Topics in graph limits, Wasserstein isometries, and discrete harmonic functions

Aranka Hrušková

Supervisor: Dr habil. Dávid Kunszenti-Kovács

Department of Mathematics and Its Applications
Central European University

This dissertation is submitted in partial fulfilment
of the requirements for the degree of
Doctor of Philosophy

Budapest, Hungary
June 2024

Contents

Chapter 1

Introduction

The thesis consists of three parts. The first, developing various topics in the theory of graph limits, forms the bulk of the work. It is based on five papers, of which two are published, two are about to be submitted, and one is a draft in preparation. The two published papers prove that certain combinatorial constructions on infinite graphs with bounded maximum degree can be obtained by local algorithms, meaning in particular that their approximate versions can be built along any graph sequence converging to such an infinite graph. The subsequent two papers focus on graph convergence in the intermediate regime, that is in the case when the number of edges grows superlinearly but subquadratically in the number of vertices along a given graph sequence. The draft that closes the first part approaches graph limits through the language of exchangeability and proves [...] .

The second part, based on a paper under review, addresses the question of existence of isometries of the set of probability measures on a metric space which do not arise from an isometry of the metric space itself. We focus on the particular case when the metric space is a sphere and show that in that scenario, there are no such exotic isometries. The third and newest part investigates phenomena related to random walks on countable discrete groups. We ask for which bounded functions on Γ there is a probability measure on the group with respect to which the function is harmonic, conjecture that changing an originally μ-harmonic function at exactly one point produces a non-harmonisable function, and prove the conjecture in a special case. After that, we also extend the classical theorem that the values of a harmonic function converge almost surely along the trajectory of a random walk to say that the random harmonic function obtained by shifting our coordinates by the trajectory of the random walk converges pointwise almost surely to a constant function.

Chapter 2

Part I: Graph limits

Let us suppose that we have a sequence $(G_n)_{n=1}^{\infty}$ of finite simple graphs with growing number of vertices. Is there a way to tell whether as n goes to infinity, the graphs look more and more alike? And if yes, what is it?

The answer depends on what exactly it is that we consider important about graphs. In recent years, when we often store and organise large datasets as graphs, a multitude of approaches to the questions above has been emerging. In general, we divide graph sequences to *dense* and *sparse*, where the former are those for which

$$\liminf_{n\to\infty} \frac{e(G_n)}{\binom{v(G_n)}{2}} > 0,$$

that is a positive fraction of the $\binom{v(G_n)}{2}$ possible edges is always present, while under sparse sequences we understand those in which the edge density converges to zero.

Whatever graph sequence $(G_n)_{n=1}^{\infty}$ we have at hand, we can always fix a positive integer k and for each n, sample uniformly at random k vertices from G_n. These induce a subgraph of G_n, meaning we just obtained a probability distribution on the finite set of simple graphs on k vertices. If for every fixed k, these sampling distributions converge, we say that the sequence converges in the dense sense. The reason for such name is that applying this process to a sparse sequence gives us no structural information other than exactly the fact that the sequence is sparse. In other words, this convergence notion trivializes for sparse sequences since for every k, the distribution on induced k-vertex subgraphs converges to the Dirac mass at the empty graph.

The upside of the theory of dense convergence is the neat limit object called *graphon*. Introduced by Lovász and Szegedy in [3], graphons are measurable functions

$$W\colon [0,1]^2 \longrightarrow [0,1]$$

such that $W(x,y) = W(y,x)$ for almost all $(x,y) \in [0,1]^2$.

On the other side of the density spectrum, we have graph sequences with uniformly bounded maximum degree, for example the sequence $(C_n)_{n=3}^{\infty}$ of cycles or the sequence $(P_n)_{n=2}^{\infty}$ of paths. We can employ a different sampling procedure now, in which k will become the depth of exploration. Concretely, we pick one vertex uniformly at random and consider the subgraph of G_n induced by all the vertices in the ball of radius k around our random vertex. Since there is a uniform upper bound, say Δ, on the maximum degree of any vertex, we again arrive to a scenario in which we have, for every k, probability distributions on the same finite set of graphs, and are asking whether these converge. This is called *local* or *Benjamini-Schramm* convergence, for which probabilistic objects called *graphings* serve as the limit objects – see Chapter 18 in [2] for their precise definition.

In both the cases described above, once we have a limit object, we can deduce from it facts about any graph sequence converging to them. A particular example of this phenomenon in the sparse regime are so-called local algorithms: if we can construct a combinatorial structure on a graphing by a local algorithm, then we can build at least nearly optimal such structures also along any graphs sequence converging to this graphing. Examples of structures that we might have in mind include independent sets witnessing a particular independence ratio, perfect matchings and similar. In the first two appendices, we tackle the question of constructing Schreier decorations, that is, we want to find decorations of the edges of infinite $2d$-regular graphs with arrows and d colours which specify an action of the free group $\mathbb{F}_d$ on the vertices of the graph which respects the adjacency relation.

