



Faculté des Sciences
Institut de mathématiques
Rue Emile-Argand 11, 2000 Neuchâtel

Large scale geometry of graphs

Thèse

présentée à la Faculté des Sciences
pour l'obtention du grade de docteur ès Sciences en mathématiques

par

Tom Kaiser

Acceptée sur proposition du jury :

Miklós Abert	MTA Alfréd Rényi	Rapporteur
Ana Khukhro	University of Cambridge	Rapporteuse
Alexander Kolpakov	Université de Neuchâtel	Rapporteur
Alain Valette	Université de Neuchâtel	Directeur de thèse

Soutenue le 29 juin 2020

IMPRIMATUR POUR THESE DE DOCTORAT

La Faculté des sciences de l'Université de Neuchâtel
autorise l'impression de la présente thèse soutenue par

Monsieur Tom KAISER

Titre:

“Coarse properties of graphs”

sur le rapport des membres du jury composé comme suit:

- Prof. Alain Valette, directeur de thèse, Université de Neuchâtel, Suisse
- Prof. Aleksandr Kolpakov, Université de Neuchâtel, Suisse
- Dr Ana Khukhro, University of Cambridge, United Kingdom
- Prof. ass. Miklos Abért, MTA Alfréd Rényi Institute of Mathematics, Budapest, Hongrie

Neuchâtel, le 7 septembre 2020

Le Doyen, Prof. A. Bangerter



Remerciements

Of all nerve-racking duties the Gods of Thunder have imposed on the shoulders of the humble PhD student, the most essential, by far, must be the indictment of the acknowledgements. And so, as many before him, the author will attempt the last of the Herculean tasks, in a desperate attempt to make his Alma Mater proud. He knows, he will have to tread over the corpses of the fallen. But he has faith since those before him, those who have taught him and shared with him, in a stroke of absolute generosity, all they know. They, heroes who walked through the glorious halls of Unimail, of a long lost past, have set the path to victory. They are still here, for those who are willing to lend them their ears (and willing to make an offering of Quöllfrisch and hyperbolic cheese).

Dear Alain, heavenly Father, thank you for welcoming me under your benign wings, for the opportunity to pursue a PhD, for your trust and advice. You understand very well that a mathematician thrives in an amicable and warm environment, which must be the reason you welcomed Alexandre, Laurent, and me at Portiragnes, institute of mathematics, where, redeemed by the Mediterranean and propelled by the golden combo of Petanque-Pastis, we thrust ourselves into the sea and the world of spectral measures.

I would like to thank Miklós Abert, Ana Khukhro and Sasha Kolpakov for agreeing to be part of my thesis committee, for their remarks on the text, their interesting questions and kind comments.

If the institute of mathematics at Neuchâtel should be known for something, let it not be mathematics, but the kind embrace of its inhabitants. In particular, I (and the university in general) am indebted to Neuchâtel's own Castor and Pollux (Carsten and Ana) for their welcoming hearts.

Let's start with Ana, the original homo universalis (Soviet, turned German, turned Welsh, soon to be Swiss (?)), thank you for probably being the nicest person on earth. Thank you for introducing me to everyone when I first arrived, and making me feel at home. Thank you for having great mathematical taste and for being on my thesis committee. Polyglot as you are, I expect your Dutch to have fully blossomed when I return for a visit.

The first time Carsten visited my apartment, I threw a bottle opener at him. Being protective of the (my) windows behind him, he caught it with his head. This is Carsten, always ready to help, and with a firm skull of weaponised concrete. Thank you for being a more than wonderful friend, understanding and caring. For liking the simple pleasure of a beer on a sunny afternoon or drizzly evening. Let's hope Marie-Ange does not put too much sense into you and you will still appreciate these things when we see each other again.

Then of course Geert Vanpaemel, a mysterious man with surprising friendships and a turbulent love life, who evades comprehension. May you live a long and fruitful life. Thanks for

your silent support.

Pamela, hidden flower of the D-wing, apart from your friendship (which I am truly very grateful for), thank you for your lakeside BBQs, the evenings at your place, and your helping hands during the apéro of my thesis defence (without which it wouldn't even have been half as good). I wish you the very best for your last thesis year and more importantly many happy moments with Tobias # AdorableTogether. Don't move too far away please!

Tobias, most kind-hearted person alive? If not at least the sexiest underwaterhockeyplayer Neuchâtel has ever seen. Thank you for introducing me to Laurent Gilbert, even though we never met in person, and your general everyday kindness. Also, I liked your Yankee hat, please get a new one!

Dear Edouard, in your acknowledgements you compared Valentin to Christ. Last time I saw you, apparently you also took a liking to His image, with long blond locks and a prophetic understanding of my admiration for Marine Le Pen¹. Let's hope you don't become a second Andre Agassi, even though your charming smile would make up for even the shiniest of heads. I remember the week you slept in my kitchen, and the stories we shared about le saint prépuce. You have always been a great influence on my 'Vocabularius Gallicus'.

Dear Valentin, for two years we shared an office. During this time you insulted me multiple times, decorated the office with onions and had inappropriate (and sexually explicit) conversations with your plants. I think it's clear I liked you from the start. Thank you for introducing me to every good bar/beer in Switzerland, taking me on excellent hikes and for teaching me how to organise said hikes myself. I wish you the very best for the future, be it in Switzerland or Portugal. "Tout est bon avec excès"!

First there was Carl-Erik, then there was the Carl-Erik Express. Because of you visiting, BFM has become a lot easier. Thx for our 10 o'clock breaks with Carsten and for making the excellent choice of moving to the pivot of my travels home or to and from Montreal. I hope you will not lose the charming way you smile when you make a joke. See you someday in Paris. Or you could visit Belgium again; this would be the easiest solution for me :)

Dear Sasha, thank you for the apéros, your stories, and Geometric Group Theory.

Dear Félix, thank you for your fascinatingly weird sense of humour and Algebraic Topology.

There once was an Iranian girl in Neuchâtel who did some mysterious and very non-trivial mathematics. You always impressed me, Sanaz, so good luck with everything, with Sven and in your mathematical career of course.

Thiebout, Tibi, Tiblidoo en Tiblidaddle, mijn favoriete Wevelgemenaar :) Bedankt voor onze vele interessante politieke discussies en diametraal verschillende opinies. Ik zeg je geen vaarwel maar tot ziens, aangezien je in mijn zog terugkeert naar ons gemeenschappelijk vaderland.

¹Having said this, the author wishes to express his revulsion for the politics of the FN.

If some years ago someone would have told me that I would end up having an utmost respect for some Dutch guy, I would have laughed and judged them mentally impaired. However David, somehow it just happened, David. As Ana (rightly so) likes to say, one of your prime qualities is honesty. So, let me return the favour: when I first met you, I didn't like you, having decided that you were arrogant. I could not have been more wrong. You turned out to have many qualities in fact: charming, socially intelligent, funny, compassionate etc. In short, when faced with a problem I ended up asking the question "what would David do"? Anyway, thank you for your funny stories, interesting opinions and honesty. Let's hope that next year we can discuss Dumoulin's Tour de France win.

In the end it seems as though many people are vying for the same title of kindest person alive, as are you Katie. The fact that I actually awarded it multiple times already should not be seen as a way of reducing the value of the prize, since it is in a sense the kindness equivalent of the Fields Medal, which I hope you will also win together. Katie, it was a pleasure to be your slave for "Analyse Vectorielle", and I believe I learned a lot from working with you. Good luck in Durham! May the penguins be with you.

Dear H  l  ne, I tested your patience, but good things come to those who wait: Preheat the oven to 175 degrees. Cover your baking pan with aluminium foil. We start with the crust, mix together 125 grams of flour, 70 grams of brown sugar, and 110 grams of (softened) butter. Add the mixture to the baking pan and distribute it evenly. Bake for 15 minutes until the crust is set. Then for the filling. Mix together one large egg, 70 grams of brown sugar, 100 grams of maple syrup, 120 grams of chopped-up pecans, and a pinch of salt. Add this to the baking pan and pop into the oven for another 17 minutes or until the mixture is just set. Congratulations, you made a pecan-maple syrup pie!

This guy from Bienne arrived on a cold winter day in January. He had grace, certainly. A vague scent of misunderstood genius, clearly. A love for books and categories, surprisingly so. This mysterious man captivated my mind for some time and I guess he still does. R  mi, thank you for always finding the perfect moment to suggest a coffee break, being exactly the time I just found the moment to get some work done. We ended up collaborating quite a lot and it has mostly been an excellent experience. Mostly wonderful in such a way that I will not mention the times that I wanted to strangle you. We tend to admire other people for the qualities we lack and this is certainly the case with you. You're an excellent mathematician with a vast general knowledge and fine-tuned writing skills (independent of what autistic Danny might say). Thank you for everything, I really hope you keep that promise of visiting me.

Then, of course, R  mi's office mate, when she is present, Laura. Thank you for the party at your place and the French touch (excuse me, Genevois) that you brought to the maths institute. Good luck in life, I'm sure you have plenty of qualities to get you where you want to end up.

ZumZum19, the little brother with the big mouth and the kind heart. Thank you for representing the math institute with me during the confinement and for the many smoke-dates we had during that time. Thanks for helping me move (twice). Next time I'm back here, I'll finally come watch one of your rugby matches to see you "d  foncer la gueule de ces fils de pute de NSRC".

Then Laurent, the other little brother, an impressive man. My good Lord, quel homme! You're a difficult person to describe, so multi-faceted. Thank you for your hiking suggestions and your climbing invitations, which I didn't take you up on, because –no other way to put it– I'm stupid. Thank you for your great talks, midnight petanque and skinny-dipping in the Mediterranean. It's an honour to have collaborated on an article with you. Bonne chance avec la suite. I'll be coming to you for cycling tips, be prepared!

We come to the most wonderful intern the country of Switzerland has ever known. Zoé, in retrospect I noticed that you really did not stay that long, but you made a great impression on me while you were here. Joyful, charming and always motivated to spend an evening at the lake. You had a great influence on me, given that you turned my disgust of swimming into what one could only describe as an "obsessive need to bathe in every single lake in Switzerland". I hope you wait a bit longer before going back to Nouméa, since up to now there is not yet a high-speed train that goes there. Best of luck to you and David.

Aisha, you have suffered enough under the tyranny of the malevolent Besançonian, as such you have deserved the right to your own paragraph, which will also be positioned in front of the one of he, who must not be named. Thank you for your friendship and death threats. I really enjoyed our conversations.

Luc, you have suffered almost insurmountable fears for your life, by the hands of the lioness of Senegal. Weaker men would have perished, but you, jewel of France – and not just any France: La France d'Emmanuel Macron! – have survived. Impressive. Oh and of course the latex document that will make you immortal among Neuchâtel PhD students: thank you for the `modele_these_final_scrbook` !

Luxembourgian, sexy, and intelligent... Everybody knows who I'm talking about. Your talks almost convinced me that symplectic geometry is an interesting subject to dedicate your life to. You are unlucky since I am writing this moments after you walked away when we were playing the best game in the world. So I will take this moment to express my disappointment in your persona, but don't worry, this just means that you have the possibility to still improve as a human being. Hugs and kisses, Nestor... Hash-tag Heart Emoji.

Audrey, a Belgian arriving in Neuchâtel at the same time as me, of course I liked you from the start. But independent of this twist of fate, you're also a human being of superior class. Thank you for your friendship and your help at my thesis defence. I expect you to invite me to yours! ;)

Antoine, that metal guy every institute of mathematics craves. Thank you for that huge warming smile when I forced you to take a coffee break with me. Please keep the production of hydromel on point, best I've ever had. And I hope I can keep consuming it in the not so far future.

Difficult to categorise William's sense of humour, but it is exquisite. It was more than an honour to admire the way you solved the daily crossword puzzles. Best of luck in Nancy and with your very pleasant and esteemed spouse.

The “statistician” lost in the mathematics department. Well, lost, in the end it is the place Radu belongs. A man with great stories, a profound cultural knowledge and a love for beers one can only admire. And I am in awe when presented with this man, for sure. Thank you for your existence and kindness, Radu.

When we are lucky we can spot a wild Stefano roaming the halls. Full of Italian wisdom and with a formidable accent, he makes your day a bit better every moment you spend in his presence. I hope I can some day take you up on your invitation to come see one of the best cycling races in the world!

Léonard, thank you for taking care of plant. Remember that Valentin will haunt your dreams if anything happens to her. Thanks for your joyful presence at the department. It’s a shame you only started last month.

Mael, thank you for allowing me to hold my defence apéro at the scouts buildings.

Thanks to Lucas, Cristina, Elisa, Matt, Emma, Antoine, Jérémy, Bruno, Alessio, Dimitri, Moe, Marine, Stella, Farshid, Cengiz, Joontae and Christine.

Finally, all the people who were in my tutorials. Ignoring sporadic exceptions, you were all amazing. Thank you for your motivation and pleasant personalities. Good luck in whatever you’ll end up doing.

And then as a post scriptum: The corona virus and the Swiss government for the confinement, which allowed me to write this thesis “quickly” and “efficiently”.

De-acknowledgements

A huge middle finger to the corona-virus and the Swiss government for banning Belgian visitors to Switzerland for two weeks, perfectly centred around the moment of my public defence.

Abstract

The objective of this work is to study large scale properties of graphs, graph sequences and groups. Firstly we consider the combinatorial cost and prove that having cost equal to 1 and hyperfiniteness are coarse invariants of graph sequences. We show that cost is multiplicative with respect to taking finite index subgroups. For an amenable group we investigate the properties of their Farber sequences, sofic approximations and vice versa.

Secondly we consider the first uniformly finite homology group of graphs with coefficients in $\mathbb{Z}$. We show that its non-vanishing depends on the ends, large circuits and (higher-dimensional) non-expansion of the graph. When the graph is transitive this is a full description.

Finally we take a look at the Baumslag-Solitar group $BS(2, 3)$. This group is non-Hopfian, meaning it has a quotient isomorphic to itself. We give a visual interpretation of this on the Cayley graph level.

Résumé

Comme objectif de cette thèse, on étudie les propriétés grossières des graphes, des suites de graphes et des groupes. Dans un premier temps on s'intéresse au coût combinatoire. On démontre qu'un coût égal à 1, ou encore l'hyperfinitude, sont des invariants grossiers pour les suites de graphes. De plus, le coût jouit d'une sorte de propriété de multiplicativité lorsqu'on considère des sous-groupes d'indice fini. Quand le groupe est moyennable, on traite les propriétés de leurs suites de Farber et de leurs approximations sofiqes.

Dans un deuxième temps, on analyse le premier groupe d'homologie uniformément fini avec coefficients entiers pour les graphes. On montre que si le graphe a plusieurs bouts, grands circuits, et pas d'expansion (de dimension supérieure), alors le groupe ne se restreint pas à zéro. De plus, pour un graphe transitif, cette description est exhaustive.

Finalement on considère le groupe de Baumslag-Solitar $BS(2, 3)$. Ce groupe n'est pas Hopfien, i.e. il est isomorphe à un quotient propre. On donne une interprétation visuelle de ce fait en étudiant son graphe de Cayley.

Contents

Introduction	1
1 Preliminaries	3
1.1 Graphs	3
1.2 Groups, Schreier and Cayley graphs	4
1.3 Large scale geometry	6
1.4 Large scale properties of groups, graphs and graph sequences	7
1.4.1 Interplay between groups and their box spaces	12
1.4.2 Ends of graphs	14
1.5 Benjamini-Schramm convergence and sofic groups	15
1.5.1 Interplay between groups and their sofic approximations	19
1.6 Cost and combinatorial cost	20
1.6.1 Ultrafilters and ultraproducts	20
1.6.2 Cost	20
1.6.3 Hyperfinite equivalence relations and treeable graphings	23
1.6.4 Combinatorial cost and hyperfiniteness	23
1.6.5 Limit graphings of graph sequences	24
1.7 Uniformly finite homology	25
1.7.1 Uniformly finite homology for metric spaces	25
1.7.2 Simplicial Uniformly Finite Homology H^{suf}	26
1.7.3 Uniformly finite homology and expansion	27
1.7.4 A theorem and a question	30
2 Overview	31
2.1 On combinatorial cost: a coarse setting	31
2.2 On the first uniformly finite homology group for graphs	32
2.3 On $BS(2, 3)$	33
3 On combinatorial cost: a coarse setting	35
3.1 Cost 1 is a coarse invariant	35
3.2 Hyperfiniteness	39
3.3 Cost-1 behaves multiplicatively for subgroups	40
3.4 Property A implies hyperfiniteness	43
3.5 Elek's theorem for Farber sequences	46
4 On the first uniformly finite homology group for graphs	51
4.1 "Hands-on" definition of H_1^{uf}	51
4.2 Trees and End-Defining Trees	56
4.2.1 Homology of Trees	56
4.2.2 Construction of a set of bi-infinite paths	57
4.2.3 Ends	61

4.3	Coefficients in $\mathbb{Z}_2$	68
4.3.1	$H_1^{\text{uf}}(X, \mathbb{Z}) \twoheadrightarrow H_1^{\text{uf}}(X, \mathbb{Z}_2)$	68
4.3.2	Decomposition of $H_1^{\text{uf}}(X, \mathbb{Z}_2)$	68
4.3.3	Large circuits and dimensions	70
4.4	$\mathbb{Z}$ -large circuits	79
4.5	Expansion	80
4.5.1	Basics	81
4.5.2	Dimension 1	82
5	On $BS(2, 3)$	85
5.1	The quotient graph	85
5.2	A free group of infinite rank	90
5.3	Another interpretation of non-Hopfianness and a continuous map	94
	List of Figures	97
	Bibliography	99

Introduction

Groups are algebraic! “But are they really?” Says the voice in your head. Of course, groups are abstract algebraic objects. Certainly to the mathematicians who have had the privilege to enjoy the beauty of Galois theory. But aren’t some of the first groups we encounter as baby-mathematicians sets of symmetries of a given object? Is it not true that, more often than not, groups appear naturally in a context where they act on some nice geometric space? And is it that wrong to claim that in these particular instances, the considered group reveals something of itself? No of course not. So it is only a little leap of faith to go a tiny step further and notice that a group itself is a geometric object. Having chosen a generating set S , the curious mathematician is able to define the norm of a group element as the length of the shortest representing word on the alphabet $S \cup S^{-1}$. This turns the group into a metric space, more precisely into a graph: the Cayley graph. This object depends on the choice of S , but will turn out to be unique up to quasi-isometry, which is an equivalence relation that allows an affine distortion of the metric. From this viewpoint, a coarse environment is the natural setting to talk about geometric group theory. On this playground there are interesting invariants to be discovered, such as ends of a group, amenability and growth. We will try to give the reader a taste of this world by describing many of these concepts. Whenever a group wanders into the neighbouring world of graphs it always brings along a court of explicit examples which makes a graph theorist shiver with anticipation.

We have set our goal to give an ‘extensive’ overview of the concepts we work with and the general setting from where they originate. The choice what to present, is always arbitrary. Certainly, there are many results we do not even mention that deserve to be talked about. Nevertheless we hope to be able to spark the reader’s interest in certain topics. We start with an introduction to large scale geometry and its invariants for groups and graphs. After this we move to graph convergence, which gives the subject a very light probabilistic flavour. Subsequently we introduce cost theory. Although very interesting, for us this is merely a means to an end. Ultimately we want to present some results in Chapter 3 around combinatorial cost (and its coarse properties). This is an invariant of graph sequences and a direct analogue of cost for probability measure preserving actions. In order to justify its definition and claim its naturality we give an unnecessarily extensive overview of cost theory. Well, unnecessary... only to someone who does not appreciate beauty.

Finally we review some homology theory on graphs. Uniformly finite homology is a coarse homology theory of which not so much is known. We introduce it and explain how its zeroth homology group can be used to characterise expansion. In Chapter 4 we then present some results that characterise the first uniformly finite homology group, which is joint work with Bottinelli.

As a dessert, in Chapter 5, we give a geometric interpretation of the non-Hopfianess of the Baumslag-Solitar group $BS(2, 3)$ and show that a certain subgroup is free. Hence in short, this thesis is organised as follows.

Chapter 1 recalls large scale properties of groups and graphs, introduces graph convergence, cost and associated concepts.

Chapter 2 provides a short summary of the results in this thesis.

Chapter 3 presents the results related to combinatorial cost. One also finds a description of sofic approximations of amenable groups.

Chapter 4 contains a description of the first uniformly finite homology group for graphs.

Chapter 5 presents some results associated to the Baumslag-Solitar group $BS(2, 3)$.

1 Preliminaries

1.1 Graphs

A comprehensive reference is [Bol98]. Graphs (sometimes called networks in more applied domains) are elementary geometric structures consisting of points (vertices) and connections between these points known as edges. In what follows we will only be interested in simple graphs, which do not contain multiple edges (more than one edge between two given vertices), nor loops (edges that start and end in the same vertex). Since we will only be working with simple graphs, we directly exclude non-simple graphs in the definition.

Definition 1.1.1. A graph G is a tuple (V, E) where

- V is the set of vertices, and;
- E is the set of edges, a subset of the set of unordered tuples $V^{[2]} = \{\{x, y\} \mid x, y \in V \text{ and } x \neq y\}$.

A labelled rooted graph is a graph (G, o, ϕ) with a distinguished vertex o and a map $\phi : E \rightarrow \mathbb{Z}$ of labels. Two labelled graphs (G, o, ϕ) and $(\tilde{G}, \tilde{o}, \tilde{\phi})$ are isomorphic whenever there is a graph isomorphism $f : G \rightarrow \tilde{G}$ such that $f(o) = \tilde{o}$ and $\tilde{\phi} \circ f = \phi$.

For a given graph G , we set $V(G)$ its set of vertices and $E(G)$ its set of edges. One also defines the degree of a vertex to be the number of adjacent edges. I.e. $d(x) = |\{e \in E(G) \mid x \in e\}|$. A graph has a natural metric induced by the shortest path distance. One says that a metric space X has bounded geometry whenever the quantity $|B_r(X, x)|$ is bounded by a constant, which only depends on r . It is easy to see that a graph has bounded geometry whenever the degrees $d(x)$ are bounded.

One can also introduce the notion of a directed graph.

Definition 1.1.2. A directed graph G is a tuple (V, E) where

- V is the set of vertices, and;
- $\vec{E}$ is the set of edges, a subset of $(V \times V) \setminus \Delta$, where $\Delta = \{(x, x) \mid x \in V\}$.

Notice that if for constructions the direction of edges in an undirected graph G is important, we can see G as a directed graph with symmetric edge set, i.e. $(x, y) \in \vec{E}(G)$ if and only if $(y, x) \in \vec{E}(G)$.

A useful concept later on will be the boundary of a set in G .

Definition 1.1.3. Given a graph G and a finite subset $A \subset G$, the r -boundary of A is given by

$$\partial_r(A) = \{x \in G \mid x \notin A \text{ and } d(x, A) \leq r\}.$$

For the 1-boundary of a set, one drops the index.

We note that we will mostly be interested in the coarse structure of a graph. This will be made precise later in Section 1.3, but for now we invite the reader to accept that this means that ‘local structure can be ignored’. The reader could as an intuitive example take a look at the two graphs in Figure 1.2: they are locally very different, but from ‘far away’ they just look like a line. So from a coarse frame of reference these two graphs, or even $\mathbb{R}$ are all the same object. Specifically this means that from this viewpoint ‘finite graphs are points’. Hence, either one restricts to infinite graphs, or one considers infinitely many finite graphs that grow in size at the same time. This motivates the introduction of graph sequences.

Definition 1.1.4. A graph sequence $\mathcal{G} = \{G_n\}_{n \in \mathbb{N}}$ is a sequence of finite connected graphs G_n of uniformly bounded degree d , such that $\lim_{n \rightarrow \infty} |V(G_n)| = \infty$.

The principal idea of graph sequences is to be able to compare the structure of all graphs at the same time, more specifically as metric spaces. We define coarse disjoint unions in order to obtain one metric space instead of infinitely many individual ones. This is very useful since this allows us to define maps for all graphs at the same time instead of individually.

Definition 1.1.5. Let $(X_n)_n$ be a sequence of bounded metric spaces. The coarse disjoint union is then defined as $\sqcup X_n$ where distances between points of X_n are given by the distance of X_n and the distance between X_n and X_m is greater or equal than $\max(\text{diam}(X_n), \text{diam}(X_m))$.

1.2 Groups, Schreier and Cayley graphs


Groups, by definition, are sets that have a binary operation associated to them. All our groups are discrete. We will consider them from a geometric viewpoint. In order to do this we introduce the Cayley graph associated to a presentation of the group.

Definition 1.2.1. Let Γ be a finitely generated group with finite generating set S , such that $e_\Gamma \notin S$, then the Cayley graph $\text{Cay}(\Gamma, S)$ is a graph such that


- its vertices are the group elements of Γ , and
- its edges are given by the set $\{\{x, xs\} \mid x \in \Gamma, s \in S\}$.

We notice that by applying s to the element xs^{-1} , we see that $\{x, xs^{-1}\}$ is an edge of $\text{Cay}(\Gamma, S)$, so S can be chosen symmetric. One could consider Cayley graphs as labelled graphs, i.e. there is a colouring of the edges by generators. For this it is important to choose S to be symmetric. The colouring $c : \vec{E} \rightarrow S$ is given by mapping the directed edge (x, xs) to s . In this case (xs, x) is mapped to s^{-1} . The advantage of this definition is that the group of graph automorphisms $\text{Aut}(\text{Cay}(\Gamma, S))$ that preserve the colouring is isomorphic to Γ . We will only view the Cayley graph as a coloured graph when important for our constructions. It is clear that this construction depends on the choice of generating set. We will see that such graphs are quasi-isometric. If we choose to work in a coarse setting, this choice is unimportant. This will soon be made precise in Section 1.3.

Consider a finitely generated group Γ with generating set S and a subgroup $\Lambda \leq \Gamma$. One can look at the set of right cosets $\Lambda \backslash \Gamma = \{\Lambda g \mid g \in \Gamma\}$. The Schreier graph $\text{Sch}(\Gamma, \Lambda, S)$ on these cosets is defined as follows. The vertices are exactly the right cosets and there is an edge from Λg_1 to Λg_2 if there is a generator $s \in S$ such that $\Lambda g_1 s = \Lambda g_2$. One can label the edges by

Figure 1.1: The Cayley graph of $\mathbb{F}_2$ with standard generators.

their defining generator in S . Note that this graph is connected. Also note that $\text{Sch}(\Gamma, \{e\}, S)$ is just the Cayley graph of Γ with respect to S , written as $\text{Cay}(\Gamma, S)$. Cayley graphs are transitive, but Schreier graphs are not in general. However there is a natural action by Γ on the cosets, which turns them into interesting objects from the group point of view. Moreover one sees that $\text{Sch}(\Gamma, \Lambda, S) = \text{Cay}(\Lambda \backslash \Gamma, \bar{S})$ if Λ is normal in Γ . Using this one can construct interesting graph sequences originating from Γ .

Figure 1.2: Two Cayley graphs of $\mathbb{Z}$ (labels omitted).

Residually finite groups are of particular interest to us.

Definition 1.2.2. Let Γ be a finitely generated group with generating set S . A descending sequence $N_1 \geq N_2 \geq \dots$ of finite index normal subgroups, where $\bigcap_{k>0} N_k = \{e\}$, is called a filtration and Γ is called residually finite if such a filtration exists.

The corresponding box space is then defined as the graph sequence

$$\square_{N_k} \Gamma = \{\text{Cay}(N_k \backslash \Gamma, \bar{S})\}_{k \in \mathbb{N}}.$$

The world of residually finite groups is sufficiently rich and diverse to justify our interest in them and their box spaces. Some examples are finitely generated nilpotent groups, free groups, finitely generated linear groups and mapping class groups. Also note that box spaces are locally in-differentiable from the actual group, in the following sense. Since $\bigcap_{k>0} N_k = \{e\}$, for all $r > 0$

there is some $K > 0$ such that for $k > K$ all $B_r(\text{Cay}(\Gamma, S), e) \simeq B_r(\text{Cay}(\Gamma/N_k, \bar{S}), \bar{e})$.

1.3 Large scale geometry

In this section we introduce large scale geometry. The idea is to investigate the coarse structure of metric spaces and construct coarse invariants, i.e. properties which are preserved under taking coarse maps. In certain settings the ‘natural’ equivalence between metric spaces is not that of isometries. In the setting of Cayley graphs for example. Different presentations of the same group can give rise to different Cayley graphs and thus different metrics on the group. However we cannot ever hope to obtain geometric invariants of our group, if we do not relax the notion of what it means to be equivalent as metric spaces. A first relaxation is that of a bilipschitz map, allowing for stretching.

Definition 1.3.1. Let (X, d_X) and (Y, d_Y) be two metric spaces, then $f : X \rightarrow Y$ is a bilipschitz if $\exists L \in \mathbb{R}^+$ such that $\forall x, x' \in X$

$$\frac{1}{L}d_X(x, x') \leq d_Y(f(x), f(x')) \leq Ld_X(x, x').$$

Whenever f is also bijective it is said to be a bilipschitz equivalence.

Consider two distances d_S and $d_{S'}$ on the group Γ induced by two different generating sets S and S' . Then it is easy to show that the ‘identity map’ $\text{Id} : (\Gamma, d_S) \rightarrow (\Gamma, d_{S'})$ is a bilipschitz equivalence. Consider also Figure 1.2 for two bilipschitz equivalent Cayley graphs of $\mathbb{Z}$. Note however that any bilipschitz map is still necessarily an injection. There is no reason we should not relax further. For sure our intuition tells us that from a large scale point of view, finite metric spaces must be equivalent to points. Our philosophy being that we wish to ignore local restrictions, we introduce a constant term to resolve the issue. This gives the large scale variant of bilipschitz maps: quasi-isometries.

Definition 1.3.2. Let (X, d_X) and (Y, d_Y) be two metric spaces, then $f : X \rightarrow Y$ is a quasi-isometric embedding if $\exists A \in \mathbb{R}^+$ such that $\forall x, x' \in X$

$$\frac{1}{A}d_X(x, x') - A \leq d_Y(f(x), f(x')) \leq Ad_X(x, x') + A.$$

If moreover $\text{Im}(f)$ is a net in Y , i.e. $\exists R \in \mathbb{R}^+$ such that $\forall y \in Y, \exists x \in X$ such that $d_Y(f(x), y) < R$, then f is a quasi-isometry.

One sees that quasi-isometries allow for an affine transformation of the metric. The reader may wonder why we stop here. Indeed there is no reason to, one can from the same viewpoint introduce coarse equivalences.

Definition 1.3.3. Let (X, d_X) and (Y, d_Y) be two metric spaces, then $f : X \rightarrow Y$ is a coarse embedding if $\exists \rho^+, \rho^- : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ increasing monotone functions with $\lim_{t \rightarrow \infty} \rho^-(t) = \lim_{t \rightarrow \infty} \rho^+(t) = +\infty$ such that $\forall x, x' \in X$

$$\rho^-(d_X(x, x')) \leq d_Y(f(x), f(x')) \leq \rho^+(d_X(x, x')).$$

If moreover $\text{Im}(f)$ is a net in Y , i.e. $\exists R \in \mathbb{R}^+$ such that $\forall y \in Y, \exists x \in X$ such that $d_Y(f(x), y) < R$, then f is a coarse equivalence.

Suppose now we have a coarse equivalence between two graphs $f : G \rightarrow H$. Consider two vertices $x, y \in G$, then there exists a path of vertices $(x = x_0, x_1, \dots, x_{k-1}, x_k = y)$ in G such that the $(x_i, x_{i+1}) \in E(G)$ for $0 \leq i \leq k-1$. Clearly $d(x, y) = d(x, x_1) + \sum_{i=1}^{k-1} d(x_i, x_{i+1}) + d(x_k, y)$. However applying f we see that this implies that

$$d(f(x), f(y)) \leq d(f(x), f(x_1)) + \sum_{i=1}^{k-1} d(f(x_i), f(x_{i+1})) + d(f(x_k), f(y)),$$

and thus, since f is a coarse equivalence, $\rho^+(d(x, y)) \leq k \cdot \rho^+(1)$, making ρ^+ affine. A similar argument holds for ρ^- . We see thus that in the setting of graphs, the notion of coarse equivalence reduces to the one of quasi-isometry. A similar argument holds for coarse disjoint unions by the following lemma of Khukhro and Valette.

Lemma 1.3.4 (Lemma 1, [KV17]). *Let $X = \bigsqcup_{n>0}^{+\infty} X_n$ and $Y = \bigsqcup_{n>0}^{+\infty} Y_n$ be coarse disjoint unions of graphs such that the diameter tends to infinity as n tends to infinity and let $\Phi : X \rightarrow Y$ be a coarse equivalence between these metric spaces. Then there exists a constant A and an almost permutation ϕ between the components of X and the components of Y such that $\Phi|_{X_i}$ is an (A, A) -quasi-isometry between X_i and $\phi(X_i)$.*

An almost permutation between sets B and B' is a bijection between a co-finite subset of B and a co-finite subset of B' . We conclude that the notion of quasi-isometry is the ‘right’ notion to talk about large scale geometry of Cayley graphs and coarse disjoint unions. This allows us to treat coarse disjoint unions and graph sequences as if they were the same object. An interesting remark to make here is that box spaces of the same group are not necessarily quasi-isometric. Of course it is easy to construct such examples, starting with an arbitrary box space, throwing out every second Cayley graph will do the job. A much less (and maybe surprising) example is given in a paper of Delabie and Khukhro. We cite the theorem here, for a specific example we refer the reader to [DK, p. 10-11].

Theorem 1.3.5. *There exist two box spaces $\square_{N_k} \Gamma \not\approx_{CE} \square_{M_k} \Gamma$ of the same group Γ , such that $\Gamma/N_k \twoheadrightarrow \Gamma/M_k$ with $[M_k : N_k]$ bounded.*

1.4 Large scale properties of groups, graphs and graph sequences

A first large scale invariant that will be of interest is amenability. Picking a particular moment when concepts were first studied, is always somewhat arbitrary. Our pseudo-randomly chosen starting point is the 1924 Banach-Tarski paradox [BT24], which is a problem about the existence of paradoxical sets. Consider two bounded sets A, B in the Euclidean space $\mathbb{E}^n$, with non-empty interior. When can A and B be decomposed into $\sqcup_{i=1}^k A_i$ and $\sqcup_{i=1}^k B_i$, such that A_i is isometric to B_i ? The answer given by Banach and Tarski: always, whenever $n \geq 3$. Von Neumann in 1929 interested by paradoxicality, introduced the class of amenable groups to explain this difference in behaviour (i.e. $n \leq 2$ and $n \geq 3$). He identified the existence left-invariant measures on the acting group (the isometries) as highly relevant.

Definition 1.4.1. A discrete group Γ is amenable when there exists a finitely additive left-invariant probability measure $\mu : \mathcal{P}(G) \rightarrow [0, 1]$. I.e.

- $\forall g \in \Gamma, A \in \Gamma, \mu(g \cdot A) = \mu(A),$
- $\forall \sqcup_{i=1}^k A_i \subset \Gamma, \sum_{i=1}^k \mu(A_i) = \mu(\sqcup_{i=1}^k A_i).$

The finite additivity constraint gives a hint of why this is very closely linked to the non-existence of paradoxical sets. Von Neumann called these groups ‘messbar’ (measurable). These groups are nowadays known as amenable groups (a-mean-able) since there is a more modern translation of amenability which uses means instead of measures. The English name is due to Day.

In the remainder of this text we will not be using this definition, but a more combinatorial one.

Definition 1.4.2. A group (Γ, S) is amenable when for every $\epsilon > 0$, for every $r > 0$, there is a finite subset $F_{r,\epsilon} \subset \Gamma$, such that $|\partial_r F_{r,\epsilon}| < \epsilon |F_{r,\epsilon}|$. Such sets are called Følner sets.

Some remarks are to be made. The most important one is that this definition seems to depend on the chosen generating set while von Neumann’s definition did not. We will soon see that amenability is a quasi-isometric invariant, hence the choice of S is indeed unimportant. Note that it is possible to fix $r \geq 1$ and obtain an equivalent definition. Allowing r to take any value in $\mathbb{N}$, gives us more liberty in proofs, see for example Proposition 1.4.3.

As a last remark one notices that the statement of Definition 1.4.2 does not require us to consider groups, one can state that a graph is amenable when it contains Følner sets. We now show that amenability of graphs (and thus groups) is quasi-isometry invariant.

Proposition 1.4.3. *Amenability is a quasi-isometric invariant.*

Proof. Consider two graphs G and H , and an A -quasi-isometry $f : G \rightarrow H$ between them. We show that G is amenable whenever H is. Consider Følner sets $F_{r,\epsilon}$ in H . Take r at least $2A$. Note that when $d(x, f^{-1}(F)) = 1$, then $d(f(x), F) \leq 2A$. Since f is a quasi-isometry, there exists a constant C such for all $y \in H$, $|f^{-1}(y)| \leq C$. We conclude that

$$\begin{aligned} |\partial(f^{-1}(F))| &\leq |f^{-1}(\partial_{2A}(F))| \\ &\leq C |\partial_{2A}(F)| \\ &\leq C \epsilon |F| \\ &\leq C^2 \epsilon |f^{-1}(F)|. \end{aligned}$$

□

It is time for some examples.

Example 1.4.4. • Finite graphs are amenable.

- $\mathbb{Z}^n$ is amenable by taking $F_{1/k} = \{-k, \dots, k\}^n$.
- $\mathbb{F}_2$ is not amenable because of Definition 1.4.1. Consider the sets $A(a)$, $A(a^{-1})$, $A(b)$, $A(b^{-1})$ and $\{e\}$ in Figure 1.3. Now $\mathbb{F}_2 = \{e\} \sqcup A(a) \sqcup A(a^{-1}) \sqcup A(b) \sqcup A(b^{-1})$, but also $\mathbb{F}_2 = a^{-1}A(a) \sqcup A(a^{-1}) = b^{-1}A(b) \sqcup A(b^{-1})$. These are paradoxical sets that force any left-invariant finitely additive measure on $\mathbb{F}_2$ to be zero, which is of course a contradiction.

Moreover we have the following stability properties.




Figure 1.3: The domains $A(a)$, $A(a^{-1})$, $A(b)$, $A(b^{-1})$ and $\{e\}$ in red, green, yellow, blue and brown respectively.

Theorem 1.4.5. *Amenability of discrete groups is preserved under*

- *taking subgroups;*
- *quotients;*
- *extensions;*
- *upward directed unions;*
- *direct limits.*

Proof. See [Pat88, p.14]. □

Note that all these properties are only true in the situation of groups. When discussing amenable graphs, one can no longer hope that any subgraph is still amenable. An easy example can be found by taking a rooted binary tree and attaching an infinite line graph to one of its points as in Figure 1.4. The reason for this is that Cayley graphs are transitive. Hence if we can find Følner sets around one point, we can find them around any point. So we would like to generalise this property to non-transitive spaces in a ‘more intelligent’ way. Yu’s property A is such a popular weakening for discrete metric spaces (thus including both graphs, as box spaces!).

Definition 1.4.6. Let (X, d) be a discrete metric space. Then X has property A if for all $r > 0$, $\epsilon > 0$, there exists $S > 0$ and a function $\xi : X \rightarrow \ell^1(X)$, denoted $x \mapsto \xi_x$, such that

1. $\|\xi_x\|_1 = 1$ for all $x \in X$;
2. ξ_x is supported in the ball of radius S around x for all $x \in X$;
3. $\|\xi_x - \xi_y\|_1 < \epsilon$ for all $x, y \in X$, with $d(x, y) < r$.

The function ξ is said to have variation at most ϵ and support S .




