_\Gamma \mid S \subseteq H, g \notin H\}.$$

The sets $\{H \in \mathrm{Sub}_\Gamma \mid S \subseteq H, g \notin H\}$ are cylinder sets in the topology of Sub_Γ, so the above expression shows that X is measurable. ■

Proof of Lemma 1.2.2.

To prove that the map is measurable, it is enough to show that any preimage of a ball is measurable in Sub_Γ. So let $P \in \mathcal{O}$ and $n \in \mathbb{N}$ be fixed. Then the ball

$$B_{1/2^n}(P) = \{Q \in \mathcal{O} \mid Q_n = P_n\},$$

therefore its preimage is

$$X = \{H \in \mathrm{Sub}_\Gamma \mid (\mathcal{O}_H)_n = P_n\}.$$

We say that a finite subset $S \subseteq \Gamma$ witnesses P_n, if the subgroup they generate induces the same orbits on $\mathcal{L}_n$, that is $\mathcal{O}_{\langle S\rangle,\mathcal{L}_n} = P_n$. Clearly every P_n has a witness of cardinality at most $|\mathcal{L}_n|$. Let W_{P_n} be the set of possible witnesses of P_n of size at most $|\mathcal{L}_n|$:

$$W_{P_n} = \left\{S \subseteq \Gamma \;\middle|\; |S| \le |\mathcal{L}_n| \text{ and } \mathcal{O}_{\langle S\rangle,\mathcal{L}_n} = P_n\right\}.$$

Let us define $F_{P_n} \subseteq \Gamma$ to be the set of forbidden group elements, which do not preserve P_n. In other words these are the elements that cannot be in any $H \in X$.

Observe that both W_{P_n} and F_{P_n} are countable, since Γ is countable, and X can be obtained as

$$X = \bigcup_{S \in W_{P_n}} \bigcap_{g \in F_{P_n}} \{H \in \mathrm{Sub}_\Gamma \mid S \subseteq H, g \notin H\}.$$

The sets $\{H \in \mathrm{Sub}_\Gamma \mid S \subseteq H, g \notin H\}$ are cylinder sets in the topology of Sub_Γ, so the above expression shows that X is measurable. ■

1.6.2 Technical assumption in Theorem 1.1.5

First we prove a lemma on intersection probabilities.

Lemma 1.6.1. Let $B_1, \dots B_r$ be measurable subsets of the standard probability space (X, μ) with $\mu(B_j) = p$ for all j, and $r = \lceil 2/p \rceil$. Then there is some pair (j, l) such that $\mu(B_j \cap B_l) \ge p^3/6$.

Proof. Let χ_B denote the characteristic function of the measurable set B. Let D_l denote the set of points in X that are covered by at least l sets from $B_1, \dots B_r$. Then

$$\sum_{j=1}^{r} \chi_{B_j} = \sum_{l=1}^{r} \chi_{D_l},$$

$$\int_X \sum_{j=1}^{r} \chi_{B_j} \, d\mu = \sum_{j=1}^{r} \mu(B_j) = rp,$$

$$rp = \int_X \sum_{l=1}^{r} \chi_{D_l} \, d\mu = \sum_{l=1}^{r} \mu(D_l).$$

We have $D_1 \supseteq D_2 \dots \supseteq D_r$, so $1 \ge \mu(D_1) \ge \mu(D_2) \dots \ge \mu(D_r)$.

$$rp = \sum_{l=1}^{r} \mu(D_l) \le 1 + (r-1)\mu(D_2).$$

$$\mu(D_2) \ge \frac{rp - 1}{r - 1}.$$

The set D_2 is covered by the $B_j \cap B_l$, so

$$\max_{j,l} \mu(B_j \cap B_l) \geq \frac{\mu(D_2)}{\binom{r}{2}} \geq \frac{rp-1}{\binom{r}{2}(r-1)} \geq \frac{1}{\left(\frac{\left(\frac{2}{p}+1\right)\left(\frac{2}{p}\right)}{2}\right)\left(\frac{2}{p}\right)} \geq \frac{p^3}{2(p+2)} \geq \frac{p^3}{6}.$$

■

We also prove that one can find a lot of elements of order at least 3 in weakly branch groups.

Lemma 1.6.2. Let G be a weakly branch group. Then for any $v \in T$ there is a $g \in \mathrm{Rst}_G(v)$ of order at least 3.

Proof. As G is weakly branch, $\mathrm{Rst}_G(v)$ is not the trivial group. By contradiction let us assume that any nontrivial element of $\mathrm{Rst}_G(v)$ has order 2. Let a nontrivial element $g \in \mathrm{Rst}_G(v)$ be fixed. We can find descendants $u_1 \neq u_2$ of v, such that $u_1^g = u_2$. Let us choose a nontrivial element $h \in \mathrm{Rst}_G(u_1)$. As $h \in \mathrm{Rst}_G(v)$, by our assumption it has order 2.

We claim that $hg \in \mathrm{Rst}_G(v)$ has order at least three. To prove this we will find a vertex which has an orbit of size at least 3. Let $w_1 \neq w_2$ descendants of u_1, such that $w_1^h = w_2$. Since g maps descendants of u_1 to descendants of u_2, we have $w_1^{hg} = w_2^g = t_2 \neq w_1, w_2$. Then $w_1^{(hg)^2} = t_2^{hg} = t_2^g = w_2 \neq w_1$. We see that w_1, w_1^{hg} and $w_1^{(hg)^2}$ are pairwise distinct, therefore the order of hg is at least 3. ■

Now we will prove that the technical assumption we assumed in the proof of Theorem 1.1.5 can be satisfied. We remind the reader that in the setting of Theorem 1.1.5 the following were established:

(1) The random sets $(V_1^{k_0}, \dots V_j^{k_0})$ are orbits of H on $\mathcal{L}_{k_0}$;

(2) $\mathrm{Sh}(V_1^{k_0}), \dots, \mathrm{Sh}(V_j^{k_0})$ are orbit-closures of H on ∂T and their union is C_{H,k_0} almost surely;

(3) $S \leq \Gamma_k$ is a finite subgroups with $p = \mathbb{P}[S \subseteq H]$ positive;

(4) $(U_1^{k_0}, \dots U_j^{k_0})$ are a realization of $(V_1^{k_0}, \dots V_j^{k_0})$, and S acts transitively on the $U_i^{k_0}$.

(5) $V_i^n = \mathrm{Sh}_{\mathcal{L}_n}(V_i^{k_0})$ and $U_i^n = \mathrm{Sh}_{\mathcal{L}_n}(U_i^{k_0})$.

Lemma 1.6.3. By possibly replacing k, S and p we can assume that for every $u \in U_i^k$ we can find some $s \in S$ such that u, u^s and u^{s^2} are distinct.

Remark. In the case when d is not a power of 2 it can be shown that Lemma 1.6.3 is implied by the earlier properties, simply because a transitive permutation group with all nontrivial elements being fixed point free and of order 2 can only exist on 2^k points. For the case when d is a power of 2 however we can only show Lemma 1.6.3 by a probabilistic argument and by increasing k and S if necessary.

Proof. Assume that there is an $s \in S$ which admits a long cycle – that is a cycle of length at least 3 – on U_i^k for some i. In this first case we define k' such that $H_{k'}$ acts transitively on all the V_i^k with probability $1 - p/2$. Then

$$\mathbb{P}[S \subseteq H_{k'} \text{ and } H_{k'} \text{ is transitive on the } V_i^k] \geq \frac{p}{2} > 0.$$

If $S \subseteq H_{k'}$ then the V_i^k are realized as the U_i^k. Now we enlist all subsets S'_l in $\Gamma_{k'}$ that contain S and act transitively on the U_i^k. There are finitely many, so we can find some S' with

$$\mathbb{P}[S' \subseteq H_{k'}] \geq p' > 0.$$

We can assume S' to be a subgroup, and by having $s \in S'$ we will show that long cycles of S' cover U_i^k. Indeed, by conjugating s one can move the cycle around in U_i^k, and by the transitivity of S' we get that the whole of U_i^k is covered. This in turn implies that long cycles of S' cover $U_i^{k'}$ as well.

If on the other hand S acts on U_i^k by involutions, we will increase k and S while keeping p positive such that the first case holds.

Let $r = \lceil 2/p \rceil$. Furthermore let $k' > k$ such that the shadow of a vertex $v \in U_i^k$ on $\mathcal{L}_{k'}$ contains at least r vertices, namely $\{v_1, v_2, \ldots, v_r, \ldots\} \subseteq U_i^{k'}$. Let $\gamma_1, \gamma_2, \ldots, \gamma_r \in \Gamma_{k'+t}$ such that $\gamma_i \in \mathrm{Rst}_{\Gamma_{k'+t}}(v_i)$ and γ_i has order at least 3 by Lemma 1.6.2, i.e. γ_i has a long cycle on $\mathcal{L}_{k'+t}$.

As H is an IRS we have

$$\mathbb{P}[S^{\gamma_j} \subseteq H] = \mathbb{P}[S \subseteq H] = p.$$

By Lemma 1.6.1 we can find some j, l such that

$$\mathbb{P}\big[(S^{\gamma_j} \cup S^{\gamma_l}) \subseteq H\big] \geq \frac{p^3}{6}.$$

Set $S' = \langle S^{\gamma_j} \cup S^{\gamma_l} \rangle$. Pick some $s \in S$ which moves $v \in \mathcal{L}_k$. It is easy to check that $s^{\gamma_j} \cdot (s^{\gamma_l})^{-1} \in S' \cap \mathrm{Rst}_{\Gamma_{k'+t}}(\mathcal{L}_k)$ has nontrivial sections only at v_j, v_l, v_j^s and v_l^s, and these sections are some conjugates of γ_j and γ_l, and therefore $s^{\gamma_j} \cdot (s^{\gamma_l})^{-1}$ has a long cycle on $\mathcal{L}_{k'+t}$. So replacing S by S', k by $k' + t$ and p by $p^3/6$ we get to the first case.

Repeating the argument for the first case at most j times we make sure that all U_k^i are covered by long cycles, which finishes the proof. ∎

Remark. For the proof of Theorem 1.1.7 one can modify this proof such that instead of Γ_n, $H_k = \Gamma_n \cap H$ and $S \subseteq \Gamma_n$ we use $G_n = \pi_n(\Gamma)$, $H_n = \pi_n(H)$ and $\pi_n(S) \subseteq G_n$. Another difference is that H_n automatically acts transitively on all the V_i^n, so there is no need to distinguish between k_0 and k.

2

Uniform rank gradient, cost and local-global convergence

This chapter is based on the preprint [AT17]. It is joint work with Miklós Abért.

Abstract

We analyze the rank gradient of finitely generated groups with respect to sequences of subgroups of finite index that do not necessarily form a chain, by connecting it to the cost of p.m.p. actions. We generalize several results that were only known for chains before. The connection is made by the notion of local-global convergence. In particular, we show that for a finitely generated group Γ with fixed price c, every Farber sequence has rank gradient $c-1$. By adapting Lackenby's trichotomy theorem to this setting, we also show that in a finitely presented amenable group, every sequence of subgroups with index tending to infinity has vanishing rank gradient.

Contents

2.1 Introduction

For a finitely generated group Γ let $d(\Gamma)$ denote the minimal number of generators (or rank) of Γ. For a subgroup $H \leq \Gamma$ of finite index let

$$r(\Gamma, H) = (d(H) - 1)/ |\Gamma : H|.$$

The *rank gradient* of Γ with respect to a sequence (Γ_n) of finite index subgroups is defined to be

$$\mathrm{RG}(\Gamma, (\Gamma_n)) = \lim_{n\to\infty} r(\Gamma, \Gamma_n)$$

when this limit exists. This notion has been introduced by Lackenby [Lac05] and further investigated in the literature, mainly for *chains* of subgroups. Recall that a chain in Γ is a decreasing sequence $\Gamma = \Gamma_0 > \Gamma_1 > \ldots$ of subgroups of finite index in Γ. In this case, it is easy to see that $r(\Gamma, \Gamma_n)$ is non-increasing and so the limit exists.

The main goal of this chapter is to offer a general framework for understanding the rank gradient of an arbitrary sequence of subgroups in Γ using the cost of probability measure preserving (p.m.p.) actions of Γ. For chains this has been done by the first author and Nikolov in [AN12]. In arbitrary sequences were analyzed for a special class of groups called right angled groups.

Let (G_n) be a sequence of finite graphs with an absolute degree bound. We define the *edge density* as

$$e(G_n) = \lim_{n\to\infty} \frac{|E(G_n)|}{|V(G_n)|}$$

when this limit exists and the *lower edge density* $\underline{e}$ to be the lim inf of the same sequence. A *rewiring* of (G_n) is another sequence of graphs H_n on the same vertex set as of G_n, such that the bi-Lipshitz distortion of the maps $\mathrm{id}_{V(G_n)}$ stay bounded in n. The *combinatorial cost* $\mathrm{cc}(G_n)$ is defined as the infimum of the lower edge densities of possible rewirings of (G_n). This notion has been introduced by Elek [Ele07] as a discrete analogue of the notion of cost.

Our first result shows that actually combinatorial cost is more than an analogue and, when making an additional convergence assumption, it can be expressed as the cost of a limiting graphing.

Local-global convergence of graphs has been introduced by Bollobás and Riordan [BR11] under the name partition metric, while the limiting object and most of the known results were obtained by Hatami, Lovász and Szegedy [HLS14]. On the group theory side, the notion is related to the work of Kechris on weak containment, see [Kec10] and [AE12]. We postpone its definition to Section 2.2. For now, it suffices to know that every sequence has a convergent subsequence and that the limit is a graphing in the sense of Gaboriau [Gab00].

Theorem 2.1.1. Let G_n be a local-global convergent graph sequence. Then we have

$$\mathrm{cost}(\lim G_n) = \mathrm{cc}(G_n).$$

Moreover, one can choose rewirings such that the limit defining the edge density exists.

Note that similar results to Theorem 2.1.1 and its consequences (until Corollary 2.1.5) have been obtained independently by A. Carderi, D. Gaboriau and M. de la Salle [CGS17, p. 2017].

This result can be effectively used to give direct proofs of results on combinatorial cost using the already established theory of cost. For instance, Theorem 2.1.1 immediately implies the following theorem of Elek [Ele07].

Corollary 2.1.2. Let G_n be a graph sequence with girth tending to infinity such that $e(G_n)$ exists. Then we have

$$\mathrm{cc}(G_n) = e(G_n).$$

Indeed, by the girth assumption, any subsequential local-global limit of G_n will be a so called treeing and by Gaboriau [Gab00], the cost of a treeing equals its expected degree divided by two.

In Theorem 2.1.1, our graphs a priori have nothing to do with groups. When they do come from a sofic approximation of Γ, the limiting graphing gives rise to an essentially free probability measure preserving action of Γ, that is unique up to weak equivalence in the sense of Kechris [Kec10]. This case of local-global convergence has been analyzed by the first author and Elek [AE11].

Following Gaboriau, we say that Γ has fixed price c, if every essentially free probability measure preserving action of Γ has cost c. Applying Theorem 2.1.1 gives us the following new result.