Outside of the two extreme worlds of the dense sequences on one hand and sequences with uniformly bounded maximum degree on the other, we are left with an abundance of sequences of intermediate edge density. The flagship example of these is the sequence of hypercubes, where $V(H_n) = \{0,1\}^n$ and two vertices are adjacent if and only if they differ at exactly one coordinate. H_n is thus an n-regular graph with

$$\frac{e(H_n)}{\binom{v(H_n)}{2}} = \frac{n \cdot 2^{n-1}}{\binom{2^n}{2}} = \frac{n}{2^n - 1} \to 0.$$

A number of attempts have recently emerged that try to capture the essence of sequences of intermediate density in a convergence notion and a corresponding limit object. Of these, we explore *action convergence* and *logarithmic convergence* in Appendices C and D, respectively.

Chapter 3

Part II: Isometries of Wasserstein spaces

Back in the 18th century, Gaspard Monge considered the following problem. Suppose we have a pile of building material that we want to transfer to a construction site. The transfer of each piece of the material from a point $x \in \mathbb{R}^3$ to a particular place $y \in \mathbb{R}^3$ in the construction site accrues a certain cost $c(x, y) \geq 0$, and we want to find a map $T\colon \mathbb{R}^3 \to \mathbb{R}^3$ which minimizes the overall cost

$$\int c(x, T(x))\, d\mu(x),$$

where μ is the probability distribution describing the location of the pile of the building material.

Let also ν be the probability measure describing the desired location of the construction. Then we denote by

$$C_M(\mu, \nu) := \inf \int c(x, T(x))\, d\mu(x)$$

the infimum of the overall cost over all maps $T\colon \mathbb{R}^3 \longrightarrow \mathbb{R}^3$ under which the pushforward $T_{\#}\mu$ of the measure μ is equal to ν, that is all maps that indeed transport our pile to the construction site.

However, for some pairs of probability measures, no such maps exist – for example, if μ is the Dirac mass δ_x at some $x \in \mathbb{R}^3$, while ν is not one. That is why in the 20th century, Leonid Kantorovich introduced the more general notion of transport *plans*, under which mass coming from one place can be split to go to multiple locations. Formally, a transport plan between two probability measures μ and ν is any coupling π between them, that is a probability measure on $\mathbb{R}^3 \times \mathbb{R}^3$ such that for any measurable subset A of $\mathbb{R}^3$, we have that

$$\pi\left(A \times \mathbb{R}^3\right) = \mu(A)$$

and

$$\pi\left(\mathbb{R}^3 \times A\right) = \nu(A).$$

The infimum of the overall cost of the transfer when splitting mass is allowed is then denoted by

$$C_K(\mu, \nu) := \inf_{\Pi} \int c(x, y)\, d\pi(x, y),$$

where we run over all couplings of μ and ν.

From today's point of view, there is of course nothing special about the particular space $\mathbb{R}^3$, and we will replace it with a general metric space (M, d) from now on.

One of the first questions to spring to mind is of course under what conditions is the infimum in the definition of C_K realised, that is, under what conditions we can replace it with a minimum. The answer covers a surprisingly wide range of scenarios, the underlying reason for which is the tightness of the set of couplings of two given probability measures. We call a set S of probability measures ν_n *tight* if for every $\varepsilon > 0$, there is a compact set K_ε such that for every $\theta \in S$, we have $\theta\left(K_\epsilon\right) \geq 1 - \varepsilon$. One can then utilise Prokhorov's theorem on weak convergence of probability measures to show that there exists an optimal coupling for C_K whenever (M, d) is separable and c is lower semi-continuous.

The most natural costs are of course derived from the metric d. Moreover, for $p \in [1, \infty)$,

$$d_{\mathcal{W}_p}(\mu, \nu) := \inf_{\pi \in \Pi} \left(\int_{M \times M} d(x, y)^p \, d\pi(x, y) \right)^{1/p}$$

is a metric, called the *p-Wasserstein distance* on the set

$$\mathcal{W}_p := \left\{ \mu \mid \exists x \text{ such that } \int_M d(x,y)^p \, d\mu(y) < \infty \right\}.$$

The family of the Wasserstein distances is a prime, but not the only example of a metric on probability measures on (M, d) which is derived from the underlying distance d. We already encountered the Lévy-Prokhorov metric d_{LP} in the setting of action convergence in Part I of the thesis, and several more similar metrics exist. How much more structure or freedom is there then in $\left(\mathcal{W}_p, d_{\mathcal{W}_p}\right)$ than in (M, d)?

One approach to this question is to examine the isometries, that is distance-preserving maps, of these two spaces. Owing to the space-homogeneity of the definition of $d_{\mathcal{W}_p}$, every isometry ϕ of (M, d) is easily seen to also be an isometry of $\left(\mathcal{W}_p, d_{\mathcal{W}_p}\right)$. But can $\left(\mathcal{W}_p, d_{\mathcal{W}_p}\right)$ have any more isometries than that? The answer wildly differs depending on the metric space in question and on the parameter p. For example, Kloeckner showed in [CITE] that for $M = \mathbb{R}^n$, exotic isometries of $(\mathcal{W}_2, d_{\mathcal{W}_2})$ exist, but for example the spheres $\mathbb{S}^n$ equipped with the geodesic distance are isometrically rigid, meaning not admitting exotic isometries of $\left(\mathcal{W}_p, d_{\mathcal{W}_p}\right)$ for any $p \in [1, \infty)$. Somewhere in between these two scenarios lies the one of the spheres $\mathbb{S}^n$ equipped with the distance inherited from $\mathbb{R}^{n+1}$. We prove in Appendix E that these, too, are rigid.