Figure 1.4: An infinite line graph added to a rooted binary tree.

Also property A has many equivalent definitions, we refer the reader to the overview paper by Willett for more information [Wil11]. Notice that, as one would like, for groups being amenable implies having property A. Whenever $F_{r,\epsilon}$ is a Følner set for Γ (say rooted at e for simplicity), clearly $\xi : \Gamma \rightarrow \ell^1(\Gamma)$, mapping g to the function

$$\xi_g : x \mapsto \begin{cases} 0 & \text{if } x \notin gF_{r,\epsilon} \\ \frac{1}{|F|} & \text{if } x \in gF_{r,\epsilon} \end{cases}$$

For a group it is ‘quite easy’ to have property A, as shows Theorem 1.4.8. We show that $\mathbb{F}_2$ has A.

Example 1.4.7. Any tree has property A. We show this for $\mathbb{F}_2$ using a ‘visual’ proof in Figure 1.5. The adjustment of the proof to the general case is immediate. Fix a oriented bi-infinite geodesic path in the standard Cayley graph of $\mathbb{F}_2$. For example the blue line in the figure. Fix $\epsilon, r > 0$, then take n such that $\frac{2r}{n} < \epsilon$. Then, for $g \in \mathbb{F}_2$, let P_g be the unique path starting at g , following the geodesic to the blue line, and then following the blue line to infinity. Let $P_{g,n}$ be the restriction of P_g to the first n points of the path. Define ξ_g to be $\frac{1}{n} \mathbb{1}_{P_{g,n}}$. Then clearly for x, y at a distance at most r the functions ξ_x and ξ_y are equal on at least $n - 2r$ points. Hence we see that

$$\|\xi_x - \xi_y\|_1 \leq \frac{2r}{n}.$$

This concludes the proof.

Also having property A is a coarse invariant. It turns out spaces with property A have many interesting properties.

Theorem 1.4.8. • *Having property A is coarse invariant.*

- *Subspaces of property A spaces have property A.*
- *The stability properties from Theorem 1.4.5 hold for groups with property A. Amalgamated products can be added to the list.*
- *A discrete metric space with property A coarsely embeds into a Hilbert space.*




Figure 1.5: A proof that $\mathbb{F}_2$ has property A. In red P_{b^2} and in green $P_{bab^{-1}a}$

Proof. See [Wil11]. □

Let us have another look at amenability. Consider a single graph G . If it is non-amenable we say it is an expander. In this case it has no Følner sets and it makes sense to quantify non-amenable in the following way.

Definition 1.4.9. The Cheeger constant of a graph G is given by

$$h(G) := \inf \left\{ \frac{|\partial A|}{|A|} \mid A \subset V(G), 0 < |A| < \infty \right\}.$$

One sees that G is an expander if and only if $h(G)$ is bounded away from zero. We also have the following interesting description. Note that when $e = (x, y)$ is a directed edge, then $\partial(x, y) = y - x$. The map ∂ can be extended linearly.

Proposition 1.4.10. A graph G is an expander if and only if the sum $\mathbb{1}_G := \sum_{x \in V(G)} x$ in $\ell^\infty(V(G), \mathbb{Z})$ can be written as the boundary of a sum of edges. I.e. there exists some $\sum_{e \in \vec{E}(G)} a_e \cdot e$ in $\ell^\infty(\vec{E}(G), \mathbb{Z})$, such that $\partial \left(\sum_{e \in \vec{E}(G)} a_e \cdot e \right) = \mathbb{1}_G$.

Proof. This description was obtained in the context of uniformly finite homology groups, which we will describe later. For a proof see Proposition 1.7.4. □

Note that the ‘direct application’ of Definition 1.4.9 to graph sequences is not a good one. Since we can obtain ‘Følner’ sets of arbitrary size and empty boundary. However this is not a satisfactory definition, since being an expander is very much a local property and we are ignoring the possible expansion that can happen inside the finite graphs, we offer the following adaptation of the Cheeger constant and apply it to graph sequences.

Definition 1.4.11. For a finite graph G , the Cheeger constant is defined as

$$h(G) := \min \left\{ \frac{|\partial A|}{|A|} \mid A \subset V(G), 0 < |A| \leq \frac{1}{2}|V(G)| \right\}.$$

A graph sequence $\mathcal{G} = \{G_n\}_{n \in \mathbb{N}}$ is an expander sequence if there is some $c > 0$ such that for every $n \in \mathbb{N}$ the Cheeger constant $h(G_n) > c$.

We note that this is an extension of non-amenability in the following sense.

Proposition 1.4.12. *An expander sequence does not embed into a Hilbert space. As a consequence, the properties of having property A and being an expander are mutually exclusive for graph sequence.*

We introduce a last coarse property of groups: property (T). The notion will not be relevant later on, but it is an important and widely studied property, so we do take the time to introduce it. Defined by Kazhdan in 1967 [Kaz67], we give its original definition. See also [BdlHV08] for a reference. A unitary representation of a group Γ is a group homomorphism π into the group of unitary operators $\mathcal{U}(\mathcal{H})$ on a complex Hilbert space $\mathcal{H}$. An example of such a unitary representation is the left-regular representation $\lambda : \Gamma \rightarrow \mathcal{U}(\ell^2(\Gamma))$, defined by the action $\lambda(g)(f)(x) = f(xg^{-1})$.

Definition 1.4.13. For a unitary representation $\pi : \Gamma \rightarrow \mathcal{U}(\mathcal{H})$ of Γ , finitely generated by S , we say that π has almost-invariant vectors if $\forall \epsilon > 0$, there is a vector $v \in \mathcal{H}$, such that $\|v - \pi(s)v\| < \epsilon\|v\|$ for all $s \in S$.

Definition 1.4.14. A group Γ has property (T) if every unitary representation with almost-invariant vectors, has a non-zero invariant vector.

Examples of such groups are given by special linear groups $SL_n(\mathbb{Z})$ where $n \geq 3$. These are also residually finite groups by taking modulo-reductions. In Section 1.4.1 we will see that this will give rise to explicit constructions of expander sequences. Hence it is to be expected that amenability and property (T) are incompatible.

Theorem 1.4.15. *Any discrete group Γ , being amenable and having property (T), is finite.*

Proof. It is easy to see that any finite group is amenable and has property (T). We suppose Γ is infinite and consider the left-regular unitary representation $\lambda : \Gamma \rightarrow \mathcal{U}(\ell^2(\Gamma))$. Consider a Følner sequence F_ϵ . Then the characteristic functions ξ_{F_ϵ} give a sequence of almost invariant vectors. However λ only has an invariant vector if and only if Γ is finite. $\square$

1.4.1 Interplay between groups and their box spaces

When we look at the box spaces of a given residually finite group, we have seen that these are locally indistinguishable. Hence we might hope that local properties are preserved (for appropriate translations of the considered properties). And indeed, this is the case! We base our proof on [Wil11, Lemma 2.5 and 3.2]

Theorem 1.4.16. *Let Γ be residually finite with filtration $(N_k)_k$, then Γ is amenable if and only if $\square_{N_k}\Gamma$ has property A.*

Proof. For the first implication we do the same as when we showed Cayley graphs of amenable groups have property A. Let $\mathcal{G}$ be a box space of Γ . Take a Følner set F around e for $\epsilon > 0$. Let r be such that $F \subset B_r(\Gamma, e)$. Then define the following function $\xi : \mathcal{G} \rightarrow \ell^1(\mathcal{G})$ by mapping $y \in N_k \backslash \Gamma$ to

$$\xi_y : x \mapsto \begin{cases} 0 & \text{if } x \notin yF \\ \frac{1}{|F|} & \text{if } x \in yF \end{cases},$$

whenever $B_r(\Gamma, e)$ is isomorphic to $B_r(N_k \backslash \Gamma, e)$. When this is not the case, map y to

$$\xi_y : x \mapsto \begin{cases} 0 & \text{if } x \notin \Gamma/N_k \\ \frac{1}{|N_k \backslash \Gamma|} & \text{if } x \in \Gamma/N_k \end{cases}.$$

For the reverse implication, fix $\epsilon > 0$. Since $\square_{N_k} \Gamma$ has property A, there is some $\xi : \square_{N_k} \Gamma \rightarrow \ell^1(\square_{N_k} \Gamma)$ with variation ϵ and support $r > 0$. Take $K > 0$ such that the projection $B_r(\Gamma, e) \twoheadrightarrow N_k \backslash \Gamma$ is injective. We write X for the Cayley graph of $N_k \backslash \Gamma$ and d for its degree. The following inequalities are easy to check:

$$\begin{aligned} \sum_{(x,y) \in E(X)} \|\xi_x - \xi_y\|_1 &\leq d\epsilon |X| \\ &\leq d\epsilon \sum_{x \in X} \sum_{z \in X} |\xi_x(z)| \end{aligned}$$

Developping the norm on the left hand side, we notice that there exists some $\tilde{z} \in X$ such that

$$\sum_{(x,y) \in E(X)} |\xi_x(\tilde{z}) - \xi_y(\tilde{z})| \leq d\epsilon \sum_{x \in X} |\xi_x(\tilde{z})|.$$

We define the function $\phi : X \rightarrow \mathbb{R} : x \mapsto \xi_x(\tilde{z})$. We may suppose ϕ to be normalised and positive. Then since the support of ϕ is finite, there are $a_i \in \mathbb{R}^+$ and non-empty $F_i \subset \text{supp}(\phi)$ such that $F_i \subset F_{i-1}$, $\sum_i a_i = 1$ and

$$\sum_i \frac{a_i}{|F_i|} \mathbb{1}_{F_i} = \phi.$$

Hence we have that

$$\begin{aligned} \sum_{(x,y) \in E(X)} |\phi(x) - \phi(y)| &= \sum_{(x,y) \in E(X)} \left| \sum_i \frac{a_i}{|F_i|} (\mathbb{1}_{F_i}(x) - \mathbb{1}_{F_i}(y)) \right| \\ &= \sum_i \frac{a_i}{|F_i|} \sum_{(x,y) \in E(X)} |\mathbb{1}_{F_i}(x) - \mathbb{1}_{F_i}(y)| \\ &= \sum_i \frac{a_i}{|F_i|} |\partial(F_i)|. \end{aligned}$$

Note that the first sum is $\leq d\epsilon$. Since all a_i sum to one, there is some F_i such that $\frac{|\partial(F_i)|}{|F_i|} \leq d\epsilon$. Moreover by choice of $N_k \backslash \Gamma$, this F_i lifts to Γ and forms a $d\epsilon$ -Følner set there. $\square$

The first big application of property (T) groups goes in the same direction. This is a result by Margulis [Mar73].

Theorem 1.4.17. *Given a property (T) group Γ , then any box space of Γ is an expander sequence.*

One might hope for an equivalence, alas this is not true, property (T) is much stronger than having expanding box spaces. Free groups for example have box spaces that are expanders. Although not all of them are. Yu and Willett introduced geometric property (T) to save this equivalence. We do not introduce the concept since this would lead us too far, but the reader is referred to [WY14] for more information.

Theorem 1.4.18. *Let Γ be residually finite with filtration $(N_k)_k$, then Γ has property (T) if and only if $\square_{N_k}\Gamma$ has geometric property (T).*

1.4.2 Ends of graphs

The number of ends of a graph is intuitively the number of ways one can approach infinity. Our definition is slightly different from the one usually found in the literature. An infinite geodesic ray in G is an isometric embedding $r : \mathbb{Z}^+ \rightarrow G$. We say that two rays r_1, r_2 are equivalent when for any compact $K \subset V(G)$, the infinite parts of $r_1|_{G \setminus K}$ and $r_2|_{G \setminus K}$ are in the same connected component of $G \setminus K$.

Definition 1.4.19. Let R be the set of infinite geodesic rays in G , then the ends of G are given by $R / \sim$.

We also give an alternative non-equivalent definition. The only difference being that the first definition can distinguish between a countable and uncountable number of ends. This definition is purely informative and will not be used later on.

Definition 1.4.20. The number of ends of a graph is given by

$$\text{Ends}(G) := \sup\{\|G \setminus K\| \mid K \text{ a compact subset of } V(G)\},$$

where $\|G \setminus K\|$ counts the number of infinite connected components of $G \setminus K$.

Example 1.4.21. • Finite graphs have zero ends. The implication goes both ways.

- $\mathbb{Z}^n$ has one end whenever $n \geq 2$. If $n = 1$ it has two.
- $\mathbb{F}_n$ has uncountably many ends, whenever $n \geq 2$.
- The comb C_n with n teeth in Figure 1.6 has n ends. A comb with infinitely many teeth, has countably many ends.


Due to the highly symmetric nature of Cayley graphs, the different number of ends possible is highly restricted. This result can be extended to vertex-transitive graphs.

Theorem 1.4.22 (Freudenthal-Hopf theorem). *Every finitely generated group has either 0, 1, 2 or an (uncountably) infinite amount of ends.*

Proof. See [Mei08, Theorem 11.27]. □

Most groups are one-ended, as shows the following theorem.

Theorem 1.4.23. 1. *A group has two ends if and only if it is virtually $\mathbb{Z}$.*

Figure 1.6: The comb C_{10} with 10 teeth.

2. Let G be a torsion-free group, it has infinitely many ends if and only if it is a free product of non-trivial torsion-free groups.

The second item is a consequence of Stallings's theorem, which is able to describe all infinite-ended groups. The reader is referred to [Sta68] for the statement and proof.

1.5 Benjamini-Schramm convergence and sofic groups

The introduction of Benjamini-Schramm convergence is technically unnecessary to introduce sofic groups. However, the personal preferences of the author make it impossible for him not to give an overview of the more general setting. In the following all graphs are labelled, isomorphisms respect the labels and the roots in case the graphs are labelled and/or rooted.

Note that given a residually finite group Γ , we have derived graph sequences from it (box spaces) that locally very much resemble the Cayley graph of Γ . Since we are working with a graph sequence of growing graphs that resemble in some sense more and more to the $\text{Cay}(\Gamma)$, one would be tempted to think of this as a limiting process. This can indeed be made precise. Consider the following first attempt. Fix an integer D . Let $\mathcal{RG}_D$ be the set of all isomorphism classes of coloured rooted graphs of degree D , with edges labelled by $\{1, \dots, D\}$. Endow $\mathcal{RG}_D$ with the metric $d_{loc}((G, o), (\tilde{G}, \tilde{o})) = 1/\sup\{r \mid B_r(G, o) \simeq B_r(\tilde{G}, \tilde{o})\}$. One can easily show we obtain a compact (hence complete) metric space. We say that a graph sequence of finite coloured rooted graphs $\mathcal{G} = \{G_n, o_n\}_n$, rootedly converges to a coloured rooted graph (G, o) if it converges with respect to d_{loc} . A more practical and equivalent description is the following.

Definition 1.5.1. The graph sequence $\mathcal{G} = \{G_n, o_n\}_n$, rootedly converges to a coloured rooted graph (G, o) whenever $\forall r$, the sequence $\{(B_r(G_n, o_n), o_n)\}_n$ is eventually constant and its limit is isomorphic to $(B_r(G, o), o)$.

Example 1.5.2. Consider the graph sequence $\{\text{Cay}(\mathbb{Z}_n), 0\}_n$ in Figure 1.7, where the Cayley graph is taken with respect to the generating set $\{-1, 1\}$. Whenever $r < \frac{n}{2}$, $B_r(\text{Cay}(\mathbb{Z}_n), 0)$ will look like a line of length $2r$ rooted in the middle. So the rooted limit is just $\text{Cay}(\mathbb{Z})$ with natural generators.

Example 1.5.3. Consider the graph sequence $\{P_n, o_n\}_n$, where P_n is a line graph of length n and o_n is the any one of the two end points. Then its limit is the infinite line graph with one leaf,




Figure 1.7: A sofic approximation of $\mathbb{Z}$, as defined in Example 1.5.2

rooted at the leaf. If o_n is chosen to be a point in the middle of P_n , then the rooted limit is the infinite line graph with two ends, rooted at some point.

Note that this convergence is unsatisfactory in multiple ways. Firstly, we are obliged to work with rooted graphs, which is a huge restriction, secondly it only cares about what happens around one specific point and completely ignores the rest of the graph. The problem with this is that the rest of the graph can have a vastly different properties than the area close to the root. This has as a consequence that we will lose continuity of certain properties when taking limits. Remember that our hope was to capture the local behaviour of our graph sequence in our limit.

We will now introduce so-called Benjamini-Schramm convergence (also local convergence or weak convergence). Some nice references are [ATV13] and [Sab]. Our exposition follows the first reference rather closely. See [BS01] for the original paper.

Note that one of the problems of rooted convergence was the obligation to choose a root, however we still want to work locally, so this is somewhat contradictory. The solution is to pick the root uniformly at random. Consider the set of isomorphism classes of finite coloured graphs of bounded degree D , denoted $\mathcal{FG}_D$, then it embeds into the set of Borel probability measures on $\mathcal{RG}_D$ in the following natural way

$$em : \mathcal{FG}_D \rightarrow \mathcal{P}(\mathcal{RG}_D) : G \mapsto \mu_G := \frac{1}{|V(G)|} \sum_{x \in V(G)} \delta_{(G,x)},$$

where $\delta_{(G,x)}$ is the Dirac measure on (G, x) . Now instead of working with graphs we are working with probability measures. Since $\mathcal{RG}_D$ is compact, $\mathcal{P}(\mathcal{RG}_D)$ endowed with the weak topology is too. For the reader's convenience we recall the definition of weak convergence of measures. The sequence μ_n converges to μ whenever

$$\lim_{n \rightarrow \infty} \int_{\mathcal{RG}_D} f d\mu_n = \int_{\mathcal{RG}_D} f d\mu, \quad (1.5.1)$$

and this for all continuous functions $f : \mathcal{RG}_D \rightarrow \mathbb{R}$. Hence a graph sequence $\{G_n\}_n$ Benjamini-Schramm converges whenever $\{\mu_{G_n}\}$ converges in the weak topology.

Now, this calls for a more intuitive definition. Define, for any finite rooted graph $\alpha \in \mathcal{FRG}_D$ and $r > 0$, the clopen set $U_{\alpha,r} := \{(G, o) \in \mathcal{RG}_D \mid (B_r(G, o), o) \simeq \alpha\}$. Note that all these sets together form a clopen basis $\mathcal{B}$ of $\mathcal{RG}_D$. Using a stepfunction-like construction on sets of $\mathcal{B}$, we can construct all continuous functions on $\mathcal{RG}_D$. Hence convergence for all f reduces to convergence of characteristic functions $\mathbb{1}_{U_{\alpha,r}}$. Let us introduce some notation. For any finite graph G and for all $r > 0$, we define the probability measure $p_r(G, \cdot) : \mathcal{RG}_D \rightarrow [0, 1]$ in the following way. For $\alpha \in \mathcal{RG}_D$, let $V'_\alpha(G) = \{x \in V(G) \mid (B_r(G, x), x) \simeq \alpha\}$ be the vertices of G such that the r -ball around such a vertex is isomorphic to α . Next we define $p_r(G, \alpha) = \frac{|V'_\alpha(G)|}{|V(G)|}$. So we look at the probability that a uniformly chosen vertex of G has its r -ball isomorphic to




Figure 1.8: As an unrooted tree, this is the basis object in constructing the Benjamini Schramm limit of $\{B_r(\mathbb{F}_2, e)\}_{r \in \mathbb{N}}$. Depicted is the rooted tree T_1 . (see Example 1.5.5)

α . Taking as our continuous f some characteristic function $\mathbb{1}_{U_{\alpha,r}}$, one sees that Equation 1.5.1 reduces to

$$\lim_{n \rightarrow \infty} p_r(G_n, \alpha) = \mu(U_{\alpha,r}).$$

We insist to the reader that the limit is not a graph, but a probability distribution. However in the case of (infinite) vertex transitive graphs, there is an obvious associated probability distribution, so the following definition is natural.

Definition 1.5.4. Let Γ be a group with generating set S . A labelled graph sequence $\mathcal{G} = \{G_n\}_{n \in \mathbb{N}}$ with labels in S converges in a Benjamini-Schramm sense to the Cayley graph Γ if for every $r \in \mathbb{N}$, the limit $\lim_{n \rightarrow \infty} p_r(G_n, B_r(\text{Cay}(\Gamma), e)) = 1$, where $B_r(\text{Cay}(\Gamma), e)$ is the r -ball in $\text{Cay}(\Gamma)$ around the identity element. So in particular this limit exists.

Example 1.5.5. We give an example of a graph sequence that does not converge to a graph, but a probability measure. In most cases taking balls of increasing radius in the Cayley graph of the considered group does not give a sofic approximation of this group. Such is the case for $\mathbb{F}_2$. We describe the limit as an illustration. Using the definition above the reader can retrace our steps and verify our calculations. Consider the tree in Figure 1.8. Let the bottom line of points be level 1, then those above level 2 etc. Define $a_j = \frac{2}{3^j}$ and let T_j be the tree in the figure, with as a root some point on level j , then the limit of the graph sequence $\{B_r(\mathbb{F}_2, e)\}_{r \in \mathbb{N}}$ is the sum $\sum_{j>0} a_j T_j$. One should think of the coefficients a_j as probabilities.

An immediate consequence of this definition is that box spaces converge in a Benjamini-Schramm sense to the Cayley graph of the defining group. However, inspired by this definition, we can do a bit better. Let us take look at an arbitrary sequence of finite index subgroups (not necessarily normal, not necessarily with trivial intersection, not necessarily a descending chain). This induces a sequence of Schreier graphs associated to the group Γ .

Definition 1.5.6. A sequence of subgroups of a group Γ is a Farber sequence if the associated graph sequence of Schreier graphs is Benjamini-Schramm convergent to the Cayley graph of Γ .

One notices that we can still go a step further and even forget about subgroups altogether.

Definition 1.5.7. A group Γ is sofic if there exists a labelled graph sequence that converges to the Cayley graph of Γ in a Benjamini-Schramm sense. The sequence is then called a sofic approximation.

There are some nice lecture notes by Juschenko [Jus]. For more about sofic groups and their generalisations/cousins, see [Pes08]. The class of sofic groups was first introduced by Gromov [Gro99], but named by Weiss [Wei00]. Sofic comes from the Hebrew word **יָפוּי**, which means finite. Originally Gromov called these groups initially subamenable. We define it here:

Definition 1.5.8. The Cayley graph $\text{Cay}(\Gamma, S)$ is initially subamenable if $\forall r, \epsilon > 0$ and every ball $B_r(\text{Cay}(\Gamma, S), e)$ there exists an edge-coloured finite graph G and a set $W \subseteq V(G)$ such that

- $\forall w \in W$, there is an isomorphism $B_r(\text{Cay}(\Gamma, S), e) \simeq B_r(G, w)$, and
- $|W| > (1 - \epsilon)|V(G)|$.

One easily sees that Gromov's definition coincides with Definition 1.5.7 in the restrictive setting of Benjamini-Schramm convergence of Definition 1.5.4. It is often said that sofic groups are defined to be generalisations of both amenable groups and residually finite groups. One could see Definition 1.5.8 as a generalisation of Følner sets. There is an alternative definition which generalises residually finite groups in the sense that they mimic their characteristic of projecting onto finite groups in an exhaustive manner. The Hamming distance between two permutations $\sigma, \tilde{\sigma} \in S_n$ is given by

$$d_{\text{Ham}}(\sigma, \tilde{\sigma}) = \frac{1}{n} |\{1 \leq i \leq n \mid \sigma(i) \neq \tilde{\sigma}(i)\}|,$$

then we can define the following.

Definition 1.5.9 ([ES06]). A discrete group Γ is sofic if for every finite set $F \subset \Gamma$ and every $\epsilon > 0$, there exists some $n \in \mathbb{N}$ and a map $\phi : F \rightarrow S_n$ such that

- for every $g \in F \setminus \{e\}$ the distance $d_{\text{Ham}}(\phi(g), \iota) > 1 - \epsilon$, where ι is the identity element of S_n ,
- for all $g, h \in F$, such that $gh \in F$, $d_{\text{Ham}}(\phi(gh), \phi(g) \circ \phi(h)) < \epsilon$, and
- whenever $e \in F$, $\phi(e) = \iota$.

In this case ϕ is called an (F, ϵ) -almost homomorphism.

This alternative definition of sofic groups is due to Elek and Szabó. Finally a word about why sofic groups are so popular, which is also a good excuse to cite some famous conjectures. They are (in a very informal way) the largest class of groups for which one can prove theorems, before proving them for all groups. Gromov for example introduced them to solve the Gottschalk surjunctivity conjecture in the most general way until now.

Conjecture 1.5.10 ([Got73]). *Let Γ be a countable group and X a finite set. Consider the compact metrisable space X^Γ of X -valued functions on Γ equipped with the product topology. Let $A : X^\Gamma \rightarrow X^\Gamma$ be a continuous map that commutes with the natural right Γ -action. Then if A is injective, it is surjective as well.*

Theorem 1.5.11 ([Gro99]). *The Gottschalk surjunctivity conjecture holds for sofic groups.*

Another example of a well-known conjecture solved in the class of sofic groups is Kaplanski's finiteness conjecture by Elek and Szabó [ES04].

Conjecture 1.5.12. *Let Γ be a discrete group. For every field K and $a, b \in K[\Gamma]$ the identity $ab = e$ implies $ba = e$.*

Furthermore it is still unknown whether non-sofic groups exist. A surprising fact that might push the balance towards an affirmative answer, is that it is still unknown whether degree two extensions of sofic groups are still sofic.

Question 1.5.13. Are all countable groups sofic?

1.5.1 Interplay between groups and their sofic approximations

There have been very nice and successful attempts to generalise, among others, Theorems 1.4.16, 1.4.17 and 1.4.18. The idea is the following. Sofic approximations of residually finite groups are locally indistinguishable from their box spaces. In a philosophical and unprecise way, one could say that apart from a measure zero subspace, it ‘is’ a box space. Because the similarity between the two is a measurable one, this asks for a measurable interpretation of the properties given to box spaces. A very nice example is due to Gabor Kun, generalising Margulis’ Theorem on expanders.

Theorem 1.5.14. *[Kun16] Let Γ be a countably infinite Property (T) group and $\{G_n\}_n$ a sofic approximation of Γ . Then there exists a $\gamma > 0$ and a sequence of finite d -regular graphs $\{G'_n\}_n$ such that*

- $V(G_n) = V(G'_n)$
- $\lim_{n \rightarrow \infty} \frac{|E(G_n) \Delta E(G'_n)|}{V(G_n)} = 0$
- *There exists $\gamma > 0$ such that for every n the graph G'_n is a vertex-disjoint union of d -regular graphs with Cheeger constant at least γ .*

The second point in the theorem allows one to ignore a part of each graph G_n in the sofic approximation (a part that becomes negligible in the limit). This is essential and without it one could never hope to prove a theorem of the form ‘property of the group’ $\implies$ ‘property of the sofic approximation’. Other results in the other direction have been obtained by Alekseev and Finn-Sell. Since they prove reversed implications they do not need a measurable property.

Theorem 1.5.15. *[AFS19] Let $\mathcal{G}$ be a sofic approximation of Γ , then*

1. *If $\mathcal{G}$ has property A, then Γ is amenable.*
2. *If $\mathcal{G}$ has boundary geometry property (T), then Γ has property (T).*

Boundary geometry property (T) strongly resembles geometric property (T). We do not introduce it since we did not introduce geometric property (T). The reader can find a definition in [AFS19]. For a reversal of item (1), we refer the reader to Theorem 3.5.5. We state the following natural question.

Question 1.5.16. Let Γ be a countably infinite Property (T) group and $\mathcal{G}$ a sofic approximation of Γ . Is $\mathcal{G}$ a vertex-disjoint union of graph sequences that have (boundary) geometric property (T) in the measurable sense of Kun in Theorem 1.5.14?

1.6 Cost and combinatorial cost

1.6.1 Ultrafilters and ultraproducts

Ultrafilters are not strictly necessary to talk about what follows but the author enjoys conversing in this language, so the reader will have no choice but to adapt. The goal is mostly to present a direct link between traditional cost theory and combinatorial cost later on. Ultrafilters can be introduced in a more general setting and are interesting in themselves, we will be very restrictive and only use them as a tool, but the interested reader may take a look at [Zel11].

Definition 1.6.1. An ultrafilter on $\mathbb{N}$ is map $\omega : \mathcal{P}(\mathbb{N}) \rightarrow \{0, 1\}$, such that

- $\omega(\mathbb{N}) = 1$;
- Whenever a subset $A \in \mathbb{N}$ can be written as a disjoint sum of subsets $\sqcup_{i=1}^n A_i$, then $\omega(A) = \sum_{i=1}^n \omega(A_i)$.

In the end this is just a finitely additive probability measure on $\mathcal{P}(\mathbb{N})$, confer Definition 1.4.1. The idea behind ultrafilters is that they constitute a choice of ‘large sets’. They will help us generalise limits and force sequences to converge whenever possible. For example in compact spaces ultralimits will always converge. We invite the reader to follow the dogma ‘ultrafilters choose convergent subsequences in a consistent way’.

Example 1.6.2. Take any $a \in \mathbb{N}$, the principal ultrafilter associated to a is given by

$$\omega_a : \mathcal{P}(\mathbb{N}) \rightarrow \{0, 1\} : A \mapsto \begin{cases} 0 & \text{when } a \notin A \\ 1 & \text{when } a \in A \end{cases}$$

We will say an ultrafilter is non-principal whenever it is not principal. In this case it is an extension of the co-finite filter which maps finite sets to 0 and co-finite ones to 1. One can prove that non-principal ultrafilters exist as a standard application of Zorn’s lemma. We now show how we can define ultralimits.

Definition 1.6.3. Let X be a topological space, and $f : \mathbb{N} \rightarrow X$ a sequence of points in X , then $x \in X$ is an ω -ultralimit of f when for every open U containing x , $\omega(f^{-1}(U)) = 1$. Notation $\lim_{\omega} f = x$.

Of course such a point is not guaranteed to exist or be unique unless X is sufficiently nice as shows the next lemma. We also note whenever we work with a principal ultrafilter one always has that $\lim_{\omega_a} f = f(a)$. For this reason we will work with ‘non-trivial’ non-principal ultrafilters. Unless when f is converging in the traditional sense, the ultralimit will depend on the chosen non-principal ultrafilter.

Lemma 1.6.4. *If X is Hausdorff and compact then any sequence has a unique ultralimit.*

Proof. See [Zel11, Lemma 2.14]. □

1.6.2 Cost

Although we will only be interested in combinatorial cost, we introduce cost as it was defined in its original setting to make the definition of its combinatorial cousin more natural. Originally cost of probability measure preserving equivalence relation was introduced by Levitt [Lev95],

but some nice results by his student Gaboriau in [Gab98] and [Gab00], popularised the concept. Since cost is an orbit equivalence invariant it is possible to use it to differentiate p.m.p. actions of groups. For example, Gaboriau showed that free p.m.p. actions of different rank free groups are never orbit equivalent, since they necessarily have different cost. This section is heavily based on Gaboriau's overview paper [Gab10]. The idea is to mimic the rank of a group, i.e. the cardinality of a minimal generating set, but in the setting of equivalence relations. Consider a standard Borel space (X, μ) with μ an atomless measure (i.e. $A \subset X$ is an atom whenever $\mu(A) > 0$ and $\forall B \subset A$ the implication $\mu(B) < \mu(A) \implies \mu(B) = 0$ holds). From a measurable perspective this is just $[0, 1]$ with the Lebesgue measure. In order to be able to understand the dynamics of a given group one can investigate its probability measure preserving actions on such a measure space.

Definition 1.6.5. A p.m.p. action α of a group Γ is an action on (X, μ) by measure preserving Borel automorphisms. We say the action is essentially free when $\mu(\{x \in X \mid \alpha(\gamma)(x) = x\}) > 0 \implies \gamma = e$. An action is ergodic whenever there is no invariant $A \subset X$ of positive measure.

Example 1.6.6. Consider the circle S^1 (of length 1) and $s_1, s_2 \in \mathbb{R} \setminus \mathbb{Q}$, rationally inequivalent. Then $\mathbb{Z}^2$ with standard generators a and b , acts on S^1 by translations over a length s_1 (resp. s_2). This action is (essentially) free and ergodic.

One can notice that every p.m.p. action α generates a p.m.p. countable standard equivalence relation $\mathcal{R}$ by simply taking the induced orbit equivalence relation on (X, μ) . More explicitly we have

$$\mathcal{R}_\alpha = \{(x, \alpha(\gamma)(x)) \mid x \in X, \gamma \in \Gamma\}.$$

Definition 1.6.7. An equivalence relation $\mathcal{R}$ on (X, μ) is a p.m.p. countable standard equivalence relation if the following properties are satisfied.

- The equivalence classes of $\mathcal{R}$ are countable;
- $\mathcal{R}$ is a Borel subset of $X \times X$;
- Every partial isomorphism whose graph is contained in $\mathcal{R}$ preserves the measure.

First we remind the reader that a partial isomorphism of X is a Borel isomorphism between two Borel subsets of X . The reader may be surprised to hear that in fact any p.m.p. equivalence relation is the orbit equivalence relation of a group action [FM77]. We now move on to define the cost of such a p.m.p. equivalence relation (and consequently of groups).

Definition 1.6.8. A graphing on (X, μ) is a graph Φ with $V(\Phi) = X$ and Borel edge set $E(\Phi) \subset X \times X$ such that the degrees are uniformly bounded by some constant d . Moreover for all measurable sets $A, B \subset X$

$$\int_A e(x, B) d\mu(x) = \int_B e(x, A) d\mu(x),$$

where $e(x, C)$ is the number of edges from x to $C \subset X$.

One can see that a graphing generates a measure equivalence relation $\mathcal{R}_\Phi$ induced by the connected components. The idea will be to measure the cost of an equivalence relation by measuring the infimal ‘density’ of a generating graphing. As density we will take the average degree.

Definition 1.6.9. The density of a graphing Φ defined on (X, μ) is given by

$$\text{Dens}(\Phi) := \frac{1}{2} \int_X \text{Deg}(x) d\mu(x).$$

Note that we added the one half in front, because the integral counts every edge twice.

Definition 1.6.10. The cost of p.m.p. equivalence relation $\mathcal{R}$ is given by

$$c(\mathcal{R}) := \inf\{\text{Dens}(\Phi) \mid \Phi \text{ generates } \mathcal{R}\}.$$

The cost of a group Γ , is

$$c(\Gamma) := \inf\{c(\mathcal{R}_\alpha) \mid \alpha \text{ is an essentially free p.m.p. action of } \Gamma\}.$$

One says Γ has fixed price when all essentially free p.m.p. actions of Γ have the same cost.

Example 1.6.11. We continue Example 1.6.6 and calculate its cost. Φ_1 is the graphing on S^1 with edge set $E_1 = \{(x, \lambda(x)) \mid \lambda = s_1, s_2 \text{ and } x \in S^1\}$. Clearly $c(\Phi_1) = 2$ since exactly 4 edges are adjacent to each vertex. We try to do better. Fix $\epsilon > 0$ and I some interval in S^1 of length ϵ . We construct a graphing Φ_ϵ on S^1 with edge set $E_\epsilon = \{(x, s_1(x)) \mid x \in S^1\} \cup \{(x, \lambda(x)) \mid \lambda = s_2, s_2^{-1} \text{ and } x \in I\}$. Using that the orbit of s_1 is dense in S^1 and the commutativity between s_1 and s_2 , the graphing Φ_ϵ still generates the relation. It is easy to calculate that $c(\Phi_\epsilon) = 1 + \epsilon$. Hence we conclude that the cost of the relation is $c(\mathcal{R}) = 1$, since one is a lower bound for the cost of infinite groups.

We illustrate the theory with some results on the cost of certain (classes) of groups.

Theorem 1.6.12. *Let Γ be a countable group*

- *If Γ is finite, then it has fixed price $1 - \frac{1}{|\Gamma|}$. ([Lev95])*
- *If Γ is a free group on n generators, then it has fixed price n . ([Gab00])*
- *If Γ is an infinite amenable group, then it has fixed price 1, moreover Γ is amenable if and only if it has an essentially free p.m.p. action such that the infimum $c(\mathcal{R})$ is attained by some graphing. ([OW80])*
- *If Γ is surface group $\pi_1(\Sigma_g)$, where Σ_g is a closed orientable surface of genus $g \geq 2$, then it has fixed price $2g - 1$. ([Gab00])*
- *If Γ has property (T), then it has cost 1. ([HP18])*
- *Let Λ be a finite index subgroup of Γ , then ([Gab00])*

$$c(\Lambda) - 1 = [\Gamma, \Lambda](c(\Gamma) - 1).$$

This important realisation by Gaboriau in [Gab98] will explain the intuition behind the definition of combinatorial cost.

Definition 1.6.13. The L -Lipschitz cost of a graphing Φ is given by

$$c_L(\Phi) := \inf\{\text{Dens}(\Psi) \mid \Psi \text{ is } L\text{-bilipschitz equivalent to } \Phi\}.$$

Lemma 1.6.14. [Gab98, étape 1] Let Φ be graphing of finite density, then the cost of the p.m.p. equivalence relation $\mathcal{R}$ generated by Φ can be calculated in its L -Lipschitz classes, i.e.

$$c(\mathcal{R}) = \inf_L \{c_L(\Phi)\}.$$

1.6.3 Hyperfinite equivalence relations and treeable graphings

The ultimate goal of introducing cost, is to justify the analogue of combinatorial cost. In this philosophy we introduce hyperfinite relations and treeable graphings in order to show that the definition of their analogues in the combinatorial setting is a natural one. Note that a p.m.p. equivalence relation is finite whenever each of its orbits is finite.

Definition 1.6.15. A p.m.p. equivalence relation $\mathcal{R}$ is hyperfinite if there exist standard finite p.m.p. subrelations $\mathcal{R}_n$ such that $\mathcal{R} = \cup_n \mathcal{R}_n$. More precisely, for any $x \in X$, the orbits give $\mathcal{R}(x) = \cup_n \mathcal{R}_n(x)$.

We note that this is not at all a trivial property. The key constraint is that the $\mathcal{R}_n$ must form Borel subspaces of $X \times X$. One can show that hyperfinite relations have cost 1. Now we turn to graphings.

Definition 1.6.16. A graphing Φ is a treeing if almost every connected component is a tree.

Treeings are very interesting tools to calculate cost since we have the following theorem by Gaboriau.

Theorem 1.6.17. If Φ is a treeing of a p.m.p. action $\mathcal{R}$, then it realises the cost, i.e. $c(\Phi) = c(\mathcal{R})$. Moreover the converse holds. If Φ realises the cost of $\mathcal{R}$, then it must be a treeing.

1.6.4 Combinatorial cost and hyperfiniteness

The combinatorial analogue of cost was introduced by Elek in [Ele07]. We define the following basic concepts.

Definition 1.6.18. The edge number of a graph sequence $\mathcal{G}$ is defined by

$$e(\mathcal{G}) = \liminf_{n \rightarrow \infty} \frac{|E(G_n)|}{|V(G_n)|}.$$

The cost is given by the infimum

$$c(\mathcal{G}) = \inf_{\mathcal{G} \simeq \mathcal{H}} e(\mathcal{H}),$$

where $\simeq$ denotes bilipschitz equivalence of graph sequences.

The choice to take the infimum over L -bilipschitz equivalent graph sequences is clearly inspired by Lemma 1.6.14. It turns out there is also a more direct link between the two which was recently discovered by both [CGdLS18] and [AT17]. We cover this in Section 1.6.5. We note that combinatorial cost is not a coarse invariant.

Hyperfiniteness was introduced by Elek as a property of graph sequences.