Theorem 2.1.3. Let Γ be a finitely generated group of fixed price c. Then

$$\mathrm{cc}(G_n) = c$$

for any sofic approximation (G_n) of Γ.

This gives an alternate proof of another result of Elek [Ele07] that for an amenable group Γ, any sofic approximation of Γ has combinatorial cost 1. Indeed, by the Ornstein-Weiss theorem [OW80], amenable groups have fixed price 1.

A sequence of subgroups is *Farber*, if the quotient Schreier graphs $\mathrm{Sch}(\Gamma, \Gamma_n, S)$ form a sofic approximation of Γ. We can now connect the cost to the rank gradient as follows.

Theorem 2.1.4. Let Γ be a finitely generated group of fixed price c. Then we have

$$\mathrm{RG}(\Gamma, (\Gamma_n)) = c - 1$$

for any Farber sequence (Γ_n) in Γ.

The same result is proved in [AN12] for Farber *chains*. Also, in [AGN17] it is proved that any Farber sequence in a right angled group has rank gradient zero. Right angled groups have fixed price 1 by Gaboriau [Gab00], so this now immediately follows from Theorem 2.1.4. Since amenable groups also have fixed price 1, we get the following.

Corollary 2.1.5. Let Γ be a finitely generated amenable group. Then we have

$$\mathrm{RG}(\Gamma, (\Gamma_n)) = 0$$

for any Farber sequence (Γ_n) in Γ.

When the sequence is not Farber, Corollary 2.1.5 is clearly not true, already for the standard lamplighter group (see [AJN11]). However, one can show that it still holds for *finitely presented* amenable groups.

Theorem 2.1.6. Let Γ be a finitely presented amenable group. Then we have

$$\mathrm{RG}(\Gamma, (\Gamma_n)) = 0$$

for any sequence (Γ_n) of distinct subgroups in Γ.

Behind this is the following extension of [AJN11] that generalized a theorem of Lackenby for normal chains [Lac05].

We call a sequence of finite graphs G_n *dispersive* if for any subsequential local-global limit $\mathcal{G}$ of G_n, $\mathcal{G}$ has no strongly ergodic component of positive measure. For the notion of strong ergodicity and a graph theoretic reformulation see Section 2.5.

Theorem 2.1.7. Let Γ be a finitely presented group generated by a finite symmetric set S. Let (Γ_n) be an arbitrary sequence of subgroups of finite index in Γ. Then at least one of the following holds:

1) the sequence $\mathrm{Sch}(\Gamma, \Gamma_n, S)$ is not dispersive;

2) $\mathrm{RG}(\Gamma, (\Gamma_n)) = 0$;

3) there exists some n such that Γ_n decomposes as a non-trivial amalgamated product.

We show that sequences in amenable groups are dispersive, and clearly they cannot decompose as a non-trivial amalgamated product. Thus Theorem 2.1.6 follows as a corollary of Theorem 2.1.7. Note that for a chain of subgroups, being dispersive is equivalent to saying that the limiting profinite action is not strongly ergodic. Hence, Theorem 2.1.7 implies [AJN11, Theorem 3]. When the Γ_n are normal in Γ, being dispersive is equivalent to saying that (Γ_n) has no subsequence with Lubotzky's property (τ). So Theorem 2.1.7 also generalizes Lackenby's trichotomy theorem [Lac05, Theorem 1.1.].

The structure of the chapter is as follows. In Section 2.2 we define the basic notions and state some lemmas that we need for our main result. In Section 2.3 we prove Theorems 2.1.1 and 2.1.3. We introduce the analogous notions and results for group actions in Section 2.4 and prove Theorem 2.1.4. We prove the results on finitely presented groups in Section 2.5. Finally, in Section 2.6 we list some open problems and suggest further directions of research.

2.2 Preliminaries

In this section we define the basic objects of investigation of this chapter and state some known results.

2.2.1 Local-global convergence

Let U_r denote the set of connected, rooted graphs with radius at most r with all degrees bounded by some integer D. For any graph G, if we pick a vertex $v \in V(G)$ and look at its r-neighborhood $B_G(r, v)$ rooted at v we get an element of U_r. Picking v uniformly at random gives us a probability measure on U_r which we will denote $P_{G,r}$, and refer to as the r-neighborhood statistics of G.

For any finite set X let $M(X)$ denote the set of probability measures on X. We say that a sequence of graphs (G_n) is *locally* (or Benjamini-Schramm) convergent, if for any r the sequence of probability measures $P_{G_n,r} \in M(U_r)$ converge to a limit distribution as $n \to \infty$.

We will work with a more refined notion of convergence, and following notation from [HLS14] we introduce a colored version of the neighborhood statistics. Let $K(k,G) = \{\varphi : V(G) \to \{1,\dots,k\}\}$ denote the set of k-colorings of the vertices of G. Let U_r^k denote the set of rooted, connected, k-colored graphs of radius at most r. For any coloring $\varphi \in K(k,G)$ we can associate a colored neighborhood statistic $P_{G,r}[\varphi] \in M(U_r^k)$ as before, by choosing a uniform random vertex v, and then considering its r-neighborhood $B_G(r,v)$, this time together with the coloring $\varphi|_{B_G(r,v)}$.

For $\eta_1, \eta_2 \in M(U_r^k)$ let

$$d_{TV}(\eta_1,\eta_2) = \sup_{A \subseteq U_r^k} |\eta_1(A) - \eta_2(A)|.$$

Note that d_{TV} is the total variation distance. As we are operating in a finite dimensional space all the usual norms are equivalent.

Intuitively, a sequence of graphs (G_n) is local-global convergent if for any $r,k \in \mathbb{N}$ and for i,j large enough the colored neighborhood distribution $P_{G_i,r}[\varphi]$ for any k-coloring φ can be approximately modeled on G_j, that is we can find some coloring ψ such that $P_{G_i,r}[\varphi]$ and $P_{G_j,r}[\psi]$ are arbitrarily close.

For a finite graph G let $Q^k_{G,r}$ denote the finite set of possible colored neighborhood statistics arising from a graph G:

$$Q^k_{G,r} = \{P_{G,r}[\varphi] \mid \varphi \in K(k,G)\} \subseteq M(U_r^k).$$

Definition 2.2.1. We say that a sequence of graphs (G_n) is *local-global convergent* if for every $r,k \in \mathbb{N}$ the compact sets $(Q^k_{G_n,r})$ converge in the Hausdorff distance on $\big(M(U_r^k), d_{TV}\big)$.

In [HLS14] the authors show that every sequence of bounded degree graphs has a locally-globally convergent subsequence, and that *graphings* can be considered as the limit objects of convergent sequences.

Definition 2.2.2 ([HLS14] Definition 3.1)**.** Let X be a Polish topological space and let μ be a probability measure on the Borel sets in X. A *graphing* (with degree bound D) is a graph $\mathcal{G}$ on $V(\mathcal{G}) = X$ with Borel edge set $E(\mathcal{G}) \subset X \times X$ in which all degrees are at most D and

$$\int_A e(x,B)\, d\mu(x) = \int_B e(x,A)\, d\mu(x) \tag{2.1}$$

for all measurable sets $A, B \subseteq X$, where $e(x,S)$ is the number of edges from $x \in X$ to $S \subseteq X$.

Every finite graph G is a graphing with $X = V(G)$ and μ the uniform distribution on $V(G)$.

The colored neighborhood statistics $P_{\mathcal{G},r}[\varphi]$ can easily be defined for a graphing $\mathcal{G}$, provided that the coloring $\varphi : X \to \{1,\dots,k\}$ is chosen to be Borel. We pick a random vertex $x \in X$ according to μ, and consider its colored r-neighborhood in $\mathcal{G}$.

As opposed to the finite case we now have to take the closure of all possible such statistics in order to obtain a compact set. Let

$$Q^k_{\mathcal{G},r} = \overline{\{\mathcal{P}_{G,r}[\varphi] \mid \varphi : V(\mathcal{G}) \to \{1,\dots,k\} \text{ Borel}\}}^{d_{TV}} \subseteq M(U^k_r).$$

The graphing $\mathcal{G}$ is a local-global limit of the sequence (G_n) if $Q^k_{G_n,r} \to Q^k_{\mathcal{G},r}$ in the Hausdorff distance for all r and k.

We say that two graphings $\mathcal{G}$ and $\mathcal{H}$ are local-global equivalent, if the sets $Q^k_{\mathcal{G},r}$ and $Q^k_{\mathcal{H},r}$ are the same. Note that the limit is unique only up to local-global equivalence. Although we will only be dealing with sequences of finite graphs, observe that the above definition of convergence makes sense for sequences of graphings as well.

2.2.2 Cost

For a graphing $\mathcal{G}$ and a vertex $x \in X$ let $[x]_\mathcal{G}$ denote the connected component of x in $\mathcal{G}$. For two graphings $\mathcal{G}$ and $\mathcal{H}$ on the same vertex set X we write $\mathcal{G} \sim \mathcal{H}$ if they have the same connected components, that is $[x]_\mathcal{G} = [x]_\mathcal{H}$ μ-almost surely.

Let $\mathcal{R}_\mathcal{G} \subseteq X \times X$ denote the measurable equivalence relation generated by $\mathcal{G}$, where two points are in the same equivalence class if they are in the same connected component of $\mathcal{G}$. Clearly $\mathcal{G} \sim \mathcal{H}$ if and only if $\mathcal{R}_\mathcal{G} = \mathcal{R}_\mathcal{H}$ up to measure zero. Note that every component of $\mathcal{G}$ is countable.

We will introduce a way of measuring edge sets of graphings. Let $\tilde{\mu}$ be the measure on $X \times X$ obtained the following way. For a measurable subset $C \subseteq X \times X$ let

$$\tilde{\mu}(C) = \int_X \#\{y \mid (x,y) \in C, y \in [x]_\mathcal{G}\}\ d\mu(x).$$

In other words on each fiber $\{x\} \times X$ we consider the counting measure concentrated on $\{x\} \times [x]_\mathcal{G}$, and integrate these with respect to μ on the first coordinate.

This measure $\tilde{\mu}$ is σ-finite, it is concentrated on $\mathcal{R}_\mathcal{G}$, and it is easy to see that in fact it only depends on the relation $\mathcal{R}_\mathcal{G}$. We can similarly define $\tilde{\mu}'$ by taking the counting measures on the fibers $[x]_\mathcal{G}$ and integrate with respect to μ over the second coordinate. A standard argument shows that condition (2.1) in Definition 2.2.2 is equivalent to $\tilde{\mu} = \tilde{\mu}'$.

The *cost* of $\mathcal{G}$ is defined to be

$$\operatorname{cost}(\mathcal{G}) = \frac{1}{2} \inf \{\tilde{\mu}(E(\mathcal{H})) \mid \mathcal{H} \sim \mathcal{G}\}.$$

The normalization factor $\frac{1}{2}$ is included to account for counting every edge twice and to ensure coherence with [Gab00]. Note that the $\tilde{\mu}$ measure of the edge set of a graphing is half the expected degree of a μ-random point. It is clear that if $\mathcal{H} \sim \mathcal{G}$, then their cost is the same, in fact the cost only depends on $\mathcal{R}_\mathcal{G}$.

The following lemma proved in [Gab00] sheds some light on the bi-Lipschitz condition used in the definition of combinatorial cost. We will also use it in the proof of Theorem 2.1.1.

Lemma 2.2.3 (Gaboriau)**.** Let $\mathcal{G}$ be a graphing. For every $\varepsilon > 0$ there exists some integer L and some $\mathcal{H} \sim \mathcal{G}$ such that $\tilde{\mu}(E(\mathcal{H})) < \operatorname{cost}(\mathcal{G}) + \varepsilon$ and $\mathcal{G}$ and $\mathcal{H}$ are L-bi-Lipschitz equivalent, that is the graph metrics they define on the connected components are within a factor of L from each other.

2.2.3 Combinatorial cost

The combinatorial analogue of cost for sequences of graphs is due to Elek [Ele07]. Let (G_n) be a sequence of graphs with $|V(G_n)| \to \infty$, and degree bounded by D. The sequence (H_n) is a *rewiring* of (G_n) – which we will denote $(H_n) \sim (G_n)$ – if they have the same vertex set, and the distances defined by the graphs are uniformly bi-Lipschitz equvialent, that is $V(G_n) = V(H_n)$ and there exists some natural number L such that for all $n \in \mathbb{N}$

$$\frac{1}{L} d_{H_n}(x, y) \leq d_{G_n}(x, y) \leq L d_{H_n}(x, y) \text{ for all } x, y \in V(G_n).$$

The *lower edge density* of a graph sequence is defined as follows.

$$\underline{e}\big((H_n)\big) = \liminf_{n\to\infty} \frac{|E(H_n)|}{|V(H_n)|}.$$

Definition 2.2.4. The *combinatorial cost* of a sequence (G_n) is the infimum of the lower edge densities of its rewirings:

$$\operatorname{cc}\big((G_n)\big) = \inf\big\{\underline{e}\big((H_n)\big) \bigm| (H_n) \sim (G_n)\big\}.$$

2.3 The cost of a local-global limit

In this section we will prove Theorems 2.1.1 and 2.1.3.

2.3.1 Proof of the main result

We aim to show that the combinatorial cost of a locally-globally convergent graph sequence is equal to the cost of its limit. The idea of the proof is that if there is a cheap rewiring of the sequence (G_n), then we can encode it into a coloring which then can be modeled with small error on the limit graphing $\mathcal{G}$. Using this coloring on the limit we can reconstruct a cheap graphing that (after some small modification) spans the same connected components as $\mathcal{G}$. In order to make this reconstruction process possible we will need to break the possible local symmetries of the graphs.

Proof of Theorem 2.1.1. Fix $\varepsilon > 0$ and suppose that (H_n) is an L-rewiring of G_n such that $\underline{e}(H_n) < \operatorname{cc}(G_n) + \varepsilon$. Set $r = L^2 + 1$, $R = 2r$.

As the degrees of the G_n are bounded by D, there is a constant k such that each G_n can be vertex colored by k colors so that no two vertices within distance $2R$ have the same color. Fix such a coloring $\eta_n : V(G_n) \to \{1, \ldots, k\}$ for each G_n. The role of these η_n is merely to break all possible symmetries of the R-neighborhoods.

For each vertex $v \in V(G_n)$ define its *type* to be the following data. Let (α_v, η_v) denote the colored R-neighborhood of v in G_n, that is $B_{G_n}(R, v)$ rooted at v, together with $\eta_n|_{B_{G_n}(R,v)}$. Let F_v denote the set of edges of H_n that connect two vertices from $B_{G_n}(R, v)$. The type of v is the triple (α_v, η_v, F_v). We think of this as a rooted, vertex colored graph with some extra distinguished edges (F_v) indicated. Note that η_v assigns distinct colors to the vertices of α_v.

Let T denote the set of all possible types. Note that $|T|$ is finite, as k and R are fixed. Now assigning each vertex its type can be considered as a coloring of $V(G_n)$ by $|T|$ colors. Let φ_n denote this coloring:

$$\varphi_n : V(G_n) \to T, \quad \varphi_n(v) = (\alpha_v, \eta_v, F_v) \text{ for all } v \in V(G_n).$$

Observe that η_n is a function of φ_n: for all $v \in V(G_n)$, $\eta_n(v)$ equals the color of the root of φ_n.