Chapter 4

Part III: Harmonic functions on discrete countable groups

Let Γ be a countable group on which we consider the σ-algebra of all its subsets. Then for every probability measure μ on Γ, we can consider the Markov chain on Γ whose transition probabilities $p(g,h)$ are given by $\mu\left(g^{-1}h\right)$ – that is, wherever we are in the state space, we sample a group element according to μ and multiply it on the right with out current state. Such a time-homogeneous Markov chain is called a *group walk*, and coming hand in hand with it are μ-harmonic functions on Γ, that is, functions f which satisfy that for every $g \in \Gamma$,

$$f(g) = \sum_{h\in\Gamma} \mu(h) f(gh).$$

Whenever we have a *bounded* μ-harmonic function, the martingale convergence theorem tells us that its value converges along almost every realisation of a μ-random walk. Informally, there seem to be some "points at infinity" towards one of which the walk is eventually heading and whose function value it is adopting. This heuristic is formulated precisely via an object known as *Poisson boundary*, or sometimes *Furstenberg-Poisson boundary*, whose existence was shown by H. Furstenberg in [CITE].

Poisson boundary of a pair (Γ, μ) is a probability space (B, ν) such that

$$H^\infty(\Gamma, \mu) \cong L^\infty(B, \nu), \tag{4.1}$$

where $H^\infty(\Gamma, \mu)$ stands for the set of bounded μ-harmonic functions on Γ and the isomorphism is that of normed vector spaces. We think of our random walk eventually "hitting" a point in the boundary B on which a function $F \in L^\infty(\nu)$ is waiting, which gives rise to a harmonic function on Γ by considering the expected value of F if we start the random walk from a given element g of Γ. Harmonicity of this expectation is simply the law of total probability used with conditioning on where the first μ-random step takes us from g, and necessarily $|\mathbb{E}_g[F]| \leq \|F\|_\infty$ for every $g \in \Gamma$, where again, this is a somewhat informal statement based on the assumption that almost every μ-random walk eventually 'hits' B, making F a random variable. One can prove that considering expectations as above is in fact the only source of bounded harmonic functions, and thus obtaining the isomorphism (4.1).

Let us observe that so far, we have not truly used that our state space is a group and indeed, the phenomena described up to now hold true for a general time-homogeneous Markov chain just as well. However, when we do have a group walk at hand, the group acts on trajectories via the diagonal action $g \cdot (x_0, x_1, x_2, \ldots) = (gx_0, gx_1, gx_2, \ldots)$ and hence it also acts on the boundary B. One can for example prove that the centre $Z(\Gamma)$ of Γ always acts trivially on B, a consequence of which is that the Poisson boundary of abelian groups is always trivial or in other words, they admit no bounded harmonic functions other than constant. The property, known as the Liouville property, of having no non-constant bounded harmonic functions, in fact holds true for all virtually nilpotent groups with respect to *any* probability measure one might try to choose on them. The converse of this result, i.e., that every countable group which is not virtually nilpotent admits a probability measure μ such that bounded non-constant μ-harmonic functions exist, was proved only recently by Frisch, Hartman, Tamuz and Vahidi Ferdowsi [1].

After the diagonal action of Γ on trajectories, one can also consider the left action on functions on Γ given by

$$(g \cdot f)(x) = f\left(g^{-1}x\right) \quad \text{for all } x \in \Gamma.$$

Intuitively, $g \cdot f$ is the function f as g sees it when it looks around and thinks that it is the origin.

Recast in this language, the fact that for any bounded, or even just positive, μ-harmonic function h and almost every μ-random trajectory $(x_n)_{n=0}^{\infty}$, the sequence $h(x_n)$ converges is expressed by saying that $x_n \cdot h(e)$ converges. It seems natural that $x_n \cdot h(g)$ should converge for every $g \in \Gamma$, that is, $x_n \cdot h$ should converge pointwise, and indeed, this is what we prove in the following proposition. In Appendix F, a joint work with Omer Segev, we leverage this fact to show that if probability measures θ and μ on Γ satisfy that

$$\|\mu^{*n} * \theta - \theta * \mu^{*n}\|_{\mathrm{TV}} \to 0 \quad \text{as } n \to \infty,$$

then $H^{\infty}(\Gamma, \theta) \subseteq H^{\infty}(\Gamma, \mu)$, where μ^{*n} are the convolution powers giving the distribution of the n-th step of the μ-random walk started from the origin. This in particular includes the case when μ and θ commute, but there do exist measures μ, θ such that the total variation distance of $\mu^{*n} * \theta$ and $\theta * \mu^{*n}$ goes to zero while $\mu * \theta \neq \theta * \mu$ (see Appendix F for an example). The inclusion $H^{\infty}(\Gamma, \theta) \subseteq H^{\infty}(\Gamma, \mu)$ then provides an elegant alternate way of showing some known results, e.g. the aforementioned triviality of the action of $Z(\Gamma)$ on the boundary.

Proposition 1.