Definition 1.6.19. A graph sequence $\mathcal{G} = \{G_n\}_{n \in \mathbb{N}}$ is hyperfinite if $\forall n \in \mathbb{N}$ and $\forall \epsilon > 0$ there exist sets $E_n^\epsilon \subset E(G_n)$ and a constant $K > 0$ such that

1. $\lim_{n \rightarrow \infty} \frac{|E_n^\epsilon|}{|V(G_n)|} < \epsilon$;
2. $\forall n \in \mathbb{N}$ one has that the connected components of $G_n \setminus E_n^\epsilon$ contain at most K vertices.

It is very similar to the definition of hyperfiniteness for p.m.p. equivalence relations. As an unprecise philosophical statement one could say that in a hyperfinite graph sequence, the graphs G_n are asymptotically equal to disjoint unions of finite graphs of uniformly bounded diameter. Also it is easy to show that the combinatorial cost of such a graph sequence is one. Finally we introduce the equivalent of treeings. Define $g(G)$ the girth of a graph as the length of the smallest cycle. A graph sequence is of large girth if $\lim_{n \rightarrow \infty} g(G_n) = \infty$. We have the following result by Elek, similar to Theorem 1.6.17.

Theorem 1.6.20. *If $\mathcal{G}$ is a large girth sequence, then $c(\mathcal{G}) = e(\mathcal{G})$.*

1.6.5 Limit graphings of graph sequences

The goal of this section is to state a nice result linking together cost and combinatorial cost. As such the goal is *not* to introduce all constructions in a detailed way since this would lead us too far. We hope to bring across some ideas and in doing so convince the reader of the beauty and usefulness of ultraproducts, local-global convergence of graphs and the encompassing theory in general. As a last remark, we do not claim to present constructions, definitions and results in their full generality. We start by constructing the ultraproduct graphing Φ_ω of a graph sequence $\mathcal{G}$ along a non-principal ultrafilter ω . We define a pseudo-distance on the product $\prod_{n=1}^\infty G_n$ as follows: two sequences (x_n) and (y_n) are at a distance

$$d_\omega((x_n), (y_n)) := \lim_{\omega} d_n(x_n, y_n),$$

where d_n is the distance induced by the graph G_n . This function can both be zero or infinite. We say two sequences are equivalent whenever they are at a distance zero. The ultraproduct is then equal to the product above quotiented by the equivalence, i.e. $\prod_{n=1}^\infty G_n / \sim$. By d_ω we denote the induced distance on $\Phi_\mathcal{G}$, which only takes values in $\mathbb{N}$ and thus induces also a graph structure on the limit by taking an edge between points which are a distance one apart from each other. Of course to obtain a graphing one needs a probability space. The basic idea is to consider the G_n with a uniform probability measure μ_n , then for any sequence of subsets $A_n \subset G_n$, we define

$$\mu_\omega(\prod_{n=1}^\infty A_n / \sim) = \lim_{\omega} \mu_n(A_n).$$

Of course this does not define a probability measure on the limit, but there are techniques to extend this to a measure (see [Car15]). Great, let us assume that our limit graphing Φ_ω is really a graphing, then it is time for a warning: Φ_ω is *not* defined on a standard probability space! We are no longer in this nice pleasant environment as in Section 1.6.2, where we played before. But the good news is that all the concepts we defined generalise in a natural way, including cost. Let us define the following generalisation of combinatorial cost as introduced in [CGdLS18].

Definition 1.6.21. Let ω be a (non-principal) ultrafilter, then the ω -edge number of a graph sequence $\mathcal{G}$ is defined by

$$e_\omega(\mathcal{G}) = \lim_{\omega} \frac{|E(G_n)|}{|V(G_n)|}.$$

The ω -cost is given by the infimum

$$c_\omega(\mathcal{G}) = \inf_{\mathcal{G} \approx \mathcal{H}} e_\omega(\mathcal{H}),$$

where $\approx$ denotes bilipschitz equivalence of graph sequences.

The following result is due to Abert and Toth, and Carderi, de la Salle and Gaboriau independently [AT17, CGdLS18]. The two articles have different viewpoints and use a different mathematical language to (for this common result) arrive at the same conclusion.

Theorem 1.6.22. *Let ω be a (non-principal) ultrafilter, $\mathcal{G}$ a graph sequence, and Φ_ω the ultraproduct graphing along ω , then $c_\omega(\mathcal{G}) = c(\Phi_\omega)$*

Moreover considering Elek's definition of cost this becomes.

Theorem 1.6.23. *Let $\mathcal{G}$ be a graph sequence, and Φ_ω the ultraproduct graphing along a non-principal ultrafilter ω , then $c(\mathcal{G}) = \inf\{c(\Phi_\omega) \mid \omega \text{ is a non-principal ultrafilter}\}$*

The reason for these results is somewhat hidden. The key is a strengthening of Benjamini-Schramm convergence, called local-global convergence (introduced in [HLS14]). We explain the idea. Local-global convergence does not only force local neighbourhood statistics to converge, but also global colourings and this in a measurable way. However bilipschitz equivalent graph sequences (rewirings) are exactly this, colourings of the edges! This means that we can copy the colouring to the limit graphing and vice versa. Hence we can rewire the limit graphing as we have rewired the graphs in our graph sequence. This is where ultralimits come in handy. While not every graph sequence converges in a local-global way, it always contains a subsequence of graphs that converges locally-globally. The ultrafilter choses such a subsequence for us and uses it to model a limit graphing that strongly resembles it. We refer the reader to the paper of Abert and Toth for the specifics of the result.

1.7 Uniformly finite homology

Uniformly finite homology was introduced by Block and Weinberger in [BW92] to study the coarse structure of metric spaces with bounded local geometry. The chains considered are basically those of Roe's coarse homology (see [Roe03]), but with an extra finiteness boundedness condition. The obtained homology is a quasi-isometric invariant. We introduce uniformly finite homology for metric spaces, then we define simplicial uniformly finite homology and an alternative definition of uniformly finite homology for simplicial complexes. A good reference is the thesis of Diana [Dia15].

1.7.1 Uniformly finite homology for metric spaces

Let A be in $\{\mathbb{R}, \mathbb{Z}, \mathbb{Z}_2\}$. Let (X, d) be metric space and endow the product space X^{n+1} with the max distance, i.e. for $x = (x_0 \dots x_n)$ and $y = (y_0 \dots y_n)$ in X^{n+1} ,

$$d(x, y) = \max\{d(x_i, y_i) \mid 0 \leq i \leq n\}.$$

We base the following on [BW92]. For each $n \in \mathbb{N}$ we define the chains as $C_n^{\text{uf}}(X, A)$ as the vector space generated by the infinite sums

$$c = \sum a_x x,$$

where $a_x \in A$ and $x \in X^{n+1}$ such that

1. There is some K such that $|a_x| < K$, for any $x \in X^{n+1}$.
2. $\forall r \in \mathbb{R}$, there is a K_r , such that $\forall x \in X^{n+1}$

$$|B_r(x) \cap \{y \in X^{n+1} \mid a_y \neq 0\}| \leq K_r.$$

3. There is some $R > 0$, such that whenever $a_x \neq 0$, then $\max_{i \neq j} d(x_i, x_j) \leq R$. This is the propagation of c .

The boundary map is the ordinary boundary map, which is well-defined because of all the boundedness conditions. Uniformly finite homology is the homology associated to this chain complex.

1.7.2 Simplicial Uniformly Finite Homology H^{suf}

Since we will only be interested in the homology of graphs we give an alternative definition in this section, which better suits our purpose and passes through simplicial complexes. We say a simplicial complex is uniformly locally finite if there is a uniform bound on the cardinality of the links of its vertices. Let X be a such a ULF simplicial complex, take $A \in \{\mathbb{Z}, \mathbb{R}, \mathbb{Z}_2\}$. Define:

$$C_n^{\text{suf}}(X, A) = l^\infty(X_{(n)}, A),$$

where $X_{(n)}$ is the set of $n + 1$ -simplices of X . As a boundary map ∂ , we take the ordinary boundary map on a simplicial complex, which we extend to infinite sums. This is well-defined since X is ULF.

Then, the *simplicial uniformly finite homology* of X is defined as

$$H_*^{\text{suf}}(X) = H(C_*^{\text{suf}}(X)).$$

Now if X is a ULF graph and $r \in \mathbb{N}$, define the *Rips Complex* of radius r as having n -simplices:

$$(R_r(X))_{(n)} = \{(x_1, \dots, x_n) \in X_{(0)}^n : d(x_i, x_j) \leq r\},$$

where d is the graph distance (i.e. length of a shortest path). A simplex of $R_r(x)$ is seen as a “virtual” simplex in X . Then $R_1(X) = X$, $R_r X \subseteq R_{r'}(X)$ whenever $r \leq r'$, and $R_r(X)$ is itself a ULF simplicial complex. It follows that we have a directed system:

$$(C_*^{\text{suf}}(R_r(X)))_{r \in \mathbb{N}}$$

and can define:

$$C_*^{\text{uf}}(X) := \lim_r C_*^{\text{suf}}(R_r(X))$$

and finally, the *uniformly finite homology* of X is defined as:

$$H_*^{\text{uf}}(X) = H(C_*^{\text{uf}}(X))$$

As a remark, note that $H_*^{\text{uf}}(X)$ can also be defined as the limit of the system

$$(H_*^{\text{suf}}(R_r(X)))_r.$$

In the following, we will use

$$Z^{\text{uf}}, B^{\text{uf}}, Z^{\text{suf}}, B^{\text{suf}},$$

for the kernels and image of the boundary maps in the chain complexes C^{uf} and C^{suf} respectively.

The following two important facts are proved in [Mos03].

Proposition 1.7.1 (H^{uf} is QI-invariant.). H^{uf} is a quasi-isometry invariant.

Proof. See [Mos03, Step 3 in the proof of Theorem 12]. □

Proposition 1.7.2. If X is ULF, UC, then $H^{\text{suf}}(X) \cong H^{\text{uf}}(X)$.

Proof. See [Mos03, Step 2 in the proof of Theorem 12]. □

1.7.3 Uniformly finite homology and expansion

As promised in Section 1.4, we show that H_0^{uf} detects expansion. We base our proof on [BW92] and for the implication (5) $\implies$ (1) we follow [BS97]. For the latter we introduce a restricted version of the very classical (and very beautiful) Max-Flow Min-Cut theorem. For a more in depth description see [LP16]. Consider a finite graph G with two distinguished vertices, a source s and a sink t . One also has a capacity function on the edges $c : E(G) \rightarrow \mathbb{R}^+$. A flow on such graph is a function on directed edges $f : \vec{E}(G) \rightarrow \mathbb{R}^+$, such that $f(\vec{e}) \leq c(e)$, and $\forall x \in G \setminus \{s, t\}$, the sum of the values of incident vertices is zero, i.e. $\forall x$

$$\sum_{y \sim x} f(y, x) - f(x, y) = 0.$$

The flow rate of a given f , is equal to $\dagger(f) := \sum_{y \sim s} f(y, s) - f(s, y)$, being the value of current coming into the network. The goal of the Max-Flow Min-Cut theorem is to describe the maximal flow rate that can be obtained on a given graph. A cut C is a subset of edges, such that in $G \setminus C$ the source and sink are in different connected components. The capacity of a cut is the sum of the capacities on the edges, i.e. $c(C) := \sum_{e \in C} c(e)$. This value gives a bound on the possible flow rate on G since all current going from s to t must pass through one of the edges of C . The surprising result is that the maximal flow rate attainable on G is exactly the value of the minimal cut.

Theorem 1.7.3 (Max-Flow Min-Cut). Let G be a finite graph with distinguished vertices s and t , let c be a capacity function on its edges, then the maximal flow rate equals the minimal cut, i.e.

$$\max\{\dagger(f) \mid f \text{ is a flow on } G\} = \min\{c(C) \mid C \text{ is a cut on } G\}$$

Proposition 1.7.4. Let G be a graph, then the following are equivalent:

1. $H_0^{\text{uf}}(G, \mathbb{Z}) = 0$;
2. $H_0^{\text{uf}}(G, \mathbb{R}) = 0$;
3. There is some $c = \sum a_x x$, with $a_x > 0$, such that $[c] = 0$ in $H_0^{\text{uf}}(G, \mathbb{Z}) = 0$;
4. There is some $c = \sum a_x x$, with $a_x > \epsilon > 0$, such that $[c] = 0$ in $H_0^{\text{uf}}(G, \mathbb{R}) = 0$;
5. G is not amenable.

Proof. We show the equivalence of (1), (3) and (5). For a complete proof see [BW92]. Note that point (1) trivially implies point (3). To see that (3) also implies (1) we apply Lemma 1.7.5. Fix any $c' \in C_0^{\text{uf}}(G, \mathbb{Z})$ of the form $c' = \sum_{x \in G} a'_x x$. Then $\sum_{x \in G} a'_x t_x$ is well-defined and bounded with boundary c' .

We show that (1) implies (5). Write $\mathbb{1}_G$ for the characteristic function on the vertices of G . Now there is some $\phi \in C_1^{\text{uf}}(G, \mathbb{Z})$, such that $\partial(\phi) = \mathbb{1}_G$ with propagation r . Note that we can suppose ϕ to be of propagation one. The reason is the following. For each edge (x, y) with non-zero coefficient in ϕ , choose a path $p_{y,x}$ of length $\leq r$ in G from y to x , then $(x, y) + \sum_{e \in p_{x,y}} e$ is a boundary. Note that the sum of all these paths is well-defined and bounded, since G has bounded geometry. Hence consider ϕ of propagation one. Then for any finite $U \subset_f G$ we have

$$\begin{aligned}
|U| &= \sum_{x \in U} \mathbb{1}_G(x) \\
&= \sum_{x \in U} (\partial(\phi))(x) \\
&= \sum_{x \in U} \sum_{y \in G} \phi(y, x) - \phi(x, y) \\
&= \sum_{x \in U} \sum_{y \in U} \phi(y, x) - \phi(x, y) + \sum_{x \in U} \sum_{y \notin U} \phi(y, x) - \phi(x, y) \\
&\leq \sum_{x \in U} \sum_{y \notin U} |\phi(y, x) - \phi(x, y)| \\
&\leq \sum_{x \in U} \sum_{y \notin U} 2\|\phi\|_\infty \\
&\leq 2|\partial(U)|\|\phi\|_\infty,
\end{aligned}$$

which shows expansion.

Finally we show (5) implies (3). See also Figure 1.9. We will use the Max-Flow Min-Cut principle. Let $\alpha > 0$ be the Cheeger constant of G . Consider an exhaustion of G , by finite sets K_n , i.e. $K_n \subset K_{n+1}$ and $\bigcup_n K_n = G$. Then construct the finite graph G_n . Start with the vertex-induced subgraph on $K_n \cup \partial(K_n)$, to which we add two vertices, a source s and a sink t . Link the vertices of K_n to s and the vertices of $\partial(K_n)$ to t . We define a conductance function on the edges by

$$c : E(G_n) \rightarrow \mathbb{R}^+ : e \mapsto \begin{cases} \alpha & \text{if } e = (s, k), \\ 1 & \text{otherwise.} \end{cases}$$

Then we calculate a lower bound for the value of a cut C on G_n , separating s and t . Define U to be the connected component of s in $G_n \setminus C$, constraint to K_n . We remark two things:

1. whenever $k \in K_n \setminus U$, then $(s, k) \in C$,
2. whenever $k \in \partial(U)$, there is some $\tilde{k} \in K_n$ such that $(\tilde{k}, k) \in C$.

Hence $c(C) \geq \alpha|K_n \setminus U| + |\partial(U)| \geq \alpha|K_n|$. By the Max-Flow Min-Cut theorem, we obtain that there is a flow f_n with $\dagger(f_n) = \alpha|K_n|$ flowing from s to t in G_n . Let $m \in \mathbb{N}$ be such that $\frac{1}{m} \leq \alpha$.




Figure 1.9: Construction of the graph G_n on the vertex set $K_n \cup \partial(K_n)$.

Note that we may suppose values of f_n to be of the form $\frac{l}{m}$, with $l \in \mathbb{N}$. Then forgetting about the points s and t , the function $m \cdot f_n$ runs from $E(K_n \cup \partial(K_n))$ to $\mathbb{Z}$, with

1. $\partial(m \cdot f_n)$ equal to the identity on K_n , and
2. $\|m \cdot f_n\|_\infty \leq m$.

Hence letting f be an accumulation point of the sequence $\{m \cdot f_n\}_n$ gives a function on C_1^{uf} , such that $\partial(f) = \mathbb{1}_G$. $\square$

Lemma 1.7.5. *Suppose there exists $c = \sum a_x x$, with $a_x > 0$, such that $[c] = 0$ in $H_0^{uf}(G, \mathbb{Z}) = 0$. Then for each $x \in G$, there is some tail $t_x \in C_1^{uf}(G, \mathbb{Z})$, such that $\partial(t_x) = x$ and $\sum_{x \in G} t_x$ is locally finite and bounded.*

Proof. Firstly there is some $\phi \in C_1^{uf}(G, \mathbb{Z})$, such that $\partial(\phi) = c$. Note that we may suppose the coefficients of ϕ to be positive (i.e. ≥ 0). Now fix x . Since $\partial(\phi) = c$, there is some edge (x_1, x) in ϕ of non-zero coefficient. Now consider the positive sum $\phi - (x_1, x)$, then the coefficient of x_1 in $\partial(\phi - (x_1, x))$ is strictly positive. So there is some edge (x_2, x_1) with non-zero coefficient. We continue inductively and obtain a tail $t_x = (x_1, x) + (x_2, x_1) + (x_3, x_2) \dots$ which is uniformly bounded since the value of its coefficients is bounded by the coefficients of ϕ . Now to find a tail $t_{x'}$ for $x' \neq x \in G$, we can apply the same algorithm to $\phi - t_x$ and inductively for any vertex of G . This concludes the proof. $\square$

1.7.4 A theorem and a question

The technique used in Lemma 1.7.5 is called the Eilenberg swindle technique. Diana and Nowak use this approach to calculate the higher degree uniformly finite homology of certain spaces, such as products of trees in [DN17]. In general, however, few results exist. We also refer the reader to the overview paper by Block and Weinberger [BW97].

Theorem 1.7.6. *Let T_i be uniformly locally finite infinite trees, in which each vertex has degree at least 3 and let $R = \mathbb{Z}, \mathbb{R}$. Let $T_1 \times T_2 \times \dots \times T_k$ be their Cartesian product endowed with the maximum metric. Then*

$$H_n^{\text{uf}}(T_1 \times T_2 \times \dots \times T_k, R) = \begin{cases} 0 & \text{if } k \neq n, \\ \text{infinite dimensional} & \text{if } k = n. \end{cases}$$

The attentive reader may notice that for another homology theory this theorem would simply be the result of a Künneth formula. Such a tool is not available for uniformly finite homology, even though it exists for its cousin, Roe's coarse cohomology. We present a conjecture stated by Diana.

Conjecture 1.7.7. *Let X, Y be metric spaces of coarse bounded geometry and let $A \in \{\mathbb{R}, \mathbb{Z}\}$. Let $n_1, n_2 \in \mathbb{N}$. Assume that $H_k^{\text{uf}}(X, A) = 0$ for all $k \leq n_1$ and $H_k^{\text{uf}}(Y, A) = 0$ for all $k \leq n_2$. Then $H_k^{\text{uf}}(X \times Y, A) = 0$ for all $k \leq n_1 + n_2 + 1$.*

2 Overview

2.1 On combinatorial cost: a coarse setting

The main inspiration for this chapter is a paper by Elek [Ele07] where he introduces combinatorial cost for graph sequences as a combinatorial version of cost in the setting of probability measure preserving actions. We consider cost from a coarse setting and show that having cost equal to 1 and hyperfiniteness are coarse invariants. We show a combinatorial version of Gabo-riau's theorem 1.6.12, which says that 'cost-1' for Farber sequences behaves multiplicatively when taking subgroups.

Theorem 2.1.1. *Take Γ a residually finite group and $\Lambda \leq \Gamma$ with $[\Gamma : \Lambda] < \infty$. Consider a Farber sequence $(N_k)_{k \in \mathbb{N}}$ of Γ that is contained in Λ . Define $\mathcal{G} = \{\text{Sch}(\Gamma, N_k, S)\}_{k \in \mathbb{N}}$ and $\mathcal{H} = \{\text{Sch}(\Lambda, N_k, S')\}_{k \in \mathbb{N}}$ then*

$$[\Gamma : \Lambda](c(\mathcal{G}) - 1) = c(\mathcal{H}) - 1.$$

We show that graph sequences coming from Farber sequences of a group have property A if and only if the group is amenable. The same is true for hyperfiniteness. This generalises the following theorem by Elek who stated it for box spaces. The equivalence between 1 and 2 is originally due to Guentner and proven by Roe [Roe03]. The equivalence between 1 and 3 is due to Elek.

Theorem 2.1.2 (Theorem 1.4, [Ele07] and Proposition 11.39, [Roe03]). *Consider a residually finite group Γ and a filtration $(N_k)_{k \in \mathbb{N}}$. If $\mathcal{G}$ is the associated graph sequence of Cayley graphs then the following are equivalent:*

1. Γ is amenable;
2. $\mathcal{G}$ has property A;
3. $\mathcal{G}$ is hyperfinite.

Furthermore we optimise the result when Farber sequences are replaced by sofic approximations. The following is already known.

Theorem 2.1.3. *Consider a group Γ and a sofic approximation $\mathcal{G}$ of Γ . The following holds*

1. *If $\mathcal{G}$ has property A, then Γ is amenable (Theorem 4.5, [AFS19]).
The reverse arrow does not hold.*
2. *If $\mathcal{G}$ is hyperfinite, then Γ is amenable.
The reverse arrow holds (Theorem 2, [Ele06] and Proof of Theorem 1.4, [Ele07]).*

Sadly enough this generalisation to the sofic approximation case breaks the symmetry. This motivates us to introduce property almost-A in Definition 3.5.4. This 'saves' the theorem.

Theorem 2.1.4. *Consider a group Γ and a sofic approximation $\mathcal{G}$ of that group, then the following are equivalent:*

1. Γ is amenable;
2. $\mathcal{G}$ has property almost-A;
3. $\mathcal{G}$ is hyperfinite.

Finally we show the following result for general graph sequences. It is a generalisation of a result in [AFS19], where it is proved for sofic approximations.

Theorem 2.1.5. *A graph sequence $\mathcal{G}$ that has property A is also hyperfinite.*

2.2 On the first uniformly finite homology group for graphs

Uniformly finite homology, introduced in [BW92], is a coarse invariant of well-behaved metric spaces. In [BW92], it is shown that vanishing of zeroth uniformly finite homology with coefficients in $\mathbb{Z}$ or $\mathbb{R}$ is equivalent to expansion (non-amenability) of the space at hand. Higher uniformly finite homology groups do not enjoy, as far as we know, such clear cut descriptions. Partial results have been obtained, e.g. in [BD15, DN17]. This is joint work with Rémi Bottinelli.

In Chapter 4 we focus on the first homology group, with coefficients in $\mathbb{Z}$ of (tame enough) graphs: $H_1^{\text{uf}}(\cdot, \mathbb{Z})$. Informally, three relatively independent phenomena are responsible for the appearance of homology classes in dimension one:

1. Graph theoretical ends.
2. The existence of circuits of arbitrarily large girth. More precisely, the non-existence of a uniform constant R such that any circuit can be written as a (possibly infinite) sum of circuits of length $\leq R$.
3. A lack of 2-simplices in the clique complex. This should be understood quantitatively, and likened to amenability.

We give a classification of the (first) homology of trees in terms of their ends, we then show that the homology of an arbitrary graph is bounded from below by that of an embedded tree that “realises the ends”.

Proposition 2.2.1. *Let X be ULF and $T \leq X$ an end-respecting tree. Then:*

$$H_1^{\text{uf}}(T, \mathbb{Z}) \leq H_1^{\text{uf}}(X, \mathbb{Z}).$$

The same holds for $\mathbb{Z}_2$.

Then, motivated by the observation that there is a surjection:

$$H_1^{\text{uf}}(\cdot, \mathbb{Z}) \rightarrow H_1^{\text{uf}}(\cdot, \mathbb{Z}_2),$$

we study $H_1^{\text{uf}}(\cdot, \mathbb{Z}_2)$ (note that the “uniformly finite” condition becomes vacuous in $\mathbb{Z}_2$ — this may sound like a bad omen, but turns out to clarify a lot of constructions). In particular, we get a

full description of $H_1^{\text{uf}}(\cdot, \mathbb{Z}_2)$ as a direct sum of two factors corresponding to exactly the first and second phenomena, that is, ends and “large circuits” respectively. Moreover, it is shown that the existence of “large circuits” actually implies infinite-dimensionality of $H_1^{\text{uf}}(\cdot, \mathbb{Z}_2)$. Moreover we define a dimension of ‘thin’ sums which is able to distinguish between graphs with “large circuits” and those without.

Theorem 2.2.2. • $\dim_{\text{thin}}(H_1^{\text{uf}}(X, \mathbb{Z}_2))$ is countable if and only if X has no large circuits. In this case it is equal to the number of ends -1 (where we do not make the difference between countably many or uncountably many ends).

• $\dim_{\text{thin}}(H_1^{\text{uf}}(X, \mathbb{Z}_2))$ is uncountable if and only if X has large circuits.

Finally, once again in for $\mathbb{Z}$, we make precise the claim that “large circuits” imply non-zero homology. Since “large circuits” is a notion which depends on the coefficient rings, this result is *not* a consequence of the previous one. Furthermore inspired by a construction of [BS97], we define a notion of “expansion” in higher dimension, which aims at capturing the third phenomenon. It is easily verified that our “expansion” implies vanishing of homology. As a (very partial) converse, we show that in dimension one, and for vertex transitive graphs, vanishing implies “expansion”.

Finally, refining the notion of expansion, we reach, for (infinite) transitive graphs (or more generally, infinite graphs with a cocompact action by their automorphism group), a characterisation of the non-vanishing of homology in terms of the three phenomena of (the refined) non-expansion, ends, and large circuits.

Theorem 2.2.3 (H_1^{uf} triad). *Let X a transitive graph, and consider the following three conditions:*

1. X has more than one end.
2. X has $\mathbb{Z}$ -large circuits; that is, $C_r(X, \mathbb{Z}) \neq C_\infty(X, \mathbb{Z})$ for any r .
3. X does not have pure H_1^{uf} -expansion.

Then $H_1^{\text{uf}}(X, \mathbb{Z}) \neq 0$ iff any of the above holds.

2.3 On $BS(2, 3)$

The Baumslag Solitar group $BS(2, 3)$, is a so-called non-Hopfian group, meaning that it has an epimorphism onto itself, that is not injective. In particular this is equivalent to saying that $BS(2, 3)$ has a quotient that is isomorphic to itself. This is truly a magnificent property to have. As a consequence some Cayley graph of $BS(2, 3)$ has a quotient that is isomorphic to itself up to change of generators. We describe this quotient on the graph-level and take a closer look at the specific epimorphism

$$\phi : BS(2, 3) \rightarrow BS(2, 3) : \begin{cases} a \mapsto a^2 \\ b \mapsto b \end{cases}.$$

We show that its kernel is a free group of infinite rank, by showing that its action on the Bass-Serre tree is free.

3 On combinatorial cost: a coarse setting

The main inspiration for this chapter is a paper by Elek where he introduces combinatorial cost for graph sequences [Ele07]. In Section 3.1 and 3.2 we show that having cost equal to 1 and hyperfiniteness are coarse invariants. In Section 3.3 we show that ‘cost-1’ for box spaces behaves multiplicatively when taking subgroups. In Section 3.5 we show that graph sequences coming from Farber sequences of a group have property A if and only if the group is amenable. The same is true for hyperfiniteness. This generalises a theorem by Elek. Furthermore we optimise this result when Farber sequences are replaced by sofic approximations. In doing so we introduce a new concept: property almost-A. This Chapter is based on [Kai19].

3.1 Cost 1 is a coarse invariant

We show that having cost 1 is a coarse invariant. First consider at a specific case, where the two graph sequences only differ by degree one vertices. We introduce the following terminology to make this precise.

Definition 3.1.1. Consider two graph sequences $\mathcal{G}$ and $\mathcal{G}'$.

- $\mathcal{G}'$ is a subgraph sequence of $\mathcal{G}$ if for every $n \in \mathbb{N}$, the graph G'_n is a vertex induced subgraph of G_n .
- When $\mathcal{G}'$ is a subgraph sequence of $\mathcal{G}$, a vertex of $\mathcal{G}'$ will be called an internal vertex of $\mathcal{G}$.
- When $\mathcal{G}'$ is a subgraph sequence of $\mathcal{G}$, a vertex in $V(G_n) \setminus V(G'_n)$ will be called an external vertex of $\mathcal{G}$.
- A vertex x is a terminal vertex of $\mathcal{G}$ with respect to the subgraph sequence $\mathcal{G}'$ if it is external and has degree one in G_n for some n .
- Given a terminal vertex $x \in G_n$, its base vertex is the unique vertex y of G_n such that $(x, y) \in E(G_n)$.

Consider the situation where $\mathcal{G}'$ is a subgraph sequence of $\mathcal{G}$, and any external vertex is terminal. Note that this means that the base vertex of any terminal vertex is internal. We show that $c(\mathcal{G}') = 1$ if and only if $c(\mathcal{G}) = 1$. In order to do this we first show that we can minimise the edge number of $\mathcal{G}$ by only rewiring the subgraph sequence $\mathcal{G}'$ (hence leaving the terminal vertices and their edge untouched).

Lemma 3.1.2. *Given a graph sequence $\mathcal{G}$ and a subgraph sequence $\mathcal{G}'$ where any external vertex is terminal, then there exists for every $\epsilon > 0$ a graph sequence $\mathcal{G}_\epsilon$ such that*

1. $\mathcal{G}_\epsilon \simeq \mathcal{G}$;
2. $e(\mathcal{G}_\epsilon) \leq c(\mathcal{G}) + \epsilon$;

3. if (x, y) is an edge in $\mathcal{G}_\epsilon$, but not in $\mathcal{G}$ (or vice versa), then x and y are internal vertices.

Proof. There trivially exist graph sequences $\mathcal{H}_\epsilon$ (these are L_ϵ -bilipschitz to $\mathcal{G}$) that satisfy point 1 and 2. It is enough to show that from such sequences we can construct new ones that have property 3 but keep properties 1 and 2. Notice that we have the following types of edges in $\mathcal{H}_\epsilon$. In what follows, vertices are terminal/internal with respect to the graph sequences $\mathcal{G}$ and $\mathcal{G}'$.

1. edges between terminal vertices;
2. edges between internal vertices;
3. edges between terminal vertices and internal vertices.

We show that $\mathcal{H}_\epsilon$ can be chosen such that it does not contain edges of type one. Suppose an edge (y_1, y_2) in $\mathcal{H}_\epsilon$ is an edge between terminal vertices such that y_1 is connected to an internal vertex x , then we remove (y_1, y_2) and add the edge (x, y_2) (see Figure 3.1a). One can do this for all such edges and obtain a graph sequence that is 2-bilipschitz equivalent to $\mathcal{H}_\epsilon$. Of course it is possible that neither y_1 nor y_2 is connected to an internal vertex. However, since, between any terminal vertex and its base vertex there is a path of length at most L_ϵ in $\mathcal{H}_\epsilon$, any edge between terminal vertices is L_ϵ -close to an edge between a terminal vertex and an internal vertex. This means that repeating the process above at most L_ϵ times gives us a graph sequence that is $2L_\epsilon$ -bilipschitz to $\mathcal{H}_\epsilon$ without edges of type one. Hence we can assume that $\mathcal{H}_\epsilon$ does not have edges of type one.




Figure 3.1: Construction of edges.

Let us take a look at edges of type three. Take a terminal vertex y . If it is connected to more than one internal vertex, then we choose one of these vertices (for example x) and reconnect as follows: take any other vertex connected to y (for example x'). Then we delete the edge (x', y) and introduce the edge (x', x) . (see Figure 3.1b). Doing this for all terminal vertices gives a graph sequence that is 2-bilipschitz to $\mathcal{H}_\epsilon$. We can assume terminal vertices in $\mathcal{H}_\epsilon$ are connected to at most one internal vertex.




Figure 3.2: Construction of edges of type three (b).

Now it is possible that some terminal vertices are not linked to their base vertex x (in $\mathcal{G}'$), but to some other internal vertex x' . However since there is no edge between the terminal vertex y

and its base vertex x , there is a path in $\mathcal{H}_\epsilon$ of length at most L_ϵ between them. The first edge of this path links y to some internal vertex x' . We delete the edge (x', y) and add the edge (x, y) (see Figure 3.2). It is clear this gives a graph sequence that is L_ϵ -bilipschitz to $\mathcal{H}_\epsilon$. Note that we can do this independently for all such edges (since, apart from the first edge, these paths are in the subgraph sequence $\mathcal{G}'$). Note that this last graph sequence has all the needed properties. Replacing $\mathcal{H}_\epsilon$ by this graph sequence concludes the proof. $\square$

Corollary 3.1.3. *Given a graph sequence $\mathcal{G}$ and a subgraph sequence $\mathcal{G}'$ where any external vertex is terminal, then $\mathcal{G}$ has cost 1 if and only if $\mathcal{G}'$ has cost 1.*

Proof. If $c(\mathcal{G}') = 1$, then for every $\epsilon > 0$ we can find equivalent graph sequences $\mathcal{G}'_\epsilon$ such that $e(\mathcal{G}'_\epsilon) \leq 1 + \epsilon$. We can obtain a rewiring of $\mathcal{G}$ by rewiring the subgraph sequence $\mathcal{G}'$ inside it in the same way as the rewiring $\mathcal{G}'_\epsilon$. This yields a graph sequence $\mathcal{G}_\epsilon$ equivalent to $\mathcal{G}$ which has edge number less or equal than $1 + \epsilon$.

On the other hand, if $c(\mathcal{G}) = 1$, we can find for every $\epsilon > 0$ equivalent graph sequences $\mathcal{G}_\epsilon$ as in Lemma 3.1.2 (i.e. $\mathcal{G}$ can be rewired by only making changes to the subgraph sequence $\mathcal{G}'$). Note that, given a graph, deleting terminal vertices does not change the distance between any two remaining vertices. This means that reducing $\mathcal{G}_\epsilon$ to the vertices of $\mathcal{G}'$ gives a graph sequence $\mathcal{G}'_{\text{red},\epsilon} = \{G'_{\text{red},\epsilon,n}\}_{n \in \mathbb{N}}$ that is bilipschitz to $\mathcal{G}'$. We calculate the edge number of this graph sequence. The expression $C_{\mathbb{Q}}(G'_{\text{red},\epsilon,n})$ denotes the vector space generated by the cycles of $G'_{\text{red},\epsilon,n}$ over the field $\mathbb{Q}$. Moreover since $\mathcal{G}'_{\text{red},\epsilon}$ is just $\mathcal{G}_\epsilon$ without some terminal vertices, one sees that $\dim_{\mathbb{Q}} C_{\mathbb{Q}}(G_{\epsilon,n}) = \dim_{\mathbb{Q}} C_{\mathbb{Q}}(G'_{\text{red},\epsilon,n})$. Note that since the degree of $\mathcal{G}$ is given by d , we can bound the number of vertices of G_n from above by $d|V(G'_n)|$. This allows us to calculate:

$$\begin{aligned} e(\mathcal{G}'_{\text{red},\epsilon}) &= \liminf_{n \rightarrow \infty} \frac{|E(G'_{\text{red},\epsilon,n})|}{|V(G'_{\text{red},\epsilon,n})|} \\ &= \liminf_{n \rightarrow \infty} \frac{|V(G'_{\text{red},\epsilon,n})| + \dim_{\mathbb{Q}} C_{\mathbb{Q}}(G'_{\text{red},\epsilon,n})}{|V(G'_{\text{red},\epsilon,n})|} \\ &\leq 1 + \liminf_{n \rightarrow \infty} \frac{\dim_{\mathbb{Q}} C_{\mathbb{Q}}(G_{\epsilon,n})}{|V(G_{\epsilon,n})|/d} \\ &\leq 1 + d\epsilon. \end{aligned}$$

This finishes the proof. $\square$

Proposition 3.1.4. *Suppose $f : \mathcal{G}' \rightarrow \mathcal{H}'$ is a coarse equivalence of graph sequences, then there exist graph sequences $\mathcal{G}$ and $\mathcal{H}$ such that*

1. $\mathcal{G}'$ is a subgraph sequence of $\mathcal{G}$ such that any external vertex is terminal;
2. $\mathcal{H}'$ is a subgraph sequence of $\mathcal{H}$ such that any external vertex is terminal and;
3. $\mathcal{G} \simeq \mathcal{H}$.

Proof. Note that by throwing out graphs and reordering the sequences we can assume that restricting f gives an (A, A) -quasi-isometry $f_n : G'_n \rightarrow H'_n$ for every n with uniform A (Lemma 1.3.4). The map f is not necessarily injective, but we know that all vertices in G'_n that have the same image over f must be inside the ball $B(x, A^2)$, where x is one of these vertices. Note that by the bounded degree of the graphs, we can find a bound on the number of elements inside this

ball (say C). We construct a graph sequence $\mathcal{H}$ and a map $f' : \mathcal{G}' \rightarrow \mathcal{H}$ as follows: take $y \in H'_n$. If $|f^{-1}(y)| > 1$, then we add $|f^{-1}(y)| - 1$ points to H'_n and connect these to y . Now we construct a new map f' by bijectively mapping each of these $|f^{-1}(y)| - 1$ points to one of the newly added points in $\mathcal{H}$.




Figure 3.3: Construction of $\mathcal{H}$.

As such we obtain $\mathcal{H}$ (with new degree $d + C$ where d is the original degree of $\mathcal{H}'$) and a coarse equivalence f' with a constant A' such that $f'_n : G'_n \rightarrow H_n$ is an (A', A') -quasi-isometry for every n . Note that the new coarse equivalence $f' : \mathcal{G}' \rightarrow \mathcal{H}$ is injective.

Consider for every $n \in \mathbb{N}$ a partition $\{A_x\}_{x \in f'(V(G'_n))}$ of $V(H_n)$, indexed by the image points of f' in $V(H_n)$, such that

- if y has a unique closest image point x in H_n , then $y \in A_x$ and;
- if y has multiple closest image points $x_1, \dots, x_k$, then there is exactly one $i \in \{1 \dots k\}$ such that $y \in A_{x_i}$.

Now for every $n \in \mathbb{N}$ for every image point $x \in f'(V(G'_n))$ we add $|A_x| - 1$ vertices to G'_n and link these to $f'^{-1}(x)$. This gives us $\mathcal{G}$. We define the extension $f'' : \mathcal{G} \rightarrow \mathcal{H}$ of f' by for every n and for every $x \in f'(V(G'_n))$ injectively mapping the new $|A_x| - 1$ vertices to the vertices in $A_x \setminus \{x\}$. We now verify that f'' is an L -bilipschitz map for some $L > 0$. Note that any vertex in A_x is at a distance at most A' from x .

Take an edge (x, y) in $\mathcal{G}$. If both vertices are internal with respect to the subgraph sequence $\mathcal{G}'$, then the images $f''(x) = f'(x)$ and $f''(y) = f'(y)$ are at most a distance $2A'$ apart. If x is internal and y external, then the distance between their images is at most A' by construction of f'' .