For any vertex $v \in V(G_n)$ the edge $(v, u) \in E(G_n)$ connecting v to its neighbor u can be traversed using at most L edges of H_n. Choose a shortest path between v and u in H_n. Because the length of the steps on the edges of H_n are bounded by L in terms of the graph distance in G_n the $r = L^2 + 1$-neighborhood of v in G_n already contains this shortest path. This holds for all neighbors u. This fact will be reflected in the type of v, namely for every neighbor of the root of α_v there will be a path of length at most L using edges from F_v connecting the root to the neighbor. We will refer to this property by saying that F_v *witnesses L-bi-Lipschitz equivalence at the root.*

Since the G_n converge locally-globally to $\mathcal{G}$ if we choose n large enough, we can find a Borel coloring $\varphi : X \to T$ such that

$$d_{TV}\big(P_{G_n,r}[\varphi_n], P_{\mathcal{G},r}[\varphi]\big) < \delta,$$

that is we can model the local statistics of φ_n on $\mathcal{G}$ with at most δ error. Choose δ such that $\delta(D + \frac{1}{2}D^L) < \varepsilon$.

The type of x gives a suggestion on how to construct a cheap graphing around x, which is F_x, the collection of distiguished edges. The idea is to consider the r-neighborhood of x in $\mathcal{G}$, and choose the edges of some graphing $\mathcal{H}$ locally according to F_x. This $\mathcal{H}$ would have the same connected components as $\mathcal{G}$ because F_x witnesses bi-Lipschitz equivalence at the root, and would be cheap because the expected $\mathcal{H}$-degree of a point is close to the expceted F_v-degree of the root in $P_{G_n,r}[\varphi_n]$. The problem is that the r-neighborhood of x in $\mathcal{G}$ is a priori not the same as what the type of x suggests it is.

However, we will show that the above idea works for most of the points, and after a slight modification the resulting graphing will be a cheap generating graphing of the relation $\mathcal{R}_{\mathcal{G}}$.

First we construct a Borel coloring $\eta : X \to \{1, \ldots, k\}$ from φ imitating the way the η_n could be recovered from the φ_n. Let $x \in X$, and let $\varphi(x) = (\alpha_x, \eta_x, F_x)$ be the type assigned to x by φ. The type suggests a color for the root, namely $\eta_x(o)$ where o is the root of α_x. So we set $\eta(x) = \eta_x(o)$. Observe that η_x is a coloring of the rooted graph α_x, and a priori neither α_x nor η_x has anything to do with the structure of $\mathcal{G}$. The value $\eta(x)$ on the other hand is a concrete color from $\{1, \ldots, \text{k}\}$ that is assigned to the point $x \in X$.

It will turn out that that for most points $x \in X$, their η-colored neighborhood in $\mathcal{G}$ is the same as the colored neighborhood (α_x, η_x) suggested by thier type $\varphi(x)$, and η breaks the possible local symmetries of $\mathcal{G}$ by being injective on the neighborhood. We define Y_1 as the set of points where this does not hold up to distance r:

$$Y_1 = \Big\{ x \in X \;\Big|\; \big(B_{\mathcal{G}}(r, x), \eta|_{B_{\mathcal{G}}(r,x)}\big) \not\cong \big(B_{\alpha_x}(r, o), \eta_x|_{B_{\alpha_x}(r,o)}\big) \Big\}$$

$$\bigcup\left\{x \in X \;\middle|\; \eta|_{B_{\mathcal{G}}(r,x)} \text{ is not injective}\right\}.$$

For any x outside Y_1 we can identify $B_{\alpha_x}(r,o)$ with $B_{\mathcal{G}}(r,x)$ using their colorings. For any $v \in V(\alpha_x)$ there exists a unique $y \in B_{\mathcal{G}}(r,x)$ with $\eta(y) = \eta_x(v)$. Such a y exists because of the isomorphism, and the injectivity of the colorings implies uniqueness. Later on we will denote this unique y by $y_{x,v}$. The identification works the other way around as well, for every y in $B_{\mathcal{G}}(r,x)$ we can find a unique $v \in V(\alpha_x)$ such that $\eta_x(v) = \eta(y)$. Let us denote this unique v by $v_{x,y}$.

We use this identification to reconstruct our rewiring on $\mathcal{G}$. Define the edges of a graphing $\mathcal{H}_0$ around $x \in X \setminus Y_1$ as follows: for every edge $(o,v) \in F_x$ that connects the root o of α_x with some other point $v \in V(\alpha_x)$ we include the edge $(x, y_{x,v})$ in $\mathcal{H}_0$.

Recall that we aim to show that $\mathcal{H}_0$ (with some small later adjustments) spans the same connected components as $\mathcal{G}$.

Definition 2.3.1 (Perfect points)**.** Call a point $x \in X$ *perfect*, if the following conditions hold:

1) F_x witnesses L-bilipchitz equivalence at the root;

2) all the edges in F_x that φ suggests in $B_{\alpha_x}(r,o)$ are indeed chosen to be in the edge set of $\mathcal{H}_0$.

For a type $\varphi(x)$, which is a rooted graph of radius $R = 2r$ with some additional decorations we write $\varphi(x)|_r$ for the graph where we simply forget everthing outside radius r. Similarly when $v \in B_{\alpha_x}(r,o)$ write $\varphi(x)|_{r,v}$ for the rooted, decorated graph we get by considering v as the root, and then forgetting everything outside radius r from v.

Definition 2.3.2 (Problematic points)**.** Let Y_2 be the set of points where one of the following holds.

i) φ does not witness the bi-Lipschitz connectivity at the root;

ii) φ fails to capture the local η-colored structure (up to distance R);

iii) η is not injective up to radius R;

iv) there is some y close to x where $\varphi(y)|_r$ differs from $\varphi(x)|_{r,y_x}$.

We call these points *problematic.*

$$\begin{aligned} Y_2 \;=\; & \left\{x \in X \;\middle|\; F_x \text{ does not witness bi-Lipschitz equivalence at the root}\right\} \\ & \bigcup\left\{x \in X \;\middle|\; (B_{\mathcal{G},R}(x), \eta|_{B_{\mathcal{G},R(x)}}) \not\cong (\alpha_x, \eta_x)\right\} \\ & \bigcup\left\{x \in X \;\middle|\; \eta|_{B_{\mathcal{G},R(x)}} \text{ is not injective}\right\} \\ & \bigcup\left\{x \in X \setminus Y_1 \;\middle|\; \exists v \in B_{\alpha_x}(r,o) \text{ s.t. } \varphi(x)|_{r,v} \not\cong \varphi(v_x)|_r\right\}. \end{aligned}$$

The next lemma shows that because no such incoherencies happen in G_n the measure of the problematic points will be small. Also note that $Y_1 \subset Y_2$, as in Y_2 we include all points where the local structure is not captured up to distance R instead of r.

Lemma 2.3.3. The points in $(X \setminus Y_2)$ are perfect and $\mu(Y_2) < \delta$.

Proof. The $P_{G_n,r}[\varphi_n]$ and $P_{\mathcal{G},r}[\varphi]$ are probability distributions on the set U_R^T of rooted, T-colored graphs of radius at most R, where T is the set of all possible types.

Let (β, o_β, ψ) denote such a graph with root o_β and coloring $\psi : V(\beta) \to T$. For a vertex $u \in V(\beta)$ its type $\psi(u) \in T$ is the rooted k-colored graph (α_v, η_v) with the additional distingushed edges F_v. There is some root o_{α_v} of α_v, and this way we define the coloring $\eta_\psi : V(\beta) \to \{1, \ldots, k\}$ by $\eta_\psi(v) = \eta_v(o_{\alpha_v})$. This η_ψ is defined from ψ the same way as η (on X) is defined from φ. Now let (β, o_β, ψ) be random with distribution $P_{G_n,r}[\varphi_n]$, then

$$\mathbb{P}_{P_{G_n,R}[\varphi_n]}\left[(\alpha_{o_\beta}, \eta_{o_\beta}) \cong (\beta, \eta_\psi)\right] = 1.$$

The above equality just restates that the type encodes the local colored structure (specifically the color of the root), but this formulation shows that this property will be inherited with small error when the total variation distance is small.

This isomorphism again enables us to identify α_{o_β} with β. For any $v \in V(\alpha_{o_\beta})$ write $y_{o_\beta,v}$ for the unique $y \in V(\beta)$ for which $\eta_{o_\beta}(v) = \eta_\psi(y)$.

We restate that (H_n) is a rewiring by saying that the distinguished edges witness the L-bi-Lipschitz equivalence at the root:

$$\mathbb{P}_{P_{G_n,R}[\varphi_n]}\left[F_{o_\beta} \text{ witnesses } L\text{-bi-Lipschitz equivalence at the root}\right] = 1.$$

We also restate the fact that the type of v, which is all the information up to distance $R = 2r$, includes all the information in the r-neighborhood of some other point u, provided that u is within distance r from v.

$$\mathbb{P}_{P_{G_n,R}[\varphi_n]}\left[\psi(o_\beta)|_{r,v} \cong \psi(y_{o_\beta,v})|_r \text{ for all } v \in V\big(B_{\alpha_{o_\beta}}(r, o_{\alpha_{o_\beta}})\big)\right] = 1.$$

Finally we restate that η distinguishes all points in the R-neighborhoods.

$$\mathbb{P}_{P_{G_n,R}[\varphi_n]}\left[\eta_\psi \text{ is injective}\right] = 1.$$

We see that the four events together hold with probability 1 with respect to $P_{G_n,R}[\varphi_n]$. Since $P_{\mathcal{G},r}[\varphi]$ is close to $P_{G_n,R}[\varphi_n]$ we get that the same holds for φ and $\mathcal{G}$ with probability at least $1 - \delta$, which implies $\mu(Y_2) < \delta$.

If $x \in X \setminus Y_2$ and $y \in B_{\mathcal{G}}(r, x)$ then $y \in X \setminus Y_1$, which means all the distinguished edges starting from y suggested by $\varphi(y)$ are indeed in $\mathcal{H}_0$, and by the definition of Y_2 we know that these are exactly the ones that $\varphi(x)$ would suggest. It is also clear that F_x has to witness generation, as otherwise x would be in Y_2. This implies that x is perfect. $\square$

Adding all the edges leaving the points in Y_2 we get $\mathcal{H}$:

$$\mathcal{H} = \mathcal{H}_0 \cup \big\{(x, y) \in E(\mathcal{G}) \;\big|\; x \in Y_2, y \in X\big\}.$$

Lemma 2.3.4. $\mathcal{H}$ has the same connected components as $\mathcal{G}$, and

$$\tilde{\mu}\big(E(\mathcal{H})\big) \le \frac{|E(H_n)|}{|V(H_n)|} + \varepsilon.$$

Proof. $\mathcal{H}$ will have the same connected components as $\mathcal{G}$, because for any edge $(x,y) \in E(\mathcal{G})$ where x is perfect the connection is witnessed by the r-neighborhood of x. If x is not perfect, then $x \in Y_2$, so $(x,y) \in E(\mathcal{H})$ by definition.

We now aim to show that $\mathcal{H}$ is indeed a cheap generator for the equivalence relation.

$$\tilde{\mu}\big(E(\mathcal{H})\big) \leq \tilde{\mu}\big(E(\mathcal{H}_0)\big) + \tilde{\mu}\big(\{(x,y) \in E(\mathcal{G}) \mid x \in Y_2, y \in X\}\big) \leq$$

$$\leq \tilde{\mu}\big(E(\mathcal{H}_0)\big) + \delta D.$$

For every point $x \in X$ let $\deg_{F_x}(o)$ denote the number of edges in F_x leaving the root of α_x. It also makes sense to talk about the expectation of this F-degree with respect to colored neighborhood statistics, as the F-degree of the root can be determined from its type.

$$\tilde{\mu}\big(E(\mathcal{H}_0)\big) = \frac{1}{2}\int_X \deg_{\mathcal{H}_0}(x)\ d\mu \leq \frac{1}{2}\int_X \deg_{F_x}(o)\ d\mu = \frac{1}{2}\mathbb{E}_{P_{\mathcal{G},r}[\varphi]}\Big[\deg_F(o)\Big] \leq$$

$$\leq \frac{1}{2}\left(\mathbb{E}_{P_{G_n,r}[\varphi_n]}\Big[\deg_F(o)\Big] + \delta D^L\right) = \frac{|E(H_n)|}{|V(H_n)|} + \frac{1}{2}\delta D^L.$$

Here we used the fact that there can be no more than D^L edges leaving the root in F_v (because of the bi-Lipschitz condition), and that the two distributions are close in total variation. Putting all this together and using that we chose δ to ensure that $\delta(D + (1/2)D^L) < \varepsilon$ we get

$$\tilde{\mu}\big(E(\mathcal{H})\big) \leq \frac{|E(H_n)|}{|V(H_n)|} + \varepsilon.$$

□

By the choice of (H_n) we can assume that

$$\frac{|E(H_n)|}{|V(H_n)|} \leq \mathrm{cc}(G_n) + 2\varepsilon,$$

which implies $\tilde{\mu}(E(\mathcal{H})) < \mathrm{cc}(G_n) + 3\varepsilon$. This shows the inequality $\mathrm{cost}(\mathcal{G}) \leq \mathrm{cc}(G_n)$.

The other inequality is proved exactly the same way. The condition that (H_n) is a rewiring was only used to ensure that the bi-Lipschitz constant L does not depend on n, only on ε. To prove that $\mathrm{cost}(\mathcal{G}) \geq \mathrm{cc}(G_n)$ we start by picking a cheap L-bi-Lipschitz generator for the single graphing $\mathcal{G}$ using Lemma 2.2.3, and by local-global convergence we know that for n large enough we can copy it to G_n with small error.

As for any large enough n and m the graphs G_n and G_m are arbitrarily close in the local-global topology we can do the same copying argument between the two. We fix the constant L first, and then choose n and m accordingly. This shows that (for all L) the rewirings (H_n) can indeed be choosen such that the limits defining the edge densities exist. This finishes the proof of Theorem 2.1.1. □

2.3.2 Sofic approximations

Using Theorem 2.1.1 we will show that sofic approximations of a group with fixed price c have combinatorial cost c as well.

Proof of Theorem 2.1.3. The sequence G_n of S-edge-labeled graphs converges to $\mathrm{Cay}(\Gamma, S)$ in the Benjamini-Schramm sense, so any subsequential local-global limit will be a graphing of an essentially free action of Γ, which by the fixed price assumption implies that it has cost c.