Proof. □

While the research of random walks and the corresponding harmonic functions has mostly been centred on finding the harmonic functions for a given measure in the past decades, one can also turn the question around and ask with respect to which probability measures is a given (bounded) function harmonic. We close the thesis with a conjecture and a partial result supporting it, which is based on the fact that for any generating measure μ on the free group $\mathbb{F}_2$, the Gromov boundary $\partial\mathbb{F}_2$ is a factor of the Poisson boundary of $(\mathbb{F}_2, \mu)$.

Conjecture 2. *Let Γ be a countable group and μ a probability measure on Γ. Let h be a μ-harmonic function on Γ. Then if f is obtained from h by changing its value at the origin, there is no probability measure ν on Γ with respect to which f is harmonic. We say that f is* non-harmonisable.

Theorem 3.

Proof. □

References

[1] Joshua Frisch, Yair Hartman, Omer Tamuz, and Pooya Vahidi Ferdowsi. Choquet-deny groups and the infinite conjugacy class property. *Annals of Mathematics*, 190(1):307–320, 2019.

[2] László Lovász. *Large Networks and Graph Limits*, volume 60. American Mathematical Society, 2012.

[3] László Lovász and Balázs Szegedy. Limits of dense graph sequences. *Journal of Combinatorial Theory, Series B*, 96(6):933–957, 2006.

Appendices

Appendix A

Factor-of-iid balanced orientation of non-amenable graphs

by Ferenc Bencs, Aranka Hrušková, László Márton Tóth

European Journal of Combinatorics 115 (2024) 103784

Contents lists available at ScienceDirect

European Journal of Combinatorics

journal homepage: www.elsevier.com/locate/ejc

Factor-of-iid balanced orientation of non-amenable graphs

Ferenc Bencs [a], Aranka Hrušková [b,a], László Márton Tóth [c]

[a] Alfréd Rényi Institute of Mathematics, Budapest, Hungary
[b] Central European University, Budapest, Hungary
[c] École Polytechnique Fédérale de Lausanne, Switzerland

ARTICLE INFO

Article history:
Received 2 February 2023
Accepted 11 July 2023
Available online 9 August 2023

ABSTRACT

We show that if a non-amenable, quasi-transitive, unimodular graph G has all degrees even then it has a factor-of-iid balanced orientation, meaning each vertex has equal in- and outdegree. This result involves extending earlier spectral-theoretic results on Bernoulli shifts to the Bernoulli graphings of quasi-transitive, unimodular graphs.

As a consequence, we also obtain that when G is regular (of either odd or even degree) and bipartite, it has a factor-of-iid perfect matching. This generalizes a result of Lyons and Nazarov beyond transitive graphs.

1. Introduction

Let G be a simple connected graph with all degrees even. An orientation of the edges of G is *balanced* if the indegree of any vertex is equal to its outdegree. When G is finite, the term *Eulerian orientation* is often used, as such an orientation can be obtained from an Eulerian cycle. Our interest lies in infinite graphs, so we prioritize the term balanced. Our main result is the following.

Theorem 1. *Every non-amenable, quasi-transitive, unimodular graph G with all degrees even has a factor-of-iid orientation that is balanced almost surely.*

The precise definitions of these notions are given in Section 2. Non-amenable means that all finite subsets of G expand, quasi-transitive means G has finitely many types of vertices, and

E-mail addresses: bencs.ferenc@renyi.mta.hu (F. Bencs), aranka@renyi.mta.hu (A. Hrušková), toth.laszlo.marton@renyi.hu (L.M. Tóth).

https://doi.org/10.1016/j.ejc.2023.103784
0195-6698/

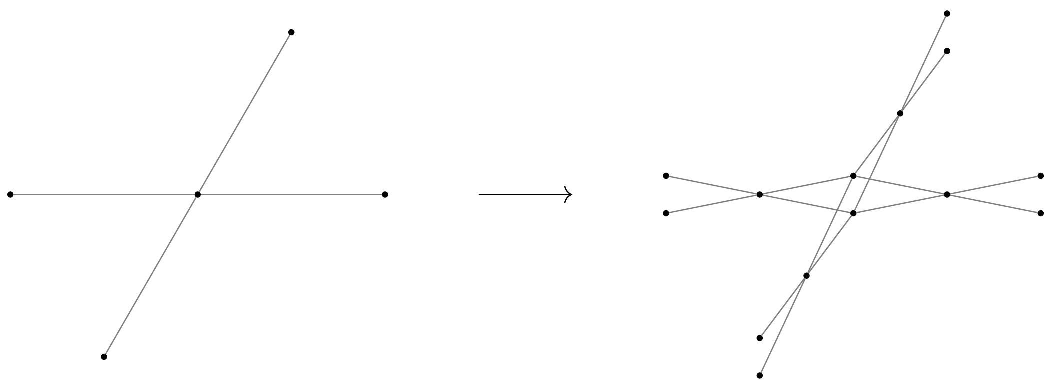

Fig. 1. Obtaining G^* from G — the combinatorial transformation around a vertex of degree 4.

unimodularity is a reversibility condition of the simple random walk on G. Informally speaking, a balanced orientation is a factor of iid if it is produced by a randomized "local" algorithm. To start with, each vertex of G gets a random label from $[0, 1]$ independently and uniformly. Then it makes a deterministic measurable decision about the orientation of its incident edges, based on the labeled graph that it sees from itself as a root. Neighboring vertices must make a consistent decision regarding the edge between them. To make the statements of our results less cumbersome, instead of saying "a factor-of-iid orientation of the edges that is balanced almost surely" we will simply say "factor-of-iid balanced orientation". (The naming is analogous for other decorations of vertices or edges.)