Take an edge $(x, y) \in \mathcal{H}$. If both are image points of f' , then their pre-images are at most a distance $A'(A' + 1)$ apart. Suppose at least one of the two is not an image point of f' . Then there exist z_1 and z_2 , image points of f' such that $x \in A_{z_1}$ and $y \in A_{z_2}$. Note that $d(z_1, z_2)$ is at most $2A'$. Hence their pre-images are at most a distance $3A'^2$ apart. By construction of f'' one sees that $d(f''^{-1}(x), f''^{-1}(y)) \leq 3A'^2 + 2$. Taking L the maximum of all these values finishes the proof. $\square$

We now move on to show that cost 1 is an invariant of coarse equivalence. We do this by constructing coarse equivalences between the graph sequences of the form we described above.

Theorem 3.1.5. *If two graph sequences $\mathcal{G}'$ and $\mathcal{H}'$ are coarsely equivalent and one of them has cost 1, then the other one also has cost 1.*

Proof. Consider the graph sequences $\mathcal{G}$ and $\mathcal{H}$ given by Proposition 3.1.4. Note that $\mathcal{G} \simeq \mathcal{H}$. So $\mathcal{G}$ has cost 1 if and only if $\mathcal{H}$ has cost 1. Using Corollary 3.1.3 we see that $\mathcal{G}'$ has cost 1 if and only if $\mathcal{G}$ has cost 1. The same is true for $\mathcal{H}'$ and $\mathcal{H}$. This finishes the proof. $\square$

3.2 Hyperfiniteness

We show that hyperfiniteness is a coarse invariant. First we show it is preserved by bilipschitz equivalence of graph sequences.

Lemma 3.2.1. *If $\mathcal{G} \simeq \mathcal{G}'$ and $\mathcal{G}$ is hyperfinite, then $\mathcal{G}'$ is as well.*

Proof. We look at the ‘identity map’ $\text{Id} : \mathcal{G} \rightarrow \mathcal{G}'$ which is an L -bilipschitz map for some $L > 0$. Take $\epsilon > 0$. Since $\mathcal{G}$ is hyperfinite there exist for every n sets E_n^ϵ and a uniform constant K as in Definition 1.6.19. For every n the connected components induced by the set E_n^ϵ provide a partition $\{A_i^n\}_{i=1}^{k_n}$ of the set of vertices $V(G_n)$.

The graph sequence $\mathcal{G}'$ is a rewiring of $\mathcal{G}$, hence going from one to the other can be seen as first deleting some edges and afterwards adding some edges. Suppose one starts with a graph sequence and then deletes all edges in some set E_n^{del} . If the former graph sequence is hyperfinite, then the latter will be too with respect to the edge sets $E_n^\epsilon \setminus (E_n^\epsilon \cap E_n^{\text{del}})$ and the constant K .

Suppose on the other hand, we add some edges. Note that we will not just add an edge between any two vertices since $\mathcal{G}'$ is L -bilipschitz to $\mathcal{G}$. Hence we can only add edges between vertices that are at most a distance L apart in $\mathcal{G}$. We can only add an edge (x, y) between two different components (i.e. $x \in A_i^n$ and $y \in A_j^n$, where $i \neq j$) if there is a path in G_n between the vertices x and y of length less than or equal to L . Hence the number of added edges is bounded by the number of such paths.

Within a distance $\leq L$ of our starting vertex x , there must be an edge going to another component. The number of paths is bounded by how many vertices are L -close to edges between components. Our graph sequence $\mathcal{G}$ has bounded degree, let's say d . So given one edge, there are at most $\sum_{i=1}^L d^i$ vertices close to one of its end vertices. This means that in the graph G_n there

are at most $|E_n^\epsilon| \left(\sum_{i=1}^L d^i \right)^2$ paths of length $\leq L$ that start in one component and end in another.

Hence, if we consider the components $\{A_i^n\}_{i=1}^{k_n}$ as subsets of G'_n , then the number of edges between them might increase, but at most by a factor $\left(\sum_{i=1}^L d^i \right)^2$. So by taking $E_n^{\epsilon'}$ equal to the set of edges in between the partitioning sets $\{A_i^n\}_{i=1}^{k_n}$ as a partition of $V(G'_n)$ we see that $\liminf_{n \rightarrow \infty} \frac{|E_n^{\epsilon'}|}{|V(G'_n)|} \leq \epsilon \left(\sum_{i=1}^L d^i \right)^2$. This shows that $\mathcal{G}'$ is also hyperfinite. $\square$

We now show that hyperfiniteness is a coarse invariant.

Lemma 3.2.2. *If $\mathcal{G}'$ and $\mathcal{H}'$ are coarsely equivalent and $\mathcal{G}'$ is hyperfinite, then $\mathcal{H}'$ is as well.*

Proof. Proposition 3.1.4 provides us with equivalent graph sequences $\mathcal{G}$ and $\mathcal{H}$, such that $\mathcal{G}'$ (resp. $\mathcal{H}'$) is a subgraph sequence of $\mathcal{G}$ (resp. $\mathcal{H}$) where all external vertices are terminal. We show that if $\mathcal{G}'$ is hyperfinite, then $\mathcal{G}$ is too. For every $\epsilon > 0$, for every $n \in \mathbb{N}$ there exist E_n^ϵ and a constant K that bounds the size of the induced connected components as in Definition 1.6.19.

Denote the degree of $\mathcal{G}$ by d . If we interpret the sets E_n^ϵ as subsets of $E(G_n)$, then we have that $\lim_{n \rightarrow \infty} \frac{|E_n^\epsilon|}{|V(G_n)|} < \epsilon$. Note that the connected components induced in G_n will look the same as those in G'_n with possibly at most d terminal vertices attached to each vertex. Hence the size of the connected components in G_n can be bounded by $K + dK$.

On the other hand if $\mathcal{G}$ is hyperfinite, we obtain once again, for every $n \in \mathbb{N}$, these sets E_n^ϵ , now as subsets of $E(G_n)$. Denote the degree of $\mathcal{G}$ by d . Now

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{|E_n^\epsilon \cap E(G'_n)|}{|V(G'_n)|} &\leq \lim_{n \rightarrow \infty} \frac{|E_n^\epsilon|}{(1/d)|V(G_n)|} \\ &\leq d\epsilon \end{aligned}$$

and the constant K still bounds the size of the connected components. This together with Lemma 3.2.1 ends the proof. $\square$

3.3 Cost-1 behaves multiplicatively for subgroups

Since box spaces originating from filtrations of residually finite groups are in fact graph sequences, one can consider the cost of such a group with respect to a given filtration. We show that cost-1 is multiplicative with respect to taking finite index subgroups. Here we consider cost with respect to a fixed filtration. It is not known whether the cost originating from a certain group depends on the chosen filtration.

Theorem 3.3.1. *Take Γ a residually finite group and $\Lambda \leq \Gamma$ with $[\Gamma : \Lambda] < \infty$. Consider a Farber sequence $(N_k)_{k \in \mathbb{N}}$ of Γ that is contained in Λ . Define $\mathcal{G} = \{\text{Sch}(\Gamma, N_k, S)\}_{k \in \mathbb{N}}$ and $\mathcal{H} = \{\text{Sch}(\Lambda, N_k, S')\}_{k \in \mathbb{N}}$ then*

$$[\Gamma : \Lambda](c(\mathcal{G}) - 1) = c(\mathcal{H}) - 1.$$

Proof. Since Γ is finitely generated there exist elements such that $\Lambda = \langle h_1, \dots, h_n \rangle$ and $\Gamma = \langle h_1, \dots, h_n, g_1, \dots, g_r \rangle$. We denote $S' = \{h_1, \dots, h_n\}$ and $S = \{h_1, \dots, h_n, g_1, \dots, g_r\}$. Now since $\Lambda \leq \Gamma$ one can take a look at $\text{Sch}(\Gamma, \Lambda, S)$. Note that this graph is connected and thus we can find a path (or equivalently a group element) that takes any coset to any other coset in less than $[\Gamma : \Lambda]$ steps. There are $[\Gamma : \Lambda]$ right-cosets of Λ and we write them as $C_1 \dots C_{[\Gamma : \Lambda]}$ where $C_1 = \Lambda$. Each coset C_j contains an element c_j of the form $\prod_{i=1}^s z_i$ where each $z_i \in S$ and $s \leq [\Gamma : \Lambda]$.

Fix a coset C_j and a generator $h \in \{h_i\}_{i=1}^n$ which fixes the coset C_j (with respect to the right action). So $C_j h c_j^{-1} = C_1$ and specifically $c_j h c_j^{-1} := \tilde{c}_j \in \Lambda$. Hence it can be written in the form $\prod_{l=1}^{t_j} \tilde{z}_l$ with $\tilde{z}_l \in \{h_i\}_{i=1}^n$. Define $R_h^j = t_j + 2[\Gamma : \Lambda]$. Now consider the graph $N_k \backslash \Gamma$ and take an h -labelled edge $(\bar{\omega}, \bar{\omega}h)$ where $\bar{\omega} \in C_j$. Now

$$N_k \omega h c_j^{-1} \tilde{c}_j^{-1} c_j = N_k \omega$$

and

$$\begin{aligned} N_k \omega h c_j^{-1} &= N_k \omega c_j^{-1} \\ &\in N_k \backslash \Lambda. \end{aligned}$$

Hence any such h -labelled edge can be replaced by a path of length at most $[\Gamma : \Lambda]$, followed by a path of length t_j in the subgraph $\text{Sch}(\Lambda, N_k, S')$ (only running through edges with labels in $\{h_i\}_{i=1}^n$) and a path of length at most $[\Gamma : \Lambda]$. Note that this combined path has length at most R_h^j (see Figure 3.4). If on the other hand h acts non-trivially on C_j , then we define $R_h^j = 0$.





Figure 3.4: Sketch of what happens in $N_k \setminus \Gamma$.

We define the graph sequence $\mathcal{G}'$. Starting from $\mathcal{G}$, for all k , delete in $\text{Sch}(\Gamma, N_k, S)$ all edges of the form $(\bar{x}, \bar{x}h)$, where $h \in \{h_i\}_{i=1}^n$ and $x, xh \in C_j$ for some $j \neq 1$. We call the obtained graph G'_k and together they form the graph sequence $\mathcal{G}'$. Clearly $\text{Sch}(\Lambda, N_k, S')$ is a subgraph of $\text{Sch}(\Gamma, N_k, S)$ and since this procedure never deletes an edge of the form (x, y) , with $x, y \in N_k \setminus \Lambda$, one sees that $\text{Sch}(\Lambda, N_k, S')$ is still a subgraph of G'_k . Then by the previous computations it is clear that $\mathcal{G}' \simeq \mathcal{G}$ with constant $R_{\max} := \max_{h \in \{h_i\}_{i=1}^n, 1 \leq j \leq [\Gamma : \Lambda]} R_h^j$.

Now we construct another graph sequence $\mathcal{G}''$ by also deleting some edges that have endpoints in two different cosets. We do this by constructing tree-like structures suspended from the subgraphs $\text{Sch}(\Lambda, N_k, S')$ where we make sure that any vertex not in the subgraph sequence $\mathcal{H}$ is the same distance from $\mathcal{H}$ as before.

The edges to be deleted are chosen as follows. We say an arc is a path in G'_k that starts in one of the cosets $N_k \setminus \Lambda$ and ends in another one, only using edges that pass from elements in one right-coset of Λ to another. Consider the set consisting of all arcs and take a (not necessarily unique) shortest arc in this set. Suppose it contains $2n + 1$ edges (resp. $2n$). We delete the middle edge $(\bar{x}, \bar{y})$ (resp. one of the middle edges). Suppose by doing this another point $\bar{z}$ becomes further away from the subgraph $\text{Sch}(\Lambda, N_k, S')$ than it was before. Clearly the shortest path must contain the deleted edge. Hence there exists a path from either $\bar{x}$ or $\bar{y}$ to $\text{Sch}(\Lambda, N_k, S')$ of length less than or equal to $n - 2$. As can be seen in Figure 3.5 this induces an arc of length strictly less than $2n - 1$ (resp. $2n - 2$), which is a contradiction to the minimality of the chosen arc.

Figure 3.5: Deleting g -edges.

We show that $\mathcal{G}'' \simeq \mathcal{G}'$. Consider an edge $(\bar{x}, \bar{y})$ in G'_k , labelled by $g \in S$, that we have removed. Hence there is a path in G'_k of length less than or equal to $[\Gamma : \Lambda]$ from the subgraph $\text{Sch}(\Lambda, N_k, S')$ to $\bar{x}$ by construction. Say the element $p_x \in \Gamma$ is the element obtained by multiplying the labels of the edges on this path. Similarly there is such a path from $\bar{y}$ to $\text{Sch}(\Lambda, N_k, S')$ and p_y is the corresponding element in Γ . Hence $p_x g p_y \in \Lambda$ and can thus be written by an element of the form $\prod_{i=1}^s z_i$, where each $z_i \in \{h_i\}_{i=1}^n$. So we see that this edge can be replaced by a path of length at most $s + 2[\Gamma : \Lambda] + 1$. Note that s only depends on the labels of these paths and thus the elements $p_x g p_y$, which are bounded in length by $2[\Gamma : \Lambda] + 1$. Since there are a finite number of group elements of length at most $2[\Gamma : \Lambda] + 1$, we obtain a finite amount of elements $\prod_{i=1}^s z_i$. Hence we can take a maximum and this shows that indeed $\mathcal{G}'' \simeq \mathcal{G}'$.

We make one last change to the graph sequence. By construction of $\mathcal{G}''$ there are no more arcs and thus any vertex that is not in the subgraph $\text{Sch}(\Lambda, N_k, S')$ has a unique closest point in $\text{Sch}(\Lambda, N_k, S')$. We construct $\mathcal{G}'''$ by linking all vertices not in $\text{Sch}(\Lambda, N_k, S')$ to their closest points in $\text{Sch}(\Lambda, N_k, S')$ and delete all other edges that have these vertices as their end points. Clearly $\mathcal{G}'' \simeq \mathcal{G}'''$.

Note that $\mathcal{G}'''$ has $\mathcal{H}$ as a subgraph sequence where any external vertex is terminal. Hence Lemma 3.1.2 applies. Thus for every $\epsilon > 0$ we get a graph sequence $\mathcal{G}'''_\epsilon$ such that $e(\mathcal{G}'''_\epsilon) \leq c(\mathcal{G}''') + \epsilon$ and $\mathcal{G}'''_\epsilon \simeq \mathcal{G}'''$ and the only changes made to $\mathcal{G}'''$ are inside the subgraph sequence $\mathcal{H}$. We construct $\mathcal{H}_\epsilon$ by taking the rewiring inside $\mathcal{G}'''_\epsilon$ of the subgraph sequence $\mathcal{H}$. One easily sees that $\mathcal{H} \simeq \mathcal{H}_\epsilon$. We claim that the infimum of the edge numbers of the graph sequences $\mathcal{H}_\epsilon$ equals the cost of $\mathcal{H}$. Indeed, otherwise we can take a graph sequence $\mathcal{H}' \simeq \mathcal{H}$ such that $e(\mathcal{H}') < \inf_{\epsilon > 0} e(\mathcal{H}_\epsilon)$. Then construct $\tilde{\mathcal{G}}'''$ by rewiring inside $\mathcal{G}'''$ the subgraph sequence $\mathcal{H}$ as $\mathcal{H}'$. Now $\tilde{\mathcal{G}}''' \simeq \mathcal{G}'''$, but by construction

$$\begin{aligned} e(\tilde{\mathcal{G}}''') &< \inf_{\epsilon > 0} e(\mathcal{G}'''_\epsilon) \\ &\leq \inf_{\epsilon > 0} (c(\mathcal{G}''') + \epsilon) \\ &= c(\mathcal{G}''') \end{aligned}$$

– which is a contradiction. We conclude that $\inf_{\epsilon > 0} e(\mathcal{H}_\epsilon)$ realises the cost of $\mathcal{H}$. Moreover the

following calculations give a relation between $e(\mathcal{G}_\epsilon''')$ and $e(\mathcal{H}_\epsilon)$:

$$\begin{aligned} e(\mathcal{G}_\epsilon''') &= \liminf_{k \rightarrow \infty} \frac{|E(G_{\epsilon,k}''')|}{|V(G_{\epsilon,k}''')|} \\ &= \liminf_{k \rightarrow \infty} \frac{|E(H_{\epsilon,k})| + ([\Gamma : \Lambda] - 1)|V(H_{\epsilon,k})|}{[\Gamma : \Lambda]|V(H_{\epsilon,k})|} \\ &= \frac{[\Gamma : \Lambda] - 1}{[\Gamma : \Lambda]} + \frac{e(\mathcal{H}_\epsilon)}{[\Gamma : \Lambda]}. \end{aligned}$$

This implies that

$$[\Gamma : \Lambda](e(\mathcal{G}_\epsilon''') - 1) = e(\mathcal{H}_\epsilon) - 1$$

and taking the infimum on both sides for ϵ we obtain the result

$$[\Gamma : \Lambda](c(\mathcal{G}) - 1) = c(\mathcal{H}) - 1.$$

□

In particular we have the following

Corollary 3.3.2. *Take Γ a residually finite group and $\Lambda \leq \Gamma$ with $[\Gamma : \Lambda] < \infty$. Consider a filtration $(N_k)_{k \in \mathbb{N}}$ of Γ that is contained in Λ , then*

$$[\Gamma : \Lambda](c(\square_{N_k} \Gamma) - 1) = c(\square_{N_k} \Lambda) - 1.$$

3.4 Property A implies hyperfiniteness

We show that property A implies hyperfiniteness for graph sequences. We recall the following definition for property A and borrow Lemma 3.4.2 and Lemma 3.4.3 from Rufus Willett [Wil11].

Definition 3.4.1. Let $\mathcal{G}$ be a graph sequence. Then $\mathcal{G}$ has property A if for all $\epsilon > 0$, there exists $S > 0$ and a function $\xi : \mathcal{G} \rightarrow \ell^1(\mathcal{G})$, denoted $x \mapsto \xi_x$, such that

1. $\|\xi_x\|_1 = 1$ for all $x \in \mathcal{G}$;
2. ξ_x is supported in the ball of radius S around x for all $x \in \mathcal{G}$;
3. $\|\xi_x - \xi_y\|_1 < \epsilon$ for all $(x, y) \in E(\mathcal{G})$.

The function ξ is said to have variation at most ϵ and support S .

Lemma 3.4.2 (Lemma 3.2, [Wil11]). *Let G be a graph with all vertices of degree at most some constant d , and let $\epsilon > 0$. Say $\phi \in \ell^1(G)$ is a finitely supported function of norm one such that*

$$\sum_{(x,y) \in E(G)} |\phi(x) - \phi(y)| < \epsilon.$$

Then there exists a subset F of the support of ϕ such that $|\partial F| < \epsilon|F|$.

Lemma 3.4.3 (Proof of Proposition 2.5, [Wil11]). *Let $\mathcal{G}$ be a graph sequence with degree d and ξ as in Definition 3.4.1 with variation at most ϵ/d and support S . Then there exists an $N \in \mathbb{N}$ such that for $n \geq N$, there are $F_n \subset G_n$ and $x_n \in G_n$ such that $F_n \subset B(x_n, S)$ and $|\partial F_n| < \epsilon|F_n|$.*

The following is taken entirely from [Wil11] where it is part of a larger proof. We reproduce this to improve readability.

Proof. By possibly taking absolute values, we may assume $\xi_x(y) \geq 0$ for all $x, y \in \mathcal{G}$. For large n (i.e. n greater than or equal to some N) one can restrict ξ to G_n and obtain $\xi^n : G_n \rightarrow \ell^1(G_n)$. We have the following inequality:

$$\sum_{(x,y) \in E(G_n)} \|\xi_x - \xi_y\|_1 < d|G_n|\epsilon/d. \quad (3.4.1)$$

On the other hand we have

$$|G_n| = \sum_{x \in G_n} \|\xi_x^n\|_1 = \sum_{x \in G_n} \sum_{z \in G_n} |\xi_x^n(z)|$$

and introducing this into (3.4.1) gives

$$\sum_{z \in G_n} \sum_{(x,y) \in E(G_n)} |\xi_x^n(z) - \xi_y^n(z)| < \sum_{z \in G_n} \epsilon \sum_{x \in G_n} |\xi_x^n(z)|.$$

So there must exist some $z_0 \in G_n$ such that

$$\sum_{(x,y) \in E(G_n)} |\xi_x^n(z_0) - \xi_y^n(z_0)| < \epsilon \sum_{x \in G_n} |\xi_x^n(z_0)|.$$

Hence if we define $\Psi^n(x) = \xi_x^n(z_0)$, we see that $\Psi^n / \|\Psi^n\|_1$ has the necessary properties to apply Lemma 3.4.2. $\square$

Before proving the theorem we introduce a lemma which will also be useful later on.

Lemma 3.4.4. *Consider a graph sequence $\mathcal{G}$ of degree d , a subgraph sequence $\mathcal{G}'$, an $\epsilon > 0$, a constant S and a function $\xi : \mathcal{G}' \rightarrow \ell^1(\mathcal{G})$, such that*

1. $\|\xi_x\|_1 = 1$ for all $x \in \mathcal{G}'$;
2. ξ_x is supported in the ball of radius S around x in $\mathcal{G}$ for all $x \in \mathcal{G}'$;
3. $\|\xi_x - \xi_y\|_1 < \epsilon$ for all $(x, y) \in E(\mathcal{G}')$.

Then there is a function $\xi' : \mathcal{G}' \rightarrow \ell^1(\mathcal{G}')$ of variation at most ϵ as in Definition 3.4.1 where the support of ξ'_x is bounded by $2d^S$ for every $x \in \mathcal{G}'$. In particular $\mathcal{G}'$ has property A.

Proof. We construct $\xi' : \mathcal{G}' \rightarrow \ell^1(\mathcal{G}')$ by doing the following. For $a \in \mathcal{G}$ Consider the set $R_a = \{x \in \mathcal{G}' \mid \xi_x(a) \neq 0\}$. Note that R_a as a subset of $\mathcal{G}'$ can be decomposed into its connected components (i.e. $R_a = \sqcup_{i=1 \dots k_a} A_i^a$). Next one takes elements $\mu_1^a \dots \mu_{k_a}^a \in R_a$, one for each connected component A_i^a .

Define for $x \in \mathcal{G}'$ the set $A_x = \{\mu_i^a \mid \exists a \in \mathcal{G} \text{ and } \exists i \in \{1 \dots k_a\} \text{ such that } x \in A_i^a\}$ and for $x, y \in \mathcal{G}'$ the set $A_x^y = \{a \in \mathcal{G} \mid \exists i \in \{1 \dots k_a\} \text{ such that } y = \mu_i^a \text{ and } x \in A_x\}$. Note that $\text{supp}(\xi_x) = \sqcup_{y \in \mathcal{G}'} A_x^y$. We define the map $\xi' : \mathcal{G}' \rightarrow \ell^1(\mathcal{G}')$ by setting $\xi'_x(y) = \sum_{a \in A_x^y} |\xi_x(a)|$. Now

$$\begin{aligned} \|\xi'_x\|_1 &= \sum_{z \in \mathcal{G}'} |\xi'_x(z)| \\ &= \sum_{z \in \mathcal{G}'} \left| \sum_{a \in A_x^z} |\xi_x(a)| \right| \\ &= \sum_{z \in \mathcal{G}'} |\xi_x(z)| \\ &= \|\xi_x\|_1. \end{aligned}$$

Consider also $(x, y) \in E(\mathcal{G}')$, then

$$\begin{aligned} \|\xi'_x - \xi'_y\|_1 &= \sum_{z \in \mathcal{G}'} |\xi'_x(z) - \xi'_y(z)| \\ &\leq \sum_{z \in \mathcal{G}'} \left| \sum_{a \in A_x^z} |\xi_x(a)| - \sum_{a \in A_y^z} |\xi_y(a)| \right| \\ &\leq \sum_{z \in \mathcal{G}'} \left| \sum_{a \in A_x^z \cap A_y^z} |\xi_x(a) - \xi_y(a)| \right| + \sum_{a \in A_x^z \setminus A_y^z} |\xi_x(a)| + \sum_{a \in A_y^z \setminus A_x^z} |\xi_y(a)| \\ &\leq \sum_{z \in \mathcal{G}'} \sum_{a \in A_x^z \cap A_y^z} |\xi_x(a) - \xi_y(a)| + \sum_{a \in A_x^z \setminus A_y^z} |\xi_x(a) + \xi_y(a)| \\ &\quad + \sum_{a \in A_y^z \setminus A_x^z} |\xi_x(a) - \xi_y(a)| \end{aligned} \tag{3.4.2}$$

$$\begin{aligned} &\leq \sum_{z \in \mathcal{G}'} \sum_{a \in A_x^z \cup A_y^z} |\xi_x(a) - \xi_y(a)| \\ &= \|\xi_x - \xi_y\|_1. \end{aligned} \tag{3.4.3}$$

One can see inequality (3.4.2) holds by the following argument: if $\xi_x(a)$ and $\xi_y(a)$ are both non-zero, then x and y are in the same connected component of R_a (say A_i^a) and thus our construction of ξ' moves the values $\xi_x(a)$ and $\xi_y(a)$ from a to the same point μ_i^a . Hence the added terms are zero. This also shows that in (3.4.3) the union $\bigcup_{z \in \mathcal{G}'} (A_x^z \cup A_y^z)$ is in fact a disjoint union equal to $\text{supp}(\xi_x) \cup \text{supp}(\xi_y)$.




Figure 3.6: Construction of R_a for a vertex that is in $\mathcal{G}$, but not in $\mathcal{G}'$.

Let us consider the support of an arbitrary ξ'_x . Suppose $\xi'_x(z) \neq 0$, then z is equal to μ_i^a for some $a \in \mathcal{G}$ and some $i \in \{1 \dots k_a\}$. But by construction both x as μ_i^a are inside A_i^a , which is connected in $\mathcal{G}'$. Moreover Since A_i^a is a subset of R_a , which is inside $B_S(\mathcal{G}, a)$, we have that the distance between x and z is at most $\max_{y \in \mathcal{G}} |B_S(\mathcal{G}, y)|$. This value can be bounded by $2d^S$, so this finishes the proof. $\square$

We are now ready to show that any graph sequence that has property A is also hyperfinite.

Theorem 3.4.5. *A graph sequence $\mathcal{G}$ that has property A is also hyperfinite.*

Proof. Fix $\epsilon > 0$. Since $\mathcal{G}$ has property A we can also fix some $\xi : \mathcal{G} \rightarrow \ell^1(\mathcal{G})$ with variation at most ϵ/d and support S as in Definition 3.4.1. We say a good partition of order r is a set of non-empty sets $\{A_i^n\}_{i=1 \dots k_n}^{n \in \mathbb{N}}$ such that

1. For every $n \in \mathbb{N}$ the sets $\{A_i^n\}_{i=1 \dots k_n}$ either form a partition of $V(G_n)$ or form a partition of a subset of $V(G_n)$ with $k_n \geq r$.
2. If we define $E(n) = \{(x, y) \in E(G_n) \mid \exists i \in \{1 \dots k_n\} \text{ such that } x \in A_i^n \text{ and } y \notin A_i^n\}$, then for every $n \in \mathbb{N}$, we have the inequality $|E(n)| < \epsilon |\bigcup_i A_i^n|$.
3. Any set in $\{A_i^n\}_{i=1 \dots k_n}$ has at most d^{2d^S} elements.

Clearly the empty partition is a good partition of order zero. We show that if we have a good partition of order r , there is also a good partition of order $r + 1$. Consider the graph sequence $\mathcal{G}'$ consisting of the graphs G'_n which are the vertex induced subgraphs of G_n on the vertex set $V(G_n) \setminus \bigcup_i A_i^n$. Now restrict ξ to $\mathcal{G}'$ and apply Lemma 3.4.4 to obtain ξ' of variation ϵ/d and support $2d^S$. To this we may apply Lemma 3.4.3, to obtain new sets F_n with $|\partial F_n| < \epsilon |F_n|$ and $|F_n| < d^{2d^S}$. These sets F_n , together with the existing partition, now give a partition of order $r + 1$.

Hence if we fix n , then there exists a good partition of order $|G_n|$. And restricting the partition to a partition of G_n we can find a partition of the set of vertices such that the number of edges between the partitioning sets is smaller than $\epsilon |G_n|$. This shows that $\mathcal{G}$ is indeed hyperfinite. $\square$

3.5 Elek's theorem for Farber sequences

Elek shows in [Ele07] that amenability of a residually finite group is equivalent to hyperfiniteness of box spaces originating from filtrations of the group. Roe shows in [Roe03] that this is

equivalent to those box spaces having property A. It is believed that what holds for filtrations also holds for Farber sequences and their associated graph sequences of Schreier graphs. This motivates us to prove Elek's theorem for Farber sequences.

A filtration approximates a given group by taking a sequence of normal subgroups with trivial intersection, more specifically it provides a Benjamini-Schramm convergent sequence of Cayley graphs. For a Farber sequence we allow any finite index subgroup and ask that the associated graph sequence of Schreier graphs is Benjamini-Schramm convergent. We use this characterisation later on.

Lemma 3.5.1. *If a finitely generated group Γ is amenable, then any graph sequence consisting of Schreier graphs of this group has property A.*

Proof. Fix $\epsilon > 0$. Since Γ is amenable there exists a Følner set $F \subset \Gamma$ such that $|\partial F| < \epsilon|F|$, where ∂F denotes the boundary of F .

Now we will show that we can find $\bar{\xi} : \Lambda \backslash \Gamma \rightarrow \ell^1(\Lambda \backslash \Gamma)$ that has the properties as in 3.4.1 for our previously chosen variation ϵ . We define $\bar{\xi}$ by

$$\bar{x} \rightarrow \left(\bar{y} \mapsto \sum_{z \in \bar{y} \cap xF} \frac{1}{|F|} \right).$$

Here z is obviously in Γ . We check this map is well-defined (i.e. that the image of $\bar{x}$ does not depend on the chosen representative). Suppose $\bar{x} = \bar{x}'$. Hence there exists an element $h \in \Lambda$ such that $hx = x'$. One can calculate the following: $|\bar{y} \cap x'F| = |h\bar{y} \cap hx'F| = |h(\bar{y} \cap xF)| = |\bar{y} \cap xF|$. So clearly $\bar{\xi}_{\bar{x}} = \bar{\xi}_{\bar{x}'}$. Also observe that $\|\bar{\xi}_{\bar{x}}\| = 1$.

Let us take a look at $\|\bar{\xi}_{\bar{x}} - \bar{\xi}_{\bar{y}}\|$, where $\bar{x} = \bar{y}g$ with g some generator of Γ . Note, that here we are looking at points that are linked by an edge in the Schreier graph of $\Lambda \backslash \Gamma$. There exist representatives x of $\bar{x}$ and y of $\bar{y}$ such that $x = yg$ in Γ . Hence

$$\begin{aligned} \|\bar{\xi}_{\bar{x}} - \bar{\xi}_{\bar{y}}\| &= \frac{1}{|F|} \sum_{\Lambda z \in \Lambda \backslash \Gamma} ||\Lambda z \cap xF| - |\Lambda z \cap yF|| \\ &\leq \frac{1}{|F|} \sum_{\Lambda z \in \Lambda \backslash \Gamma} |(xF \Delta yF) \cap \Lambda z| \\ &= \frac{|xF \Delta yF|}{|F|} \end{aligned} \tag{3.5.4}$$

$$\leq \epsilon. \tag{3.5.5}$$

Here equality (3.5.4) follows from the Λz forming a partition of Γ and inequality (3.5.5) follows from the choice of x and y (i.e. being connected by an edge in the Cayley graph of Γ). This finishes the proof. $\square$

We show that hyperfiniteness of a sofic approximation implies amenability of the group.

Theorem 3.5.2. *Take a finitely generated group Γ and a sofic approximation $\mathcal{G} = \{G_n\}_{n \in \mathbb{N}}$. If $\mathcal{G}$ is hyperfinite, then Γ is amenable.*

Proof. We show that hyperfiniteness of a sofic approximation implies amenability for Γ . Now since $\mathcal{G}$ is hyperfinite there exists an $R > 0$ and for every n a partition $\{A_i^n\}_{i=1 \dots k_n}$, where $|A_i^n| < R$.

Moreover if E_n^ϵ is the collection of edges between partitioning sets in G_n , then we have that $\liminf_{n \rightarrow \infty} \frac{|E_n^\epsilon|}{|V_n|} < \epsilon$. From here on we take such a partition with respect to $\epsilon/4$. Now since $\mathcal{G}$ is a sofic approximation (and thus the graph sequence is equivalently Benjamini-Schramm convergent) the number of vertices in G_n for which the $2R$ -ball centered at this vertex is different from the $2R$ -ball in the Cayley-graph of Γ will proportionally go to zero. Take V_n^{2R} the vertices in G_n where the $2R$ -ball is not isomorphic to the ball in the Cayley graph. So we can now choose $N > 0$ such that $\frac{|V_N^{2R}|}{|V_N|} < 1/2$.

We say a vertex is nice if the ball $B_R(G_n, x)$ is isomorphic to the ball $B_R(\text{Cay}(\Gamma), e)$. We say a partitioning set A_i^N is particularly nice if all its vertices are nice. This is interesting because then the set can be lifted to the Cayley graph of Γ . Note that if A_i^N contains one point of $V_N \setminus V_N^{2R}$, then all its points are nice. Hence the union of all particularly nice partitioning sets contains at least $|V_N \setminus V_N^{2R}|$ vertices. We can calculate the following

$$\frac{|E_N^{\epsilon/4}|}{|V_N \setminus V_N^{2R}|} \leq \frac{|E_N^{\epsilon/4}|}{|V_N|(1 - 1/2)} < \frac{\epsilon}{4} \cdot 2 = \frac{\epsilon}{2}.$$

Hence there exists some particularly nice A_i^N that has boundary smaller than ϵ times its number of points. As it is particularly nice this set is situated in the part of the graph for which the R -balls are isomorphic to the R -balls of the Cayley graph of Γ and thus this set can be lifted to Γ , defining a Følner set with respect to ϵ . If we do this for all ϵ , we obtain a Følner sequence for Γ . This concludes the proof. $\square$

Now the hard work is done and our main theorem is an immediate corollary.

Theorem 3.5.3. *Consider a finitely generated group Γ and a Farber sequence $(N_k)_{k \in \mathbb{N}}$. If $\mathcal{G}$ is the associated graph sequence of Schreier graphs then the following are equivalent:*

1. Γ is amenable
2. $\mathcal{G}$ has property A
3. $\mathcal{G}$ is hyperfinite.

Proof. We prove $1 \Rightarrow 2$, $2 \Rightarrow 3$, $3 \Rightarrow 1$. The three implications are particular cases of respectively 3.5.1, 3.4.5 and 3.5.2. $\square$

Our results are optimal in the following sense. Alekseev and Finn-Sell show in [AFS19] that if a sofic approximation has property A, then the group must be amenable. They do not expect the converse to hold in full generality and ask what the suitable notions for the converse statement should be. We first give an easy counterexample, using expanders. After this we answer their question in Theorem 3.5.5.

It is a well-known fact that expanders do not have property A (see for example Proposition 7.2.5 in [Wil09]). To answer our question, consider the following scenario. Take an amenable residually finite group Γ and a box space $\square_{N_k} \Gamma$ of this group with respect to a generating set S . This Benjamini-Schramm converges to the Cayley graph of Γ . Now take an expander sequence $\{H_n\}_{n \in \mathbb{N}}$ and fix a labeling in S on each H_n . Call this graph sequence $\mathcal{H}$. For each H_n , there exists a graph $\text{Cay}(N_k \setminus \Gamma, \bar{S})$ in the box space of $\square_{N_k} \Gamma$ such that $\frac{|V(H_n)|}{|V(N_k \setminus \Gamma)|} < \frac{1}{n}$. Now we choose one point of H_n and one point of $\text{Cay}(N_k \setminus \Gamma, \bar{S})$ and identify them, obtaining the graph G'_n . Now we construct the following graph sequence $\mathcal{G}' = \{G'_n\}$. One sees that the expander sequence

that we added to the box space is negligible with respect to the box space. As such this still Benjamini-Schramm converges and still forms a sofic approximation for Γ . However this no longer has property A since it contains an expander. Hence there is no hope to show Theorem 3.5.3 for arbitrary sofic approximations.

Another question asked is whether hyperfiniteness implies property A for sofic approximations. The previous example also gives a negative answer in this case. For given ϵ , for each G'_n we define a partition of the vertices as follows: Theorem 2.1.2 provides a suitable partition for $\text{Cay}(N_k \setminus \Gamma, \bar{S})$, which we extend to a partition of the vertices of G'_n , by adding the vertices of H_n as singletons. Since $\mathcal{H}$ has bounded degree this shows that $\mathcal{G}'$ is hyperfinite. So also here we have found the limits of the implication. In Theorem 3.5.2 we showed that hyperfiniteness of the sofic approximation implies amenability of the group. The other implication also holds. This is due to Elek in [Ele07]. The original theorem only mentions that this is true for box spaces, however the exact same proof can be used to obtain the same result for sofic approximations.

We now introduce a new concept property almost-A in order to 'fix' the theorem.

Definition 3.5.4. Let $\mathcal{G}$ be a graph sequence, then $\mathcal{G}$ has property almost-A if for all $n \in \mathbb{N}$, there exists sets $V'_n \subset V(G_n)$ such that:

1. $\lim_{n \rightarrow \infty} \frac{|V'_n|}{|V(G_n)|} = 0$;
2. $\mathcal{G}' := \{G_n \setminus V'_n\}_{n \in \mathbb{N}}$ has property A.

The following theorem saves the equivalences for sofic approximations and in doing so, it gives an equivalent definition for amenability.

Theorem 3.5.5. Consider a group Γ and a sofic approximation $\mathcal{G}$ of that group, then the following are equivalent:

1. Γ is amenable;
2. $\mathcal{G}$ has property almost-A;
3. $\mathcal{G}$ is hyperfinite.

Proof. Firstly if $\mathcal{G}$ has property almost-A, then the graph sequence $\mathcal{G}'$ in Definition 3.5.4 is still a sofic approximation of the group Γ . Hence by the theorem of Alekseev and Finn-Sell (see Theorem 1.1, [AFS19]), Γ is amenable.

On the other hand for a sofic approximation $\mathcal{G}$ of an amenable group Γ we define $\phi : \mathbb{N} \rightarrow \mathbb{N}$ by setting $\phi(s)$ to be the smallest natural number such that for all $n \geq \phi(s)$ the inequality

$$p_s(G_n, B_s(\text{Cay}(\Gamma), e)) \geq 1 - \frac{1}{s}$$

holds. This function is well-defined by Benjamini-Schramm convergence of $\mathcal{G}$ and is increasing. We define the following for $\phi(s) \leq n < \phi(s+1)$

$$V'_n = \{x \in V(G_n) \mid B_s(G_n, x) \neq B_s(\text{Cay}(\Gamma), e)\}.$$

We show $\mathcal{G}$ has property almost-A with respect to the sets V'_n and thus the induced graph sequence $\mathcal{G}' := \{G_n \setminus V'_n\}_{n \in \mathbb{N}}$ must have property A. Fix $\epsilon > 0$. Since Γ is amenable we have a Følner set F such that $|\partial F| < \epsilon|F|$. We define $\xi : \mathcal{G}' \rightarrow \ell^1(\mathcal{G})$ for $x \in G'_n$ where $n \geq \phi(|F|)$ by

$$x \mapsto \left(y \mapsto \begin{cases} \frac{1}{|F|} & \text{if } y \in xF \\ 0 & \text{otherwise} \end{cases} \right),$$

and otherwise for $x \in G'_n$ and $n < \phi(|F|)$

$$x \mapsto \left(y \mapsto \begin{cases} \frac{1}{|V(G'_n)|} & \text{if } y \in G'_n \\ 0 & \text{otherwise} \end{cases} \right).$$

Note that the expression xF is well-defined by the isomorphism $B_{|F|}(G_n, x) \simeq B_{|F|}(\Gamma, e)$. We can extract from ξ some $\xi' : \mathcal{G}' \rightarrow \ell^1(\mathcal{G}')$ by applying Lemma 3.4.4. This ends the proof. $\square$

We finish with this proposition.