First pick a locally-globally convergent subsequence G_{n_k} with limit $\mathcal{G}_1$.

$$\mathrm{cc}(G_n) \leq \mathrm{cc}(G_{n_k}) = \mathrm{cost}(\mathcal{G}_1) = c.$$

Now assume that $\mathrm{cc}(G_n) < c$. We pick an L large enough such that there is some L-bi-Lipschitz rewiring (H_n) with $\underline{e}(H_n) < c$. As $\underline{e}$ is defined by a liminf we can choose a subsequence n_l such that

$$\lim \frac{|E(H_{n_l})|}{|V(H_{n_l})|} < c.$$

Now by passing to a further subsequence we can assume that the (G_{n_l}) converge locally-globally to some $\mathcal{G}_2$. The H_{n_l} witness that $\mathrm{cc}(G_{n_l}) < c$, while local-global convergence implies $\mathrm{cc}(G_{n_l}) = \mathrm{cost}(\mathcal{G}_2) = c$ by Theorem 2.1.1. This is clearly a contradiction, hence $\mathrm{cc}(G_n) = c$. $\square$

2.4 Group actions

The same notions and results exist in the world of measure preserving group actions, where convergence with respect to the weak containment topology takes the place of local-global convergence. The analogous definitions and statements will be introduced in this section.

2.4.1 Groupoid cost

Let Γ be a finitely generated group, generated by the finite symmetric set $S = S^{-1}$. Let (X, μ) be either a standard Borel probability space or X a finite set with μ the uniform measure on X. A probability p.m.p. action f of Γ is a homomorphism from Γ to the group of measure preserving transformations of (X, μ). The image of some $\gamma \in \Gamma$ under this homomorphism will be denoted by f_γ.

Any such p.m.p. action gives rise to a groupoid denoted M_f: endow Γ with the discrete topology and counting measure, consider $M_f = X \times \Gamma$ with the product Borel structure and product measure $\tilde{\mu}$. We also define a partial product on $X \times \Gamma$: $(x_1, \gamma_1) \cdot (x_2, \gamma_2) = (x_1, \gamma_1\gamma_2)$ whenever $x_2 = f_{\gamma_1}(x_1)$. The inverse is defined by $(x, \gamma)^{-1} = (f_\gamma(x), \gamma^{-1})$, and so $X \times \Gamma$ becomes a groupoid with respect to this partial product. We think of the element (x, γ) as an arrow pointing from x to $f_\gamma(x)$, with the arrow labeled by γ.

The notion of a generating subset of the groupoid is just as one would expect it: a subset generates, if all elements of M_f can be written as a product of elements and their inverses chosen from the subset.

For $A, B \subseteq M_f$ we will write

$$A \cdot B = \{a \cdot b \mid a \in A, b \in B, \text{ and } a \cdot b \text{ is defined}\}.$$

Also let $E = X \times \{e\}$, where e is the identity element of Γ. Using our notation A generates M_f if and only if

$$M_f = \bigcup_{n=1}^{\infty} (A \cup A^{-1} \cup E)^n.$$

The groupoid cost of f is

$$\mathrm{gcost}(f) = \inf\{\tilde{\mu}(A) \mid A \text{ generates } M_f\}.$$

The generators S of the group give rise to a specific generating subset of M_f, namely $X_S = X \times S$. We will say that a generating subset A is L-bi-Lipschitz, if all the elements of X_S can be generated by using at most L arrows from A and vice versa. More precisely we require that

$$X_S \subseteq (A \cup A^{-1} \cup E)^L, \text{ and } A \subseteq (X_S \cup X_S^{-1} \cup E)^L.$$

Note that while the actual value of the bi-Lipschitz constant L may depend on the choice of S, the property of A being a bi-Lipschitz generating subset (with some bi-Lipschitz constant) does not.

In this setting Lemma 2.2.3 was stated by Abért and Nikolov [AN12]. It says that by paying an arbitrarily small amount, we can choose the generating subset to be bi-Lipschitz. That is, for any $\varepsilon > 0$ there exists some integer L and an L-bi-Lipschitz generating subset $A \subseteq M_f$ such that $\tilde{\mu}(A) < \mathrm{gcost}(f) + \varepsilon$.

2.4.2 The weak containment topology

The notion of weak containment of actions was introduced by Kechris [Kec10]. The topology described below on the weak equivalence classes was defined by Abért and Elek in [AE11] and then studied further by Carderi in [Car15]. They showed that the topology of local-global convergence is a compact topology on the weak equivalence classes of actions.

For an action f of the group Γ and a point $x \in X$ the Γ-orbit of x admits a Schreier graph structure: for two points $y, z \in \Gamma x$ in the orbit draw an oriented edge from y to z labeled by some $s \in S$ if $f_s(y) = z$. Denote this graph by $\mathrm{Sch}(\Gamma, f, x)$.

The only difference compared to the local-global convergence of graph sequences and graphings is that in this case we consider the neighborhoods in the Schreier graphs together with the edge labeling by the generators S.

To an action f we again associate a set $Q^k_{f,r}$ that is the closure of all local statistics arising from Borel k-colorings with respect to the total variation distance. We say that an action f *weakly contains* another action g (denoted $f \succeq g$) if $Q^k_{g,r} \subseteq Q^k_{f,r}$ for all r and k. This means that all colorings of g can be modeled on f with arbitrarily small error. The actions are *weakly equivalent* if they both weakly contain the other, that is $Q^k_{g,r} = Q^k_{f,r}$

Convergence with respect to the weak containment topology is defined by the convergence of $Q^k_{f_n,r}$ for all r, k as compact sets with respect to the Hausdorff distance. The intuitive meaning of this convergence is the same as the one for local-global convergence. Abért and Elek showed that the topology induced by this convergence notion is compact, in particular every convergent sequence has a limit [AE11].

Kechris showed that if f and g are free p.m.p. actions and $f \succeq g$, then $\mathrm{cost}(\mathcal{R}_f) \leq \mathrm{cost}(\mathcal{R}_g)$ [Kec10, Corollary 10.14]. Here $\mathcal{R}_f$ denotes the orbit equivalence relation generated by the action f. Abért and Weiss extended this beyond free actions in [AW13]: for any actions with $f \succeq g$ the groupoid cost satisfies $\mathrm{gcost}(f) \leq \mathrm{gcost}(g)$. This implies that the groupoid cost is well defined on weak equivalence classes, and studying the continuity properties of the groupoid cost with respect to the weak containment topology makes sense.

2.4.3 The groupoid cost of weak containment limits

Following the proof of Theorem 2.1.1 we get a result for group actions.

Proposition 2.4.1. Suppose that the sequence $f_1, f_2, \dots$ of p.m.p. actions is convergent in the weak containment topology to the p.m.p. action f. Then

$$\limsup_{n\to\infty} \mathrm{gcost}(f_n) \leq \mathrm{gcost}(f). \tag{2.2}$$

This is a semicontinuity result for the groupoid cost with respect to the weak containment topology. The proof follows exactly the same steps as in Theorem 2.1.1: we choose a cheap bi-Lipschitz generating set for the groupoid M_f, record all local information into a coloring of X, modell this coloring on the f_n when n is large enough with some small error and build a cheap generating set for M_{f_n} by decoding the coloring.

The whole process is actually slightly easier in this setting, because there is no need to break the local symmetries of the graphs as the Schreier edge labeling already takes care of that. As we are not imposing a uniform bound on the complexity of generation by talking about "combinatorial groupoid cost", we only get an inequality.

However, for the inequality we only need that colorings of f can be modeled with small error on the f_n, and we don't have to require it the other way around. That is, if the sequence "asymptotically weakly contains" f, then we have (2.2). This can be thought of as an asymptotic version of the monotonicity results by Kechris [Kec10], and Abért-Weiss [AW13].

Remark (The ultraproduct technique). These results, together with Theorem 1 for graphings of free p.m.p. actions can be obtained by using the ultraproduct techniques introduced in [AE11] and [Car15], Carderi's result on ultraproduct actions being weakly equivalent to some standard action and the monotonicity results of Kechris and Abért-Weiss.

If one modifies (the somewhat arbitrary) choice of lower edge density in the definition of the combinatorial cost to *edge density along an ultrafilter* ω by taking an ultralimit instead of a liminf, then this modified combinatorial cost of the sequence will equal the cost of the ultraproduct graphing.

2.4.4 Rank gradient in groups with fixed price

We need one further tool to prove Theorem 2.1.4. The following lemma is stated in [AGN17, Lemma 21].

Lemma 2.4.2. Let Γ be a countable group, and H a subgroup of finite index in Γ. Let f be the right coset action of Γ on Γ/H. Then we have

$$r(\Gamma, H) = \frac{\mathrm{rank}(H) - 1}{|\Gamma : H|} = \mathrm{gcost}(f) - 1.$$

Proof of Theorem 2.1.4. First we show that $\liminf r(\Gamma, \Gamma_n) \geq c-1$. We select a subsequence Γ_{n_k} such that $r(\Gamma, \Gamma_{n_k})$ converges to the liminf. Taking the diagonal product of the corresponding group actions f_{n_k} we get an action f of Γ that factors onto each f_{n_k}, which implies $\operatorname{gcost}(f) \leq \operatorname{gcost}(f_{n_k})$ for every n_k. The measure of the set of fixed points of a group element can only increase for factros, which implies that f is essentially free because of the Farber condition. Using Lemma 2.4.2 and $\operatorname{gcost}(f) = c$ we get

$$\lim_{k\to\infty} r(\Gamma, \Gamma_{n_k}) = \lim_{k\to\infty} (\operatorname{gcost}(f_{n_k}) - 1) \geq \operatorname{gcost}(f) - 1 = c - 1.$$

Similarly we can choose a subsequence such that

$$\limsup r(\Gamma, \Gamma_n) = \lim r(\Gamma, \Gamma_{n_l}).$$

By passing to a further subsequence we can also assume that the actions f_{n_l} converge in the weak containment topology to some action $\hat{f}$. This limit is essentially free by the Farber condition, so $\operatorname{gcost}(\hat{f}) = c$. Using Proposition 2.4.1 we get

$$\lim_{l\to\infty} r(\Gamma, \Gamma_{n_l}) = \lim_{l\to\infty} (\operatorname{gcost}(f_{n_l}) - 1) \leq \operatorname{gcost}(\hat{f}) - 1 = c - 1.$$

□

Remark (alternative proof)**.** The second part of the proof can be obtained without Proposition 2.4.1 – which we only sketched – by using a result from [AGN17].

After choosing a subsequence such that $\limsup r(\Gamma, \Gamma_n) = \lim r(\Gamma, \Gamma_{n_l})$, [AGN17, Theorem 8] states that $\lim r(\Gamma, \Gamma_{n_l}) \leq \operatorname{cc}\big(\operatorname{Sch}(\Gamma, \Gamma_{n_l}, S)\big) - 1$. Using Theorem 2.1.3 we get

$$\lim_{l\to\infty} r(\Gamma, \Gamma_{n_l}) \leq \operatorname{cc}\big(\operatorname{Sch}(\Gamma, \Gamma_{n_l}, S)\big) - 1 = c - 1.$$

2.5 The trichotomy theorem

In this section we introduce strong ergodicity, and prove the results on finitely presented groups.

2.5.1 Strong ergodicity

Let f be a p.m.p. action of the countable group Γ on a standard Borel space (X, μ). A sequence A_n of measurable subsets is called *almost invariant*, if

$$\lim_{n\to\infty} \mu(f_\gamma A_n \triangle A_n) = 0, \text{ for all } \gamma \in \Gamma.$$

The action f is *strongly ergodic*, if for any almost invariant sequence A_n we have

$$\lim_{n\to\infty} \mu(A_n)\big(1 - \mu(A_n)\big) = 1.$$

We will make use of the following result of Abért and Weiss [AW13, Theorem 3].

Theorem 2.5.1. Let f be an ergodic p.m.p. action of a countable group Γ on a standard Borel space (X,μ). If f is not strongly ergodic, then f is weakly equivalent to $f \times I$,which is the diagonal action on $(X,\mu) \times [0,1]$ with Γ acting trivially on the second coordinate.

2.5.2 Dispersive actions

Let $\mathrm{Sch}(\Gamma,\Gamma_n,S)$ be a sequence of Schreier graphs, and let f_n denote the corresponding finite actions. We call the sequence $\mathrm{Sch}(\Gamma,\Gamma_n,S)$ *dispersive* if for any subsequential weak containment limit f of f_n, f has no strongly ergodic, ergodic component of positive measure.

Lemma 2.5.2. Let Γ be a group generated by the finite set S and Γ_n a sequence of subgroups such that the corresponding Schreier graphs $\mathrm{Sch}(\Gamma,\Gamma_n,S)$ form a dispersive sequence. Then for every $\varepsilon > 0$ and $k \in \mathbb{N}$ we can find some n such that the vertex set V of $\mathrm{Sch}(\Gamma,\Gamma_n,S)$ can be partitioned into k sets $A_1,\dots,A_k$ such that

1. $\frac{1}{k} - \varepsilon \le \frac{|A_i|}{|V|} \le \frac{1}{k} + \varepsilon$ for all A_i,
2. $\sum |SA_i \setminus A_i| < \varepsilon|V|$.

Proof. Pick a subsequence Γ_{n_l} such that the $\mathrm{Sch}(\Gamma,\Gamma_{n_l},S)$ converge in the weak containment topology to some Γ action f on a standard Borel space (X,μ). As the sequence is dispersive we know that f has no strongly egrodic, ergodic components of positive measure.

We claim that X can be partitioned into k disjoined Borel sets $B_1,\dots,B_k$ of approximately equal measure that are almost invariant, namely

$$\sum_{s\in S}\sum_{1\le i\le k} \mu(sA_i \setminus A_i) < \varepsilon/2.$$

Assume first that f is ergodic. Then it is not strongly ergodic, and we can use Theorem 2.5.1. The base space of $f \times I$ can be easily partitioned into k invariant subsets, specifically $B_i' = (X,\mu)\times[(i-1)/k, i/k]$. Weak equivalence guarantees that this partition can be modeled with arbitrarily small error (e.g. $\varepsilon/2$) for f on the set X, and thus we have the desired B_i.

When f is not ergodic the we can divide X into two invariant sets X_1 and X_2 of positive measure. If f is ergodic on one of the X_i, then it is not strongly ergodic on that part and hence we can use the above argument to partion that X_i. If f is not ergodic on X_i, then we can again divide it into invariant subsets of positive measure.

Since all positive measure ergodic components are not strongly ergodic, we can continue this procedure and find a partition into some tiny (measure $\varepsilon/100$) invariant sets and some non-strongy-ergodic components, which we can partition into almost invariant sets. Putting these blocks together into k sets we can get an almost invariant partition into pieces with measure in $[1/k - \varepsilon/2, 1/k + \varepsilon/2]$.

Now since f is a limit, for some n_l large enough we can model the partition $B_1,\dots,B_k$ with error $\varepsilon/2$ on $\mathrm{Sch}(\Gamma,\Gamma_{n_l},S)$, and get the desired A_i. $\square$

2.5.3 Finitely presented groups

We now briefly discuss how we will present finite index subgroups of finitely presented groups. Notation and general framework follows [AJN11].

Let $\Gamma = \langle S \mid R\rangle$ be a finitely presented group, and $H \subseteq \Gamma$ a finite index subgroup. Let $\mathcal{T}$ denote a spanning tree of the Schreier graph $\mathrm{Sch}(\Gamma,H,S)$. We select a transversal T for the subgroup H as follows: for each coset γH we consider the unique path in $\mathrm{Sch}(\Gamma, H, S)$ from the root H to γH, and select the corresponding S-word to be in T. This T is called the left *Schreier transversal* corresponding to $\mathcal{T}$ with respect to S. For a group element $\gamma \in \Gamma$ let $\tilde{\gamma}$ denote the unique element in T such that $\gamma H = \tilde{\gamma} H$.