Obtaining combinatorial structures or certain models in statistical mechanics as factors of iid is a central topic in ergodic theory. See [1] and the references therein for an overview in the non-amenable setting.

All Cayley graphs, in particular regular trees are unimodular. For $d > 1$, the $2d$-regular tree T_{2d} is also non-amenable, so it is covered by Theorem 1. Note that on T_{2d} there is a unique invariant random balanced orientation, which by Theorem 1 is a factor of iid. Moreover, this result cannot be obtained by measurable versions of the Lovász Local Lemma, see Remark 8.

Our interest in balanced orientations is due to the fact that on a $2d$-regular graph a balanced orientation is a partial result towards a Schreier decoration. A *Schreier decoration* of G is a coloring of the edges with d colors together with an orientation such that at every vertex, there is exactly one incoming and one outgoing edge of each color. It is a combinatorial coding of an action of the free group F_d on the vertex set of the graphs. Every Schreier decoration gives a balanced orientation by forgetting the colors. In [2], the third author proved that all $2d$-regular unimodular random rooted graphs admit an invariant random Schreier decoration, and the current authors show in a parallel work [3] that such invariant random Schreier decorations can be obtained as a factor of iid in Euclidean grids in all dimensions greater than 1 as well as on all Archimedean (planar) lattices of even degree.

It remains an open question whether there indeed is a factor-of-iid Schreier decoration of T_{2d}. In [4], Thornton studies when graphs have factor-of-iid Cayley diagrams. Finding a Cayley diagram of a fixed group as a random decoration comes with (compared to a Schreier decoration) additional local restrictions on how the decoration should behave on loops. Nevertheless, the question of finding a Cayley diagram of F_d overlaps with our interest in Schreier decorations on T_{2d}. Thornton has a result on factor-of-iid Cayley diagrams on non-amenable graphs ([4, Theorem 1.7]) that provides *approximate* Cayley diagrams, but we do not allow for a small-probability local error here.

The proof of Theorem 1 relies on two main ingredients. First, we reduce the question of finding a balanced orientation of G to finding a perfect matching in an auxiliary graph G^*. Fig. 1 illustrates the construction, which already appears in works of Schrijver [5] and Mihail–Winkler [6]. The precise formulation is given in Section 5.

Second, we apply earlier matching results of Lyons and Nazarov [7], who proved that bipartite, non-amenable Cayley graphs have a factor-of-iid perfect matching. Csóka and Lippner extended this to all non-amenable Cayley graphs in [8].

In order to deduce Theorem 1 from these matching results, we have to establish vertex expansion in the appropriate Bernoulli graphing (see Section 2.7 for the definition). We do this via spectral theory, and state our spectral-theoretic result here because we believe it is of interest in itself.

Theorem 2. *Let G be a connected, unimodular, quasi-transitive graph. If G is non-amenable then its Bernoulli graphing $\mathcal{G}$ has positive spectral gap.*

The interpretation of spectral gap is slightly different depending on the Bernoulli graphing being measurably bipartite or not. See Theorems 17 and 18 for exact statements.

Our proof of Theorem 2 requires more sophistication than simply repeating earlier arguments in a more general setting. To emphasize this, we point out that (unlike in the transitive case) -1 can indeed be part of the spectrum. Also our proof does not bound $\|\mathcal{M}\|$ above by $\|M_G\|$, where $\|M_G\|$ is the operator norm of the Markov operator M_G on $\ell^2(\mathrm{V}(\mathrm{G}))$, while for Cayley graphs, one has $\|\mathcal{M}\| \leq \|M_G\|$.

As a consequence of Theorem 2, we also obtain the following generalization of the result of Lyons and Nazarov.

Corollary 3. *Let G be a connected, unimodular, quasi-transitive non-amenable regular bipartite graph. Then G has a factor-of-iid perfect matching.*

The bipartite assumption in Corollary 3 cannot be dropped because there are unimodular, quasi-transitive regular graphs that have no perfect matching at all, see Remark 21. Regularity cannot be dropped either, as for example bi-regular trees (of two different degrees of regularity) have no factor-of-iid perfect matching.

We also discuss how far Theorem 1 goes towards obtaining Schreier decorations on the regular tree T_{2d}.

Proposition 4. *Regarding Schreier decorations of T_{2d} we observe the following.*

(i) If T_d has a factor-of-iid proper edge coloring with d colors then T_{2d} has a factor-of-iid Schreier decoration.
(ii) T_{2d} has a factor-of-iid Schreier decoration with the last two colors unordered.
(iii) Let $\vec{T}_4$ denote the tree T_4 with edges oriented in a balanced way. ($\vec{T}_4$ is unique up to isomorphism.) There is no $\mathrm{Aut}(\vec{T}_4)$-factor-of-iid Schreier decoration of $\vec{T}_4$ with the additional property that after forgetting the colors, it coincides with the original orientation of $\vec{T}_4$.
(iv) For every positive integer d, if T_{2d} has a factor-of-iid Schreier decoration then so does T_{2d+2}.