Proposition 3.5.6. *A sofic approximation $\mathcal{G}'$ of an amenable group Γ consisting only of Følner sets has property A.*

Proof. Since G'_n is some Følner set, we can embed it into $\text{Cay}(\Gamma)$. Fix $\epsilon > 0$ and take $N \in \mathbb{N}$, such that for every $n \geq N$, $|\partial G'_n| < \epsilon|G'_n|$ as subgraphs of $\text{Cay}(\Gamma)$. Now define G_n to be G'_n to which we add its $|G'_n|$ -boundary in $\text{Cay}(\Gamma)$. Then define $\xi : \mathcal{G}' \rightarrow \ell^1(\mathcal{G})$ by

$$x \mapsto \left(y \mapsto \begin{cases} \frac{1}{|F|} & \text{if } y \in xF \\ 0 & \text{otherwise} \end{cases} \right),$$

and otherwise for $x \in G_n$ and $n < N$

$$x \mapsto \left(y \mapsto \begin{cases} \frac{1}{|V(G_n)|} & \text{if } y \in G_n \\ 0 & \text{otherwise} \end{cases} \right).$$

Lemma 3.4.4 now provides a $\xi : \mathcal{G}' \rightarrow \ell^1(\mathcal{G}')$ with the properties as in Definition 3.4.1. $\square$

4 On the first uniformly finite homology group for graphs

This chapter is joint work with Rémi Bottinelli [BK20]. We aim to describe the first uniformly finite homology group for graphs and find three phenomena that are responsible for its non-vanishing: graph-theoretical ends, large circuits and non-expansion. After providing a more practical description of all actors in play in Section 4.1, we proceed in four parts.

In Section 4.2, we study the first of the above phenomena. Starting with a classification of the (first) homology of trees (Theorem 4.2.6) in terms of their ends, we then show that the homology of an arbitrary graph is bounded from below by that of an embedded tree that “realises the ends” (Proposition 4.2.12).

In Section 4.3, motivated by the observation that there is a surjection:

$$H_1^{\text{uf}}(\cdot, \mathbb{Z}) \rightarrow H_1^{\text{uf}}(\cdot, \mathbb{Z}_2),$$

we study $H_1^{\text{uf}}(\cdot, \mathbb{Z}_2)$ (note that the “uniformly finite” condition becomes vacuous in $\mathbb{Z}_2$ — this may sound like a bad omen, but turns out to clarify a lot of constructions). In particular, we get, in Proposition 4.3.2, a full description of $H_1^{\text{uf}}(\cdot, \mathbb{Z}_2)$ as a direct sum of two factors corresponding to exactly the first and second phenomena, that is, ends and “large circuits” respectively.

In Section 4.4, again in $\mathbb{Z}$, we make precise the claim that “large circuits” imply non-zero homology. Since “large circuits” is a notion which depends on the coefficient rings, this result is *not* a consequence of the previous section.

In Section 4.5, inspired by a construction of [BS97], we define a notion of “expansion” in higher dimension, which aims at capturing the third phenomenon. It is easily verified (Lemma 4.5.3) that our “expansion” implies vanishing of homology. As a (very partial) converse, we show (Theorem 4.5.6) that in dimension one, and for vertex transitive graphs, vanishing implies “expansion”.

Finally, refining the notion of expansion, we reach, for (infinite) transitive graphs (or more generally, infinite graphs with a cocompact action by their automorphism group), a characterisation of the non-vanishing of homology in terms of the three phenomena of (the refined) non-expansion, ends, and large circuits.

4.1 “Hands-on” definition of H_1^{uf}

The definition of uniformly finite homology is rather abstract; the goal in this section is to find a more concrete description of H_1^{uf} .

Let X be a ULF graph, and A either $\mathbb{Z}$ or $\mathbb{Z}_2$.

An element of $Z_1^{\text{uf}}(X, A)$ can be seen as a (uniformly bounded) *flow* on X ; that is, each (directed) edge takes a certain value, which we view as flowing through the edge. The condition of being a homological cycle states that at any vertex, the sum of flows of edges directed to the vertex is the same as the flow of edges directed out of the vertex (“in=out”). The two building blocks for such flows are circuits and bi-infinite paths: Indeed, any oriented, graph-theoretical

circuit in X defines an element of $Z_1^{\text{su}}(X, A)$ by taking the sum of its (oriented) edges. Similarly, a bi-infinite path (that is, a morphism of graphs from $\mathbb{Z}$ to X , injective on edges for convenience) also defines an element of $Z_1^{\text{su}}(X, A)$. Now, any flow $f \in Z_1^{\text{su}}(X, A)$ can be decomposed into a sum of circuits and bi-infinite paths, in such a way that there is a uniform bound on the number of circuits/bi-infinite paths passing over any given edge. Indeed, say that a circuit or bi-infinite path g is *below* f if for all edges $e \in EX$:

$$\begin{aligned} f(e) \geq 0 &\Rightarrow f(e) \geq g(e) \geq 0 \\ f(e) \leq 0 &\Rightarrow f(e) \leq g(e) \leq 0. \end{aligned}$$

It is easy to see that it is always possible to find a bi-infinite path or circuit below a flow f . Hence if one takes a maximal multiset of circuits below f and subtracts them, then takes a maximal multiset of bi-infinite paths below the result and subtracts it, nothing can remain. A flow f can thus be decomposed into a well-behaved sum of circuits and bi-infinite paths.

Since we will be working with infinite sums of flows, it is important to know when they are well-defined:

Definition 4.1.1 (Circuits and thinness). Let $\mathfrak{C}(X)$ define the set of oriented circuits in X , and $\mathfrak{C}_r(X)$ the set of oriented circuits of length at most r ($r > 0$). Let $\mathfrak{e} : \mathfrak{C}(X) \rightarrow Z_1^{\text{su}}(X, A)$ the function assigning to any oriented circuit its decomposition as a sum of edges (with orientation and weight; $\mathfrak{e}$ stands for weighted edges).

A set C of circuits is *locally finite* if, for each $e \in EX$, the set $\{c \in C : e \in \text{supp } \mathfrak{e}(C)\}$ is finite.

A function $\lambda : \mathfrak{C}(X) \rightarrow A$ is *thin* if its support is locally finite, and there is some constant K such that, for each $e \in EX$

$$\sum_{c \in \mathfrak{C}(X) : e \in \text{supp } \mathfrak{e}(c)} i(e : \mathfrak{e}(c)) \lambda(c) \leq K,$$

where $i(e : \mathfrak{e}(c))$ is 1 if $\mathfrak{e}(c)(e) > 0$ and -1 otherwise. In this case λ defines an element of $Z_1^{\text{su}}(X, A)$ by extending $\mathfrak{e}$ linearly:

$$\mathfrak{e}(\lambda) := \sum_{c \in \mathfrak{C}(X)} \lambda(c) \mathfrak{e}(c).$$

Informally, a set of circuits is locally finite if any edge has only a finite number of circuits in this set lying on it. A function $\lambda : \mathfrak{C}(X) \rightarrow A$ is interpreted as an assignment of coefficients to circuits; it is thin when summing the circuits with these coefficients defines an element of $Z_1^{\text{su}}(X, A)$.

In particular, if $\lambda : \mathfrak{C}(X) \rightarrow A$ is thin, then $\text{supp } \lambda$ is locally finite. Furthermore, any element of $\lambda \in l^\infty(\mathfrak{C}_r)$, is necessarily thin, since, by ULF, there is some uniform bound on the number of circuits of length $\leq r$ sharing any given edge.

See [BG11] for some motivation on “thinness”, but note however that we have somewhat diverged from their definitions!

Definition 4.1.2. For X a ULF graph and A either $\mathbb{Z}_2$ or $\mathbb{Z}$ set

$$\begin{aligned} C_r(X, A) &:= \left\{ \sum_{c \in \mathfrak{C}_r} \lambda(c) e(c) \mid \lambda : \mathfrak{C}_r \rightarrow A \text{ is thin} \right\} \\ C_\infty(X, A) &:= \cup_r C_r(X, A) \\ C(X, A) &:= \left\{ \sum_{c \in \mathfrak{C}} \lambda(c) e(c) \mid \lambda : \mathfrak{C} \rightarrow A \text{ is thin} \right\} \end{aligned}$$

It is easy to see that all of those objects are actually vector subspaces of $Z_1^{\text{uf}}(X, A)$. Note that in the case $A = \mathbb{Z}_2$, C_r can be described simply as the space of sums of elements of $\mathfrak{C}_r$, and C as the space of locally finite sums of elements of $\mathfrak{C}$. Furthermore, by our remarks:

$$C_r(X, A) = e(l^\infty(\mathfrak{C}_r)).$$

Also, we have the inclusions:

$$C_1(X, A) \subseteq C_2(X, A) \subseteq \dots \subseteq C_\infty(X, A) \subseteq C(X, A).$$

We say that X has *A-large circuits* if there does not exist some r such that $C_r(X, A) = C_\infty(X, A)$.

Proposition 4.1.3. *Having A-large circuits is a quasi-isometry invariant.*

Proof. Let $g : X \rightarrow Y$ be a quasi-isometry between two graphs, of which the first has A -large circuits and the other does not (say $C_R(Y) = C_\infty(Y)$). Assume for now that g is actually bijective on vertices. For each choice of $x_1, x_2 \in VX$ (resp. VY), let p_{x_1, x_2} be a choice of shortest path between x_1 and x_2 in X (resp. Y). Since X has A -large circuits, for each $r \in \mathbb{N}$ there is some flow $f \in C_\infty(X) - C_r(X)$. Define the image by g to be $\sum_{(x_1, x_2) \in \text{supp}(f)} p_{g(x_1), g(x_2)}$. Since Y has A -large circuits, this can be rewritten as a sum of cycles of length $\leq R$, say $\sum_i c_i$. Since g is bijective, we can also define the inverse image $\sum_i \sum_{(y_1, y_2) \in \text{supp}(c_i)} p_{g^{-1}(y_1), g^{-1}(y_2)}$, which is necessarily equal to f . Since the cycles $\sum_{(y_1, y_2) \in \text{supp}(c_i)} p_{g^{-1}(y_1), g^{-1}(y_2)}$ are of length at most $R \cdot L$, where L is the quasi-isometry constant of f^{-1} , this finishes the proof when g is bijective.

If g is not bijective, X and Y can be extended by connecting leaves to their vertices, yielding, say, X' and Y' , and g can then be modified to get a bijective quasi-isometry $g' : X' \rightarrow Y'$ (See [Kai19, Proof of Proposition 2.4] for more details). Since X' is obtained from X by adding leaves, it is easy to see that $C_r(X') = C_r(X)$ for all r and $C(X') = C(X)$ (and similarly for Y and Y'). In particular, X has A -large circuits iff X' does, and by the argument above X has A -large circuits iff Y does. $\square$

We start with a description of $C_\infty(X, A)$ in terms of our chain complexes.

Lemma 4.1.4. *If X is a ULF graph, then, for any $r > 0$:*

$$\begin{aligned} C_r(X, A) &\subseteq Z_1^{\text{uf}}(X, A) \cap B_1^{\text{uf}}(R_{r/2+1}(X), A), \\ Z_1^{\text{uf}}(X, A) \cap B_1^{\text{uf}}(R_r(X), A) &\subseteq C_{3r}(X, A). \end{aligned}$$

Consequently:

$$C_\infty(X, A) = Z_1^{\text{uf}}(X, A) \cap B_1^{\text{uf}}(X, A).$$

Proof. If $f \in C_r(X)$ then $f = e\phi$ for some $\phi : \mathfrak{C}_r \rightarrow A$ uniformly bounded. Obviously we have $f \in Z_1^{\text{uf}}(X)$. For any $c \in \text{supp } \phi$, one can “triangulate” c : that is, if c is the circuit $(v_1, \dots, v_s)$ with $s \leq r$, we can consider the sum

$$\Delta c := \sum_{i=1}^s (v_1, v_i, v_{i+1}) \in C_2^{\text{uf}}(R_{2r+1}(X))$$

(with addition in the indices modulo s). One sees that $\partial(\Delta c) = ec$, and letting

$$\Delta \phi := \sum_c \phi(c) \cdot (\Delta c),$$

we get $\Delta \phi \in C_2^{\text{uf}}(R_{r/2+1}(X))$ by ULF, and $\partial(\Delta \phi) = e\phi = f$. Thus, $f \in B_2^{\text{uf}}(R_{r/2+1}(X))$.

Conversely, assume $f \in Z_1^{\text{uf}}(X) \cap B_1^{\text{uf}}(R_r(X))$. Let $g \in C_2^{\text{uf}}(R_r(X))$ such that $\partial g = f$.

For any two vertices u, v at distance $\leq r$, fix a shortest path $p_{u,v} = (u = u_0, u_1, \dots, u_s = v)$ from u to v . If $t := (u, v, w)$ is a triangle in $R_r(X)$, consider the circuit Ot obtained by concatenating the paths $p_{u,v}$, $p_{v,w}$ and $p_{w,u}$. By construction, $Ot \in \mathfrak{C}_{3r}$. Let also

$$Og := \sum_{t \in (R_r(X))_{(3)}} g(t) \cdot (Ot).$$

Then $Og \in l^\infty(\mathfrak{C}_{3r})$, again by ULF, and by our assumption that $\partial g \in Z_1^{\text{uf}}(X)$, we get

$$e(Og) = \partial g.$$

It follows that $f = e(Og) \in C_{3r}$.

The last equality of the lemma follows simply by taking unions. □

We can now describe $H_1^{\text{uf}}(X)$ in a way that only deals with circuits and paths, and does not deal with Rips complexes:

Proposition 4.1.5. *If X is a ULF graph, then, for any $r > 0$, the composite homomorphism*

$$\Phi_r : Z_1^{\text{uf}}(X, A) \hookrightarrow Z_1^{\text{uf}}(R_r(X), A) \twoheadrightarrow H_1^{\text{uf}}(R_r(X), A)$$

is surjective and its kernel satisfies

$$C_{2r} \leq \ker \Phi_r \leq C_{3r}.$$

Consequently, the composite homomorphism

$$\Phi_\infty : Z_1^{\text{uf}}(X, A) \hookrightarrow Z_1^{\text{uf}}(X, A) \twoheadrightarrow H_1^{\text{uf}}(X, A)$$

is surjective and has kernel

$$\ker \Phi_\infty = C_\infty(X, A).$$

In particular:

$$H_1^{\text{uf}}(X, A) \cong \frac{Z_1^{\text{uf}}(X, A)}{C_\infty(X, A)}.$$

Proof. Fix $r > 0$. We first verify surjectivity of Φ_r . Fix any $f \in Z_1^{\text{uf}}(R_r(X))$. Choose for any two u, v in X , a unique shortest path $p_{u,v} = (u = u_0, u_1, \dots, u_s = v)$ from u to v , and let

$$\tilde{f} := \sum_{(u,v) \in (R_r(X))_{(2)}} f(u,v) \cdot p_{u,v},$$

that is, $\tilde{f}$ is obtained by replacing “virtual edges” (i.e. edges in $(R_r(X))_{(2)}$) by paths. By ULF, $\tilde{f} \in C_1^{\text{uf}}(X)$ and $\partial \tilde{f} = \partial f$, so that $\partial \tilde{f} = 0$, hence $\tilde{f} \in Z_1^{\text{uf}}(X)$. We claim that $\tilde{f} - f \in B_1^{\text{uf}}(R_r(X))$, which will clearly imply surjectivity of Φ_r . Indeed, if for any (u, v) at distance $\leq r$, one considers

$$\Delta(u, v) := \sum_{i=1}^{s-1} (u_0, u_i, u_{i+1}) \in C_2^{\text{uf}}(R_r(X)),$$

where $(u = u_0, u_1, \dots, u_s = v)$ is a path from u to v , then $\partial \Delta(u, v) = (u, v) - p_{u,v}$. Summing over the virtual edges (u, v) in the support of f , we get

$$\partial \left(\sum_{(u,v) \in (R_r(X))_{(1)}} f(u, v) \cdot \Delta(u, v) \right) = f - \tilde{f}.$$

This shows that, Φ_r is surjective, assuming the sum in the parentheses is in $C_2^{\text{uf}}(R_r(X))$. This holds as usual by ULF. Now, the inclusions related to $\ker \Phi_r$ remain, but these follow at once from the first part of Lemma 4.1.4.

The properties of Φ_∞ follow from the second part of Lemma 4.1.4 and the fact that

$$H_1^{\text{uf}}(X) = \frac{\bigcup_r Z_1^{\text{uf}}(R_r(X))}{\bigcup_r B_1^{\text{uf}}(R_r(X))}.$$

□

For later use, we note also that:

Corollary 4.1.6. *Let X a ULF graph and assume that $C_r(X, A) = C_\infty(X, A)$ for some $r > 0$. Then, for any s large enough, the map Φ_s descends to an isomorphism*

$$\tilde{\Phi}_s : Z_1^{\text{uf}}(X, A)/C_r(X, A) \rightarrow H_1^{\text{uf}}(R_s(X), A),$$

and so does Φ_∞ :

$$\tilde{\Phi}_\infty : Z_1^{\text{uf}}(X, A)/C_r(X, A) \rightarrow H_1^{\text{uf}}(X, A).$$

Furthermore, the following diagram (of isomorphisms) commutes:

$$\begin{array}{ccc} & \frac{Z_1^{\text{uf}}(X, A)}{C_r(X, A)} & \\ \swarrow \tilde{\Phi}_s & & \searrow \tilde{\Phi}_\infty \\ H_1^{\text{uf}}(R_s(X), A) & \xrightarrow{[i]} & H_1^{\text{uf}}(X, A). \end{array}$$

Proof. Consider the maps Φ_s, Φ_∞ of Proposition 4.1.5. Since for all $s \geq r$, $C_\infty(X) = C_s(X) = C_r(X)$, the equations involving the respective kernels give us the desired isomorphisms. Com-




Figure 4.1: Every ULF tree is quasi-isomorphic to a trivalent one through a simple rewiring at each vertex of high valency.

mutativity of the triangles follows, since all maps involved are (appropriate quotients of) inclusions. $\square$

4.2 Trees and End-Defining Trees

Expanding a result of [Dia15], we describe the first homology of any (ULF) tree T as $\ell^\infty(B)$, for some set B of bi-infinite paths in T , of cardinality the number of ends of T minus 1 (in the finite case).

We will then show that the first homology of a sufficiently well-behaved subtree of a graph X injects into the first homology of X .

From now on, by a trivalent tree, we mean a tree with vertices of degree either 2 or 3. We first verify that in terms of uniformly finite homology, leaves (that is, vertices of degree 1) can be safely forgotten.

Lemma 4.2.1. *Fix $A \in \{\mathbb{Z}, \mathbb{Z}_2\}$ and T' the maximal leafless (i.e. with no vertex of degree 1) subtree of T . Then*

$$H_1^{\text{uf}}(T, A) = H_1^{\text{uf}}(T', A).$$

Proof. Since T has no non-trivial circuits, we have that $H_1^{\text{uf}}(T, A) = Z_1^{\text{uf}}(T, A)$ (Proposition 4.1.5). Recall that an element $c \in Z_1^{\text{uf}}(T, A)$ is a flow on T , and since leaves correspond to dead-end for the flow, the edges adjacent to leaves must have zero flow. Arguing inductively implies then that c must have support in T' , and the equality is achieved. $\square$

To reduce to trivalent trees, it is now enough to note that any ULF tree is quasi-isometric to a trivalent one. Indeed, suppose such a tree T is given. A trivalent “unfolding” of T can be obtained by replacing each star around a vertex as in Figure 4.1. The uniform bound on degrees ensures that the result, $\tilde{T}$, is quasi-isometric to T . Since H_1^{uf} is QI-invariant, it is therefore sufficient to study trivalent trees.

4.2.1 Homology of Trees

We want a concrete description of the first uniformly finite homology of trees. Recall that we essentially have two kinds of flows on a graph X : circuits and bi-infinite paths. When X is a tree, non-trivial circuits are never reduced, thus define zero flow. The “building blocks” for a

flow on X are therefore bi-infinite paths. We will extract from X a specific set P of bi-infinite paths which suffices to fully describe flows on X . The construction of P will be explicit but only some abstract “tameness” properties of P will actually be used.

In order to ease notation we consider the following construction of a trivalent tree (with no leaves).

- Start with a single vertex v_0 , and call that graph T_0 .
- Let T_1 be a tree isomorphic to $\mathbb{Z}$, containing T_0 as $v_0 = 0$. T_1 can be viewed as two copies of $\mathbb{N}$ glued on v_0 .

And now inductively:

- In T_n , choose a set B_n of vertices that does not intersect T_{n-1} ; call its elements *branching points* of depth n . To each $v \in B_n$, glue a copy called R_v of $\mathbb{N}$ at zero (a new branch). For $i \in \mathbb{N}$, write i_v for the vertex i of $R_v \simeq \mathbb{N}$, and ∞_v for its end (this is not a vertex). Let T_{n+1} be the result. Formally:

$$T_{n+1} = (T_n \sqcup \bigsqcup_{v \in B_n} R_v) / \{0_v \sim v \mid v \in B_n\}.$$

Obviously, $T_n \leq T_{n+1}$.

Let us call each R_v the *ray* at v and call n its depth if it is in $T_{n+1} - T_n$. The construction $(T_n, B_n)_{n \in \mathbb{N}}$ has depth N if it ends at step T_N , i.e. if for $k > N$ we have $B_k = \emptyset$. The construction has infinite depth otherwise: $T = \bigcup_n T_n$.

Note that for each n , the set B_n of branching points of depth n can be partitioned as:

$$B_n = \bigsqcup_{v \in B_{n-1}} B_n \cap (R_v - \{v\}),$$


since each new branching point has to lie on $T_n - T_{n-1} = \bigsqcup_{v \in B_{n-1}} R_v - \{v\}$. Let $B_{n,v} := B_n \cap R_v - \{v\} = B_n \cap R_v$ denote the branching points lying on R_v . There is an obvious ordering on $B_{n,v}$ as a subset of $\mathbb{N}$: write $v_{v,i}$ the i -th element for this ordering (where $i < \alpha$, and $\alpha \in \mathbb{N} \cup \{\infty\}$ is the cardinal of $B_{n,v}$).

Proposition 4.2.2. *All (infinite) trivalent leafless trees are constructed in this fashion, modulo one edge.*

Proof. Start with a vertex $v_0 \in T$ of degree two (if such a vertex does not exist, split an edge to obtain one). Follow any two paths starting from v_0 : this defines T_1 . Let $B_1 := \{v \in T_1 : \deg(v) > 2\}$, and from each $v \in B_1$, follow any new ray to get T_2 . Proceed inductively like that, and let $T_\omega := \bigcup_n T_n$. It remains to see that $T_\omega = T$, but this follows from the fact that if $v \in VT$ is at distance d from v , then necessarily $v \in T_d$. $\square$

4.2.2 Construction of a set of bi-infinite paths

We now describe how to construct a set P of bi-infinite paths inductively, by following the construction of T . Fix T a trivalent tree, and consider a sequence $(T_n, B_n)_{n \in \mathbb{N}}$ that matches the construction above. We construct P by induction: P_{n+1} extends P_n and is a set of bi-infinite paths for T_{n+1} .



(a) Decomposition of a tree into sets T_n . The black vertex is T_0 , and cumulatively adding the blue, red, and orange segments correspond respectively to T_1, T_2 , and T_3 .

(b) Construction of the set of bi-infinite paths associated to said tree. The blue one makes up P_1 , and cumulatively adding the red and orange ones yield P_2 and P_3 respectively.

Figure 4.2: Example of the decomposition of a tree and the the associated sets of bi-infinite paths.

T_0 : No bi-infinite path, i.e. $P_0 = \emptyset$.

T_1 : P_1 contains the unique bi-infinite path in T , i.e. $P_1 = \{T_1\}$.

T_{n+1} : Assume now P_i has been constructed for $i \leq n$.

We will define $P_{n+1} = P_n \sqcup P'_{n+1}$, with P'_{n+1} the set of *new* bi-infinite paths.

First fix $u \in B_{n-1}$, and consider the ray R_u in T_n . We recall that the branching points of B_{n+1} can be partitioned by the rays on which they lie. Likewise, we define new bi-infinite paths for each ray of depth n :

$$P'_{n+1,u} := \begin{cases} \{R_{v_{u,i}}^{-1} + [v_{u,i}, v_{u,i+1}] + R_{v_{u,i+1}} \mid i \in \mathbb{N}\} & \text{if } |B_{n,u}| = \infty \\ \{(R_{v_{u,i}}^{-1} + [v_{u,i}, v_{u,i+1}]) + R_{v_{u,i+1}} \mid i < k\} \cup \{R_{v_{u,k}}^{-1} + [v_{u,k}, \infty_u]\} & \text{if } |B_{n,u}| = k + 1 \end{cases}$$

where $(v_{u,i})_{i < |B_{n,u}|}$ is the ordering of $B_{n,u}$, and obtain

$$P'_{n+1} := \bigsqcup_{u \in B_{n-1}} P'_{n+1,u}, \quad P_{n+1} := P_n \sqcup P'_{n+1}.$$

If $(T_n, B_n)_{n \in \mathbb{N}}$ has infinite depth, let $P = \bigcup_n P_n$. We then define a (partial) order on P as follows. If $p \in P_n$ and $q \in P_{n+1}$ share an edge, let $p \leq q$, and take the transitive closure. Say that $p \in P$ is the *last* bi-infinite path over an edge e if p lies over e and for any other $q \in P$ lying over e , we have $q < p$. Write

$$M_p := \{e \in ET : p \text{ is last over } e\}.$$

We can now describe some “tameness” properties of P that will be used later on:

Proposition 4.2.3. *Let $(T_n, B_n)_{n \in \mathbb{N}}$ be a tree construction as above, and $(P_n)_{n \in \mathbb{N}}$ the corresponding construction of bi-infinite paths. Let also $\leq$ be the partial order on P described above.*

1. Each edge is covered by at most three bi-infinite paths.
2. All elements of P'_n lie in $T_{n+1} - T_{n-1}^+$ (where T_k^+ is the 1-neighbourhood of T_k).
3. P_n covers T_n .
4. Each $p \in P$ has finitely many predecessors (with respect to $\leq$), and is the last bi-infinite path over at least one edge (i.e. $M_p \neq \emptyset$).
5. For any edge e where $p \in P_{n+1}$ is last over e , there exists $e' \in T_n^+$ such that:

$$\{q \in P : e' \in q\} = \{q \in P : e \in q < p\} = \{q \in P_n : e \in q\} =: Q_p.$$

6. Each set M_p is connected.

Proof. By construction. □

We can actually endow P with a well-ordering (extending $\leq$), which will facilitate the proof in Theorem 4.2.6.

Lemma 4.2.4. *If $(P, \leq)$ is a countable poset such that $\{q \in P : q \leq p\}$ is finite for all p , there exist a well ordering $\preccurlyeq$ on P extending $\leq$ (i.e. $p \leq q \Rightarrow p \preccurlyeq q$) and such that $\{q \in P : q \preccurlyeq p\}$ is still finite for any p .*

Proof. Define

$$P_0 = \left\{ p \in P : \{q \in P : q < p\} = \emptyset \right\}$$

$$P_{n+1} = \left\{ p \in P : \{q \in P : q < p\} \subseteq \bigcup_{j=1}^n P_j \right\}$$

and note that, assuming P is non-empty:

- $P_0 \neq \emptyset$ and $P = \bigcup P_j$, since $\{q \in P : q \leq p\}$ is finite for all p ;
- The elements of P_j are all pairwise incomparable (endow each P_j with any well-ordering);
- Each P_j is countable (since P is).

Assume that P is infinite countable (the finite case is treated similarly). It remains then to construct a bijection $\sigma : \mathbb{N} \rightarrow P$ such that $\sigma n < \sigma m \Rightarrow n < m$.

The algorithm in Algorithm 1 constructs σ . The idea is to “diagonally” (as in the usual bijection $\mathbb{N}^2 \leftrightarrow \mathbb{N}$) exhaust all P_i s, while preserving the existing order.

□

Corollary 4.2.5. *The set of bi-infinite paths P can be assumed to be well-ordered without breaking any of the properties in Proposition 4.2.3. In particular, any edge has a last bi-infinite path on it:*

$$ET = \bigsqcup_p M_p$$

Algorithm 1 Extending partial order to total order.

```

1:  $i \leftarrow 0, j \leftarrow 0, k \leftarrow 0, \sigma \leftarrow \emptyset$ 
2: while  $\exists p \in P_j - \text{im } \sigma$  such that  $\forall q < p : q \in \text{im } \sigma$  do
3:    $\sigma \leftarrow \sigma \cup \{i \mapsto p\}$ 
4:    $i \leftarrow i + 1$ 
5:   if  $j \leq k$  then
6:      $j \leftarrow j + 1$ 
7:   else
8:      $j \leftarrow 0, k \leftarrow k + 1$ 

```

We prove that if T is a tree (with uniformly bounded vertex degree) and P is as constructed above, it can be used to describe $H_1^{\text{uf}}(X)$. We invite the reader to have a quick peak at the theorem. Before proceeding with the statement, let us first motivate the properties of P stated in Proposition 4.2.3. The imposed order and the finiteness condition help in making an inductive argument. The non-emptiness of the sets M_p guarantees that the chosen bi-infinite paths are independent, that is, there is no bounded sum that is zero in homology. Once we have independence of bi-infinite paths, the connectedness of M_p will guarantee generation of all homology classes. The reader might have the feeling that we are constructing some kind of infinite basis — this would be correct, as the theorem states. Finally the two remaining conditions just make sure that we only have to consider uniformly bounded sums of bi-infinite paths.

Theorem 4.2.6. *Given a trivalent tree T , a set P of bi-infinite paths on T as constructed above and $A \in \{\mathbb{Z}, \mathbb{Z}_2\}$, the following map:*

$$l^\infty(P) \rightarrow H_1^{\text{uf}}(T, A) \cong H_1^{\text{uf}}(T, A)$$

$$f \mapsto c_f := [e \mapsto \sum_{p \in P : e \in p} f(p)]$$

is an isomorphism of modules (continuous when $A = \mathbb{Z}$ and $H_1^{\text{uf}}(T, A)$ has the l^∞ norm).

Proof. By Corollary 4.2.5, we can assume that P is well-ordered.

The map above is easily seen to be linear, and of norm ≤ 3 (since $|\{q \in P : e \in q\}| \leq 3$ for all $e \in E$).

Injectivity: By induction. If $f \neq 0 \in l^\infty(P)$, there exists a bi-infinite path p such that $f(p) \neq 0$; take a least such p . Consider then $e \in M_p$:

$$c_f(e) = \sum_{q \in P} f(q) = \sum_{q < p} f(q) + f(p)$$

Since by assumption, for each $e \in q < p$, $f(q) = 0$, it follows that $c_f(e) \neq 0$.

Surjectivity: Take a cycle $c \in Z_1^{\text{uf}}(T, A)$. Then, define inductively for $p \in P$:

$$f(p) := c(e) - \sum_{e \in q < p} f(q),$$

where $e \in M_p$. We check that this is independent of the choice of $e \in M_p$. By connectedness of

M_p , it is enough to verify that adjacent edges e_1, e_2 in M_p yield the same value. Locally, there are only two possible configurations as in Figure 4.3.




Figure 4.3: Local configurations of bi-infinite paths.

Horizontally one can see the bi-infinite path p , for which we are defining the value $f(p)$. The green (G), red (R) and blue (B) lines represent groups of bi-infinite paths running over these vertices (possibly empty). Note that all these bi-infinite paths are in order strictly smaller than p (because p is last over e_1, e_2). Hence f is defined for all of them. We check that $f(p)$ does not depend on whether we choose e_1 or e_2 . There are two possible situations, as depicted above. For the first picture the claim clearly holds, we check the second picture:

$$\begin{aligned}
 c(e_1) - \sum_{q \in G \text{ or } q \in R} f(q) &= c(e_2) + c(e_3) - \sum_{q \in G \text{ or } q \in R} f(q) \\
 &= c(e_2) + \sum_{q \in R} f(q) - \sum_{q \in B} f(q) - \sum_{q \in G \text{ or } q \in R} f(q) \\
 &= c(e_2) - \sum_{q \in G \text{ or } q \in B} f(q).
 \end{aligned}$$

The second equality holds by the induction hypothesis.

Now f is uniformly bounded. Indeed, from

$$f(p) := c(e) - \sum_{e \in q < p} f(q),$$

we see that

$$|f(p)| \leq \|c\|_\infty + \left\| \sum_{e \in q < p} f(q) \right\| \leq 2\|c\|_\infty,$$

since the elements of $\{q \in P : e \in q < p\}$ are assumed to share an edge (so that the sum over $\{q \in P : e \in q < p\}$ is exactly the value $c(e')$ of some edge e'). It is now easily seen that $c_f = c$, and since $\|f\|_\infty \leq 2\|c\|_\infty$, the desired isomorphism is reached. $\square$

As an example of necessity of the “tameness” conditions, we give Figure 4.4; the set of bi-infinite paths chosen satisfy each condition but the fifth one — the existence of an edge e' covered by exactly the predecessors of a given bi-infinite path p . As a result the flow defined on the teeth of the comb cannot be described as a bounded sum of the given bi-infinite paths.

4.2.3 Ends

Now that the homology of trees is understood, we can relate it to the homology of arbitrary graphs via their ends. The core idea is that, when working with ends of graphs, working with nicely embedded trees is enough.

We start with a definition:




Figure 4.4: A bi-infinite comb with bi-infinite paths as shown in green and red. The cycle we consider is defined by alternating positive and negative values on the teeth of the comb. The reader can check that the obtained function on the bi-infinite paths is unbounded.

Definition 4.2.7 (End-respecting (defining) trees). Let X be a ULF graph. An end-respecting (resp. end-defining) tree for X is a leafless tree $T \leq X$ such that any end of X is represented by at most one (resp. a unique) infinite branch of T . Furthermore, T is said to *have separators* if there exists a sequence $(K_n)_n$ of increasing finite subsets of VX satisfying:

- The vertex-induced subgraph of K_n in X is connected for all n ;
- The vertex-induced subgraph of $T \cap K_n$ in T is connected for all n ;
- $\bigcup_n K_n = VX$;
- For each infinite connected component L of $X - K_n$, there is exactly one edge of $\partial^e(T \cap K_n) \cap ET$. This defines a bijection between infinite connected components of $X - K_n$ and $\partial^e(T \cap K_n) \cap ET$.

Remark 4.2.8. End-defining trees with separators can always be constructed. Indeed, fix X and some $x \in VX$. Consider the smallest $R > 0$ such that removing $B(x, R)$ from X disconnects the graph in multiple infinite components. Now choose as many edges in the boundary of the ball as there are infinite components, and a geodesic for each of these starting in x . Combining these geodesics gives a subgraph, and deleting edges as necessary yields a tree. Repeat this approach inductively: suppose we have a finite tree defined up to $B(x, R)$, with boundary edges l_i in correspondence to infinite connected components L_i (of $X - B(x, R)$). Take $\tilde{R} > R$ least such that $X - B(x, \tilde{R})$ has more components than $X - B(x, R)$. If a connected component L_i of $X - B(x, R)$ is further separated in components L_i^j in $X - B(x, \tilde{R})$, where $j \in \{1 \dots k\}$, with distinguished boundary edges l_i^j for each component in the boundary of $B(x, \tilde{R})$, then as before take geodesics from l_i to each l_i^j and delete possible cycles.

On the other hand, it is not clear to us whether one can always find separators for a given end-respecting tree.

Definition 4.2.9. Given $\phi, \psi : EX \rightarrow \mathbb{Z}$, say that $\psi \leq \phi$ if, for any $e \in EX$: $\phi e \geq 0$ implies $0 \leq \psi e \leq \phi e$ and $\phi e \leq 0$ implies $\phi e \leq \psi e \leq 0$.

The following technical result is close to [BW92, Lemma 2.4].

Lemma 4.2.10. Let $\phi \in Z_1^{\text{suf}}(X, \mathbb{Z})$ and $r : EX \rightarrow \mathbb{Z}$ with $r \leq \phi$, $r \notin Z_1^{\text{suf}}(X, \mathbb{Z})$, and $\partial r \geq 0$. Then there exists $\hat{r} \in Z_1^{\text{suf}}(X, \mathbb{Z})$ satisfying $r \leq \hat{r} \leq \phi$. Furthermore, $\hat{r}$ and r are different on infinitely many edges.

Here is an intuitive interpretation: view ϕ as closed (in the sense that in=out) flow on the graph X . Then r is a non-closed flow that is no greater than ϕ : any flow on the edges is in the same direction as ϕ , and not stronger. Then r can be made into a closed flow by simply extending it in the directions allowed by ϕ (so, no reduction of what r had and not exceeding what ϕ allows).

Proof. We extend r step by step, and see that the sequence of extensions converges:

1. For any v with $\partial r(v) > 0$, we have:

$$\sum_{u \sim v} \phi(u, v) = 0,$$

so, there must exist $e_v = (u, v)$ with $\phi(e_v) < r_v(e) \leq 0$.

2. Consider now $r^1 := r + \sum_{v \in V_0} e_v$, where V_0 is the set of v s with $\partial r(v) > 0$. We notice that this sum is well-defined since no e_v appears in two distinct v s, and $r \leq r^1 \leq \phi$ by construction. Also, $\partial r^1 \geq 0$ since if $v \in V_0$, then v stays (non-strictly) positive in ∂r^1 , and the other v s can only increase, by the direction of the added edges.
3. Redo the argument with r replaced by r^1 , and construct inductively a sequence $r = r^0 \leq r^1 \leq \dots \leq \phi$.
4. Since the sequence $(r^i)_i$ is monotonous and below ϕ , it must converge to some $r \leq \hat{r} \leq \phi$, in $Z_1^{\text{uf}}(X, \mathbb{Z})$, by boundedness of ϕ .
5. The statement that r and $\hat{r}$ differ on an infinite set can be deduced from the fact that the process of extension $r^i \rightsquigarrow r^{i+1}$ cannot end after finitely many steps. Indeed, if $v \in V_i$, and $(u, v) = e_v$ is the edge chosen to “cancel” v , then u becomes positive in ∂r^{i+1} , so that there are always positive edges.

□

A similar result holds for $\mathbb{Z}_2$. In $EX \rightarrow \mathbb{Z}_2$, say $r \leq \phi$ iff $\phi(e) = 0$ implies $r(e) = 0$ for all $e \in EX$.

Lemma 4.2.11. *Let X a (ULF) graph, and $\phi \in Z_1^{\text{uf}}(X, \mathbb{Z}_2)$, $r \in C_1^{\text{uf}}(X, \mathbb{Z}_2)$ with $r \leq \phi$, and $\partial r = \delta_v$ for some vertex v . Then, there exists $r \leq \hat{r} \leq \phi$ with $\hat{r} \in Z_1^{\text{uf}}(X, \mathbb{Z}_2)$ and such that $\hat{r}$ differs from r on infinitely many edges.*

Proof. By mimicking the $\mathbb{Z}$ case. Since $\sum_{e \sim v} \phi(e) = 0 \neq \sum_{e \sim v} r(e)$, there exists e with $\phi(e) = 1, r(e) = 0$. Thus, define $r^1 := r + e$ and continue by induction. □

We can now prove the following.