For every edge $e = (\gamma H, s\gamma H)$ of $\mathrm{Sch}(\Gamma, H, S)$ we put $T(e) = (\tilde{s\gamma})^{-1} s\tilde{\gamma}$. It is known that the $\{T(e)\}$ belong to and generate H. Note that if $e \in E(\mathcal{T})$, then $T(e) = 1$.

For a relation $r = s_l \dots s_1 \in R$ and group element $t \in T$ let $r_t = t^{-1} r t$. This r_t is an element of H, and can be considered as a word in the $T(e)$: $r_t = T(e_l) \dots T(e_1)$, where $e_i = (s_{i-1} \dots s_1 t H, s_i \dots s_1 t H)$. We are going to use the fact that these relations give a presentation of H:

$$H = \big\langle \{T(e)\}_{e \in E(\mathrm{Sch}(\Gamma, H, S)) \setminus E(\mathcal{T})} \mid \{r_t\}_{r \in R, t \in T} \big\rangle.$$

Suppose a group H has subgroups $H_i \subseteq H$ ($1 \leq i \leq k$) which all contain a fixed subgroup L and $H \cong *_L H_i$. We say that this decomposition is non-trivial, if L has index at least 3 in at least two of the subgroups.

Proof of Theorem 2.1.7. Assume that $\mathrm{Sch}(\Gamma, \Gamma_n, S)$ is dispersive and $\mathrm{RG}(\Gamma, (\Gamma_n)) > 0$. We will show that some Γ_n decomposes as a non-trivial amalgamated product. We can pass to a subsequence and assume that

$$\frac{d(\Gamma_n) - 1}{[\Gamma : \Gamma_n]} > c > 0 \text{ for all } n.$$

Choose an integer k such that

$$\left(\frac{3}{2}|S| + 1\right) \frac{1}{k} \leq c/2.$$

Let M be the sum of the lengths of the relations in R. As the sequence is dispersive, using Lemma 2.5.2 we can choose some n such that the vertex set $V\big(\mathrm{Sch}(\Gamma, \Gamma_n, S)\big)$ can be split into the disjoint union of k sets $A_1, \dots, A_k$ such that

1. $\frac{[\Gamma:\Gamma_n]}{k} - \frac{[\Gamma:\Gamma_n]}{2k} < |A_j| < \frac{[\Gamma:\Gamma_n]}{k} + \frac{[\Gamma:\Gamma_n]}{2k}$ for all $j \in \{1, \dots, k\}$ and
2. $|\partial(A_1, \dots, A_k)| < \frac{1}{k(1+M^2)} [\Gamma : \Gamma_n]$, where

$$\partial(A_1, \dots, A_k) = \big\{ e \in E\big(\mathrm{Sch}(\Gamma, \Gamma_n, S)\big) \;\big|\; e = (x, y), x \in A_j, y \in A_l, j \neq l \big\}.$$

As we have $[\Gamma : \Gamma_n] \to \infty$ we can also make sure that we choose n large enough so that

$$\frac{k-1}{[\Gamma : \Gamma_n]} \leq c/2.$$

Put $H = \Gamma_n$ and follow the above construction for a presentation of H. Note that $\big|V\big(\mathrm{Sch}(\Gamma, H, S)\big)\big| = [\Gamma : H]$.

Define Y to be the collection of generators $T(e)$ that share a relation with an inbetween edge. More precisely let $T(e) \in Y$ if either $e \in \partial(A_1, \dots, A_k)$ or there exists a relation $r_t = T(e_1)^{\pm 1} \dots T(e_l)^{\pm 1}$ for which some $e_j \in \partial(A_1, \dots, A_k)$ and some $e_m = e$. Let X_i be the set of generators $T(e)$ for which both endpoints of e are in A_i.

Define the subgroups $L = \langle Y \rangle$ and $H_i = \langle Y \cup X_i \rangle$. Clearly $L \leq H_i$ for all i.

Lemma 2.5.3. *H decomposes as the amalgamated product of the H_j over L, $H \cong *_L H_i$.*

We postpone the proof of the lemma, and show that this decomposition is non-trivial. Suppose that L has index at most 3 in $H_1, H_2, \dots, H_{k-1}$. Then each of these H_i ($1 \le i \le k-1$) are generated by L and at most 1 other element. Thus H is generated by H_k (which includes L) and at most $k-1$ other elements, $d(H) \le d(H_k) + k - 1$.

It is easy to bound the cardinality of $|X_k|$:

$$|X_k| \le |S||A_k| \le |S| \left(\frac{[\Gamma : H]}{k} + \frac{[\Gamma : H]}{2k} \right).$$

Let $r = s_l \dots s_1$ be a relation of Γ of lenght l. Note that there are at most $l|\partial(A_1, \dots, A_k)|$ different lifts $r_{\tilde{g}} = T(e_1)^{\pm 1} \dots T(e_l)^{\pm 1}$ of r for which some $e_j \in \partial(A_1, \dots, A_k)$. Also for each such relation of H we have at most l generators $T(e)$ of H that are getting into Y. Thus, if $\{l_j\}$ is the set of lenghts of the relations of Γ (so $M = \sum l_j$), then we can bound the cardinality of Y.

$$|Y| \le |\partial(A_1, \dots, A_k)|(1 + \sum l_j^2) \le |\partial(A_1, \dots, A_k)|(1 + M^2) \le \frac{[\Gamma : H]}{k}.$$

Putting our bounds together we get

$$d(H) \le d(H_k) + k - 1 \le |X_k| + |Y| + k - 1 \le \left(\frac{3}{2}|S| + 1 \right) \frac{[\Gamma : H]}{k} + k - 1.$$

This however gives an upper bound on the rank quotient at H:

$$\frac{d(H) - 1}{[\Gamma : H]} \le \left(\frac{3}{2}|S| + 1 \right) \frac{1}{k} + \frac{k-1}{[\Gamma : H]} \le c/2 + c/2 = c.$$

This contradicts our assumption that each such quotient is more than c, hence the decomposition is non-trivial. □

Proof of Lemma 2.5.3. The argument follows the one in [AJN11, Section 3].

Consider the following sets of relations. Let R_i be the set of all the $r_t = T(e_1)^{\pm 1} \dots T(e_l)^{\pm 1}$ where either all e_j have both endpoints in A_i or some e_j is in $\partial(A_1, \dots, A_k)$. Now $R_i \cup R_j$ is the same set $\bar{R}$ for all pairs (i, j), that is the relations having an inbetween edge. Define the groups T_i by the presentations $\langle X_i \cup Y \mid R_i \rangle$, let $\bar{T} = \langle Y \mid \bar{R} \rangle$. We have a homomorphisms $\phi_j : \bar{T} \to T_i$ by the inclusion of Y into $X_i \cup Y$. From the presentations we see that $H \cong *_{\bar{T}} T_i$.

Each T_i surjects onto H_i (and $\bar{T}$ surjects onto L) by mapping the abstract generators to their counterparts in H. By the universal property of the amalgamated product one can see $H \cong *_L H_i$. □

Proposition 2.5.4. *For a countable amenable group Γ all sequences (Γ_n) of distinct finite index subgroups are dispersive.*

Proof. It is a result of Schmidt [Sch81, Theorem 2.4] that amenable groups admit no strongly ergodic actions. If any subsequential limit would have an ergodic component of positive measure that is strongly ergodic, then restricting the action to that component would contradict [Sch81]. □

As amenable groups cannot decompose as non-trivial amalgamated products this proves Theorem 2.1.6.

2.6 Open problems

Note that an even stronger form of Theorem 2.1.4 would simply express the rank gradient of a local-global convergent sequence as the cost of the limiting graphing of the sequence. As of this moment, we do not know how to prove or disprove this. The obstacle is that combinatorial cost handles sequences of generating sets with a bounded complexity with respect to some standard generating set. A priori, it could happen that actual small generating sets over the sequence need a very aggressive growth of complexity, and the local-global metric is too weak to connect such generating sets over the sequence. This is the same obstacle that makes the proof of Theorem 2.1.7 somewhat tricky.

The following problem connects two well-known unsolved problems, one in ergodic theory, the other in 3-manifold theory.

Problem 2.6.1. Let Γ be a finitely generated group. Does there exist c such that for any Farber sequence (Γ_n) in Γ, we have $\mathrm{RG}(\Gamma, (\Gamma_n)) = c$?

Equivalently, one can ask whether $\mathrm{RG}(\Gamma, (\Gamma_n))$ exists for any Farber sequence (Γ_n) in Γ. Indeed, an advantage of Farber sequences over Farber chains is that they are closed to merging.

By Theorem 2.1.4, a negative answer to Problem 2.6.1 would immediately give a negative answer to the Fixed Price problem of Gaboriau [Gab00], that asks whether for an arbitrary countable group Γ, all essentially free p.m.p. actions of Γ have the same cost. A positive answer, on the other hand, would specifically show that in a finitely generated group, any two normal chains with trivial intersection have the same rank gradient, which by [AN12] would then solve the strong Rank vs Heegaard genus problem on hyperbolic 3-manifolds.

One possible approach to Problem 2.6.1 is through graph theory as follows. Let G be a finite, connected graph with maximal degree D. For $L \geq 1$ let

$$c_L(G) = \min_H \frac{|E(H)|}{|V(G)|}$$

where H runs through all rewirings of G with bi-Lipschitz constant at most L. It is easy to see that

$$1 - o(1) \leq c_L(G) \leq D.$$

The following problem is related to the Fixed Price problem of Gaboriau.

Problem 2.6.2. Let (G_n) be a Benjamini-Schramm convergent sequence of graphs of bounded degree. Does $c_L(G_n)$ converge for every $L \geq 1$?

The connection is one sided.

Proposition 2.6.3. An affirmative solution of Problem 2.6.2 implies an affirmative solution of Problem 2.6.1.

Note, however, that this problem seems to be a real strenghtening. Indeed, it could happen that for two large graphs G_1 and G_2 that are very close in the Benjamini-Schramm topology, one can find a cheap rewiring of G_1 with a bi-Lipschitz constant L_1 but only do the same to G_2 with a much bigger constant.

Problem 2.6.4. Let Γ be a finitely presented group generated by a finite symmetric set S. Let (Γ_n) be a sequence of subgroups of finite index in Γ and let $G_n = \mathrm{Sch}(\Gamma, \Gamma_n, S)$. Assume that no Γ_n decomposes as a non-trivial amalgamated product and that the sequence (G_n) is dispersive. Is it true that $\mathrm{cc}(G_n) = 1$?

In Theorem 2.1.7 we show that the rank gradient of (Γ_n) must vanish.

3

The distortion function

Abstract

We study the distortion function of p.m.p. actions. We investigate its behavior with respect to weak containment, and observe monotonicity outside a countable set of discontinuity points. We compute the distortion function of free actions of $\mathbb{Z}^d$ up to constant multiple, and conclude that the distortion function distinguishes them. For lamplighter groups we establish a logarithmic bound on the distortion function in case the lamps have finitely many states and the lamplighter walks on a finitely generated infinite group.

Contents

3.1 Preliminaries

The distortion function of p.m.p. actions was introduced in [AGN17]. It is an invariant of group actions that is related to the cost, so our setting will be the same as in Subsection 2.4.1.

Let f be a p.m.p. action of the finitely generated group Γ on the standard Borel probability space (X, μ), and let M_f denote the associated measured groupoid. Fix a finite generating set S, and write

$$X_S = \{(x, s) \in M_f \mid x \in X, s \in S\}$$

and

$$E = \{(x, e) \in M_f \mid x \in X\}.$$

Recall that we say a generating subset $A \subseteq M_f$ is L-bi-Lipschitz if

$$X_S \subseteq (A \cup A^{-1} \cup E)^L \text{ and } A \subseteq (X_S \cup X_S^{-1} \cup E)^L.$$

The distortion function measures how large one has to chose the bi-Lipschitz constant as we are getting closer and closer to the cost.

Definition 3.1.1 (distortion function)**.** For all positive x define $\delta_f(x)$ to be the least integer L such that there is an L-bi-Lipschitz generating set $A \subseteq M_f$ with $\tilde{\mu}(A) \leq \operatorname{gcost}(f) + x$.

To make notation more compact we introduce the bi-Lipschitz distance of two generating sets $A, B \subseteq M_f$:

$$d_L(A, B) = \inf\{k \mid B \subseteq (A \cup A^{-1} \cup E)^k, \text{ and } A \subseteq (B \cup B^{-1} \cup E)^k\}.$$

This way we can write the distortion function as

$$\delta_f(x) = \inf_{\substack{A \subseteq M_f \text{ generating} \\ \tilde{\mu}(A) \leq \operatorname{gcost}(f) + x}} d_L(A, X_S).$$

The distance of two generating subsets might be infinite, but as we discussed in Subsection 2.4.1 one can always assume bi-Lipschitz generation by paying an arbitrarily small amount. In other words $\delta_f(\varepsilon)$ is finite for any $\varepsilon > 0$.

To sum up, we see that the distortion function is actually defined on $\mathbb{R}^+$, takes positive integer values and it is clearly monotone decreasing.

It is also clear that δ_f depends on the choice of the generating set S, but only up to a constant multiple. For this reason we suppress S from the notation. It will either be fixed, or the results will not depend on its choice.

3.1.1 Monotonicity with respect to weak containment

As we have seen in the proof of Theorem 2.1.1 and Proposition 2.4.1, weak containment of group actions allows one to copy generating subsets with little error. Notice that during this procedure the bi-Lipschitz constant does not increase, which we express in the following proposition.

Proposition 3.1.2. Let Γ be a group with a fixed finite generating set S. Let f and g be actions of Γ on the standard Borel probability space (X, μ). Let $A \subseteq M_f$ be an L-bi-Lipschitz generating set, and assume $f \preceq g$. Then for any $\varepsilon > 0$ there is a $B \subseteq M_g$ L-bi-Lipschitz generating set with $\tilde{\mu}_g(B) \leq \tilde{\mu}_f(A) + \varepsilon$.

This means we can relate the distortion functions of free actions of a fixed price group.

Corollary 3.1.3. Assume $\Gamma = \langle S \rangle$ has fixed price. If f, g are free actions with $f \preceq g$ then $\delta_g(x + \varepsilon) \leq \delta_f(x)$ for all $x, \varepsilon \in \mathbb{R}^+$.

The fixed price assumption ensures that f and g have the same cost, so the distortion functions measure the necessary complexity as we are approaching the same number.

It would be inviting to state that if f and g are weakly equivalent then their distortion functions (with respect to a fixed S) are the same. However we can see in Corollary 3.1.3 that there is a slight issue preventing this: we are only able to bound $\delta_g(x + \varepsilon)$ instead of $\delta_g(x)$.

Nonetheless our functions are integer valued and monotone decreasing. This means that apart from a countable set of discontinuity points we can still relate δ_f and δ_g.