It is an open question whether T_d (for $d > 2$) has a factor-of-iid proper edge coloring by d colors. Part (ii) utilizes the partial result towards such a factor-of-iid proper edge coloring presented in [1]; see Section 6.2 for further comments. Note, however, that by part (iii), obtaining a factor-of-iid Schreier decoration of T_4 cannot be achieved by selecting a balanced orientation first and then choosing the colors without modifying the orientation. This observation is unique to degree 4 because it relies on the 2-regular tree, otherwise known as the bi-infinite path, not having a factor-of-iid proper edge coloring with two colors. For higher degree, a construction might be finished this way, as pointed out in part (i).

Finally, regarding the auxiliary graph G^* we show that existences of different factors of iid are equivalent.

Proposition 5. *For every $2d$-regular graph G, the bipartite graph G^* is also $2d$-regular, and the following are equivalent.*

1. *G^* has a factor-of-iid proper edge $2d$-coloring.*
2. *G^* has a factor-of-iid perfect matching.*
3. *G^* has a factor-of-iid Schreier decoration.*

Moreover, if any of these is a finitary factor, the others are too.

For the definition of *finitary* factors see Section 2.4.

The structure of the paper is as follows. In Section 2, we introduce the necessary notions and existing results. In Section 3, we prove our spectral-theoretic result, Theorem 2. We deduce Corollary 3 in Section 4, and in Section 5, we prove our main result, Theorem 1. Our results on other types of decorations are collected in Section 6. Section 7 lists some open questions.

Addendum. After our manuscript was made available online, Riley Thornton brought to our attention that his Theorem 2.8 in [9] provides a measurable balanced orientation in $2d$-regular graphings with expansion. Since Backhausz, Szegedy, and Virág show in [10, Theorem 2.2] that the Bernoulli graphing of T_{2d} does have expansion, a factor-of-iid balanced orientation of T_{2d} can also be obtained by combining these two results.

2. Notation and basics

Some of the descriptions in this section are identical to the ones in our parallel work [3].

2.1. Graphs

A graph G is given by its vertex set $V(G)$ and edge set $E(G)$, where $E(G) \subset V(G)^{(2)}$ is a collection of 2-element subsets of $V(G)$ and we write uv for the subset $\{u, v\}$. For any subset $A \subset V(G)$, we denote by $N_G(A)$ the neighborhood of A, that is $\{u \in V(G) : \exists v \in A \text{ such that } uv \in E(G)\}$. We use the calligraphic $\mathcal{G}$ for graphs that have a probability measure associated to them that makes them a graphing (see Section 2.6 for precise definition).

2.2. Amenability

Let G be a locally finite connected graph, and let $p_n(x, y)$ denote the probability of the simple random walk started from x reaching y in n steps. Then the value $\limsup_{n\to\infty} \sqrt[n]{p_n(x, y)}$ is independent of the choice of x and y, and is in fact equal to the norm of the Markov operator $M : \ell^2(V(G), m_{\mathrm{st}}) \to \ell^2(V(G), m_{\mathrm{st}})$. Here m_{st} is the degree-biased version of the counting measure, i.e. $m_{\mathrm{st}}(X) = \sum_{v\in X} \deg(v)$, which is a stationary measure with respect to the random walk. The operator M is defined by

$$(M(f))(v) = \frac{1}{\deg(v)} \sum_{uv\in E(G)} f(u).$$

M is self-adjoint and has norm at most 1 for any G. We will denote its norm (and spectral radius) by ρ:

$$\rho = \|M\| = \limsup_{n\to\infty} \sqrt[n]{p_n(x, y)}.$$

We say G is *amenable* if $\rho = 1$ and *non-amenable* if $\rho < 1$.

This characterization of amenability, due to Kesten [11], is of course only one of many. In particular non-amenability is equivalent to the positivity of the *Cheeger constant* of G. In Section 5 we also work with Cheeger constants, but we do so on graphings, not on countably infinite graphs.

2.3. Schreier graphs

Given a finitely generated group $\Gamma = \langle S\rangle$ and an action $\Gamma \curvearrowright X$ on some set X, the *Schreier graph* $\mathrm{Sch}(\Gamma \curvearrowright X, S)$ of the action is defined as follows. The set of vertices is X, and for every $x \in X$, $s \in S$, we introduce an oriented s-labeled edge from x to $s.x$.

Rooted connected Schreier graphs of Γ come from pointed transitive actions of Γ, which are in one-to-one correspondence with subgroups of Γ. Trivially, a graph with a Schreier decoration is a Schreier graph of the free group F_d on d generators. A special case is the (left) *Cayley graph* of Γ, denoted $\mathrm{Cay}(\Gamma, S)$, which is the Schreier graph of the (left) translation action $\Gamma \curvearrowright \Gamma$.

2.4. Factors of iid

Let Γ be a group. A *Γ-space* is a measurable space X with an action $\Gamma \curvearrowright X$. A map $\Phi : X \to Y$ between two Γ-spaces is a *Γ-factor* if it is measurable and Γ-equivariant, that is $\gamma.\Phi(x) = \Phi(\gamma.x)$ for every $\gamma \in \Gamma$ and $x \in X$.

A measure μ on a Γ-space X is *invariant* if $\mu(\gamma.A) = \mu(A)$ for all $\gamma \in \Gamma$ and all measurable $A \subseteq X$. We say an action $\Gamma \curvearrowright (X, \mu)$ is *probability-measure-preserving* (p.m.p.) if μ is a Γ-invariant probability measure.