Proposition 4.2.12 (End-respecting trees inject in H_1^{uf}). *Let X be ULF and T an end-respecting tree with separators. Then, the inclusion $T \subseteq X$ induces an embedding of homology groups:*

$$H_1^{\text{uf}}(T, \mathbb{Z}) \leq H_1^{\text{uf}}(X, \mathbb{Z}).$$

The same holds for $\mathbb{Z}_2$.

Proof. We first consider $\mathbb{Z}$. Let $(K_i)_i$ the sequence of separators associated to T .

1. First, recall that $H_1^{\text{uf}}(X, \mathbb{Z}) \cong Z_1^{\text{uf}}(X, \mathbb{Z})/C_\infty$, so, in particular, $H_1^{\text{uf}}(T, \mathbb{Z}) \cong Z_1^{\text{uf}}(T, \mathbb{Z})$, T being a tree. From now on, drop the $\mathbb{Z}$. We consider the function:

$$Z_1^{\text{uf}}(T) \xrightarrow{\subseteq} Z_1^{\text{uf}}(X) \xrightarrow{\not\subseteq} Z_1^{\text{uf}}(X)/C_\infty$$

so that showing injectivity amounts to showing that no non-zero function in $Z_1^{\text{uf}}(T)$ can be written as an element of C_∞ . We actually show more: the intersection of $Z_1^{\text{uf}}(T)$ and C is trivial.

2. Assume, towards a contradiction, that there exists $0 \neq f \in Z_1^{\text{uf}}(T) \cap C$. In particular there exists a thin set of circuits C such that:

$$f = \sum C.$$

(Formally, one should add integer coefficients to the sum, but we will stick to the shorthand).

3. Since T is end-respecting, there exists some $K = K_n$ and an edge $e \in \text{supp } f$ so that K disconnects the space in multiple infinite components, one of which – call it L – is such that $\text{supp } f \cap \partial^e L$, consists of exactly one edge e .
4. Up to taking $-f$, we may assume that e is directed “away from L ”, and that f is positive on e .
5. Following f backwards starting from e , we get a set of $k := f(e)$ rays $r_1, \dots, r_k$ supported in $T \cap (L \cup \{e\})$ and all ending with e .
6. Let C_L the set of circuits in C that lie in L and do not intersect K .

Let r be the sum of initial segments of each r_i , so that r has support not intersecting the circuits of C not in C_L (only a finite number of them has to be taken into account since C is thin and K finite). Note that $r \leq f$ by construction.

Let f_L be the restriction of $f - (\sum r_i)$ to L .

7. Remark that $f_L \in Z_1^{\text{uf}}(L)$, $r \leq C_L - f_L$ and, far enough from K , we have:

$$r = C_L - f_L$$

8. By Lemma 4.2.10, there exists $\hat{r} \in Z_1^{\text{uf}}(L)$ with $r \leq \hat{r} \leq C_L - f_L$, $\hat{r}$ distinct from r at infinitely many edges.
9. Then $\hat{r}$ must be distinct from r away from K . Since $r \leq \hat{r} \leq C_L - f_L$ and, away from K , $r = C_L - f_L$, we reach a contradiction.

The case of $\mathbb{Z}_2$ is a simple adaptation. We follow the same procedure up to step 5. There, we can just follow *one* ray r_1 along f . Then, at step 8, we use Lemma 4.2.11 instead of Lemma 4.2.10 to extend r to some bi-infinite path $\hat{r}$, with $\hat{r} \leq \sum C_L - f_L$, and differing from r arbitrarily far from K , from which we again deduce a contradiction.

□

We extract the core property used in the proof for future use:

Corollary 4.2.13 (of the proof). *With the assumptions of Proposition 4.2.12:*

$$Z_1^{\text{uf}}(T) \cap C(X) = 0.$$

Corollary 4.2.14. *The dimension of $H_1^{\text{uf}}(X, \mathbb{Z})$ is at least the number of ends of X minus 1.*

Note that the number of ends has already been seen to appear in terms of cohomological dimension; see e.g. [SW79, Proposition 5.1] and [Pes90, Theorem 9], in which the number of ends is characterised as the dimension of the zeroth cohomology “modulo finite support” with coefficients in $\mathbb{Z}_2$.

Recall that any finite subset in a graph induces a partition of its ends given by the connected component in the complement of the set in which they lie.

Lemma 4.2.15. *Let T_1, T_2 be two end-defining trees in the ULF graph X . For any edge e in T_1 , there exists a compact $K \subset X$, such that the partition $\mathcal{E}$ induced by K on the ends of T_2 is a subpartition of the one induced by e on T_1 . In particular, the shadow topology induced on the end spaces $\text{ends}(X) = \text{ends}(T_1) = \text{ends}(T_2)$ by T_1 is the same as the topology induced by T_2 .*

Proof. Write $\{E^+, E^-\}$ for the partition of $\text{ends}(X)$ induced by e in T_1 . Fix a vertex x_1 on T_1 and x_2 on T_2 . Let B_i be the ball of radius i and center x_1 in X . The partition of $\text{ends}(T_2)$ induced by B_i is indirectly given by the connected components of $T_2 - B_i$. Suppose it is not a subpartition of $\{E^+, E^-\}$, then there are ends $\varepsilon_i, \varepsilon'_i$ such that they are in the same connected component of $T_2 - B_i$, but $\varepsilon_i \in E^-$ and $\varepsilon'_i \in E^+$. Now take rays r_i, r'_i starting at x_1 (resp. $\tilde{r}_i, \tilde{r}'_i$ starting at x_2) in T_1 (resp. T_2), such that their limits are the ends ε_i and ε'_i . Note that since $\tilde{r}_i$ and $\tilde{r}'_i$ are in the same connected component, they overlap on their first vertices inside B_i . This means that the sequences $\{\tilde{r}_i\}_i$ and $\{\tilde{r}'_i\}_i$ have the same pointwise limit. On the other hand consider r_i and $\tilde{r}_i$. Note that any end ε of X corresponds to a choice of connected components of $X - B_i$. Note that since $r_i(\infty) = \tilde{r}_i(\infty)$, they are eventually in the same connected component of $X - B_i$. For each i there are only a finite amount of connected components of $X - B_i$. Hence we can take a subsequence of $\{r_i\}_i$ which for each i is eventually in the same connected component of $X - B_i$. Note that the subsequence of $\tilde{r}_i$ defined on the same indices is eventually in the same connected components. Note that if we restrict to the subsequences, they will converge pointwise. This means that the two limits of these sequences must have the same end. Then, because of previous arguments, the end of $(\lim r_i)(\infty) = (\lim \tilde{r}_i)(\infty)$, which is equal to $(\lim \tilde{r}'_i)(\infty)$ and thus equal to $(\lim r'_i)(\infty)$. However this is a contradiction, since the ends of r_i are in E^- and those of r'_i are in E^+ . Since both kinds of infinite rays are separated by the edge e , also the ends of the limit rays are in these respective sets. $\square$

The result in Proposition 4.2.12 can actually be generalised to end-respecting trees *without* separators.

Proposition 4.2.16. *Let X be ULF and $T \leq X$ an end-respecting tree. Then:*

$$H_1^{\text{uf}}(T, \mathbb{Z}) \leq H_1^{\text{uf}}(X, \mathbb{Z}).$$

The same holds for $\mathbb{Z}_2$.

Proof. Let $T \leq X$ be such a tree, and $\tilde{T}$ an end-defining tree for X with separators.

- We show that $H_1^{\text{suif}}(T, \mathbb{Z}) \cap C(X, \mathbb{Z}) = \{0\}$, and then the argument in the first step of the proof of Proposition 4.2.12 gives the desired embedding.
- Towards a contradiction, assume that $f \in H_1^{\text{suif}}(T, \mathbb{Z}) \cap C$ is non-zero. By the decomposition of $H_1^{\text{suif}}(T, \mathbb{Z})$ into $l^\infty(P)$ for a set of bi-infinite paths P , one can write

$$f = \sum_{b \in P} a_b \cdot b.$$

- For any bi-infinite path $b \in P$, one can consider its “parallel” bi-infinite path $\tilde{b}$ in the tree $\tilde{T}$; that is, b has two ends corresponding to ends of the graph X , which correspond also to two ends in $\tilde{T}$ (since $\tilde{T}$ is end-defining), and finally to a unique bi-infinite path. Let $\tilde{P} := \{\tilde{b} \mid b \in P\}$ and

$$\tilde{f} = \sum_{b \in P} a_b \cdot \tilde{b}.$$

Note that by “projecting”, one can find a collection of circuits C_f (and coefficients a_c) such that

$$\tilde{f} = f + \sum_{c \in C_f} a_c \cdot c$$

(the construction for C_f is similar to the one in the proof of Proposition 4.3.2).

- We claim that $\tilde{f}$ is well-defined (but potentially unbounded) and non-zero.
- Assume for now that this is the case. Since $\tilde{f}$ is non-zero, there exists a first separator $K := K_i$ and an edge e lying on $\partial^e K$ such that $\tilde{f}$ is non-zero on e .
- Let C_f^K denote those circuits in C_f that intersect the separator K_{i+1} and set

$$f' := f - \sum_{c \in C_f^K} a_c \cdot c.$$

By construction, f' is equal to $\tilde{f}$ inside of K_{i+1} , so in particular in K_i and on e ; hence $f'(e) \neq 0$.

- Following the argument of Proposition 4.2.12 starting from step 4 with function f' shows that f' cannot be non-zero, a contradiction.

It remains to verify that $\tilde{f}$ is actually well-defined and non-zero. Being well-defined is a consequence of Lemma 4.2.15.

Assume that $\tilde{f} = 0$, and let e be some edge of T on which f is non-zero.

- Let $\{E^+, E^-\}$ be the partition ∂T induced by e . By Lemma 4.2.15 there exists a large enough separator K_i in X such that the induced partition on ends is a subpartition of $\{E^+, E^-\}$

Let then K be such a separator and $\{E_1^+, \dots, E_n^+\} \cup \{E_1^-, \dots, E_m^-\}$ the subpartition induced by K . Let also $e_1^+, \dots, e_n^+, e_1^-, e_m^-$ the edges in the boundary of K , with corresponding signs. The ends in $E_i^\pm$ are then exactly those in the component of $X - K$ corresponding to $e_i^\pm$.

For a collection of bi-infinite paths B , let $a_B := \sum_{b \in B} a_b$. If we note b^{-1} for the bi-infinite path b with opposite orientation, then $a_{b^{-1}} = -a_b$. For a collection of oriented bi-infinite paths B and two disjoint sets of ends E_1, E_2 , let $[E_1, E_2]$ be $\{b \in B \mid b(-\infty) \in E_1 \text{ and } b(+\infty) \in E_2\} \sqcup \{b^{-1} \mid b \in B \text{ and } b(-\infty) \in E_2 \text{ and } b(+\infty) \in E_1\}$. In words, the bi-infinite paths having one end in E_1 and one in E_2 , taken with consistent orientation.

- Since $\tilde{f} = 0$, we have $\tilde{f}(e_i^s) = 0$ for all i and $s = \pm$. This implies

$$a_{[E_i^+, (E_i^+)^c]} = 0.$$

- We also have:

$$f(e) = a_{[E^-, E^+]},$$

and using $[E^-, E^+] = \bigsqcup_{i=1}^m [E_i^-, E^+]$, we get

$$\begin{aligned} &= \sum_{i=1}^m a_{[E_i^-, E^+]} \\ &= \sum_{i=1}^m a_{[E_i^-, E^+]} + 0. \end{aligned}$$

Now, using the fact that $a_{[E_j^-, E_i^+]} = -a_{[E_i^+, E_j^-]}$ for any j , this becomes

$$\begin{aligned} &= \sum_{i=1}^m \left(a_{[E_i^-, E^+]} + \sum_{j=1, j \neq i}^m \overbrace{a_{[E_i^-, E_j^-]} + a_{[E_j^-, E_i^+]}}^{=0} \right) \\ &= \sum_{i=1}^m \left(\left(a_{[E_i^-, E^+]} + \sum_{j=1, j \neq i}^m a_{[E_i^-, E_j^-]} \right) + \sum_{j=1, j \neq i}^m a_{[E_j^-, E_i^-]} \right), \end{aligned}$$

and by the decomposition $[E_i^-, (E_i^-)^c] = [E_i^-, E_i^+] \sqcup \bigsqcup_{j=1, j \neq i}^m [E_i^-, E_j^-]$ for any i :

$$= \sum_{i=1}^m \left(\left(a_{[E_i^-, (E_i^-)^c]} \right) + \sum_{j=1, j \neq i}^m a_{[E_j^-, E_i^-]} \right).$$

The content of the inner parenthesis vanishes since $\tilde{f}(e_i^-) = 0$:

$$\begin{aligned} &= \sum_{i=1}^m \sum_{j=1, j \neq i}^m a_{[E_j^-, E_i^-]} \\ &= \sum_{1 \leq i \neq j \leq m} a_{[E_j^-, E_i^-]} + a_{[E_i^-, E_j^-]}. \end{aligned}$$

And finally, again since inverting directions negates the sign:

$$= 0.$$

It follows therefore that $\tilde{f} = 0$ implies $f(e) = 0$, which is a contradiction.

□

4.3 Coefficients in $\mathbb{Z}_2$

In this section, we study the special case of $H_1^{\text{uf}}(X, \mathbb{Z}_2)$, or “uniformly finite homology with coefficients in $\mathbb{Z}_2$ ”. In that case, the “uniformly finite” condition becomes vacuous, and one may wonder if any insight can be gained with such a restriction. This is indeed the case, as the next proposition shows, thus motivating the study of $H_1^{\text{uf}}(X, \mathbb{Z}_2)$.

4.3.1 $H_1^{\text{uf}}(X, \mathbb{Z}) \twoheadrightarrow H_1^{\text{uf}}(X, \mathbb{Z}_2)$

Proposition 4.3.1. *Let X a (ULF) graph. Then $H_1^{\text{uf}}(X, \mathbb{Z}_2)$ is a quotient of $H_1^{\text{uf}}(X, \mathbb{Z})$.*

Proof. Consider the natural map:

$$\Phi : Z_1^{\text{suf}}(X, \mathbb{Z}) \xrightarrow{q} Z_1^{\text{suf}}(X, \mathbb{Z}_2) \xrightarrow{\pi} Z_1^{\text{suf}}(X, \mathbb{Z}_2) / C_\infty(X, \mathbb{Z}_2).$$

The first arrow is induced by the quotient $\mathbb{Z} \rightarrow \mathbb{Z}_2$. We show it is surjective. Indeed, if $f : EX \rightarrow \mathbb{Z}_2$ lies in $Z_1^{\text{suf}}(X, \mathbb{Z}_2)$, one can decompose f and write

$$f = \sum C + \sum B,$$

for C a set of circuits and B a set of bi-infinite paths, such that the elements of $B \cup C$ have pairwise disjoint support. Choosing signs arbitrarily, we can lift f to $\tilde{f} \in Z_1^{\text{suf}}(X, \mathbb{Z})$ by lifting each element of C and B .

To get a surjection $H_1^{\text{uf}}(X, \mathbb{Z}) \rightarrow H_1^{\text{uf}}(X, \mathbb{Z}_2)$, it remains to check that $C_\infty(X, \mathbb{Z})$ is contained in $\ker \Phi$. But if $\sum C \in C_\infty(X, \mathbb{Z})$ is given, then the image of $\sum C$ through q is in $C_\infty(X, \mathbb{Z}_2)$.

□

4.3.2 Decomposition of $H_1^{\text{uf}}(X, \mathbb{Z}_2)$

The main result of the section is the following:

Proposition 4.3.2. *Let X be a (ULF) graph, and $T \leq X$ an end-defining subtree. Then:*

$$H_1^{\text{uf}}(X, \mathbb{Z}_2) \cong H_1^{\text{uf}}(T, \mathbb{Z}_2) \oplus \frac{C(X, \mathbb{Z}_2)}{C_\infty(X, \mathbb{Z}_2)}.$$

Proof. Drop the $\mathbb{Z}_2$. Recall that $H_1^{\text{uf}}(X) = Z_1^{\text{suf}}(X) / C_\infty$ and $H_1^{\text{uf}}(T) \cong Z_1^{\text{suf}}(T)$. Furthermore, both $Z_1^{\text{suf}}(T)$ and C lie in $Z_1^{\text{suf}}(X)$. Since $C_\infty \leq C$, it suffices therefore to show that $Z_1^{\text{suf}}(T) \cap C = 0$ and $Z_1^{\text{suf}}(T) + C = Z_1^{\text{suf}}(X)$. We already know, by Corollary 4.2.13 that $Z_1^{\text{suf}}(T) \cap C = 0$. It remains therefore to show $Z_1^{\text{suf}}(T) + C = Z_1^{\text{suf}}(X)$.

Fix $f \in Z_1^{\text{suf}}(X)$. We will find a thin set of circuits D such that $f - \sum D$ is supported on T ; in some sense “pushing” f towards T using circuits.

By “following paths”, one can write:

$$f = \sum_{\varepsilon \neq \varepsilon'} \sum_{b \in B_{\varepsilon, \varepsilon'}} b + \sum_{b \in B_{=}} b + \sum_{c \in C} c$$

with $B_{\varepsilon, \varepsilon'}$ a set of bi-infinite paths going from end ε to end ε' , $B_{=}$ a set of bi-infinite paths going to the same ends, and C a set of circuit, and so that no edge is shared between elements of $\bigcup_{\varepsilon \neq \varepsilon'} B_{\varepsilon, \varepsilon'} \cup B_{=} \cup C$ – the sum is therefore (very!) thin.

If $B_{\varepsilon, \varepsilon'}$ is finite, let, for any $b \in B_{\varepsilon, \varepsilon'}$, b^* denote the unique bi-infinite path in T with same ends. If $B_{\varepsilon, \varepsilon'}$ is infinite, let $B_{\varepsilon, \varepsilon'}^+ \sqcup B_{\varepsilon, \varepsilon'}^- = B_{\varepsilon, \varepsilon'}$ be a partition into two infinite parts, and choose a bijection $B_{\varepsilon, \varepsilon'}^+ \rightarrow B_{\varepsilon, \varepsilon'}^-$ denoted as sending $b \in B_{\varepsilon, \varepsilon'}^+$ to $b^* \in B_{\varepsilon, \varepsilon'}^-$. Denote also the inverse of this bijection with $\cdot^*$, so that $(b^*)^* = b$.

Fix some $R_0 > 0$, and let $B_{\varepsilon, \varepsilon'}^0, B_{=}^0, C^0$ be the elements of each set $(B_{\varepsilon, \varepsilon'}, B_{=}, C)$ that intersect $B(v_T, R_0)$. By ULF, each of those sets is finite, and only a finite number of pairs $\varepsilon \neq \varepsilon'$ have $B_{\varepsilon, \varepsilon'}^0$ non-empty. For any $b \in B_{\varepsilon, \varepsilon'}^0$, choose two paths $p_{b,0}^+, p_{b,0}^-$ connecting the two rays of $b - B(v_T, R_0)$ outside of $B(v_T, R_0)$ to b^* . Similarly, for any $b \in B_{=}^0$, find a path $p_{b,0}$ joining the two rays of b outside of $B(v_T, R_0)$, say $p_{b,0}$ connects $b(t_{b,0}^+)$ to $b(t_{b,0}^-)$. Let R_1 large enough that $B(v_T, R_1)$ contains each of the path used above and each element of C^0 , let $B_{\varepsilon, \varepsilon'}^1, B_{=}^1, C^1$ be the elements of each set that intersect $B(v_T, R_1)$ and repeat the above.

This defines circuits, say c_b^i such that:

$$\sum c_b^i = b - b^*$$

if $b \in B_{\varepsilon, \varepsilon'}$, and

$$\sum c_b^i = b,$$

if $b \in B_{=}$. Furthermore, at each step i , the circuits added do not intersect $B(v_T, R_{i-1})$, except for the segments corresponding to b^* . Since those segments appear in exactly one circuit in the case of infinite $B_{\varepsilon, \varepsilon'}$, and in a finite number of circuits in the case of finite $B_{\varepsilon, \varepsilon'}$, the sum of circuits is thin, and we have successfully pushed f to the tree. $\square$

Lemma 4.3.3. *If there exists R such that any circuit can be written as a $\mathbb{Z}_2$ thin-sum of circuits of length $\leq R$, then $C(X, \mathbb{Z}_2) = C_R(X, \mathbb{Z}_2)$.*

Proof. By endowing both $\mathbb{Z}_2^{EX}$ and $\mathbb{Z}_2^{\mathbb{C}}$ with the product topology, both spaces are compact, $\mathbb{Z}_2^{\mathbb{C}_R}$ is a closed subspace of $\mathbb{Z}_2^{\mathbb{C}}$, and the operator $e : \mathbb{Z}_2^{\mathbb{C}} \rightarrow \mathbb{Z}_2^{EX}$ is continuous.

For each $c \in \mathbb{C}$, fix some $\phi_c \in \mathbb{Z}_2^{\mathbb{C}_R}$ with $ec = e\phi_c$.

Fix some $\psi \in \mathbb{Z}_2^{\mathbb{C}}$ locally finite; we have to show that $e\psi = e\phi$ for some $\phi \in \mathbb{Z}_2^{\mathbb{C}_R}$. Let $(c_i \in \mathbb{C})_i$ be an enumeration of the circuits appearing in ψ , then

$$\psi = \lim_n \sum_{i=1}^n c_i.$$

Since e is continuous and linear:

$$\begin{aligned} e\psi &= \lim_n \sum_{i=1}^n e c_i \\ &= \lim_n \sum_{i=1}^n e \phi_{c_i} \\ &= \lim_n e \left(\sum_{i=1}^n \phi_{c_i} \right) \end{aligned}$$

and, taking a subsequence, one may assume that the sequence $\sum_{i=1}^n \phi_{c_i}$ converges to some $\phi \in \mathbb{Z}_2^{\mathbb{C}_R}$, so that

$$\begin{aligned} &= e \lim_n \left(\sum_{i=1}^n \phi_{c_i} \right) \\ &= e\phi. \end{aligned}$$

□

Corollary 4.3.4. *For a finitely presented group G :*

$$H_1^{\text{uf}}(G, \mathbb{Z}_2) \cong H_1^{\text{uf}}(T, \mathbb{Z}_2),$$

which has dimension $\# \text{ends}(G) - 1$.

Proof. By Proposition 4.3.2, the above lemma and the fact that any circuit in the Cayley graph of G can be written as a finite sums of relators. □

4.3.3 Large circuits and dimensions

Our goal in this section is to show that, intuitively, if a graph has circuits of increasing length that cannot be decomposed into smaller ones, then $\mathbb{Z}_2$ -homology cannot vanish, and will in fact be infinite dimensional. Recall our definition of thinness and of the spaces $C_R(X, \mathbb{Z}_2)$, $C_\infty(X, \mathbb{Z}_2)$ and $C(X, \mathbb{Z}_2)$ from Section 4.1.

In $\mathbb{Z}_2$, thinness is equivalent to local finiteness, and orientation is forgotten, so that we can safely simplify our definitions:

A collection C of elements of $\mathbb{Z}_2^{EX}$ is said to be *thin* if it is locally finite. The pointwise sum of the elements of C is then well-defined, and we have:

$$\begin{aligned} C &:= \{\text{“thin sum of circuits”}\}, \\ C_R &:= \{\text{“thin sum of circuits of length } \leq R\}, \\ C_\infty &:= \cup_R C_R. \end{aligned}$$

Fix some $R > 0$. The following are then equivalent:

1. $C = C_R$.
2. $C_\infty = C_R$

3. For all $R' > R$, $C_{R'} = C_R$.
4. Any circuit can be written as a thin-sum sum of circuits of length $\leq R$.

If those do not hold, we say that X has $\mathbb{Z}_2$ -large circuits. It is relatively clear that $1 \Rightarrow 2 \Rightarrow 3 \Rightarrow 4$, and Lemma 4.3.3 takes care of $4 \Rightarrow 1$.

In particular, the following holds.

Proposition 4.3.5. *If X does not have $\mathbb{Z}_2$ -large circuits, then $C(X, \mathbb{Z}_2)/C_\infty(X, \mathbb{Z}_2) = 0$.*

For the sake of clarity, we will settle on the following definitions:

Definition 4.3.6. Fix a *ULF* graph X .

- A thin-sum is the well-defined sum of a thin collection $A \subseteq \mathbb{Z}_2^{EX}$.
- A vector subspace of $\mathbb{Z}_2^{EX}$ is said to be a thin-sum space if it is closed under thin-sums.
- The thin-sum-span of a collection A (not necessarily thin itself) is the set:

$$\{\sum B \mid B \text{ a thin subset of } A\}.$$

- A collection A is thin-sum-independent if no two distinct thin subcollections of A yield the same sum; equivalently, no thin subcollection has sum 0.
- A thin-sum basis is a collection that is thin-sum-independent and a thin-sum basis for a given thin-sum space.

By compactness, the thin-sum span of a collection is closed under thin-sums (see [BG11]). Note that we do not require a thin-sum basis A for a thin-sum space $\mathcal{W}$ to be thin itself, but if it is, we get a bijection

$$\mathbb{Z}_2^A \leftrightarrow \mathcal{W}.$$

With those definitions in mind, we see that each C_R is a thin-sum space, as is C (as, respectively, the thin-sum span of circuits of length $\leq R$, and all circuits), but C_∞ is not necessarily. Finally, all the concepts above still make sense if we replace $\mathbb{Z}_2^{EX}$ by $\mathbb{Z}_2^W$ for any countable set W .

We will show that as soon as X has large circuits, then $H_1^{\text{uf}}(X, \mathbb{Z}_2)$ (more precisely C/C_∞) is infinite-dimensional.

To that end, we first construct a set of circuits in X that is thin *and* a thin-sum basis for C .

Proposition 4.3.7 (C has thin basis). *There exists a set $\mathcal{B}$ of circuits that is thin and a thin-sum basis for C .*

The proof uses an infinite variation of Gaussian elimination, which we will use again later in a transfinite way.

Any element of C is a function $EX \rightarrow \mathbb{Z}_2$. Order the edges EX as $EX = (e_i)_{i \in \mathbb{N}}$. For any non-zero element $f \in C$, there is then a well-defined *leading index* $l(f)$ of f :

$$l(f) := \min\{i \in \mathbb{N} : f(i) \neq 0\}.$$

If $f = 0$, we set $l(f) = \infty$ by convention, and will also set $g(\infty) = 0$ for any $g \in C$.

Proof. Let $\mathcal{S}$ be the set of all simple¹ circuits, viewed as maps $EX \rightarrow \mathbb{Z}_2$, and order $\mathcal{S}$: $\mathcal{S} = (f_j)_{j \in \mathbb{N}}$ ($\mathcal{S}$ is countable). We then apply Gaussian elimination on $\mathcal{S}$ as follows. Define the following function inductively:

$$\begin{aligned} F : \mathbb{N} \times \mathbb{N} &\rightarrow (\mathbb{N} \rightarrow \mathbb{Z}_2) \\ F_{0,j} &= f_j \\ F_{t,j} &= \begin{cases} F_{j,j} & \text{if } j < t \\ F_{t-1,j} + F_{t-1,j}(l(F_{t-1,t-1})) \cdot F_{t-1,t-1} & \text{if } j \geq t \end{cases} \end{aligned}$$

(with index notation used for readability). Note that by our convention on leading index, the definition makes sense even for $F_{t-1,t-1} = 0$. We first remark that for any $t > j$, $F(t, j) = F(j, j)$. Let then $g_j := F(j, j)$, and

$$\mathcal{B} := \{g_j \mid j \in \mathbb{N}, g_j \neq 0\}.$$

Any element of $\mathcal{B}$ is non-zero, and since for $j > i$, $g_j(l(g_i)) = 0$ by construction, it follows that they are all distinct.

It remains to verify that: $\mathcal{B}$ is thin, thin-sum independent and a thin-sum basis.

Thinness. First, note that as soon as $i = l(g_j)$ for some j , then e_i is in the support of finitely many elements of $\mathcal{B}$ by construction.

Assume now there exists a non-leading edge e_i that is in the support of infinitely many elements of $\mathcal{B}$, and without loss of generality that it is actually the first such (i.e. $i' < i$ implies $e_{i'}$ is in the support of finitely many circuits). Let t be such that

$$g_t(i) \neq 0 \text{ and } g_t(i') = 0 \quad \forall i' < i.$$

The existence of such a t is guaranteed by our assumption on i . Then $g_t(i) \neq 0$ but $g_t(i') = 0$ whenever $i' < i$, so that $l(g_t) = i$, which is a contradiction.

Independence. Given any non-empty set of elements $\{g_i\}_{i \in I}$ of $\mathcal{B}$, consider $\hat{i} := \min I$, and let $\hat{j} = l(g_{\hat{i}})$. Then $(\sum_{i \in I} g_i)(\hat{j}) = 1$ since for all $I \ni i \neq \hat{i}$, $g_i(\hat{j}) = 0$. Thus, no non-trivial sum is zero.

Basis. Fix $h : \mathbb{N} \cong EX \rightarrow \mathbb{Z}_2$ any element of $\mathcal{C}$. By definition, h is a thin sum of simple circuits $\sum_{i \in \mathcal{I}} f_i$, where $\mathcal{I} \subset \mathbb{N}$. By definition of F we see that any f_i can be written

$$f_i = F_{i,i} + \sum_{t=1}^i F_{t-1,i}(l(F_{t-1,t-1})) \cdot F_{t-1,t-1}$$

and each summand is either zero or a basis element:

$$= g_i + \sum_{t=1}^i F_{t-1,i}(l(F_{t-1,t-1})) \cdot g_{t-1}.$$

Therefore, f_i can also be seen as a map $\tilde{f}_i : \mathcal{B} \rightarrow \mathbb{Z}_2$, where each basis element is mapped to its

¹Recall that a circuit is said to be *simple* if no vertex appears twice.

coefficient in the sum above. Define now

$$\begin{aligned}\mathcal{I}_s &:= \{i \in \mathcal{I} \mid i \leq s\} \\ h_s &:= \sum_{i \in \mathcal{I}_s} f_i \in (\mathbb{Z}_2)^{EX} \\ \tilde{h}_s &:= \sum_{i \in \mathcal{I}_s} \tilde{f}_i \in (\mathbb{Z}_2)^{\mathcal{B}}.\end{aligned}$$

Note that, by construction:

$$h_s = \sum_g \tilde{h}_s(g) \cdot g.$$

Note that if we fix an edge e_r , then the sequence $\{h_s(r)\}_{s \in \mathbb{N}}$ is eventually constant (by thinness). In other words, the sequence $\{h_s\}_{s \in \mathbb{N}}$ converges (pointwise). Take an accumulation point and obtain the limit $\tilde{h}$ of the sequence $\{\tilde{h}_s\}_s$ and a subsequence $\{\tilde{h}_{s_k}\}_k$ converging to $\tilde{h}$. We now have

$$h_{s_k} \rightarrow h, \quad \tilde{h}_{s_k} \rightarrow \tilde{h}.$$

Define

$$h' := \sum_{g \in \mathcal{B}} \tilde{h}(g) \cdot g.$$

We see that for a fixed edge e_r ,

$$\begin{aligned}h'(r) &= \sum_{g \in \mathcal{B}} \tilde{h}(g) \cdot g(r) \\ &= \sum_{g \in \mathcal{B}} \left(\lim_k \tilde{h}_{s_k}(g) \right) \cdot g(r) \\ &= \sum_{g \in \mathcal{B}} \left(\lim_k \tilde{h}_{s_k}(g) \cdot g(r) \right)\end{aligned}$$

and since $g(r) = 0$ for all but a finitely number of $g \in \mathcal{B}$, we can exchange limit and sum:

$$\begin{aligned}&= \lim_k \sum_{g \in \mathcal{B}} \tilde{h}_{s_k}(g) \cdot g(r) \\ &= \lim_k h_{s_k}(r) \\ &= h(r).\end{aligned}$$

This shows that any $h \in C$ can be written as a thin-sum of elements of $\mathcal{B}$, and concludes the proof. $\square$

Let $\mathcal{B}$ denote a thin thin-sum basis consisting of circuits as constructed above. By thinness, any set of elements of $\mathcal{B}$ defines a (thin) sum, and thus, we get a bijection:

$$\mathbb{Z}_2^{\mathcal{B}} \leftrightarrow C.$$

We now want to tweak $\mathcal{B}$ so as to have a sequence of increasing thin bases $\mathcal{B}'_i$ for C_i and $\mathcal{B}'$

for C such that:

$$\mathcal{B}'_i \subseteq \mathcal{B}'_{i+1} \subseteq \mathcal{B}'.$$

This can be done in a somewhat abstract setting:

Proposition 4.3.8. *Assume a thin-sum space $W \subseteq \mathbb{Z}_2^X$ is given, along with a sequence of nested thin-sum subspaces V_i . If $\mathcal{W}$ is a countable thin thin-sum basis for W , then one can construct a countable, thin, thin-sum basis $\mathcal{B}$ for W , along with sequence of nested thin-sum bases $\mathcal{B}_i$ for V_i .*

Since $\mathcal{W}$ is thin and a basis, we have $W \leftrightarrow \mathbb{Z}_2^{\mathcal{W}}$. Given any well-ordering of $\mathcal{W}$, one can let, for any $v \in W$:

$$l(v) = \min\{i : v(i) \neq 0 \text{ when } v \text{ is viewed as an element of } \mathbb{Z}_2^{\mathcal{W}}\}.$$

If $v = 0$, we use the same convention as above: $l(v) = \infty$, and for any w , $w(\infty) = 0$. In the proof, we will assume a fixed ordering on $\mathcal{W}$ and identify $W \leftrightarrow \mathbb{Z}_2^{\mathcal{W}}$.

Proof. Well-order the set W in a way that respects the subsets V_i , i.e.

$$W = \{v_\alpha \mid \alpha < \kappa\},$$

with κ the cardinal of W and such that there exist increasing ordinals α_i with

$$V_i = \{v_\alpha \in W \mid \alpha < \alpha_{i+1}\}.$$

We show that we can construct $\mathcal{B}_i$ ($i \in \mathbb{N}$) and $\mathcal{B}$ satisfying these requirements. We define the “Gaussian elimination” map

$$\nu : \kappa \times \kappa \times \mathbb{N} \rightarrow \mathbb{Z}_2$$

inductively as follows:

$$\nu_{0,\alpha} = v_\alpha, \quad \nu_{\tau,\alpha} = \begin{cases} \lim_{\sigma < \tau} \nu_{\sigma,\alpha} & \text{if } \tau \text{ is a limit ordinal,} \\ v_{\alpha,\alpha} & \text{otherwise and } \tau > \alpha, \\ v_{\tau-1,\alpha} - v_{\tau-1,\alpha}(l(v_{\tau-1,\tau-1})) \cdot v_{\tau-1,\tau-1} & \text{otherwise and } \tau \leq \alpha, \end{cases} \quad (4.3.1)$$

where the limit is a pointwise limit (which means that at any coordinate r , the transfinite sequence, evaluated at that coordinate, is eventually constant).

From now on, we will write ν_α for the diagonal element $\nu_{\alpha,\alpha}$.

The interpretation of the map ν is as follows. An element of W is uniquely determined by a map $\mathbb{N} \rightarrow \mathbb{Z}_2$ (a row). Thus, the set W , once ordered, can be represented as a map $\kappa \times \mathbb{N} \rightarrow \mathbb{Z}_2$ (a matrix). Gaussian elimination, in the finite case, is an iterative process where one iterates over the rows of a matrix, each the “pivot”, while modifying the matrix at each step, hence a map $\kappa \times \kappa \times \mathbb{N} \rightarrow \mathbb{Z}_2$. In short, the first coordinate represents the iteration/pivot, the second the element in question, and the third its value at a given base coordinate.

At time zero, the matrix has not been modified. At time τ , if τ is a successor ordinal, we pivot around the element $\nu_{\tau-1,\tau-1}$. Note that if the row α “lies before the pivot” or “is the pivot” (i.e. $\alpha < \tau$), it is not changed. If row α “lies after the pivot” ($\alpha \geq \tau$), we must subtract the pivot to

what row α “just was”. In case τ is a limit ordinal, then we will just re-iterate all changes that happened before time τ and restart pivoting in its successor.

We start by showing by induction on τ that:

1. For any $n \in \mathbb{N}$, the number of ordinals $\sigma \leq \tau$ such that $v_\sigma(n) \neq 0$ is finite.
2. The limit in Equation (4.3.1) is actually well-defined.
3. For any $\sigma < \tau, \alpha$, $v_{\tau,\alpha}(l(v_\sigma)) = 0$.
4. The value of $v_{\tau,\alpha}$ is

$$v_{\tau,\alpha} = v_{0,\alpha} + \sum_{\sigma < \min\{\alpha, \tau\}} v_{\sigma,\alpha}(l(v_\sigma)) \cdot v_\sigma. \quad (4.3.2)$$

Indeed, assume the above holds for any $\theta < \tau$.

If $\tau = 0$: Then Items 1 to 4 hold trivially.

If τ is successor:

Item 1 is easy. By the induction hypothesis the number of $\sigma \leq \tau - 1$ with $v_\sigma(n) \neq 0$ is finite. Hence it must be finite for τ too.

Item 2 does not apply successor ordinals.

For Item 3, we see that if $\tau > \alpha$, then

$$v_{\tau,\alpha}(l(v_\sigma)) = v_\alpha(l(v_\sigma))$$

which is zero by induction hypothesis (i.e. Item 3 holds at $\theta = \alpha$). If $\tau \leq \alpha$, then

$$v_{\tau,\alpha} = v_{\tau-1,\alpha} - v_{\tau-1,\alpha}(l(v_{\tau-1,\tau-1})) \cdot v_{\tau-1,\tau-1}, \quad (4.3.3)$$

which, evaluated in $l(v_{\tau-1,\tau-1})$ is zero by construction. If $\sigma < \tau - 1$, both $v_{\tau-1,\alpha}(l(v_\sigma))$ and $v_{\tau-1,\tau-1}(l(v_\sigma))$ are zero, hence so is $v_{\tau,\alpha}$.

Finally, Item 4. One sees that if $\tau \leq \alpha$, we get $\min\{\tau, \alpha\} = \tau$, and:

$$\begin{aligned} v_{\tau,\alpha} &= v_{\tau-1,\alpha} + v_{\tau-1,\alpha}(l(v_{\tau-1,\tau-1})) \cdot v_{\tau-1,\tau-1} \\ &= v_{0,\alpha} + \sum_{\sigma < \tau-1} v_{\sigma,\alpha}(l(v_\sigma)) \cdot v_\sigma + v_{\tau-1,\alpha}(l(v_{\tau-1,\tau-1})) \cdot v_{\tau-1,\tau-1}, \end{aligned}$$

by definition of $v_{\tau,\alpha}$. If $\tau > \alpha$, we get $\min\{\tau, \alpha\} = \alpha$, and

$$\begin{aligned} v_{\tau,\alpha} &= v_{\tau-1,\alpha} = v_{0,\alpha} + \sum_{\sigma < \min\{\alpha, \tau-1\}} v_{\sigma,\alpha}(l(v_\sigma)) \cdot v_\sigma \\ &= v_{0,\alpha} + \sum_{\sigma < \min\{\alpha, \tau\}} v_{\sigma,\alpha}(l(v_\sigma)) \cdot v_\sigma \end{aligned}$$

Now moving on to the limit case.