Corollary 3.1.4. Let Γ be as before, assumed to have fixed price. Let f, g be weakly equivalent free actions of Γ. Assume that $D \subseteq \mathbb{R}^+$ is the set of discontinuity points of f. Then the set of discontinuity points of g is also D, and $\delta_f(x) = \delta_g(x)$ for all $x \in R^+ \setminus D$.

If one could prove continuity of the distortion function from right or left the above corollary would of course extend to D.

Problem 3.1.5. Is the distortion function of a free p.m.p. action always right (always left) continuous?

The problem with the discontinuity points is not a significant issue, as we will generally only consider distortion functions up to constant multiple. For real valued functions h, g we will denote equivalence up to constant multiple by $h \asymp g$. That is

$$h \asymp g \Leftrightarrow \exists C > 0 \text{ s.t. } \frac{1}{C} g \leq h \leq Cg.$$

3.2 Distortion of $\mathbb{Z}^d$

In this section we will investigate the distortion function of free actions of $\mathbb{Z}^d$. For amenable groups there is only one weak equivalence class of free actions. In fact this property characterizes amenability, see [Kec10, Proposition 13.2]. Our result is the following.

Theorem 3.2.1. For any free action f of $\mathbb{Z}^d$ we have

$$\delta_f(x) \asymp \left(\frac{1}{x}\right)^{\frac{1}{2d}}.$$

Before proceeding to the proof of Theorem 3.2.1 we state a quick corollary.

Corollary 3.2.2. The distortion function distinguishes the actions of $\mathbb{Z}^d$ for different d. In particular the distortion function is not an orbit equivalence invariant.

In the remainder of this section we prove Theorem 3.2.1. The upper bound is established in Proposition 3.2.4. By the characterization of amenability mentioned above it suffices to find one free p.m.p. action with the desired distortion. The lower bound is proved in Proposition 3.2.9.

3.2.1 Connection to invariant processes

For our construction proving the upper bound it will be more convenient to think of *invariant random rewirings* on the Cayley graph of $\mathbb{Z}^d$ as opposed to measurable subsets of groupoids associated to free actions.

Definition 3.2.3. An invariant random L-rewiring of a group Γ with fixed generating set S is a random L-rewiring H of the countable graph $\mathrm{Cay}(\Gamma, S)$ whose distribution is invariant under the translation action of Γ. We say H is aperiodic, if the rewiring almost surely has no translational symmetries.

It is a standard argument to associate essentially free p.m.p. actions to aperiodic invariant random processes.

In our situation we can define $\mathrm{Rew}(\Gamma, S, L)$ to be the space of all L-rewirings of $\mathrm{Cay}(\Gamma, S)$. This space carries a rooted distance and an action of Γ by translation of the set of edges constituting the rewiring. This action is by homeomorphisms.

An invariant random L-rewiring is simply a Borel probability measure on $\mathrm{Rew}(\Gamma, S, L)$ that is Γ-invariant. If the random rewiring is aperiodic the action will be essentially free with respect to the measure.

3.2.2 Construction for the upper bound

Using the connection above the following proposition will suffice to find an upper bound in Theorem 3.2.1. By the *average degree* of an invariant random rewiring we mean the expected degree of the identity element.

Proposition 3.2.4. Let $L \in \mathbb{N}$ arbitrary. There exists an invariant random aperiodic rewiring H of $\mathbb{Z}^d$ with bi-Lipschitz constant $(4d + 1)L$ such that H has average degree

$$2 + \frac{2d - 2}{L^{2d}}.$$

Proof. First we construct an aperiodic random rewiring H_0 which will be invariant only with respect to the finite index subgroup $L^2 \cdot \mathbb{Z}^d \leq \mathbb{Z}^d$. Then taking H to be a uniform random translate of H_0 will give invariance under the whole $\mathbb{Z}^d$.

The vertex set of H_0 will be $V = \mathbb{Z}^d$. The edge set of H_0 will consist of two types of edges: *short* and *long*. The short edges will simply be edges of Z^d, while the long edges will connect certain vertices that are at distance L in a vertical or horizontal direction.

We will introduce edges in a grid-like way, and to that end we denote by $V_2 = (L^2) \cdot Z^d$ the vectors with all coordinates divisible by L^2. Similarly let $V_1 = L \cdot Z^d$ denote the vectors with all coordinates divisible by L. Clearly $V_2 \subseteq V_1$.

At first we introduce long edges between vertices in V_1 to form an "L-by-L grid" on V_2. More precisely, let $i \in \{1, \dots d\}$ be an arbitrary coordinate, and let e_i denote the standard unit vector in the i-th coordinate. For any vector $v = (v_1, \dots, v_d) \in V_1$ such that $L^2 \mid v_j$ for all other coordinates $j \neq i$ and $L \mid v_i$ we connect v to $v + Le_i$ and $v - Le_i$. We do this for all coordinates i.

The second step, still using long edges, is filling in the cubes of our grid by a spanning tree. We will do this randomly and independently for all cubes, in order to get aperiodicity.

For $v \in V_1$ let $[v]_2 \in V_2$ denote the vector we get by decreasing each coordinate v_i to the greatest number $u_i \leq v_i$ with $L^2 \mid u_i$. That is $[v]_2 = (u_1, \dots, u_d) \in V_2$ and $v_i - L^2 < u_i \leq v_i$.

Fix $u \in V_2$. We will connect the vertices $v \in V_1 \setminus V_2$ with $[v]_2 = u$ by a random spanning tree. First randomly choose an ordering of the coordinates $\{1, \dots, d\}$, say i comes before j in the ordering if $\sigma(i) < \sigma(j)$ for a uniform random permutation $\sigma : \{1, \dots, d\} \to \{1, \dots, d\}$.

Let $v \in V_1 \setminus V_2$ with $[v]_2 = u$, and we look at the first coordinate i in our random ordering such that $L^2 \nmid v_i$. Introduce an edge from v to $v - Le_i$. This way we make sure that from any such v we can reach u in at most $d(L-1)$ steps. The ordering is chosen independently for all $u \in V_2$, so H_0 becomes aperiodic.

We claim that one can get from any $v \in V_1$ to any of its L-neighbors $v' = v \pm Le_i$ using at most $(d(L-1) + L + d(L-1)) \le (2d+1)L$ long edges. Indeed from v and v' we can get to $[v]_2$ and $[v']_2$ in $d(L-1)$ steps, and either $[v']_2 = [v]_2$, or $[v']_2 = [v]_2 \pm L^2 e_i$. In the first case we have nothing to do, while in the second we can get from $[v]_2$ to $[v']_2$ using L long edges.

The last step is to repeat the same spanning tree trick for vertices $v \in V \setminus V_1$, but this time with short edges. There is no need to do this at random anymore, as we already broke all symmetries.

We connect each such v to $v - e_i$, where i is the first coordinate with $L \nmid v_i$. We denote by $[v]_1$ the vector $(u_1, \dots, u_d) \in V_1$ with $v_i - L < u_i \le v_i$. From any $v \in V$ we can reach $[v]_1$ in at most $d(L-1)$ steps, and so we can get from any v to any neighboring vector v' in at most $d(L-1) + (2d+1)L + d(L-1) \le (4d+1)L$ steps.

As the edges we used are of length at most L we have shown that this H_0 is indeed a rewiring with bi-Lipschitz constant $(4d+1)L$.

It is also clear that the distribution of H_0 is $L^2 \cdot \mathbb{Z}^d$-invariant, because of the i.i.d. choice of the spanning trees.

To compute the average degree we can look at the average of degrees in the $\{0, \dots, (L^2 - 1)\}^d$ cube. At first disregarding the edges going out we see a spanning tree, which has $(L^2)^d - 1$ edges. The contribution of these edges to the average degree is

$$\frac{2\big((L^2)^d - 1\big)}{(L^2)^d}.$$

There are $2d$ outgoing edges, because of the grid we placed in the first step. These contribute $(2d)/(L^2)^d$ to the average degree. In total the average degree is

$$2 + \frac{2d-2}{L^{2d}}.$$

□

3.2.3 The lower bound

Now we turn to proving a lower bound on the distortion of $\mathbb{Z}^d$ actions. First we show that any individual L-rewiring has to include an L^2-by-L^2 grid in a topological sense. This gives us sufficiently many edges that are inessential for connectedness, and get a lower bound on the degree when averaging over Følner sets. Then we use a pointwise ergodic theorem for amenable groups to obtain our result on invariant random rewirings.

Lemma 3.2.5. Let F_n be a Følner sequence for the amenable, bounded degree graph G. If H is an L-rewiring of G, then F_n is also a Følner sequence for H.

Proof. By our assumption
$$\lim_{n\to\infty} \frac{|\partial_E^G(F_n)|}{|F_n|} = 0.$$
We investigate the number a_k of edges of H leaving F_n that traverse distance k in G. Every such edge corresponds to a path of length k that uses an edge from $\partial_E^G(F_n)$ along the way. Thus we get
$$a_k \leq k \cdot |\partial_E^G(F_n)| \cdot D^{k-1}.$$
Here D stands for the degree bound of G. Summing for all $k \leq L$ we get
$$|\partial_E^H(F_n)| \leq \sum_{k=1}^{L} k \cdot |\partial_E^G(F_n)| \cdot D^{k-1} = C(L, D) \cdot |\partial_E^G(F_n)|.$$
The constant C does not depend on n, so we get
$$\lim_{n\to\infty} \frac{|\partial_E^H(F_n)|}{|F_n|} = 0.$$
□

Lemma 3.2.6. Let H be an L-rewiring of $\mathbb{Z}^d$. Set $K = 4L^2 + 1$, and let B be a K-by-K box in $\mathbb{Z}^d$. We claim that we can delete $(d-2)2^{d-1}+1$ edges of H inside B while preserving connectedness of H.

Proof. First we outline the proof when $d = 2$. For the convenience of the reader the main objects of the proof are outlined on Figure 1. The proof for the higher dimensional case will be similar.

Without loss of generality we assume F to be the box $\{0, \dots, 4L^2+1\} \times \{0, \dots, 4L^2+1\}$. We have to show that we can delete 1 edge of H inside F, so we need to find a cycle inside H. Consider the 4 verices $u = (L^2, L^2)$, $v = (L^2, 2L^2 + 1)$, $w = (2L^2 + 1, L^2)$ and $z = (2L^2 + 1, 2L^2 + 1)$.

The vertices u and v are of distance $2L^2 + 1$ in $\mathbb{Z}^d$, connected by a unique path $P_{u,v}$. Let $P_{u,v} = (u_0, u_1, \dots, u_{L^2+1})$, with $u_0 = u$ and $u_{L^2+1} = v$. As H is a rewiring for each (u_i, u_{i+1}) there is a path $Q_{u_i,u_{i+1}}$ in H of length at most L. Moreover, this path stays within distance L^2 of the edge (u_i, u_{i+1}). Concatenating these paths and erasing possible cycles or backtracking gives rise to a path $\tilde{Q}_{u,v}$ in H connecting u and v which stays within distance L^2 of $P_{u,v}$. We repeat this argument for the pairs (v, z), (z, w) and (w, u).

The paths $\tilde{Q}_{u,v}$, $\tilde{Q}_{v,z}$, $\tilde{Q}_{z,w}$ and $\tilde{Q}_{w,u}$ concatenate to a returning walk W in H inside the tubular L^2-neighborhood of the square $P_{u,v} \cup P_{v,z} \cup P_{z,w} \cup P_{w,u}$. This walk is clearly not nullhomotopic, hence we find a cycle after removal of backtracking. This cycle is clearly inside F.

In higher dimension we consider the edge graph C of a d-dimensional cube. The graph C has 2^d vertices and $(d-1)2^{d-1}$ edges.

In $\mathbb{Z}^d$ we also consider a d-dimensional cube with sides $L^2 + 1$, and build the paths $Q_{u,v}$ between neighboring vertices. This way we get a continuous map from C to F which is a local embedding everywhere except maybe the vertices of C. Moreover every $Q_{u,v}$ path contains an edge $e_{u,v}$ that is not in the union of the other paths, simply because all other $Q_{u',v'}$ paths stay close to their respective $P_{u',v'}$.

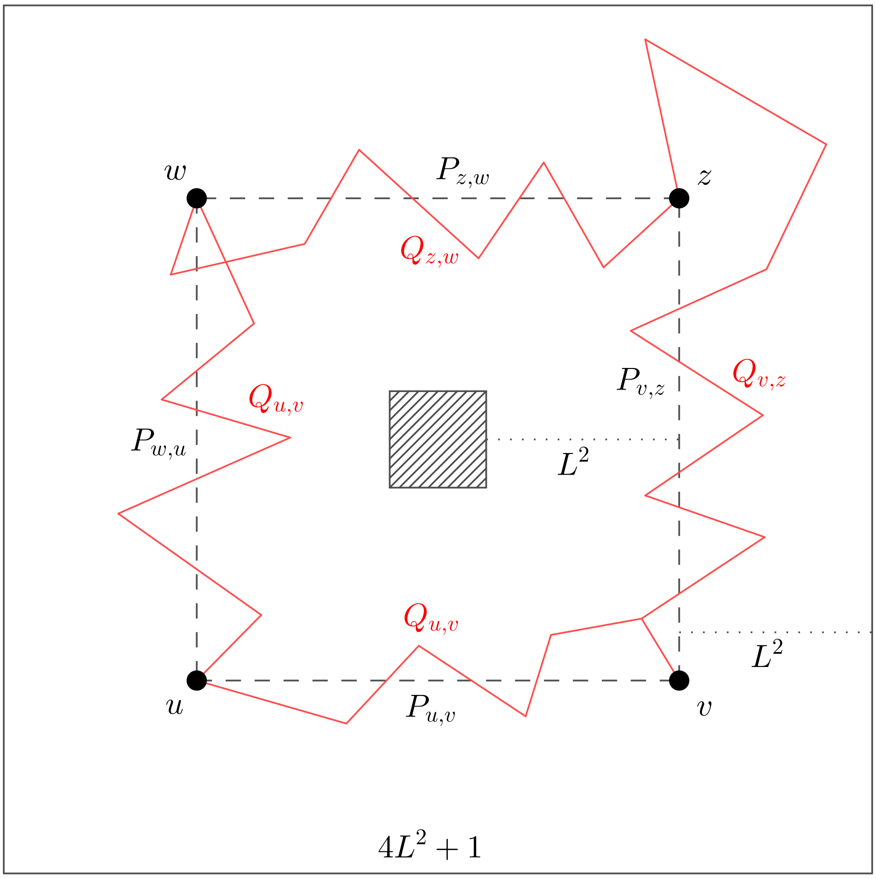

Figure 1: Finding a cycle

We select the complement of a spanning tree of C. This has $(d-1)2^{d-1} - (2^d - 1) = (d-2)2^{d-1}+1$ edges. We can delete the edge $e_{u,v}$ corresponding to all selected edges without breaking connectedness. □

Lemma 3.2.7. Let H be an L-rewiring of $\mathbb{Z}^d$. Then the average degree of H is at least

$$2 + \frac{d-2}{(2L^2+1)^d}.$$

By average degree of an individual rewiring we mean the limit of the average of degrees along a Følner sequence. When H is a random instance of an invariant random rewiring this limit exists with probability 1. For general H our proof shows that the liminf is at least $2 + (d-2)/(2L^2+1)^d$.