Let G be a countable graph and $\Gamma \leq \mathrm{Aut}(G)$. Let u denote the Lebesgue measure on $[0, 1]$. We endow the space $[0, 1]^{V(G)}$ with the product measure $\mathrm{u}^{V(G)}$. The translation action $\Gamma \curvearrowright [0, 1]^{V(G)}$ is defined by

$$(\gamma.f)(v) = f(\gamma^{-1}.v), \ \forall \gamma \in \Gamma, v \in V(G).$$

The action $\Gamma \curvearrowright ([0, 1]^{V(G)}, \mathrm{u}^{V(G)})$ is p.m.p.

An orientation of G can be thought of as a function on $E(G)$ sending every edge to one of its endpoints. Viewed like this, orientations of G form a standard Borel space in the product $E(G)^{V(G)}$. We denote this space of orientations $\mathtt{Or}(G)$, and note that it comes with a natural action of Γ. The set $\mathtt{BalOr}(G) \subseteq \mathtt{Or}(G)$ of balanced orientations is Γ-invariant and Borel, so it is a Γ-space in itself. Similarly, the set of all Schreier decorations of G forms the Γ-space $\mathtt{Sch}(G)$.

Definition 6. A Γ-factor of iid balanced orientation (respectively, Schreier decoration) of a graph G is a Γ-factor $\Phi : ([0, 1]^{V(G)}, \mathrm{u}^{V(G)}) \to \mathtt{BalOr}(G)$ (respectively, to $\mathtt{Sch}(G)$). If the subgroup $\Gamma \leq \mathrm{Aut}(G)$ is not specified, we mean an $\mathrm{Aut}(G)$-factor.

Remark 7. We allow Φ to not be defined on a $\mathrm{u}^{V(G)}$-null subset $X_0 \subseteq [0, 1]^{V(G)}$.

Let us now recall some special classes of factor of iid processes on graphs. For a fixed vertex $x \in V(G)$, let $\big(\Phi(\omega)\big)(x)$ denote the restriction of $\Phi(\omega)$ to the edges incident to x. We say Φ is a *finitary* factor of iid if for almost all $\omega \in [0, 1]^{V(G)}$, there exists an $R \in \mathbb{N}$ such that $\big(\Phi(\omega)\big)(x)$ is already determined by $\omega|_{B_G(x,R)}$. That is, if we change ω outside $B_G(x, R)$, the decoration $\Phi(\omega)$ does not change around x. This radius R can depend on the particular ω. If it does not then we say Φ is a *block factor*.

When constructing factors of iid algorithmically, one often makes use of the fact that a uniform [0, 1] random variable can be decomposed into countably many independent uniform [0, 1] random variables. In practice, this means that we can assume that a vertex has multiple labels or that a new independent random label is always available after a previous one was used.

We will use a reverse operation as well: the composition of countably many uniform [0, 1] random variables is again a uniform [0, 1] random variable.

Remark 8. Note that balanced orientations of T_{2d} have the property that fixing the orientation on all edges at distance r from some vertex u determines the orientation of edges incident to u, independently of r. Consequently, the balanced orientation constructed in Theorem 1 has no local reduction to the Lovász Local Lemma (LLL). Indeed, by [12, Section 11.1] it has randomized local complexity $\Theta(\log n)$, whereas the algorithm of [13] implies $o(\log n)$ complexity for problems that have local reductions to the LLL. So although there are measurable versions of the LLL [14,15], factor-of-iid balanced orientations of T_{2d} cannot be obtained that way.

2.5. Unimodular quasi-transitive graphs

Unimodular random rooted graphs are central objects in sparse graph limit theory because they can represent limits of locally convergent sequences of finite graphs. In this paper, however, we only deal with a special case, namely unimodular quasi-transitive graphs. For a thorough treatment of the topic and the connection to unimodular random rooted graphs, we refer the reader to [16, Chapter 8.2] and [17].

Let G be a locally finite graph, $\Gamma = \mathrm{Aut}(G)$. There is a function $\mu : V(G) \to \mathbb{R}^+$ such that for any $x, y \in V(G)$, we have

$$\frac{\mu(x)}{\mu(y)} = \frac{|\mathrm{Stab}_\Gamma(x).y|}{|\mathrm{Stab}_\Gamma(y).x|}.$$

The function μ is unique up to multiplication by a constant. We say G is *unimodular* if $|\mathrm{Stab}_\Gamma(x).y| = |\mathrm{Stab}_\Gamma(y).x|$ for any pair $x, y \in V(G)$ that are in the same Γ orbit, that is $y \in \Gamma.x$. So G is unimodular if and only if $\mu(y) = \mu(x)$ for any $y \in \Gamma.x$.

Moreover, when $\{o_i\}$ is the orbit section of G and $\sum_i \mu(o_i)^{-1} < \infty$, then we can normalize μ to obtain a probability measure on $\{o_i\}$.

In particular, when G is quasi-transitive, let $T = \{o_1, \dots, o_t\} \subset V(G)$ be a set of representatives of the orbits of $\Gamma \curvearrowright V(G)$. Let p be the normalized version of μ^{-1} as above — we think of p as a distribution of a random root in G.