If τ is limit: We first check Item 1. It suffices to show that for any n , the number of ordinals $\sigma < \tau$ with $v_\sigma(n) \neq 0$ is finite, since then the number of $\sigma \leq \tau$ with $v_\sigma(n) \neq 0$ will also be finite. If $n = l(v_\sigma)$ for some $\sigma < \tau$, then the number of $\theta < \sigma$ for which $v_\theta(n) = 0$ is finite by the induction hypothesis (Item 1), and for any $\tau > \theta > \sigma$, $v_\theta(n) = 0$, also by the induction hypothesis (Item 3). Assume then that there exists some n which is not a leading coefficient,

which is also in the support of infinitely many v_σ ($\sigma < \tau$). We can take n to be the least such, and for σ large enough, we will have:

$$v_\sigma(n) \neq 0, \quad v_\sigma(k) = 0, \quad \forall k < n,$$

so that n is actually the leading coefficient of v_σ , a contradiction.

We now check Item 2. Let α arbitrary, then

$$v_{\tau,\alpha} = \lim_{\sigma < \tau} \left(v_{0,\alpha} + \sum_{\sigma' < \min\{\alpha, \sigma\}} v_{\sigma',\alpha}(l(v_{\sigma'})) \cdot v_{\sigma'} \right)$$

which we need to verify to be well-defined. In other words, we need to check that for any n , the sequence

$$\sigma \mapsto v_{0,\alpha}(n) + \sum_{\sigma' < \min\{\alpha, \sigma\}} v_{\sigma',\alpha}(l(v_{\sigma'})) \cdot v_{\sigma'}(n)$$

stabilises starting at some ordinal. But by Item 1, $v_{\sigma'}(n)$ is zero for σ' large enough, so that the sequence indeed stabilises. Furthermore, it stabilises to:

$$v_{0,\alpha}(n) + \sum_{\sigma < \min\{\alpha, \tau\}} v_{\sigma,\alpha}(l(v_\sigma)) \cdot v_\sigma(n),$$

which also proves Item 4.

It remains to check Item 3. But if $\sigma < \tau, \alpha$, then

$$v_{\tau,\alpha}(l(v_\sigma)) = \lim_{\sigma' < \tau} v_{\sigma',\alpha}(l(v_\sigma)) = \lim_{\sigma < \sigma' < \tau} v_{\sigma',\alpha}(l(v_\sigma)) = \lim 0 = 0.$$

This closes the induction.

Let now

$$\mathcal{B} := \{v_\alpha | \alpha < \kappa, v_\alpha \neq 0\}, \quad \mathcal{B}_i := \{v_\alpha | \alpha < \alpha_{i+1}, v_\alpha \neq 0\}$$

Note that by Item 1, each non-zero v_α has a well-defined leading index $l(v_\alpha)$, and $v_\alpha \neq v_\beta$ implies $l(v_\alpha) \neq l(v_\beta)$. It follows that $\mathcal{B}$ is at most countable. We can now check the following consequences of this process:

1. By Equation (4.3.2), $\mathcal{B}_i$ generates V_i and $\mathcal{B}$ generates W . Indeed, Equation (4.3.2) yields (by setting $\alpha = \tau$)

$$v_\tau = v_{0,\tau} + \sum_{\sigma < \tau} v_{\sigma,\tau}(l(v_\sigma)) \cdot v_\sigma,$$

so that, by rewriting the sum, we get

$$v_{0,\tau} = v_\tau - \sum_{\sigma < \tau} v_{\sigma,\tau}(l(v_\sigma)) \cdot v_\sigma,$$

so that any element of W (resp. V_i) is a thin sum of elements of $\mathcal{B}$ (resp. $\mathcal{B}_i$).

2. $\mathcal{B}$ is thin, as a subset of $\mathbb{Z}_2^{\mathcal{W}}$, by the same argument as the proof of Item 1 for τ limit.
3. $\mathcal{B}$ is thin, as a subset of $\mathbb{Z}_2^X$. Every $x \in X$ appears in the support of a finite number of $v \in \mathcal{W}$, and every $v \in \mathcal{W}$ in the support of a finite number of elements of $\mathcal{B}$.

4. The set $\mathcal{B}$ is thin-sum independent. Suppose not. Then, for some $\tilde{\kappa} \subset \kappa$, we have $\sum_{\sigma \in \tilde{\kappa}} \nu_\sigma = 0$. Let $\tau_0 := \min\{\sigma \in \tilde{\kappa} \mid \nu_\sigma \neq 0\}$. Then $\sum_{\sigma \in \tilde{\kappa}} \nu_\sigma(l(\nu_{\tau_0})) = 1$, a contradiction.

□

In our case, this yields:

Corollary 4.3.9. *There exists a thin thin-sum basis $\mathcal{B}$ for C with nested subsets $\mathcal{B}_i$, each a thin-sum basis for C_i .*

Proof. By Proposition 4.3.7 and Proposition 4.3.8. □

And we can now easily get:

Theorem 4.3.10. *If X has large circuits, then $C(X, \mathbb{Z}_2)/C_\infty(X, \mathbb{Z}_2)$ is infinite dimensional (as a vector space).*

Proof. Let $\mathcal{B}, \mathcal{B}_i$ be as in Corollary 4.3.9.

If X has large circuits, the sequence

$$C_1 \subseteq \dots \subseteq C_i \subseteq C_{i+1} \subseteq \dots$$

does not stabilise, so that there exists a sequence $(i_j)_{j \in \mathbb{N}}$ satisfying:

$$C_{i_j} \subsetneq C_{i_{j+1}}.$$

In particular, this implies that $\mathcal{B}_{i_j} \subsetneq \mathcal{B}_{i_{j+1}}$. Choose then, for each j , $b_j \in \mathcal{B}_{i_{j+1}} - \mathcal{B}_{i_j}$.

Let $\mathcal{J}$ be an infinite set of infinite pairwise disjoint subsets of $\mathbb{N}$ and

$$f_J := \sum_{j \in J} b_j, \quad J \in \mathcal{J}.$$

We claim that the family $\{[f_J] \in C(X, \mathbb{Z}_2)/C_\infty(X, \mathbb{Z}_2)\}_{J \in \mathcal{J}}$ is linearly independent, which will imply infinite dimensionality.

We can actually show more: the family $\{f_J\}_{J \in \mathcal{J}}$ is thin, and no sum of the form:

$$\sum_{J \in \mathcal{J}'} f_J, \quad \mathcal{J}' \subseteq \mathcal{J},$$

lies in C_∞ . Indeed, since the f_J s have disjoint support, any such combination can be written as an infinite sum:

$$f = \sum_{j \in J'} b_j,$$

for some $J' \subseteq \mathbb{N}$. If $f \in C_\infty$, then there would exist R with $f \in C_R$, so that $f = \sum_{b \in B} b$ for $B \subseteq \mathcal{B}_R$, and $0 = \sum_{B_{J'} \Delta B} b$, where $B_{J'} = \{b_j \mid j \in J'\}$. Since J' is infinite, $B_{J'}$ contains elements of $\mathcal{B}_{R'}$ for R' arbitrarily high, so that the symmetric difference above is not empty. This is a contradiction with $\mathcal{B}$ being thin-sum independent.

□

Corollary 4.3.11. *Let X be a ULF graph, then the following are equivalent:*

1. X has $\mathbb{Z}_2$ -large circuits;
2. $C(X, \mathbb{Z}_2)/C_\infty(X, \mathbb{Z}_2)$ is infinite dimensional (as a vector space).

Proof. By the above and Proposition 4.3.5. □

In $Z_1^{\text{uf}}(X, \mathbb{Z}_2)$, infinite sums are allowed, as long as they are locally finite. However, in $H_1^{\text{uf}}(X, \mathbb{Z}_2)$, only *finite* sums are naturally permitted; there is no obvious “thin-sum” space structure on quotients. Thus, the passage to a quotient involves a major loss of information. We suggest a very simple approach that remedies this problem in our ad hoc situation. It makes sense to define a thin quotient of thin vector spaces. There might be a more abstract setting about which one could say something.

Definition 4.3.12. Fix a thin-sum space $V \leq \mathbb{Z}_2^X$, and a sub-vector space $W \leq V$ (not necessarily closed under thin-sum). In the quotient V/W , we say that a potentially infinite sum of elements $\sum a_i$ is *consistent* iff there is a collection of preimages $\tilde{a}_i \in a_i$ such that the collection $\{\tilde{a}_i\}_i$ is thin (in particular, abusing notation, any $v \in V - \{0\}$ appears as at most a finite number of $\tilde{a}_i$ s), and all such thin choices of representatives result in the same element in V/W :

$$\sum \tilde{a}_i - \sum \tilde{b}_i \in W,$$

if $\{\tilde{a}_i\}_i, \{\tilde{b}_i\}_i$ are two choices of representatives.

In order to have a cleaner theorem we also introduce the following “thin” sum.

Theorem 4.3.13. *The following are equivalent:*

1. In $H_1^{\text{uf}}(X, \mathbb{Z}_2)$, the infinite sum $\{[0]\}_{i=0}^\infty$ of the zero-class is consistent.
2. X has no large circuits.
3. Any sum of cosets that has a thin representative sum is consistent.

Proof. Clearly 3 implies 1.

We show 1 implies 2. Suppose X has large circuits, then there exists an infinite thin sum of large circuits that is non-zero in homology. However every individual circuit is finite and thus zero in homology. On the other hand the repeated infinite sum $\sum_{i=1}^\infty [0]$ gives 0. So we see that the infinite sum of the zero class has two representing thin sums that give a different solution. Hence it is inconsistent.

We show 2 implies 1. Consider an infinite representing thin sum $\sum_{i=1}^\infty \tilde{x}_i$ of the infinite zero class sum. Since there are no large circuits, there is some $R > 0$ such that each $\tilde{x}_i$ can be written as a sum of R -circuits. More specifically there is a g_i such that $\partial g_i = \tilde{x}_i$. Now take some accumulation point λ of the sequence $\left\{ \sum_{i=1}^n g_i \right\}_{n \in \mathbb{N}}$. Then $\partial \lambda = \sum_{i=1}^\infty \tilde{x}_i$. The infinite zero-class sum is as such consistent.

We show 1 implies 3. Suppose there is some inconsistent thin sum. So two representing sums $\sum_{i=1}^{\infty} \tilde{x}_i$ and $\sum_{i=1}^{\infty} \tilde{y}_i$, such that $[\tilde{x}_i] = [\tilde{y}_i]$. Then each $[\tilde{x}_i - \tilde{y}_i] = [0]$. Hence by assumption $\left[\sum_{i=1}^{\infty} \tilde{x}_i - \tilde{y}_i\right] = [0]$. Which shows that $\left[\sum_{i=1}^{\infty} \tilde{x}_i\right] = \left[\sum_{i=1}^{\infty} \tilde{y}_i\right]$, because finite sums are always well defined. $\square$

Corollary 4.3.14. *If $H_1^{\text{uf}}(X, \mathbb{Z}_2)$ has large circuits, then every thin-sum is inconsistent.*

Proof. Direct consequence of Theorem 4.3.13. $\square$

We draw the following conclusion:

Define the *thin dimension* $\dim_{\text{thin}}$ of $H_1^{\text{uf}}(X, \mathbb{Z}_2)$ as the least cardinal κ such that $H_1^{\text{uf}}(X, \mathbb{Z}_2)$ has a subset K of cardinality κ and such that any element of $H_1^{\text{uf}}(X, \mathbb{Z}_2)$ can be written as a consistent sum of elements of K .

Theorem 4.3.15. • $\dim_{\text{thin}}(H_1^{\text{uf}}(X, \mathbb{Z}_2))$ is countable if and only if X has no large circuits. In this case it is equal to the number of ends -1 (where we do not make the difference between countably many or uncountably many ends).

- $\dim_{\text{thin}}(H_1^{\text{uf}}(X, \mathbb{Z}_2))$ is uncountable if and only if X has large circuits.

Proof. First part. If there are no large circuits, then by Proposition 4.3.2 there exists a tree T inside X , such that $H_1^{\text{uf}}(T, \mathbb{Z}_2) = H_1^{\text{uf}}(X, \mathbb{Z}_2)$. The isomorphism is given by the natural inclusion map. Hence Proposition 4.3.2 gives us a thin thin-sum basis $\mathcal{B}$ in the tree. Note that all sums remain well-defined when we inject $\mathcal{B}$ into X by Theorem 4.3.13.

Second part. Suppose X has large circuits, then by Corollary 4.3.14 all consistent sums are finite. However we showed before that $H_1^{\text{uf}}(X, \mathbb{Z}_2)$ contains an uncountable number of elements. So a generating set must be uncountable too.

Note that since having large circuits and not having large circuits exhausts the universe, both implications above are actually equivalences. $\square$

4.4 $\mathbb{Z}$ -large circuits

Recall that a ULF graph X has $\mathbb{Z}$ -large circuits if

$$\forall r > 0 \ C_r(X, \mathbb{Z}) \neq C_{\infty}(X, \mathbb{Z}).$$

Proposition 4.4.1. *If X is vertex-transitive and has $\mathbb{Z}$ -large circuits, then $H_1^{\text{uf}}(X, \mathbb{Z}) \neq 0$.*

Proof. Assume X vertex-transitive with vanishing homology in dimension one and $\mathbb{Z}$ -large circuits. For any $r > 0$, take some $f_r \in C_{\infty} - C_r$.

We claim that f_r can be assumed to have norm 1. First, by assumption that $f_r \in C_{\infty}$, there exists some $s > r$ with $f_r \in C_s$, and thus $f_r = \epsilon g$ for some $g : \mathfrak{C}_s \rightarrow \mathbb{Z}$ uniformly bounded (recall the definition of ϵ from Section 4.1). One can write $g = g^+ - g^-$, so just assume for now that $g \geq 0$. Now, consider the following process: view g as a multiset, and consider a maximal subset of circuits that do not intersect pairwise, and subtract it from g ; iterate. This process has

to end eventually (indeed, fix any circuit c in the support of g ; by ULF, there exists some K , not depending on c , such that c intersects at most K other circuits in the support of g ; if during the K first steps c is never removed, this means by maximality that at each step a neighbour is removed; thus after $K + 1$ steps all neighbours are removed, and then so will c , by maximality), so that $f = \sum eg_i^+ - \sum eg_i^-$, with each $f_i^\pm := eg_i^\pm$ of norm 1. If each $f_i^\pm$ lied in C_r , so would f (C_r being a vector space). We conclude that at least one of $f_i^\pm$ does not lie in C_r , and the claim is verified.

Fix some basepoint $v_0 \in VX$, and for any r , some $f_r = eg_r$ as above, and consider the family $(f_r^i \in C_1^{\text{su}}(X, \mathbb{Z}))_{i \in \mathbb{N}}$ constructed as follows: for any i , take the restriction of g_r to circuits strictly inside the ball of radius i around v_0 , call this g_r^i and let $f_r^i = eg_r^i$. Note that since all circuits of g_r^i are disjoint, f_r^i also has norm 1; furthermore, f_r^i has support in a ball of radius i . In other words, f_r and f_r^i are both elements of Z_1^{su} of norm $(\leq) 1$ and agreeing on the ball of radius i around v_0 . The important aspect of this family is that $f_r^i \xrightarrow{i} f_r$ pointwise.

If v is a vertex in X , denote by ϕ_v some automorphism of X with $\phi_v(v_0) = v$. Then there is a set $\{v_{r,i}\}$ of vertices in X (we write $\phi_{r,i}$ for the corresponding automorphism) such that the balls $\phi_{r,i}(B_{r,i})$ are pairwise disjoint.

The sum $F := \sum_{r,i} f_r^i \circ \phi_{r,i}^{-1}$ is then well-defined, lies in Z_1^{su} and has norm 1. Since $H_1^{\text{uf}} = 0$, one has that $F \in C_\infty$, and thus, by definition, $F \in C_s$ for some s . Take therefore some $G : \mathbb{C}_s \rightarrow \mathbb{Z}$ with $eG = F$ and of uniformly bounded norm. Fix any $r > s$ and consider the sequence $(G_{r,i}) := (G \circ \phi_{r,i})_i$. Note that by construction $eG_{r,i}|_{B_{r,i}} = f_{r,i}$ and recall that $f_{r,i} \rightarrow f_r$. Up to taking a subsequence, one may assume that $G_{r,i}$ converges to some G_r (since the $G_{r,i}$ are uniformly bounded in norm). Therefore, $G_{r,i} \rightarrow G_r$, which implies $eG_{r,i} \rightarrow eG_r$, but $eG_{r,i} \rightarrow f_r$ since $eG_{r,i}|_{B_{r,i}} = f_{r,i}$ and the sequence $B_{r,i}$ exhausts X . Since $eG_r = f_r$, we get $f_r \in C_s$, thus contradicting $f_r \notin C_r$.

□

Proposition 4.4.2. *If X is vertex-transitive and has $\mathbb{Z}$ -large circuits, then $\dim(H_1^{\text{uf}}(X, \mathbb{Z})) = \infty$.*

Proof. Consider the function F constructed in the proof of Proposition 4.4.1. In the same way it is possible to construct instead countably many functions F_k each defined on balls $B_{r,i}^k$, which are the building blocks of the construction. Moreover these F_k can be chosen to not have any intersecting balls between them. By arguments of the proof before they are all non-zero in homology and finite linear combinations between them (as well as bounded infinite linear combinations for that matter) are also non-zero. Hence the result follows. □

4.5 Expansion

In this section, we follow a different trail in our search for a description of the vanishing of $H_1^{\text{uf}}(X, \mathbb{Z})$. Our motivation is twofold. On the one hand, after toying a bit with $H_1^{\text{uf}}(X, \mathbb{Z})$, one comes to the conclusion that, ends and “large circuits” aside, the main ingredient in non-vanishing of homology is a lack of triangles, compared to edges. The typical example here is (a triangulated) $\mathbb{Z}^2$: any infinite sum of parallel bi-infinite paths going in one direction will define a non-zero class in homology, since there are not enough triangles to kill all bi-infinite paths at the same time. On the other hand, (classical) graph expansion can be interpreted as the ability, given any finite subset of vertices U , and element $f \in Z_0^{\text{su}}(U)$, to push f out of U in a bounded way; that is, there exists some $g \in C_1^{\text{su}}(X)$ with $(\partial g)|_U = f$ and $\|g\| \leq \varepsilon^{-1}\|f\|$, where ε is the Cheeger constant for X . This section can then be described as an effort to try and mimic the

“homological interpretation” of expansion in higher dimensions (we call this H_n^{suf} -expansion), and fine-tune it to isolate the phenomenon of “lack of triangles” from “ends” and “large circuits” (thus getting *pure* H_n^{suf} -expansion).

Recall that a (ULF) graph X is said to be *amenable* if it satisfies the following condition:

$$\forall \epsilon > 0 \exists U \subseteq V X \text{ finite} : |\partial^e U| < \epsilon |U|.$$

If X is non-amenable, the least $\epsilon > 0$ satisfying $\forall U \subseteq V X \text{ finite}, |\partial^e U| \geq \epsilon |U|$ is called the Cheeger constant for X .

The following was proved in [BW92]:

Proposition 4.5.1 (Part of [BW92, Theorem 3.1]). *A ULF graph X is non-amenable if and only if $H_0^{\text{uf}}(X, \mathbb{Z}) = 0$.*

4.5.1 Basics

Let X be a ULF simplicial complex. If $U \subseteq V X$ is a set of vertices, we also use U for the induced simplicial subcomplex.

Definition 4.5.2 (H_n^{suf} -Expansion). We say that X has n -dimensional H_n^{suf} -expansion if there exists a function $K : \mathbb{N} \rightarrow \mathbb{N}$ satisfying the following. For any finite sets of vertices $U \subseteq U' \subseteq V X$, and any $f \in C_n^{\text{suf}}(U', \mathbb{Z})$ satisfying:

1. ∂f has support contained in $U' - U$.
2. f can be extended to some $\hat{f} \in Z_n^{\text{suf}}(X, \mathbb{Z})$ with $\|\hat{f}\| = \|f\|$ (by extended, we mean $\hat{f}|_{U'} = f$).

Then, there exists $g \in C_{n+1}^{\text{suf}}(X, \mathbb{Z})$ such that $\partial g|_{U'} = f$ and $\|g\| \leq K(\|f\|)$.

We say that a function K as above is a certificate to H_n^{suf} -expansion.

Finally, if Item 2 in the definition is replaced by

3. f can be extended to some $\hat{f} \in B_n^{\text{suf}}(X, \mathbb{Z})$ with $\|\hat{f}\| = \|f\|$.

and the rest is left as-is, then the obtained condition is called *pure* H_n^{suf} -expansion.

Intuitively, expansion can be interpreted as homological cycles locally being homological boundaries, with control. In other words, one can always kill “pieces of cycles” with “pieces of boundaries” in a bounded way. We will see that this property *always* implies that simplicial uniformly finite homology vanishes, at the corresponding dimension. Pure expansion is slightly harder to describe, but the core idea is that it only takes into account those f that “can actually be killed” by some boundary.

Note that

$$H_n^{\text{suf}}\text{-expansion} \quad \Rightarrow \quad \text{pure } H_n^{\text{suf}}\text{-expansion}.$$

Lemma 4.5.3. *If X has H_n^{suf} -expansion, then $H_n^{\text{suf}}(X, \mathbb{Z}) = 0$.*

Proof. Assume that K is a certificate. Fix $f \in Z_n^{\text{suf}}(X, \mathbb{Z})$, and $\langle U_i \subseteq V X \rangle_i$ a nested sequence of finite subsets exhausting $V X$. Let U'_i be the 2-neighbourhood of U_i . For each i , let f_i the restriction of f to U'_i , and assume that $\|f_i\| = \|f\|$ for each i (achievable by starting with U_0 large

enough). Since $\|f\| = \|f_i\|$, let $g_i \in C_{n+1}^{\text{su}}(X, \mathbb{Z})$ satisfying $\partial g_i|_{U'_i} = f_i$ and $\|g_i\| \leq K(\|f\|)$, as given by H_n^{su} -expansion.

Then the sequence $\langle g_i \rangle_i$ has an accumulation point, say g , since its elements are uniformly bounded (by $K(\|f\|)$). Since ∂g is also an accumulation point of $(\partial g_i)_i$, which converges to f , we conclude that $\partial g = f$.

Thus, f is in $B_n^{\text{su}}(X, \mathbb{Z})$ and $H_n^{\text{su}}(X, \mathbb{Z}) = 0$. $\square$

Remark 4.5.4. Consider the following condition on X : there exists some constant N such that for any $f \in Z_n^{\text{su}}(X, \mathbb{Z})$, one can write $f = \sum_{i=1}^{N\|f\|} f_i$, with each f_i of norm 1 and in $Z_n^{\text{su}}(X, \mathbb{Z})$. If this condition held, we could always take the function $K : \mathbb{N} \rightarrow \mathbb{N}$ in H_n^{su} -expansion to be linear. We do not know whether it holds, even in dimension 1 (that is: can one decompose any flow of norm k into a sum of Nk flows of norm 1, for some universal N ?).

In the zero-dimensional case, H^{su} -expansion is equivalent to non-amenability. Moreover, since $B_0^{\text{su}} = Z_0^{\text{su}}$, also pure H_0^{su} expansion coincides with H_0^{su} expansion.

Proposition 4.5.5 (H_0^{su} -expansion is just expansion.). *Let X be a ULF graph. Then X has H_0^{su} -expansion iff X is non-amenable.*

Proof. By Proposition 4.5.1 and Lemma 4.5.3, we get H_0^{su} -expansion $\Rightarrow$ non-amenability.

Now for the converse. We may take $U = U'$. By decomposing such an $f : U \rightarrow \mathbb{Z}$ into a sum of 0/1-valued functions, the condition reduces to:

- There exists $K \in \mathbb{N}$ such that for all $U \subseteq V X$ and $W \subseteq U$, there exists $g : EX \rightarrow \mathbb{Z}$ with $\partial g|_U = \chi_W$ and $\|g\| \leq K$.

Assume now that X is non-amenable, with Cheeger constant ε . By adapting the proof of [BS97, Lemma 2.1] (more precisely, their construction of a flow), we see that given $U \subseteq V X$ and $W \subseteq U$, one can construct a g satisfying $\partial g|_U = \chi_W$. Since in their argument, the “flux” (=excess flow) at vertices in W is ε and the flow at edges at most 1, scaling by ε^{-1} yields the bound of $K := \varepsilon^{-1}$ for g . $\square$

4.5.2 Dimension 1

Theorem 4.5.6. *If X is an infinite, vertex-transitive ULF simplicial complex and $H_1^{\text{su}}(X, \mathbb{Z}) = 0$, then X has H_1^{su} -expansion.*

Proof. Suppose X does not have H_1^{su} -expansion but $H_1^{\text{su}}(X, \mathbb{Z}) = 0$. Then, there exist sequences $\langle U_i \subseteq U'_i \subseteq V X \rangle$ and $\langle f_i \in C_1^{\text{su}}(X, \mathbb{Z}) \rangle_i$ as in the definition, with $\|f_i\| \leq C$ for some C , but such that if $\partial g_i|_{U'_i} = f_i$, then $\|g_i\| \geq K_i$, with $K_i \rightarrow \infty$ (the sequence $\langle f_i \rangle_i$ becomes “harder” to kill with i increasing).

Let $f'_i \in Z_1^{\text{su}}(X, \mathbb{Z})$ be an extension of f_i , as in the definition, but with finite support (their existence is proven in Lemma 4.5.7). Since X is infinite and transitive, one can assume that the support of all f'_i s are pairwise disjoint.

Let then $F = \sum_i f'_i$. Then, $F \in Z_1^{\text{su}}(X, \mathbb{Z})$ implies the existence of some $G \in C_2^{\text{su}}(X, \mathbb{Z})$ with $\partial G = F$. In particular, $\partial G|_{U'_i} = f_i$, and since $\|G\| \leq K_i$ for large enough i , we get a contradiction. $\square$

Note that we do not actually need as much as transitivity, but merely being able to move finite pieces of the graphs sufficiently far, so that, e.g. having a cocompact action by the automorphism group is enough.

Lemma 4.5.7. *Assume X is ULF and one-ended. Let $U \subseteq U' \subseteq VX$ and $f \in C_1^{\text{suf}}(U', \mathbb{Z})$ satisfying $\partial f \subseteq U' - U$ and extensible to $\hat{f} \in Z_1^{\text{suf}}(X, \mathbb{Z})$ with $\|\hat{f}\| = \|f\|$. Then $\hat{f}$ can be taken with finite support.*

Proof. That f can be extended to $\hat{f} \in Z_1^{\text{suf}}(X, \mathbb{Z})$ implies the existence of a finite number of vertices $v_1, \dots, v_n \in U' - U$, and integers $m_1, \dots, m_n$ such that $\partial f = \sum m_i v_i$. Since $\partial f \subseteq U' - U$, we must have $\sum m_i = 0$. Note that the values at these vertices must be killed. To do this, one can either connect two vertices of opposite sign outside of U' , or add an infinite ray. Because both the number of v_i s as well as their coefficients are finite, there are at most finitely many rays. Consider a ball enveloping the finite structure of $\hat{f}$, such that only the rays are escaping. Take any ray coming from a positive vertex and any ray coming from a negative vertex. Since X is one-ended, both rays can be connected outside of the chosen ball by a path p . If the path does not cross any other rays, we can reconnect locally. If it crosses other rays we can see this as sequence of basis vertices starting in one with a positive coefficient, ending in a negative one. Hence there exist two consecutive rays such that their vertices have opposite sign. This ends the proof. $\square$

The above “finite extension” property for dimension 1 is crucial in our proof that for transitive infinite graphs, vanishing of H_1^{suf} is equivalent to H_1^{suf} -expansion. Were it to also hold for higher dimensions, the proof of Theorem 4.5.6 would obviously generalise.

H^{uf}

In the above, we focused on H^{suf} , since H^{suf} -expansion is defined in terms of the simplicial complex at hand. In the (coarse) setting of H^{uf} , this is obviously problematic. Our aim now is essentially to coarsify [pure] H^{suf} -expansion.

Definition 4.5.8 (H_n^{uf} -expansion). A ULF graph X is said to have [pure] H_n^{uf} -expansion iff its Rips complexes eventually all have [pure] H_n^{suf} -expansion.

Proposition 4.5.9. *For ULF graphs, [pure] H_1^{uf} -expansion is a quasi-isometry invariant.*

Proof. Let $f : X \rightarrow Y$ be a quasi-isometry, where Y has [pure] H_1^{uf} -expansion. So there is some Rips complex $R_r Y$ that has [pure] H_1^{suf} -expansion. Since any flow is close to a flow in Z_1^{suf} , we may suppose it is actually in there. Note that for any s the induced map $f : R_s X \rightarrow R_r Y$ is still a quasi-isometry and for s sufficiently large, the pre-image of any simplex in $R_r Y$ is also a simplex in $R_s X$. Since we are allowed to choose our flow in Z_1^{suf} , it follows that also $R_s X$ has [pure] H_1^{suf} -expansion. $\square$

The following proposition justifies our interpretation of non-vanishing of $H_1^{\text{uf}}(X, \mathbb{Z})$ in terms of the three phenomena of ends, large circuits and expansion:

Proposition 4.5.10. *If X has $\mathbb{Z}$ -small circuits (i.e. $C_r(X, \mathbb{Z}) = C_\infty(X, \mathbb{Z})$ for some r) and is one-ended, then $H_1^{\text{uf}}(X, \mathbb{Z}) \neq 0$ implies that X does not have pure H_1^{uf} -expansion.*

In other words, if $H_1^{\text{uf}}(X, \mathbb{Z}) \neq 0$ then at least one of the three phenomena is responsible for it. In the non-transitive case, it is not clear that the converse holds, that is, that $H_1^{\text{uf}}(X, \mathbb{Z}) = 0$ would imply at least one of: ends, large circuits, expansion.

Proof. Assume $H_1^{\text{uf}}(X, \mathbb{Z}) \neq 0$, X one-ended and, for some r , $C_r(X, \mathbb{Z}) = C_\infty(X, \mathbb{Z})$. Then, by Corollary 4.1.6, for s big enough, $H_1^{\text{uf}}(R_s(X), \mathbb{Z}) \neq 0$. By Lemma 4.5.3, it follows that $R_s(X)$ does not have expansion, and it remains to check that it does not have *pure* expansion either.

Fix $(U \subseteq U', f)$ with $f \in C_1^{\text{uf}}(R_s(U'), \mathbb{Z})$ with $\partial f \subseteq U' - U$ and f extensible.

Note that $R_s(X)$ does not have expansion. Hence, it suffices to check that there exists some $g \in C_2^{\text{uf}}(R_s(X), \mathbb{Z})$ with $\partial g|_{R_s(U')} = f$ and $\|\partial g\| = \|f\|$. Let $\hat{f}$ a finite extension of f , as given by Lemma 4.5.7. By “tracing virtual edges”, $\hat{f}$ is, up to a finite sum of boundaries of triangles in $R_s(X)$, an element of $Z_1^{\text{uf}}(X, \mathbb{Z})$, of finite support. In particular, up to a finite sum of boundaries of triangles, $\hat{f}$ lies in $C_\infty(X, \mathbb{Z}) = C_r(X, \mathbb{Z})$, and can be decomposed into a sum of circuits of length $\leq r$. Now, if s is large enough, each of those circuits can itself be decomposed as a sum of boundaries of triangles in $R_s(X)$, so that $\hat{f}$ can be written as a sum of boundaries of triangles of $R_s(X)$. In other terms, $\hat{f} = \partial g$ for some $g \in C_2^{\text{uf}}(R_s(X), \mathbb{Z})$. □

Theorem 4.5.11 (H_1^{uf} triad). *Let X a transitive graph, and consider the following three conditions:*

1. *X has more than one end.*
2. *X has $\mathbb{Z}$ -large circuits; that is, $C_r(X, \mathbb{Z}) \neq C_\infty(X, \mathbb{Z})$ for any r .*
3. *X does not have pure H_1^{uf} -expansion.*

Then $H_1^{\text{uf}}(X, \mathbb{Z}) \neq 0$ iff either of the above holds.

Proof. By the results in Section 4.2, we know already that Item 1 implies non-zero homology. Similarly, the results in Section 4.4 give that Item 2 implies non-zero homology.

Finally, assume that Item 3 holds. If any of Items 1 and 2 hold, then obviously $H_1^{\text{uf}}(X, \mathbb{Z}) \neq 0$ by what we already know. So, one may assume now that for some r large enough, $C_r(X, \mathbb{Z}) = C_\infty(X, \mathbb{Z})$. Since $R_r(X)$ does not have (pure) expansion (r large enough), $H_1^{\text{uf}}(R_r(X), \mathbb{Z}) \neq 0$, and by Corollary 4.1.6, $H_1^{\text{uf}}(X, \mathbb{Z}) \neq 0$.

For the converse, suppose that $H_1^{\text{uf}}(X, \mathbb{Z}) \neq 0$, but X has one end and $\mathbb{Z}$ -small circuits (i.e. neither of Items 1 and 2 hold), then we can apply Proposition 4.5.10. □

Remark 4.5.12 (Finitely presented groups). We might wish to use Theorem 4.5.11 to characterise the vanishing of uniformly finite homology of finitely generated groups in terms of ends, large circuits and expansion. Ideally, each of these three phenomena would be given a group-theoretic description, thus allowing a group-theoretic description of the vanishing itself.

The theory of ends of groups is well-developed, and it is relatively clear that a finitely presented group will have $\mathbb{Z}$ -small circuits. This is as far as we could get. An obvious guess would be that a finitely generated group is finitely presented if and only if it has $\mathbb{Z}$ -small circuits, but after more thought, it does not seem unlikely that infinitely presented groups with $\mathbb{Z}$ -small circuits could be constructed. The question, then, is what exactly does $\mathbb{Z}$ -small circuits measure? The case of pure H_1^{uf} -expansion is even more opaque to us.

5 On $BS(2, 3)$

We only focus on the Baumslag-Solitar group $BS(2, 3) = \langle a, b \mid ba^2 = a^3b \rangle$ and the following homomorphism which is surjective, but not injective:

$$\phi : BS(2, 3) \rightarrow BS(2, 3) : \begin{cases} a \mapsto a^2 \\ b \mapsto b \end{cases}.$$

The group element $[a^b, a]$ is non-trivial, and inside the kernel of ϕ . As mentioned before we notice that $BS(2, 3)/\text{Ker}(\phi)$ is isomorphic to $BS(2, 3)$. And thus their Cayley graphs are isomorphic up to change of generators. We will first describe this quotient from a graph-viewpoint. Our other aim is to show that $\text{Ker}(\phi)$ is a free group of infinite rank. The proof is geometric. We will give an explicit set of generators and show that their action on the Bass-Serre tree associated to $BS(2, 3)$ is free. Finally we present how ϕ appears as a morphism on fundamental groups, induced by a continuous map.

5.1 The quotient graph

It is useful to remind ourselves of Britton's lemma. We do not give its general version, but directly apply it to $BS(2, 3)$.

Lemma 5.1.1 (Britton's lemma). *Let $\omega = a^{\alpha_0} \prod_{i=1}^n b^{\beta_i} a^{\alpha_i}$ in $BS(2, 3)$ such that $\beta_i \neq 0$. If $\omega = 1$, then either*

- $n = 0$ and $\alpha_0 = 0$;
- or $n > 0$ and $\exists i \in \{1 \dots n-1\}$ such that
 1. $\beta_i > 0, \beta_{i+1} < 0$ and $\alpha_i \in 2^{\min\{\beta_i, |\beta_{i+1}|\}} \mathbb{Z}$;
 2. $\beta_i < 0, \beta_{i+1} > 0$ and $\alpha_i \in 3^{\min\{|\beta_i|, \beta_{i+1}\}} \mathbb{Z}$.

We first calculate the kernel of ϕ , then remind the reader what the Cayley graph of $BS(2, 3)$ looks like, in order to finally describe its quotient induced by the morphism ϕ .

Definition 5.1.2. Given a word $\omega = a^{\alpha_0} \prod_{i=1}^n a^{\alpha_i} b^{\beta_i}$ in $\{a, b\}^*$, define $\rho(\omega) = \sum_{i=1}^n \beta_i$ and $\rho_a(\omega) = \sum_{i=1}^n |\beta_i|$.

Note that, as a consequence of Britton's lemma, we see that if ω represents the trivial element, then $\rho(\omega) = 0$.

Lemma 5.1.3. $\text{Ker}(\phi)$ is normally generated by $\{[a^b, a], [a^b, a^{-1}], [a^b, a^2], [a^b, a^{-2}]\}$.

Proof. Consider a word ω in the generators of $BS(2, 3)$. If $\omega \in \text{Ker}(\phi)$, then $\rho(\omega) = 0$, since $\rho(\phi(\omega)) = 0$ and ϕ does not affect b . Given that a proof is a sequence of trivialities, we can check the following:

- If $\rho_a(\omega) = 2$, then an easy calculation shows that if $\phi(\omega) = 1$, then ω represents the trivial element.
- If $\rho_a(\omega) = 4$, one can show that ω represents one of the generators given above (up to cyclic permutation of the letters).
- Consider the group element $b^\beta a^\alpha b^{\beta'}$, where $\beta, \beta' > 0$. By Bachet-Bezout, there exist $\lambda, \mu \in \mathbb{Z}$ such that $\lambda 2^\beta + \mu 3^{\beta'} = 1$. Hence $b^\beta a^\alpha b^{\beta'} = b^\beta a^{\alpha \lambda 2^\beta + \alpha \mu 3^{\beta'}} b^{\beta'}$, which equals $a^{\alpha \lambda 3^\beta} b^{\beta + \beta'} a^{\alpha \mu 2^{\beta'}}$.
- By repeatedly applying the previous item, any word ω (up to cyclic permutation) can be replaced by a word ω' that represents the same element in $BS(2, 3)$ and is of the form $a^{\alpha_1} b^{\beta_1} a^{\alpha_2} b^{\beta_2} \dots a^{\alpha_n} b^{\beta_n}$, where none of the exponents are zero and exponents of subsequent b 's have opposite sign.
- We specialise even more, in the sense that we minimise the exponents β_i . This means that $ba^{2\alpha} b^{-1}$ can be replaced by $a^{3\alpha}$ and $b^{-1} a^{3\alpha} b$ by $a^{2\alpha}$.
- So suppose ω is of the prescribed minimal form, such that $\rho_a(\omega) > 4$ and is in $\text{Ker}(\phi)$, then the image is $a^{2\alpha_1} b^{\beta_1} a^{2\alpha_2} b^{\beta_2} \dots a^{2\alpha_n} b^{\beta_n}$. We apply Britton's lemma. Suppose there is some $i \in \{1 \dots n\}$ such that $\beta_i > 0, \beta_{i+1} < 0$ and $2\alpha_i \in 2^{\min\{\beta_i, |\beta_{i+1}|\}} \mathbb{Z}$. Without loss of generality we suppose β_i is the minimum. Then $\alpha_i = 2^{\beta_i - 1} \alpha'_i$, and thus $b^{\beta_i} a^{\alpha_i} b^{\beta_{i+1}} = ba^{3^{\beta_i - 1} \alpha'_i} b^{\beta_{i+1} + \beta_i - 1}$. Note that this contradicts minimality of the word ω if $\beta_i > 1$. The reader can check that a similar contradiction is found for the second case of Britton's lemma.
- Because of the previous point, there is some b whose exponent is either 1 or -1 . Up to taking inverses and cyclicly permuting the word, we may suppose it is -1 and it appears on the second position, i.e. ω can be chosen of the form $b^{\beta_1} a^{\alpha_2} b^{-1} a^{\alpha_3} b^{\beta_3} \dots$
- Adding minimality to the argument, this can be rewritten as

$$b^{\beta_1 - 1} a^{3(\frac{\alpha_2 - 1}{2}) + \alpha_3 - t} bab^{-1} a^t ba^{-1} b^{-1} a^{-t} a^t bab^{\beta_3 - 1} \dots$$

where t is either 1 or 2, such that $\alpha_3 - t \in 3\mathbb{Z}$. Note that the first part of our rewritten word is exactly a generator described in the statement of the lemma. Suppose ω' is the remaining word, after deleting the generator. One notices that $\rho_a(\omega') = \rho_a(\omega) - 2$. Hence an induction argument based on the value of ρ_a finishes the proof.