Proof. We will choose a Følner sequence F_n in $\mathbb{Z}^d$ and compute the limit

$$\lim_{n\to\infty} \frac{\sum_{v\in F_n} \deg_H(v)}{|F_n|}.$$

Let H_n be the graph we get by collapsing the entire complement of F_n to one point v^*. This graph is connected, and

$$\frac{\sum_{v\in F_n} \deg_H(v)}{|F_n|} = \frac{2|E(H_n)| - \deg_{H_n}(v^*)}{|F_n|} = \frac{2|E(H_n)| - |\partial_E^H(F_n)|}{|F_n|}.$$

By Lemma 3.2.5 we have

$$\lim_{n\to\infty} \frac{|\partial_E^H(F_n)|}{|F_n|} = 0.$$

Let $K = 4L^2 + 1$, and let F_n be a box with sides of size $K \cdot n$. By tiling F_n with n^d K-by-K boxes and using Lemma 3.2.6 we can delete $n^d\big((d-2)2^{d-1}+1\big)$ edges preserving connectedness of H_n.

This implies that there are at least $|F_n| + n^d\big((d-2)2^{d-1}+1\big)$ edges in H_n, so the average degree is at least

$$2 + \frac{2n^d\big((d-2)2^{d-1}+1\big)}{\big((4L^2+1)n\big)^d} \geq 2 + \frac{d-2}{(2L^2+1)^d}.$$

□

3.2.4 The ergodic theorem

We will use the pointwise ergodic theorem of Lindenstrauss for amenable groups from [Lin01].

Let G act ergodically on a measure space (X,μ), and let $\varphi \in L^1(X,\mu)$. For a finite subset $F \subseteq G$ we denote the average of φ from $x \in X$ over F to be

$$A(F,\varphi)(x) = \frac{1}{|F|}\sum_{g\in F} \varphi(gx).$$

In G is amenable, the Følner sequence F_n is said to be *tempered*, if there is a $C > 0$ such that for all n

$$\left|\bigcup_{k\leq n} F_k^{-1}F_n\right| \leq C|F_{n+1}|.$$

Theorem 3.2.8 (Lindenstauss, Theorem 1.3 in [Lin01])**.** Let G be an amenable group acting ergodically on a measure space (X,μ), and let F_n be a tempered Følner sequence. Then for any $\varphi \in L^1(X,\mu)$,

$$\lim_{n\to\infty} A(F_n,\varphi)(x) = \int \varphi(x)\, d\mu(x) \quad \text{for } \mu\text{-a.e. } x \in X.$$

Using Theorem 3.2.8 component-wise we can connect the average of degrees over a Følner sequence to the expected degree of the root in an invariant random rewiring. The following proposition is a direct corollary of Proposition 3.2.7.

Proposition 3.2.9. An invariant random L-rewiring of $\mathbb{Z}^d$ has average degree at least

$$2 + \frac{2(d-1)}{(2L^2+1)^d}.$$

Remark. The theorem of Lindenstrauss assumes the Følner sequence to be tempered, but this is satisfied by the sets chosen in the proof of Lemma 3.2.7.

In fact we can get by with the weaker [Lin01, Theorem 1.1], which only states convergence in L^1. Indeed, by analyzing the proof of Lemma 3.2.7 we can see that the lower bounds on the average degrees are uniform among all L-rewirings.

3.2.5 Computing the distortion

Now we put together Propositions 3.2.4 and 3.2.9 to prove Theorem 3.2.1.

Proof of Theorem 3.2.1. Proposition 3.2.4 shows that if we aim to find a rewiring which gets $(d-1)/L^{2d}$ close to the cost we can keep the bi-Lipschitz constant as low as $(4d+1)L$, so letting $x = (d-1)/L^{2d}$ we get

$$L = \left(\frac{d-1}{x}\right)^{\frac{1}{2d}}.$$

Turning this around for the distortion we get

$$\delta_f(x) \leq (4d+1)\left(\frac{d-1}{x}\right)^{\frac{1}{2d}} + 1.$$

The $(+1)$ at the end of our bound is added because the distortion is integer valued. This is more convenient than using the ceiling function, and since we are only interested in the growth up to multiplicative constants it will not matter anyway.

On the other hand Proposition 3.2.9 shows that an L-rewiring cannot get closer to the cost than

$$\frac{d-1}{(2L^2+1)^d}.$$

For L large enough we have

$$\frac{d-1}{(2L^2+1)^d} \leq \frac{d-2}{(2L^2)^d}.$$

Letting $x = (d-2)/(2L^2)^d$ we compute what this implies for the distortion, namely

$$L = \frac{1}{\sqrt{2}}(d-2)^{\frac{1}{2d}}\left(\frac{1}{x}\right)^{\frac{1}{2d}}.$$

So for x small enough

$$\frac{1}{\sqrt{2}}(d-2)^{\frac{1}{2d}}\left(\frac{1}{x}\right)^{\frac{1}{2d}} \leq \delta_f(x) \leq (4d+1)(d-1)^{\frac{1}{2d}}\left(\frac{1}{x}\right)^{\frac{1}{2d}} + 1.$$

This concludes our proof of Theorem 3.2.1. □

3.2.6 Distortion with subgraphings

It is a natural question to ask what happens if in the definition of cost we restrict our attention to generating subgraphings, or in the group action setting sets A which are chosen as subsets of X_S. It is still an open question whether or not one can find a free p.m.p. action of a group where the two notions of cost are distinct.

For the distortion function considering only subsets of X_S for generation immediately reduces the bi-Lipschitz condition to the one-sided containment

$$X_S \subseteq (A \cup A^{-1} \cup E)^L.$$

Without repeating the details we would like to point out that the methods of Propositions 3.2.4 and 3.2.9 still work, but the results will be slightly different.

In the construction the long edges can no longer be used, so instead we can introduce an L-by-L grid of short edges. When establishing the lower bound we can use $(4L+1)$-by-$(4L+1)$ cubes because of the stronger condition. Otherwise repeating the same arguments one gets the following.

Proposition 3.2.10. If one only considers generating subsets $A \subseteq X_S$ in the definition of groupoid cost and the distortion function, then the distortion function of any free p.m.p. action f of $\mathbb{Z}^d$ becomes

$$\delta'_f(x) \asymp \left(\frac{1}{x}\right)^{\frac{1}{d}}.$$

3.3 Distortion of lamplighter groups

In this section we turn our attention to lamplighter groups.

3.3.1 Lamplighter groups

A lamplighter group is the wreath product $\Gamma = G \wr B = \bigoplus_B G \rtimes B$, where B and G are the *base* and *lamp* groups respectively. We will assume both B and G to be finitely generated, and think of $G \wr B$ as a lamplighter moving around the graph $\mathrm{Cay}(B, S_B)$ and changing the status of lamps (described by G) positioned at every vertex.

Formally elements of Γ can be written as pairs (φ, b), with $\varphi : B \to G$ having finitely many $x \in B$ with $\varphi(x) \neq e_G$, and $b \in B$. The action of b on the functions φ is by translation.

The natural way to finitely generate Γ is allowing the lamplighter to walk around on the base graph with elements of S_B and to adjust the lamps where he currently is by applying elements of S_G. That is, we embed B into Γ via the map $b \mapsto (\epsilon, b)$ where ϵ is the constant identity map on B. We also embed G into Γ via the map $g \mapsto (\varphi_g, e_B)$ where

$$\varphi_g(x) = \begin{cases} g & \text{if } x = e_B, \\ e_G & \text{if } x \neq e_B. \end{cases}$$

To ease notation from now on we will think of B and G as subgroups of Γ embedded as above. Now the natural generating set for Γ is simply $S_\Gamma = S_B \cup S_G$.

Our next theorem bounds the distortion of lamplighters when G is finite.

Theorem 3.3.1. Let G be finite and $B = \langle S_B \rangle$ finitely generated and infinite. Let $\Gamma = G \wr B$ and f be a free p.m.p. action of Γ. Then $\delta_f(x) \leq C \cdot \log(1/x)$ for some $C \in \mathbb{R}$.

Under the above assumptions Γ has fixed price 1. This is because $\bigoplus_B G$ has enough commutation to have fixed price 1, and as such Γ has an infinite fixed price 1 normal subgroup, which implies that Γ itself has fixed price 1. We will not go into more details here, as our construction for the $C \cdot \log(1/x)$ distortion will also imply fixed price 1.

Γ is not necessarily amenable, so not all free actions are weakly equivalent. Instead we will make use of a result of Abért and Weiss [AW13, Theorem 1], stating that all free actions weakly contain Bernoulli actions. Using the monotonicity of the distortion function it suffices to prove Theorem 3.3.1 for some Bernoulli action.

In the random process language this translates to constructing a *factor of i.i.d.* L-rewiring of $\mathrm{Cay}(\Gamma, S_\Gamma)$ with sufficiently small average degree.

3.3.2 Factor of i.i.d. rewirings

We will consider the Bernoulli shift of Γ with base space $\big([0,1], u\big)$ where u is the uniform measure. That is the space $[0,1]^\Gamma$ with the measure u^Γ. The natural translation action of Γ obviously preserves the product measure.

As before, let $\mathrm{Rew}(\Gamma, S, L)$ denote the space of L-rewirings of Γ. A random rewiring is a measure μ on $\mathrm{Rew}(\Gamma, S, L)$. We say that μ is a factor of i.i.d. if we can find a measurable map $\psi : [0,1]^\Gamma \to \mathrm{Rew}(\Gamma, S, L)$ that intertwines both the actions and the measures:

$$\gamma\psi(\omega) = \psi(\gamma\omega) \text{ for a.e. } \omega \in [0,1]^\Gamma, \text{ and } \psi_*(u^\Gamma) = \mu.$$

Since u^Γ is invariant with respect to the Γ action any factor of it is also invariant.

Intuitively we can think of factor of i.i.d. rewirings as having independent uniform random decorations on the vertices of $\mathrm{Cay}(\Gamma, S_\Gamma)$ and then deciding for each pair of vertices if they are connected in the rewiring. However the decision has to be deterministic, based on what the labeled graph looks like, "rooted" at the pair.

Even more informally, every possible edge just looks around. Based on what the S_Γ-edge-labeled and $[0,1]$-vertex-decorated graph looks like from there, it makes a decision whether or not it should belong to the rewiring. All edges use the same rule, so if they see the same, their decision has to be the same.

If we have a factor of i.i.d. rewiring μ, then we can use it to define a measurable rewiring of the space $\big([0,1]^\Gamma, u^\Gamma\big)$. We define a generating set A of the groupoid $[0,1]^\Gamma \times \Gamma$. For $\gamma \in \Gamma$ set $(\omega, \gamma) \in A$ if in the rewiring $\psi(\omega)$ the arrow labeled by γ starting from the identity is chosen to be in the rewiring. The measurability of this generating set is provided by the rewiring being factor of i.i.d. and half the expected degree of the identity in μ will be the measure of A.

In the rest of this chapter we focus on constructing factor of i.i.d. rewirings of lamplighter groups such that the expected degree and the distortion are both reasonably small.

3.3.3 Constructing factor of i.i.d. rewirings

We proceed to describe our construction. Fix $n \in \mathbb{N}$ as a radius around the lamplighter, and group up the elements of Γ into equivalence classes. Two elements $\gamma_1, \gamma_2 \in \Gamma$ are *equivalent* if the lamplighter is at the same position and the status of the lamps agrees outside the ball of radius n around the lamplighter. Formally $\gamma_1 \sim \gamma_2$ if $\gamma_1 \cdot \gamma_2^{-1} = (\varphi, e_B)$ with $\varphi(b) = e_G$ for all $b \notin B_{\mathrm{Cay}(B,S_B)}(n, e)$.

Let $N = |B_{\mathrm{Cay}(B,S_B)}(n, e)|$ denote the size of the n-ball in $\mathrm{Cay}(B, S_B)$. Every equivalence class contains $|G|^N$ group elements.

We subdivide the classes based on the $(n-1)$-neighborhood of the marker. We say two elements γ_1, γ_2 of Γ are *strongly equivalent*, if the lamp configurations agree outside the ball of radius $(n-1)$ around the marker. (We still assume the marker to be positioned at the same place.) Formally this means $\gamma_1 \cdot \gamma_2^{-1} = (\varphi, e_B)$ with $\varphi(b) = e_G$ for all $b \notin B_{\mathrm{Cay}(B,S_B)}(n-1, e)$.

Let $N' = |B_{\mathrm{Cay}(B,S_B)}(n-1, e)|$ denote the size of the $(n-1)$-ball in $\mathrm{Cay}(B, S_B)$.

Proposition 3.3.2. There exists a factor of i.i.d. $(2N \cdot |G| + 1)$-rewiring of $\mathrm{Cay}(\Gamma, S_\Gamma)$ with average degree at most

$$2 + \frac{4|S_B|}{|G|^{N'}}.$$

Before we turn to the proof we need a lemma on spanning trees of Cartesian powers of finite Cayley graphs. The Cartesian product of two graphs $G_i = (V_i, E_i)$, $i \in \{1, 2\}$ is defined on the vertex set $V_1 \times V_2$. The pairs (u_1, u_2) and (v_1, v_2) are connected if either $u_1 = v_1$ and $(u_2, v_2) \in E_2$, or $u_2 = v_2$ and $(u_1, v_1) \in E_1$.

Lemma 3.3.3. Let $G = \langle S_G \rangle$ be a finite group, and let C_N denote the N-fold Cartesian power of $\mathrm{Cay}(G, S_G)$. Then we can find a spanning tree T_N of C_N and a vertex $o_N \in C_N$ which can be reached from any vertex of C_N using at most $N \cdot |G|/2$ edges from T_N.

Proof. First start with any spanning tree T_1 of $C_1 = \mathrm{Cay}(G, S_G)$. Let K denote the smallest integer such that one can find a vertex $o_1 \in \mathrm{Cay}(G, S_G)$ with all other vertices being of distance at most K from o_1 in T_1.

We claim $K \leq \lfloor |G|/2 \rfloor$. We continue to remove leaves of the tree T_1 until we are left with either a single vertex o_1, or a single edge whose either endpoint suffices as o_1. Using the fact that in each step we remove at least 2 vertices shows our claim, and that our bound is sharp exactly when T_1 is a path.

When $N = 2$ we construct T_2 by copying T_1 to all *vertical* subgraphs $\{u\} \times C_1 \cong C_1$, and connecting them by a further *horizontal* copy of T_1 on the subgraph $C_1 \times \{o_1\}$. We set $o_2 = (o_1, o_1)$. From any point (u, v) we can get to (u, o_1) in T_2 with at most K steps using the appropriate vertical T_1 tree. Then we can get to (o_1, o_1) with at most K steps using the horizontal T_1. That is, we can move to o_2 with at most $2K$ steps.

For general N we define T_N and o_N recursively, and proceed by induction. We copy T_{N-1} to all subgraphs $\{v\} \times C_{N-1}$, connect them by a copy of T_1 on the graph $C_1 \times \{o_{N-1}\}$, and set $o_N = (o_1, o_{N-1})$. The above argument shows that we can get to o_N in $K + (N-1)K \leq N \cdot |G|/2$ steps. $\square$

Corollary 3.3.4. In the setting of Lemma 3.3.3 we can construct the spanning tree T_N and the vertex o_N as a factor or i.i.d. random process using uniform $[0, 1]$ labels such that o_N can be reached from any vertex of C_N using at most $N \cdot |G|/2$ edges from T_N.

Proof. We can make our construction into an invariant random process by simply listing all translates of (T_N, o_N) by $\mathrm{Aut}(C_N)$, and choosing one uniformly at random. On finite graphs any invariant process is a factor of i.i.d. process as well. $\square$

Remark. Observe that in fact the proof of Lemma 3.3.3 shows more. For each N we set K_N to be the smallest integer such that there is a spanning tree T_N of C_N, and a vertex $o_N \in C_N$ with all other vertices at most distance K_N in T_N from o_N. We argued that $K_{N+1} \leq K_N + K_1$, but in fact the same idea yields $K_{N+M} \leq K_N + K_M$. So in fact K_N as a sequence is subadditive, so Fekete's lemma implies that K_N/N converges. For our purposes though it is enough to see that it is bounded.

Proof of Proposition 3.3.2. On each equivalence class $[\gamma]$ introduce the finite graph $C_{[\gamma]}$ by connecting elements where the status of the lamps differs only at one position by some $g \in S_G$. This difference need not happen at the position of the lamplighter, but by the definition of equivalence it must happen within distance n of it. In the case when $G = Z_2$ this is simply an N-dimensional cube, and in general it is the N-fold Cartesian power of $\mathrm{Cay}(G, S_G)$.

Each edge of $C_{[\gamma]}$ travels distance at most $2n+1$ with respect to the standard generators of Γ, as we can take n steps with the lamplighter to the position we need to change, apply the change and walk back.

We will choose spanning trees of these $C_{[\gamma]}$, and then connect them by adding edges between the classes. Both will be factor of i.i.d. constructions, so their union will also be a factor of i.i.d. Indeed, a uniform $[0, 1]$ label as a measure space is isomorphic to $[0, 1]^2$, so we can assume to have two i.i.d. random labels per vertex, and use them for the first and second part of the construction respectively.

An important observation is that the equivalence classes form an invariant partition with respect to the Γ action on $\mathrm{Cay}(\Gamma, S_\Gamma)$. Generators from S_G do not move the lamplighter, and only change the lamp at the marker, so they only permute within the classes. On the other hand generators from S_B move the lamplighter and the whole lamp configuration with it, so the classes themselves are permuted, but being equivalent remains invariant.

Each element γ can identify its whole equivalence class, which is inside some bounded radius ball in $\mathrm{Cay}(\Gamma, S_\Gamma)$. Using the random labels on these vertices and Corollary 3.3.4 we can get a random spanning tree with a distinguished vertex on each equivalence class. The fact that the equivalence classes form an invariant partition ensures that the union of all the spanning trees is a factor of i.i.d. spanning forest.

For the second part of our construction in each strong equivalence class we chose a point uniformly at random – this is factor of i.i.d. – from where we include the edges labeled by s and s^{-1}. This concludes our construction.

The following two lemmas contain the computation of the bi-Lipschitz constant and the average degree, thus finishing the proof of our proposition. □

Lemma 3.3.5. The random graph constructed in Proposition 3.3.2 is connected and is a $2N \cdot |G| + 1$ rewiring of $\mathrm{Cay}(\Gamma, S_\Gamma)$.

Proof. Indeed the edges labeled by S_G connect vertices within the equivalence classes, so their endpoints can be connected using at most $N \cdot |G|$ edges from the spanning forest.

The edges labeled by S_B are a bit more tricky. Fix $s \in S_B$. From any vertex $\gamma = (\varphi, b)$ we take at most $N \cdot |G|$ steps to reach a strongly equivalent vertex (φ', b) where the edge labeled by s is included. That means we rearranged the lamps within distance n at most $N \cdot |G|$ times such that in the end we got a configuration φ' which agrees with φ within distance $(n - 1)$ of the marker b.

We move along the edge labeled s starting from (φ', b), which means we move the marker to $b \cdot s$. Then we plan to move from $(\varphi', b \cdot s)$ to $(\varphi, b \cdot s)$, so we claim the two are equivalent. Indeed, as φ and φ' agree outside the $(n - 1)$-ball around b they also agree outside the ball of radius n around $b \cdot s$. This means we can fix the difference between $(\varphi', b \cdot s)$ and $(\varphi, b \cdot s)$ using the spanning tree on the equivalence class, which we can do using at most $N \cdot |G|$ edges.

The above path used at most $2N \cdot |G| + 1$ edges. Collecting our constants we see that the random spanning subgraph has bi-Lipschitz distortion at most $\max\{2n+1, N \cdot |G|, 2N \cdot |G| + 1\} = 2N \cdot |G| + 1$. □

Lemma 3.3.6. The average degree of the random graph constructed in Proposition 3.3.2 is at most

$$2 + \frac{4|S_B|}{|G|^{N'}}.$$

Proof. The spanning forest contributes an average degree of

$$2\frac{|G|^N - 1}{|G|^N}.$$

We also include $2|S_B|$ edges for each strong equivalence class. The number of strong equivalence classes inside one equivalence class is $|G|^{N-N'}$. The total average degree is

$$2\frac{|G|^N - 1}{|G|^N} + 2\frac{2|S_B| \cdot |G|^{N-N'}}{|G|^N} \leq 2 + \frac{4|S_B|}{|G|^{N'}}.$$

□

3.3.4 Bounding the distortion in lamplighters

Proposition 3.3.2 almost suffices to show Theorem 3.3.1, as we can trivially bound N in terms of N', namely $N \leq |S_B| \cdot N$. There is a slight problem however, namely that we only constructed cheap rewirings for the bi-Lipschitz constants $2N \cdot |G| + 1$. Recall that N is the size of the ball of radius n in the Cayley graph of the base group. As we increase n the gaps between the corresponding N are typically unbounded.

In our proof however we can replace N and N' as follows. Instead of the ball $B_{\mathrm{Cay}(B,S_B)}(n-1, e)$ we consider a fixed subset $F \subseteq B_{\mathrm{Cay}(B,S_B)}(n-1, e)$, and define strong equivalence as having the same configuration outside the F-window around the lamplighter. Formally this means $\gamma_1 \cdot \gamma_2^{-1} = (\varphi, e_B)$ with $\varphi(b) = e_G$ for all $b \notin F$. We define equivalence similarly, but with the window $F \cdot S_B$ instead of using $B_{\mathrm{Cay}(B,S_B)}(n, e)$. As before we have the trivial bound on the size, $|F \cdot S_B| \leq |F| \cdot |S_B|$. For every integer L we can choose some F with $|F| = L$ within the ball of radius $L - 1$, which leads to the following strengthening of our proposition.

Proposition 3.3.7. For any $L \in \mathbb{N}$ there exists a factor of i.i.d. $(2L \cdot |S_B| \cdot |G| + 1)$-rewiring of $\mathrm{Cay}(\Gamma, S_\Gamma)$ with average degree at most

$$2 + \frac{4|S_B|}{|G|^L}.$$

Proof of Theorem 3.3.1. We use Proposition 3.3.7. The distortion is at most $2L \cdot |S_B| \cdot |G| + 1$, while the average degree is at most $2 + 4|S_B|/|G|^L$. As we get this for all L the gaps between our bi-Lipschitz constants are bounded by $2|S_B| \cdot |G|$. If we set $y = 2|S_B|/|G|^L$ we get

$$L = \log_{|G|}\left(\frac{2|S_B|}{y}\right).$$

Running L through $\mathbb{N}$, for these specific y we get

$$\delta_f(y) \leq 2 \cdot |S_B| \cdot |G| \log_{|G|} \left(\frac{2|S_B|}{y} \right) + 1.$$

In general using the monotonicity of $\delta_f(x)$ and the bounded gap between the integer values we get for all x that

$$\delta_f(x) \leq 2 \cdot |S_B| \cdot |G| \log_{|G|} \left(\frac{2|S_B|}{x} \right) + 1 + 2|S_B| \cdot |G|.$$

This concludes our proof of Theorem 3.3.1. $\square$

Bibliography

[Abé+17] Miklós Abért, Nicolas Bergeron, Ian Biringer, Tsachik Gelander, Nikolay Nikolov, Jean Raimbault, and Iddo Samet. On the growth of L^2-invariants for sequences of lattices in Lie groups. *Annals of Mathematics* 185.3 (2017), pp. 711–790.

[AE11] Miklós Abért and Gábor Elek. The space of actions, partition metric and combinatorial rigidity. *arXiv preprint arXiv:1108.2147* (2011).

[AE12] Miklós Abért and Gábor Elek. Dynamical properties of profinite actions. *Ergodic Theory and Dynamical Systems* 32.6 (2012), pp. 1805–1835.

[AGN17] Miklós Abért, Tsachik Gelander, and Nikolay Nikolov. Rank, combinatorial cost, and homology torsion growth in higher rank lattices. *Duke Mathematical Journal* 166.15 (2017), pp. 2925–2964.

[AGV14] Miklós Abért, Yair Glasner, and Bálint Virág. Kesten's theorem for invariant random subgroups. *Duke Mathematical Journal* 163.3 (2014), pp. 465–488.

[AJN11] Miklós Abért, Andrei Jaikin-Zapirain, and Nikolay Nikolov. The rank gradient from a combinatorial viewpoint. *Groups, Geometry, and Dynamics* 5.2 (2011), pp. 213–230.

[AN12] Miklós Abért and Nikolay Nikolov. Rank gradient, cost of groups and the rank versus Heegaard genus problem. *Journal of the European Mathematical Society* 14.5 (2012), pp. 1657–1677.

[AT17] Miklós Abért and László Márton Tóth. Uniform rank gradient, cost and local-global convergence. *arXiv preprint arXiv:1710.10431* (2017).

[AW13] Miklós Abért and Benjamin Weiss. Bernoulli actions are weakly contained in any free action. *Ergodic theory and dynamical systems* 33.2 (2013), pp. 323–333.

[BGŠ03] Laurent Bartholdi, Rostislav Grigorchuk, and Zoran Šunić. Branch groups. *Handbook of algebra* 3 (2003), pp. 989–1112.

[BT18] Ferenc Bencs and László Márton Tóth. Invariant random subgroups of groups acting on rooted trees. *arXiv preprint arXiv:1801.05801* (2018).

[BS01] Itai Benjamini and Oded Schramm. Recurrence of distributional limits of finite planar graphs. *Electronic Journal of Probability* 6 (2001).

[BGN15] Mustafa Gökhan Benli, Rostislav Grigorchuk, and Tatiana Nagnibeda. Universal groups of intermediate growth and their invariant random subgroups. *Functional Analysis and Its Applications* 49.3 (2015), pp. 159–174.

[BR11] Béla Bollobás and Oliver Riordan. Sparse graphs: metrics and random models. *Random Structures & Algorithms* 39.1 (2011), pp. 1–38.

[Bor+08] Christian Borgs, Jennifer T. Chayes, László Lovász, Vera T. Sós, and Katalin Vesztergombi. Convergent sequences of dense graphs I: Subgraph frequencies, metric properties and testing. *Advances in Mathematics* 219.6 (2008), pp. 1801–1851.

[Bow10] Lewis Bowen. A measure-conjugacy invariant for free group actions. *Annals of Mathematics* (2010), pp. 1387–1400.

[Bow12] Lewis Bowen. Invariant random subgroups of the free group. *arXiv preprint arXiv:1204.5939* (2012).

[BGK15] Lewis Bowen, Rostislav Grigorchuk, and Rostyslav Kravchenko. Invariant random subgroups of lamplighter groups. *Israel Journal of Mathematics* 207.2 (2015), pp. 763–782.

[Car15] Alessandro Carderi. Ultraproducts, weak equivalence and sofic entropy. *arXiv preprint arXiv:1509.03189* (2015).

[CGS17] Alessandro Carderi, Damien Gaboriau, and Mikael de la Salle. Non-standard limits of graphs and their orbit equivalence invariants. *in preparation* (2017).

[DG14] Artem Dudko and Rostislav Grigorchuk. On diagonal actions of branch groups and the corresponding characters. *arXiv preprint arXiv:1412.5476* (2014).

[DM17] Artem Dudko and Kostya Medynets. On Invariant Random Subgroups of Block-Diagonal Limits of Symmetric Groups. *arXiv preprint arXiv:1711.01653* (2017).

[Ele07] Gábor Elek. The combinatorial cost. *L'Enseignement Mathématique* 53.3-4 (2007), pp. 225–235.

[ES04] Gábor Elek and Endre Szabó. Sofic groups and direct finiteness. *Journal of Algebra* 280.2 (2004), pp. 426–434.

[Gab00] Damien Gaboriau. Coût des relations d'équivalence et des groupes. (French) [Cost of equivalence relations and of groups]. *Inventiones mathematicae* 139.1 (2000), pp. 41–98.

[Gri00] Rostislav Grigorchuk. Just infinite branch groups. In: *New horizons in pro-p groups*. Springer, 2000, pp. 121–179.

[Gro99] Mikhael Gromov. Endomorphisms of symbolic algebraic varieties. *Journal of the European Mathematical Society* 1.2 (1999), pp. 109–197.

[HLS14] Hamed Hatami, László Lovász, and Balázs Szegedy. Limits of locally–globally convergent graph sequences. *Geometric and Functional Analysis* 24.1 (2014), pp. 269–296.

[Kal02] Olav Kallenberg. *Foundations of Modern Probability*. Probability and Its Applications. Springer New York, 2002.

[Kec10] Alexander S. Kechris. *Global aspects of ergodic group actions*. 160. American Mathematical Soc., 2010.

[Lac05] Marc Lackenby. Expanders, rank and graphs of groups. *Israel Journal of Mathematics* 146.1 (2005), pp. 357–370.

[Lin01] Elon Lindenstrauss. Pointwise theorems for amenable groups. *Inventiones mathematicae* 146.2 (2001), pp. 259–295.

[Lov12] László Lovász. *Large networks and graph limits*. Vol. 60. American Mathematical Soc., 2012.

[OW80] Donald Ornstein and Benjamin Weiss. Ergodic theory of amenable group actions. I: The Rohlin lemma. *Bulletin of the American Mathematical Society* 2.1 (1980), pp. 161–164.

[Sch81] Klaus Schmidt. Amenability, Kazhdan's property T, strong ergodicity and invariant means for ergodic group-actions. *Ergodic Theory and Dynamical Systems* 1.2 (1981), pp. 223–236.

[TT14] Simon Thomas and Robin Tucker-Drob. Invariant random subgroups of strictly diagonal limits of finite symmetric groups. *Bulletin of the London Mathematical Society* 46.5 (2014), pp. 1007–1020.

[TT18] Simon Thomas and Robin Tucker-Drob. Invariant random subgroups of inductive limits of finite alternating groups. *Journal of Algebra* 503 (2018), pp. 474–533.

[Ver12] Anatolii Vershik. Totally nonfree actions and the infinite symmetric group. *Moscow Mathematical Journal* 12.1 (2012), pp. 193–212.