□

Corollary 5.1.4. $\text{Ker}(\phi)$ is normally generated by $[a^b, a]$.

Proof. Since all other generators can be normally generated by $[a^b, a]$:

- $[a^b, a^{-1}] = a^{-1} [a^b, a]^{-1} a$,
- $[a^b, a^2] = [a^b, a] a [a^b, a] a^{-1}$, and
- $[a^b, a^{-2}] = a^{-1} [a^b, a]^{-1} a^{-1} [a^b, a]^{-1} a^2$.


□

Figure 5.1: Aspects of the Cayley graph of $BS(2, 3)$.

A description of $BS(1, 2)$ and its Cayley graph is a pretty standard introduction to the world of Baumslag-Solitar groups (c.f. [Mei08]). For the purpose of keeping the exposition self-contained, we quickly describe the Cayley graph of $BS(2, 3)$, which follows a slight variation of the classical example $BS(1, 2)$. The standard building block is given in Figure 5.1a, where red edges are labeled by a and blue ones by b . This is the cycle induced by the single relator. These building blocks fit nicely together into an upper-half plane given in Figure 5.1b, based at a copy of $\mathbb{Z}$ generated by a . We may now construct the Cayley graph. Consider the base line generated by a , then we can insert three planes here such that upward arrows do not overlap. One can see this as attaching a plane based at 1 , a and a^2 . In a similar way we construct lower half planes (one can see this by interpreting Figure 5.1b with the baseline on top). Two of these can be attached at $\mathbb{Z}$ such that blue arrows do not overlap. Locally this looks like Figure 5.1c. In general consistently applying this construction to each left coset of $\langle a \rangle$, which all appear as a Cayley graphs of $\mathbb{Z}$, gives us the Cayley graph of $BS(2, 3)$. Now project such that each left coset of $\langle a \rangle$ becomes a point, we obtain a tree of degree five, as can be seen in Figure 5.1d. This is the Bass-Serre tree associated to the HNN extension $BS(2, 3) = \text{HNN}(\mathbb{Z}, 2\mathbb{Z} \sim 3\mathbb{Z})$. Choosing a specific vertex of this tree, one sees that three edges are directed upward and two downward (the halfplanes attached upward and downward respectively). This ‘side view’ will be important for our visual interpretation later. We first define the following sets.

We define the height of an element to be its image by the map

$$\alpha : BS(2, 3) \rightarrow \mathbb{Z} : \begin{cases} a \mapsto 0 \\ b \mapsto 1 \end{cases},$$

Figure 5.2: Colouring of H_λ

and we define what it means for two cosets to be neighbours.

Definition 5.1.5. Two left cosets λ and μ of $\langle a \rangle$ are same height neighbours if for $l \in \lambda$ either la^b , lad^b or la^2a^b are in μ . Notation: $\lambda \sim \mu$.

If we consider the graph defined by those cosets as vertices and edges between two cosets that are same height neighbours, then we obtain a forest (read: disjoint union) of 3-valent trees. Note that such a tree has a unique 2-colouring. This colouring can be seen in Figure 5.2 and thus also the visual interpretation of the same height neighbour relation becomes clear. Presented is a limited sideview of $BS(2, 3)$, where the cosets are coloured by the 2-colouring. Two coloured points (representing left cosets of $\langle a \rangle$) that are two edges apart, are neighbours. From the visual it is clear that a point has exactly three neighbours. Now we can define the following sets:

Definition 5.1.6. Given λ a left coset of $\langle a \rangle$, then

- H_λ is the connected component of λ of the forest induced by the same height neighbour relation;



Since H_λ is a 3-valent tree it has a unique 2-colouring.

- H_λ^+ are the cosets in H_λ , that have the same colour as λ ;
- $H_\lambda^- := H_\lambda \setminus H_\lambda^+$.

Note that if $\mu \in H_\lambda^-$, then clearly $H_\lambda^- = H_\mu^+$.

We now have the tools to show what happens when we quotient by $\text{Ker}(\phi)$. Write $G = BS(2, 3)/\text{Ker}(\phi)$, which we will not consider as an abstract group, but rather as the specific manifestation of $BS(2, 3)$ as a quotient of itself. Note that $\phi(a) = a^2$. This means that $\bar{a}$, in G , will behave like a^2 in $BS(2, 3)$. In a sense we would like to take the root of $\bar{a}$ (note that the reason for this awkward phrasing comes from the visual interpretation we will see soon). This is why we look for a pre-image of a , one can take for example the commutator $[b, a] = bab^{-1}a^{-1}$. So the specific group element $\overline{[b, a]}$ in G will behave as the abstract group element a in $BS(2, 3)$. For brevity we write $\tilde{a} := \overline{[b, a]}$.

We will construct the quotient graph by identifying points in the original graph, following a principle we will call ‘horizontal identification’. Let us see which effect, identifying two points that are at a ‘distance’ $\omega = [a^b, a^{-1}]$ has. We follow the path (in green) of ω in Figure 5.3. Suppose we start on the line $\langle a \rangle$, then we go up a level to $b\langle a \rangle$. Next we move over an edge labelled a , which means that when we take the edge b^{-1} , we do not take the same edge back (i.e. we follow the second green arrow). After this, moving by a^{-1} , we go up one of

Figure 5.3: The identification of the element $[a^b, a]$ Figure 5.4: Two lines are interleaved if at a ‘distance’ $bab^{-1}a^{-1}$.

the two other planes. Then we go down again. Note that since we identify two points of $\langle a \rangle$ and $bab^{-1}a^{-1}ba^{-1}b^{-1}\langle a \rangle$, in fact those lines are completely identified. By symmetry all cosets of $H_{\langle a \rangle}^+$, will become equivalent to $\langle a \rangle$ in G . Note that by symmetry also all lines in $H_{\langle a \rangle}^-$ are identified with one another. Let us consider the behaviour of $\tilde{a}$. Note that it is exactly equal to the the first four letters of ω . This means that we are basically following the first two green arrows of ω in Figure 5.3. Next we apply $\tilde{a}$ a second time, we arrive on the line $\omega\langle a \rangle$, which has been identified to $\langle a \rangle$. Specifically the point we arrive in is identified with a , since $\tilde{a} = \bar{a}$. One notices that viewing G wrt the new generators $\{\bar{b}, \tilde{a}\}$ will give us an intertwining of the lines $\langle a \rangle$ and $bab^{-1}\langle a \rangle$. With respect to the sets $H_{\langle a \rangle}^+$ and $H_{\langle a \rangle}^-$, which are each just one line in G , this means that the generator $\tilde{a}$ will alternate points of $H_{\langle a \rangle}^+$ and $H_{\langle a \rangle}^-$. The result can be seen in Figure 5.5. Here points on the lower line are obtained by alternating points of $\langle a \rangle$ and $[b, a^{-1}]\langle a \rangle$, where we start by 1 and $[b, a]$ respectively. Out of this new ‘baseline’ once again three planes open up. We describe one, since the other two are analogous by symmetry. Consider the plane based at $\bar{1}$ in Figure 5.5. We take one step back and look at Figure 5.4, before identification. Then we see that the intertwining respects the upper brown points, in the sense that if the first green point from the left is 1, then the first yellow point from the left is $[b, a]$. This means that the fourth point is exactly $[b, a]a$. Which is underneath the second brown point. This means that, after identification, a position opens up in between consecutive brown points. This makes sense, since by symmetry of the cayley graph, also the line $H_{b\langle a \rangle}^+$, must intertwine with $H_{b\langle a \rangle}^-$. This fills up the gap. Note that points in $H_{b\langle a \rangle}^-$ are exactly the orange points in Figure 5.5.

Note that since $\text{Ker}(\phi)$ is normally generated by $[a^b, a]$, this is the only ‘real’ identification that takes place. Other identifications are due to taking conjugates of this element. I.e. if two cosets of $\langle a \rangle$ are identified, of course the tree structure above and below it must become




Figure 5.5: A part of the Cayley graph of G , where $\tilde{a}$ is now a generator. Choices of colours are consistent with previous drawings.

identified too.

5.2 A free group of infinite rank

Let us now prove that the kernel of ϕ is a free group of infinite rank. In order to do this in an accessible way, we first introduce some terminology. For illustrations of the concepts, also see Figure 5.6.

Definition 5.2.1. Let $\mathcal{P}$ be the set of geodesic paths starting at $\langle a \rangle$ in the Bass-Serre tree in Figure 5.1d (these are the left cosets of $\langle a \rangle$). Let $\{a, b\}^*$ be all words in a, b and their inverses, and $n, m \in \mathbb{Z}$. We define:

1. Let $\omega \in \{a, b\}^*$, then the associated path in the Bass-Serre tree is the projection of the associated path in the Cayley graph of $BS(2, 3)$. This path may contain backtracking.
2. A good representative of a path $p \in \mathcal{P}$ is a word ω such that p is its associated path in the Bass-Serre tree. Equivalently: such that $\rho_a(\omega)$ is minimal and ω represents the left-coset at the end of p .
3. A path contains a tip if a good representative ω contains a subword of the form $ba^n b^{-1}$.
4. A path contains a valley if a good representative ω contains a subword of the form $b^{-1} a^n b$.
5. A path is end-essential if all good representatives do not end in a tip.
6. A path is swiss if a good representative ω contains a subword of the form $b^{-1} a^n b a^m b^{-1}$ or $ba^n b^{-1} a^m b$. A path that is not swiss is called nepalese.
7. Let ω be a good representative of $p \in \mathcal{P}$, an end-essential nepalese path with a tip. Let ω_1 be the first part of ω , let B be the tip and ω_2 the remainder. Then p has two triplets at B , namely the paths associated to $\omega_1 a B a^{-1} \omega_2$ and $\omega_1 a^{-1} B a \omega_2$.
8. Let ω be a good representative of $p \in \mathcal{P}$, an end-essential nepalese path with a valley. Let ω_1 be the first part of ω , let V be the valley and ω_2 the remainder. Then p has a twin at V , namely the path associated to $\omega_1 a^{-1} V a b^{-1} \omega_2$.




Figure 5.6: An artistic interpretation of a swiss and nepalese path. The reader should think of these paths as paths in the Bass-Serre tree 5.1d. Note that the swiss path has valleys and tips that share an edge, while for the nepalese path they are disjoint. Moreover, in this example, the nepalese path is end-essential while the swiss path is not.

9. Let $\sim$ be the sibling relation defined on end-essential nepalese paths, induced by setting $p \sim q$ if they belong to the same set of triplets or the same set of twins.
10. We denote by $G_\sim$ the graph induced by the sibling relation on end-essential nepalese paths. Two paths in the same connected component are called relatives.
11. For a path $p \in \mathcal{P}$ define the function $c : \mathcal{P} \rightarrow \mathbb{N}$ by setting $c(p)$ equal to the number of valleys and tips.

We note that the sibling relation is not an equivalence relation, since it is not transitive and not reflexive. But since it is symmetric $\sim$ clearly defines a graph on the set of end-essential nepalese paths. This graph is not connected. We show the following.

Theorem 5.2.2. $\text{Ker}(\phi) \simeq \mathbb{F}_\infty$. In particular, let O be a set of words in $\{a, b\}^*$ such that

- Each word represents an end-essential nepalese path and each such path is represented by at most one word.
- For each connected component of $G_\sim$, there is exactly one path that is represented by a word in O .

Then the symmetric set

$$\mathcal{S} = \{\omega a^i [a^b, a]^j a^{-i} \omega^{-1} \mid j \in \{-1, 1\}, i \in \{0, 1\}, \omega \in O\}$$

freely generates $\text{Ker}(\phi)$. In particular this shows that the action of $\text{Ker}(\phi)$ on the Bass-Serre tree is free.

Proof. Generation. We first show that $\mathcal{S}$ is generating. Let $\tilde{O}$ be a set containing for each geodesic path in the Bass-Serre tree exactly one word which is a good representative. Then it can easily be seen that $\text{Ker}(\phi)$ is generated by

$$\mathcal{S} = \{\omega a^i [a^b, a]^j a^{-i} \omega^{-1} \mid j \in \{-1, 1\}, i \in \{0, 1\}, \omega \in \tilde{O}\}.$$

Hence we can reduce the number of paths needed. Suppose ω represents a swiss path, then it is of the form $\omega_1 bab^{-1}ab\omega_2$ or $\omega_1 b^{-1}abab^{-1}\omega_2$ (up to possibly taking other exponents of the a 's). We treat these two cases for $i = 0$ and $j = 1$, the other cases are similar. Note that for the first calculation we consider the conjugated group element and in the second we only consider the group element representing the path. As we are conjugating the paths this leads to redundant notation. We write the first in full for the readers convenience.

$$\begin{aligned}\omega[a^b, a]\omega^{-1} &= \omega_1 bab^{-1}ab\omega_2[a^b, a]\omega_2^{-1}b^{-1}a^{-1}ba^{-1}b^{-1}\omega_2 \\ &= \omega_1[a^b, a]\omega_1^{-1}\omega_1aba\omega_2[a^b, a]\omega_2^{-1}a^{-1}b^{-1}a^{-1}\omega_1^{-1}\omega_1[a^b, a]^{-1}\omega_1^{-1}, \text{ or} \\ \omega &= \omega_1 b^{-1}abab^{-1}\omega_2 \\ &= \omega_1 a^{-1}b^{-1}[a^b, a]ba\omega_1^{-1}\omega_1 a^{-1}b^{-1}a^4\omega_2.\end{aligned}$$

Hence we have reduced the problem to the shorter paths ω_1 and $\omega_1 aba\omega_2$, and $\omega_1 a^{-1}b^{-1}$ and $\omega_1 a^{-1}b^{-1}a^4\omega_2$ respectively. An induction hypothesis on the length of the paths finishes the job.

Hence we may assume our paths are nepalese. Suppose they are not end-essential. So they are represented by a word ω of the form $\omega_1 bab^{-1}$. Note that the element $bab^{-1}a^i[a^b, a]a^{-i}ba^{-1}b^{-1}$ is in the subgroup $\mathbb{F}_2$ generated by $\{[a^b, a], a[a^b, a]a^{-1}\}$. Hence, because of its specific form, it can be written as one of the generators (or their inverses), or a product of at most four generators (or their inverses). Say g_1 or $g_1g_2g_3g_4$. In particular we can rewrite $\omega a^i[a^b, a]a^{-i}\omega^{-1}$ as $\omega_1 g_1 \omega_1^{-1}$ or $\omega_1 g_1 \omega_1^{-1} \omega_1 g_2 \omega_1^{-1} \omega_1 g_3 \omega_1^{-1} \omega_1 g_4 \omega_1^{-1}$, where ω_1 is an end-essential path.

Finally consider ω , a good representative for an end-essential nepalese path p . Let $\omega \notin O$. We will show that ω can be written as a product of one (any) of its siblings and paths q such that $c(q) < c(p)$. Consider two triplets $\omega_1 B\omega_2$ and $\omega_1 aBa^{-1}\omega_2$, where B is a tip. Then

$$\begin{aligned}\omega &= \omega_1 B\omega_2 \\ &= \omega_1 BaB^{-1}a^{-1}\omega_1^{-1}\omega_1 aBa^{-1}\omega_2 \\ &= \omega_1[a^b, a]\omega_1^{-1}\omega_1(aBa^{-1}\omega_2), \text{ or} \\ \omega &= \omega_1 B\omega_2 \\ &= (\omega_1 a^{-1})aBa^{-1}B^{-1}(a\omega_1^{-1})\omega_1 a^{-1}Ba\omega_2 \\ &= (\omega_1 a^{-1})[a^b, a]^{-1}(a\omega_1^{-1})(\omega_1 a^{-1}Ba\omega_2).\end{aligned}$$

Clearly $c(\omega_1) < c(\omega)$. On the other hand consider two twins $\omega_1 Vb^{-1}\omega_2$ and $\omega_1 a^{-1}Vab^{-1}\omega_2$. Note that the valley can either be of the form $b^{-1}ab$ or bab^{-1} . We suppose it is the former, the argument for the latter is similar. Note also Figure 5.7a for a visual interpretation. One has that

$$\begin{aligned}\omega &= \omega_1 b^{-1}ab\omega_2 \\ &= (\omega_1 a^{-1}b^{-1})bab^{-1}aba^{-1}b^{-1}a^{-1}(ba\omega_1^{-1})\omega_1 a^{-1}b^{-1}aba\omega_2 \\ &= (\omega_1 a^{-1}b^{-1})[a^b, a](ba\omega_1^{-1})\omega_1 a^{-1}b^{-1}aba\omega_2.\end{aligned}$$

The case $b^{-1}a^{-1}b$ is similar, obtaining a conjugate of $[a^b, a]^{-1}$ in the first term. Since each connected component in $G_\sim$ is represented in O by some element, each end-essential nepalese




Figure 5.7: On the left: The black path (corresponding to ω) has a valley V . Its twin at V is given by the blue path. Note that the red paths corresponds to $\omega_1 a^{-1} b^{-1}$ and $ba\omega_1^{-1}$, and the orange path to the normal generator $[a^b, a]$. On the right: An element in $\text{Ker}(\phi)$, red paths are the paths defined by ω , the orange path corresponds to $[a^b, a]$, the second edge E of the commutator path separates the Bass-Serre tree in two components.

path can be written as a product of the designated relative in $\mathcal{O}$ and elements with less tips or valleys. One ends the argument by induction on the number of tips and valleys.

Freeness. We are now ready to prove the given basis is free. Let $\omega \in \mathcal{O}$, then $\omega a^i [a^b, a] a^{-i} \omega^{-1}$ does not necessarily represent a geodesic path (we call the path p), it may have backtracking when the last edge of ω and the first of $[a^b, a]$ overlap. Similarly it might also overlap on the last edge of $[a^b, a]$ and the first of ω^{-1} . Let E be the second edge of $[a^b, a]$ as a subpath of p (see Figure 5.7b), then it separates the tree in two connected components (say E^- and E^+). We claim that applying any other element $\tilde{\omega} a^s [a^b, a] a^{-s} \tilde{\omega}^{-1}$ of $\mathcal{S}$, one will never traverse E again. There are several possibilities.

- Firstly suppose $\tilde{\omega} = \omega$, then this reduces to the fact that $\langle [a^b, a], a[a^b, a]a^{-1} \rangle$ is free on two generators.
- Consider the concatenated path $\omega a^i [a^b, a] a^{-i} \omega^{-1} \tilde{\omega} a^s [a^b, a] a^{-s} \tilde{\omega}^{-1}$. Since $\tilde{\omega}$ is nepalese and end-essential, this path remain in E^+ , once having crossed E . Let $E_{\tilde{\omega}}$ be the second edge in the second commutator $[a^b, a]$. Suppose $E_{\tilde{\omega}}$ is not part of the path defined by ω^{-1} . In particular, to cross E , we first need to cross $E_{\tilde{\omega}}$, which means we have merely displaced the problem.
- So we may suppose $E_{\tilde{\omega}}$ is part of ω^{-1} . Note that this is only possible when ω^{-1} contains a tip or a valley. In this case the composition $\omega^{-1} \tilde{\omega} [a^b, a] \tilde{\omega}^{-1}$ defines a sibling of ω^{-1} . It is impossible, using a sequence of siblings, to cross E , since for every connected component

of the sibling relation we only allowed one path to be contained in O , which is ω .

□

Remark 5.2.3. Note that if one is willing to give up the generating set, then there is a quick and clean argument to show that $\text{Ker}(\phi)$ acts freely on the Bass-Serre tree. The stabiliser of each vertex is given by a conjugate of $\langle a \rangle$. So to show that $\text{Ker}(\phi)$ acts freely, it is enough to show that the intersection with each of the stabilisers is trivial. Given that $\text{Ker}(\phi)$ is normal, it suffices to show that its intersection with $\langle a \rangle$ is trivial. Since a^n is mapped to a^{2^n} this is true.

Furthermore we look at the limit case of applying the morphism ϕ .

Proposition 5.2.4. $\text{Ker}(\phi^n)$ is normally generated by $\{[b^m ab^{-m}, a] \mid 0 < m \leq n\}$

Proof. By induction. □

Corollary 5.2.5. $BS(2, 3) / \bigcup_n \text{Ker}(\phi^n) = \langle a, b \mid ba^2b^{-1}a^{-3}, [b^m ab^{-m}, a] \text{ where } m \in \mathbb{N} \rangle$.

Call this group L . Note that $\forall n \in \mathbb{N}$ there is an element $\alpha \in L$, such that $\alpha^{2^n} = \bar{a}$. I.e $\bar{a}$ has infinitely many roots. More specifically the presentation induced by α and $\bar{b}$ is the same as the one given in Corollary 5.2.5. One sees that the group generated by $\langle \bar{b}^m ab^{-m} \mid m \in \mathbb{N} \rangle$ is isomorphic to the dyadic rationals $\mathbb{Z}[1/2]$.

5.3 Another interpretation of non-Hopfianness and a continuous map

We present another interpretation on the non-Hopfianness of $BS(2, 3)$, which was communicated to the author by Gilbert Levitt. In some sense, this is a more topological version of Remark 6.10 in [Lev15]. The paper treats generalised Baumslag-Solitar groups, which are fundamental groups of graphs of groups. This is also the setting of Remark 6.10. These groups can be retrieved as the fundamental groups of certain 2-dimensional simplicial complexes (see [DRT11, Section 7]). We will not introduce graphs of groups and associated concepts, but will rather immediately take the topological point of view adapted to $BS(2, 3)$. As a last remark note that the forthcoming example is not particular to the case of $BS(2, 3)$, but can easily be altered to apply to other Baumslag-Solitar groups. The same is true for Figure 5.8.

First as small detour. We note the following fact: in $BS(2, 3)$, the elements a^2 and a^3 are “the same” since they are conjugates. What we mean is that they behave in the same way, which is made precise in the following lemma.

Lemma 5.3.1. *The group $BS(2, 3)$ has an abstract group presentation, where the generators can be interpreted as either $\{a^2, b\}$ or $\{a^3, b\}$. This presentation is given by $\langle \lambda, \mu \mid \lambda^2 = \mu\lambda\mu^{-1}\lambda^{-1}\mu\lambda\mu^{-1} \rangle$.*

Proof. Starting from the usual presentation $\langle a, b \mid ba^2b^{-1} = a^3 \rangle$, one can apply two sequences




Figure 5.8: $BS(2, 3)$ as a fundamental group and ϕ as a morphism induced by a continuous self-map (in red) of the space B .

of Tietze transformations to obtain the presentation above. We present the calculations for a^2 :

$$\begin{aligned}
 \langle a, b \mid ba^2b^{-1} = a^3 \rangle &= \langle a, \lambda, b \mid \lambda = a^2, ba^2b^{-1} = a^3 \rangle \\
 &= \langle a, \lambda, b \mid \lambda = a^2, ba^2b^{-1} = a^3, a = b\lambda b^{-1}\lambda^{-1} \rangle \\
 &= \langle \lambda, b \mid \lambda = b\lambda b^{-1}\lambda^{-1}b\lambda b^{-1}\lambda^{-1}, b\lambda b^{-1} = b\lambda b^{-1}\lambda^{-1}\lambda \rangle \\
 &= \langle \lambda, b \mid \lambda^2 = b\lambda b^{-1}\lambda^{-1}b\lambda b^{-1} \rangle \\
 &= \langle \lambda, \mu \mid \lambda^2 = \mu\lambda\mu^{-1}\lambda^{-1}\mu\lambda\mu^{-1} \rangle.
 \end{aligned}$$

We can obtain the same presentation for a^3 by doing similar calculations. □

We define ϕ' to be the morphism

$$BS(2, 3) \rightarrow BS(2, 3) : \begin{cases} a \mapsto a^3 \\ b \mapsto b \end{cases}.$$

From a group presentation point of view the following is true: since a^2 and a^3 behave in a completely similar way the morphisms ϕ and ϕ' are in fact also the same morphism! As a consequence they have the same kernel.

We start by describing a space B such that $BS(2, 3)$ is its fundamental group. The basic building blocks are two circles S_0 and S_1 , and two annuli $A_0 := S_a^0 \times [0, 1]$ and $A_1 := S_a^1 \times [1, 2]$. Now we glue the blocks together. For A_0 we attach $S_a^0 \times 0$ to S_0 along the identity map $e^{i\theta} \mapsto e^{i\theta}$. We do the same for $S_a^0 \times 1$ to S_1 . For A_1 we attach $S_a^1 \times 1$ to S_0 along the map $e^{i\theta} \mapsto e^{3i\theta}$, and $S_a^1 \times 2$ to S_1 along the map $e^{i\theta} \mapsto e^{2i\theta}$. This space is depicted two times in Figure 5.8, being once given by the outer annuli and circles, and a second time by the inner annuli and circles.

The reader can easily verify that the fundamental group of B at $(1, \frac{1}{2})$ is generated by two loops

$$\begin{aligned} \gamma_1 : [0, 1] \rightarrow B : \theta &\mapsto \left(e^{2\pi i \theta}, \frac{1}{2} \right) \\ \gamma_2 : [0, 1] \rightarrow B : x &\mapsto \begin{cases} (1, \frac{1}{2} - 2x) & \text{if } x \in [0, \frac{1}{4}] \\ (1, 1 + 2x - \frac{1}{2}) & \text{if } x \in [\frac{1}{4}, \frac{3}{4}] \\ (1, \frac{5}{2} - 2x) & \text{if } x \in [\frac{3}{4}, 1] \end{cases}, \end{aligned}$$

where $\pi_1(B, (1, \frac{1}{2}))$ is isomorphic to $BS(2, 3)$ by mapping a to $[\gamma_1]$ and b to $[\gamma_2]$.

Finally the map $\tilde{\phi} : B \rightarrow B$ is given in red in Figure 5.8. We define the map piecewise

- $A_0 \rightarrow A_1 : (e^{i\theta}, x) \mapsto (e^{i\theta}, 2 - x)$
- $A_1 \rightarrow A_0 : (e^{i\theta}, x) \mapsto (e^{6i\theta}, 2 - x)$
- $S_0 \rightarrow S_1 : e^{i\theta} \mapsto e^{2i\theta}$
- $S_1 \rightarrow S_0 : e^{i\theta} \mapsto e^{3i\theta}$.

It is easy to verify that this map respects the gluing of the building blocks, hence is continuous on B . In particular $\tilde{\phi}$ induces a map on the fundamental group of B , being $\tilde{\phi}_* : \pi_1(B, (1, \frac{1}{2})) \rightarrow \pi_1(B, (1, \frac{3}{2}))$. One could check that $\text{Ker}(\tilde{\phi}_*)$ is exactly $\text{Ker}(\phi)$ (and thus also $\text{Ker}(\phi')$). So from a group presentation point of view, we can say that $\tilde{\phi}$ induces the group morphism ϕ . However, suppose we really want to recover ϕ , then we would need a map from $\pi_1(B, (1, \frac{1}{2}))$ to itself. Consider the path

$$p^1 : [0, 1] \rightarrow B : x \mapsto \begin{cases} (1, \frac{1}{2} + x) & \text{if } x \in [0, \frac{1}{2}] \\ (1, 2 - x) & \text{if } x \in [\frac{1}{2}, 1] \end{cases}.$$

This path imposes an isomorphism $p_*^1 : \pi_1(B, (1, \frac{3}{2})) \rightarrow \pi_1(B, (1, \frac{1}{2}))$. Now composing, we obtain that $p_*^1 \circ \tilde{\phi}_*$ is exactly ϕ . If, on the other hand, we want to recover ϕ' , we only need to compose $\tilde{\phi}_*$ with p_*^2 , where p^2 is the path given by

$$p^2 : [0, 1] \rightarrow B : x \mapsto \begin{cases} (1, \frac{1}{2} - x) & \text{if } x \in [0, \frac{1}{2}] \\ (1, \frac{1}{2} + x) & \text{if } x \in [\frac{1}{2}, 1] \end{cases}.$$

List of Figures

1.1	The Cayley graph of $\mathbb{F}_2$ with standard generators.	5
1.2	Two Cayley graphs of $\mathbb{Z}$ (labels omitted).	5
1.3	The free group of rank two: paradoxical sets	9
1.4	An infinite line graph added to a rooted binary tree.	10
1.5	A proof that $\mathbb{F}_2$ has property A	11
1.6	The comb C_{10} with 10 teeth.	15
1.7	A sofic approximation of $\mathbb{Z}$, as defined in Example 1.5.2	16
1.8	Description of the Benjamini-Schramm limit of $\{B_r(\mathbb{F}_2, e)\}_{r \in \mathbb{N}}$	17
1.9	Construction of the graph G_n on the vertex set $K_n \cup \partial(K_n)$	29
3.1	Construction of edges.	36
3.2	Construction of edges of type three (b).	36
3.3	Construction of $\mathcal{H}$	38
3.4	Sketch of what happens in $N_k \setminus \Gamma$	41
3.5	Deleting g -edges.	42
3.6	Construction of R_a for a vertex that is in $\mathcal{G}$, but not in $\mathcal{G}'$	46
4.1	Transforming a ULF tree into a trivalent one.	56
4.2	Example of the decomposition of a tree and the associated sets of bi-infinite paths.	58
4.3	Local configurations of bi-infinite paths.	61
4.4	A bi-infinite comb with a choice of bips such that the associated function on paths is unbounded.	62
5.1	Aspects of the Cayley graph of $BS(2, 3)$	87
5.2	Colouring of H_λ	88
5.3	The identification of the element $[a^b, a]$	89
5.4	Two lines are interleaved if at a ‘distance’ $bab^{-1}a^{-1}$	89
5.5	A part of the Cayley graph of G , where $\tilde{a}$ is now a generator	90
5.6	Examples of swiss and nepalese paths	91
5.7	Examples of twin paths and a general group element	93
5.8	$BS(2, 3)$ as a fundamental group and ϕ as a morphism induced by a continuous self-map of the space B	95

Bibliography

- [AFS19] Vadim Alekseev and Martin Finn-Sell. Sofic boundaries of groups and coarse geometry of sofic approximations. *Groups Geom. Dyn.*, 13(1):191–234, 2019.
- [AT17] Miklós Abért and László Tóth. Uniform rank gradient, cost, and local-global convergence. *arXiv preprint arXiv:1710.10431*, 2017.
- [ATV13] Miklós Abért, Andreas Thom, and Bálint Virág. Benjamini-Schramm convergence and pointwise convergence of the spectral measure. 2013. Preprint available on www.renyi.hu/~abert/luckapprox.pdf.
- [BD15] Matthias Blank and Francesca Diana. Uniformly finite homology and amenable groups. *Algebraic & Geometric Topology*, 15(1):467–492, 2015.
- [BdlHV08] Bachir Bekka, Pierre de la Harpe, and Alain Valette. *Kazhdan’s property (T)*, volume 11 of *New Mathematical Monographs*. Cambridge University Press, Cambridge, 2008.
- [BG11] Henning Bruhn and Agelos Georgakopoulos. Bases and closures under infinite sums. *Linear Algebra and its Applications*, 435:2007–2018, 10 2011.
- [BK20] Rémi Bottinelli and Tom Kaiser. The first uniformly finite homology group with coefficients in $\mathbb{Z}$ and a characterisation of its vanishing in the transitive case. *arXiv preprint arXiv:2001.04857*, 2020.
- [Bol98] Béla Bollobás. *Modern graph theory*, volume 184 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1998.
- [BS97] Itai Benjamini and Oded Schramm. Every graph with a positive cheeger constant contains a tree with a positive cheeger constant. *Geometric & Functional Analysis GAFA*, 7(3):403–419, Jul 1997.
- [BS01] Itai Benjamini and Oded Schramm. Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.*, 6:no. 23, 13, 2001.
- [BT24] Stefan Banach and Alfred Tarski. Sur la décomposition des ensembles de points en parties respectivement congruentes. *Fund. math*, 6(1):244–277, 1924.
- [BW92] Jonathan Block and Shmuel Weinberger. Aperiodic tilings, positive scalar curvature, and amenability of spaces. *Journal of the American Mathematical Society*, 5(4):907–918, 1992.
- [BW97] Jonathan Block and Shmuel Weinberger. Large scale homology theories and geometry. In *Geometric topology (Athens, GA, 1993)*, volume 2 of *AMS/IP Stud. Adv. Math.*, pages 522–569. Amer. Math. Soc., Providence, RI, 1997.

- [Car15] Alessandro Carderi. Ultraproducts, weak equivalence and sofic entropy. *arXiv preprint arXiv:1509.03189*, 2015.
- [CGdLS18] Alessandro Carderi, Damien Gaboriau, and Mikael de La Salle. Non-standard limits of graphs and some orbit equivalence invariants. *arXiv preprint arXiv:1812.00704*, 2018.
- [Dia15] Francesca Diana. *Aspects of uniformly finite homology*. PhD thesis, 2015.
- [DK] Thiebaut Delabie and Ana Khukhro. Coarse fundamental groups and box spaces. *Proceedings of the Royal Society of Edinburgh: Section A Mathematics*, page 1–16.
- [DN17] Francesca Diana and Piotr W. Nowak. Eilenberg swindles and higher large scale homology of products of trees. *Groups Geom. Dyn.*, 11(1):371–392, 2017.
- [DRT11] Alberto L. Delgado, Derek J. S. Robinson, and Mathew Timm. Generalized Baumslag-Solitar groups and geometric homomorphisms. *J. Pure Appl. Algebra*, 215(4):398–410, 2011.
- [Ele06] Gábor Elek. The strong approximation conjecture holds for amenable groups. *J. Funct. Anal.*, 239(1):345–355, 2006.
- [Ele07] Gábor Elek. The combinatorial cost. *Enseign. Math. (2)*, 53(3-4):225–235, 2007.
- [ES04] Gábor Elek and Endre Szabó. Sofic groups and direct finiteness. *J. Algebra*, 280(2):426–434, 2004.
- [ES06] Gábor Elek and Endre Szabó. On sofic groups. *J. Group Theory*, 9(2):161–171, 2006.
- [FM77] Jacob Feldman and Calvin C. Moore. Ergodic equivalence relations, cohomology, and von Neumann algebras. I. *Trans. Amer. Math. Soc.*, 234(2):289–324, 1977.
- [Gab98] Damien Gaboriau. Mercuriale de groupes et de relations. *C. R. Acad. Sci. Paris Sér. I Math.*, 326(2):219–222, 1998.
- [Gab00] Damien Gaboriau. Coût des relations d’équivalence et des groupes. *Invent. Math.*, 139(1):41–98, 2000.
- [Gab10] Damien Gaboriau. Orbit equivalence and measured group theory. In *Proceedings of the International Congress of Mathematicians. Volume III*, pages 1501–1527. Hindustan Book Agency, New Delhi, 2010.
- [Got73] Walter Gottschalk. Some general dynamical notions. In *Recent advances in topological dynamics (Proc. Conf. Topological Dynamics, Yale Univ., New Haven, Conn., 1972; in honor of Gustav Arnold Hedlund)*, pages 120–125. Lecture Notes in Math., Vol. 318, 1973.
- [Gro99] Mikhaïel Gromov. Endomorphisms of symbolic algebraic varieties. *J. Eur. Math. Soc. (JEMS)*, 1(2):109–197, 1999.

- [HLS14] Hamed Hatami, László Lovász, and Balázs Szegedy. Limits of locally-globally convergent graph sequences. *Geom. Funct. Anal.*, 24(1):269–296, 2014.
- [HP18] Tom Hutchcroft and Gábor Pete. Kazhdan groups have cost 1. *arXiv preprint arXiv:1810.11015*, 2018.
- [Jus] Kate Juschenko. Sofic groups. Lecture notes. <https://web.ma.utexas.edu/users/juschenko/files/soficgroups.pdf>.
- [Kai19] Tom Kaiser. Combinatorial cost: a coarse setting. *Trans. Amer. Math. Soc.*, 372(4):2855–2874, 2019.
- [Kaz67] David A. Kazhdan. On the connection of the dual space of a group with the structure of its closed subgroups. *Funkcional. Anal. i Prilozhen.*, 1:71–74, 1967.
- [Kun16] Gabor Kun. On sofic approximations of property (t) groups. *arXiv preprint arXiv:1606.04471*, 2016.
- [KV17] Ana Khukhro and Alain Valette. Expanders and box spaces. *Adv. Math.*, 314:806–834, 2017.
- [Lev95] Gilbert Levitt. On the cost of generating an equivalence relation. *Ergodic Theory Dynam. Systems*, 15(6):1173–1181, 1995.
- [Lev15] Gilbert Levitt. Quotients and subgroups of Baumslag-Solitar groups. *J. Group Theory*, 18(1):1–43, 2015.
- [LP16] Russell Lyons and Yuval Peres. *Probability on trees and networks*, volume 42 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, New York, 2016.
- [Mar73] Grigoriĭ A. Margulis. Explicit constructions of expanders. *Problemy Peredazhi Informacii*, 9(4):71–80, 1973.
- [Mei08] John Meier. *Groups, graphs and trees*, volume 73 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 2008. An introduction to the geometry of infinite groups.
- [Mos03] Lee Mosher. Homology and dynamics in quasi-isometric rigidity of once-punctured mapping class groups, 2003.
- [OW80] Donald S. Ornstein and Benjamin Weiss. Ergodic theory of amenable group actions. I. The Rohlin lemma. *Bull. Amer. Math. Soc. (N.S.)*, 2(1):161–164, 1980.
- [Pat88] Alan L. T. Paterson. *Amenability*, volume 29 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 1988.
- [Pes90] Georg Peschke. The theory of ends. *Nieuw Arch. Wisk.(4)*, 8(1):1–12, 1990.
- [Pes08] Vladimir G. Pestov. Hyperlinear and sofic groups: a brief guide. *Bull. Symbolic Logic*, 14(4):449–480, 2008.

- [Roe03] John Roe. *Lectures on coarse geometry*, volume 31 of *University Lecture Series*. American Mathematical Society, Providence, RI, 2003.
- [Sab] Mostafa Sabri. Benjamini-schramm convergence of graphs.
- [Sta68] John R. Stallings. On torsion-free groups with infinitely many ends. *Ann. of Math.* (2), 88:312–334, 1968.
- [SW79] Peter Scott and Terry Wall. Topological methods in group theory. In *Homological group theory (Proc. Sympos., Durham, 1977)*, volume 36 of *London Math. Soc. Lecture Note Ser.*, pages 137–203. Cambridge Univ. Press, Cambridge-New York, 1979.
- [Wei00] Benjamin Weiss. Sofic groups and dynamical systems. volume 62, pages 350–359. 2000. *Ergodic theory and harmonic analysis* (Mumbai, 1999).
- [Wil09] Rufus Willett. Some notes on property A. In *Limits of graphs in group theory and computer science*, pages 191–281. EPFL Press, Lausanne, 2009.
- [Wil11] Rufus Willett. Property A and graphs with large girth. *J. Topol. Anal.*, 3(3):377–384, 2011.
- [WY14] Rufus Willett and Guoliang Yu. Geometric property (T). *Chin. Ann. Math. Ser. B*, 35(5):761–800, 2014.
- [Zel11] Yevhen G. Zelenyuk. *Ultrafilters and topologies on groups*, volume 50 of *De Gruyter Expositions in Mathematics*. Walter de Gruyter GmbH & Co. KG, Berlin, 2011.