The notion of unimodularity comes hand in hand with the Mass Transport Principle. In our case, it takes the following form:

Proposition 9 (*Mass Transport Principle, Corollary 8.11. in [16]*). *Given a function $f : V(G) \times V(G) \to [0, \infty]$ that is invariant under the diagonal action of Γ, we have*

$$\sum_{i=1}^{t} p(o_i) \sum_{z \in V(G)} f(o_i, z) = \sum_{i=1}^{t} p(o_i) \sum_{z \in V(G)} f(z, o_i).$$

We immediately use the Mass Transport Principle to set up a finite state Markov chain mimicking the transitions of the random walk on G between Γ-orbits.

Lemma 10. *For any $1 \leq i \neq j \leq t$, the function p satisfies,*

$$p(o_i) \left|\{o_i v \in E \mid v \in \Gamma.o_j\}\right| = p(o_j) \left|\{v o_j \in E \mid v \in \Gamma.o_i\}\right|.$$

Proof. For fixed $i \neq j$, set up a payment function f with $f(x, y) = 1$ if $xy \in E(G)$, $x \in \Gamma.o_i$ and $y \in \Gamma.o_j$. Set $f(x, y) = 0$ otherwise. The Mass Transport Principle gives the desired equality. □

We define a Markov chain M_T with states T and transition probabilities

$$p_{M_T}(o_i, o_j) = \frac{\left|\{o_i v \in E(G) \mid v \in \Gamma.o_j\}\right|}{\deg(o_i)}.$$

Note that M_T is just the projection of the random walk on G onto $\{\Gamma.o_1, \dots, \Gamma.o_t\}$.

With slight abuse of notation, we will also denote the transition matrix by M_T. We write $\widetilde{p}$ for the degree-biased version of the root distribution p, that is

$$\widetilde{p}(o_i) = \frac{\deg(o_i)}{\Delta} \cdot p(o_i).$$

Here $\Delta = \mathbb{E}_p[\deg(o_i)]$ is the expected degree of a root picked with distribution p. Lemma 10 shows that $\widetilde{p}$ is a reversible stationary distribution for M_T.

List the eigenvalues of M_T in decreasing order, $1 = \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_t$. We say M_T is *bipartite* if $\lambda_t = -1$. M_T is bipartite if and only if we can partition T into two sets T_1 and T_2 such that whenever $o_i, o_j \in T_1$ or $o_i, o_j \in T_2$, we have $p_{M_T}(o_i, o_j) = 0$. We set $\rho_T = \max(\{0\} \cup \{|\lambda_i| \mid 1 < i \leq t \text{ and } \lambda_i > -1\})$.

We will have to treat the bipartite and non-bipartite case separately. When M_T is not bipartite, we have $\rho_T = 0$ when $t = 1$ and $\rho_T = \max\{\lambda_2, |\lambda_t|\}$ when $t \geq 2$. The following is an immediate consequence of the Convergence Theorem for finite-state Markov chains.

Lemma 11. *Assume M_T is not bipartite. Let $e_{o_i} \in \mathbb{R}^T$ denote the characteristic vector of $o_i \in T$. Then for any $v \in \mathbb{R}^T$, there exists a $C > 0$ such that for any $i \in [t]$ and $k \in \mathbb{N}$ we have*

$$|\langle M_T^k e_{o_i}, v\rangle - \langle \widetilde{p}, v\rangle| \leq C \rho_T^k.$$

When M_T is bipartite, the spectrum is symmetric, i.e. $\lambda_i = -\lambda_{t-i+1}$ for all $i \in \{1, \dots, t\}$. In particular, $\lambda_{t-1} = -\lambda_2$, so $\rho_T = \lambda_2$ whenever $t \geq 3$. (When $t = 2$ we have $\rho_T = 0$.) The reversibility of M_T ensures $\widetilde{p}(T_1) = \widetilde{p}(T_2) = 1/2$, and M_T^2 defines two disjoint Markov chains on T_1 and T_2 with stationary measures $2\widetilde{p}|_{T_1}$ and $2\widetilde{p}|_{T_2}$ respectively. The eigenvalues of M_T^2 are the squares of the eigenvalues of M_T, in particular they are non-negative, so M_T^2 is not bipartite on either T_1 or T_2. Also the second largest eigenvalue in absolute value of M_T^2 (on both T_1 and T_2) is ρ_T^2.

2.6. Graphings

Graphings play an essential role in obtaining invariant random structures on graphs as they represent a space where both the probability measure and the underlying (possibly random) countable graph are present. Their use in constructing factor-of-iid perfect matchings is well-established in [7,8]. For a more detailed introduction, see for example [18, Chapter 18].

Definition 12. Let (X, ν) be a Borel probability space. A (bounded-degree) *graphing* is a graph $\mathcal{G}$ with $V(\mathcal{G}) = X$ and Borel edge set $E(\mathcal{G})$, in which all degrees are at most $D \in \mathbb{N}$, and

$$\int_A \deg_B(x)\, d\nu(x) = \int_B \deg_A(x)\, d\nu(x) \tag{2.1}$$

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

We will now define what we mean by $E(\mathcal{G})$ being Borel. The reason we do it in a slightly convoluted way is because in this paper it will be more convenient to use $E(\mathcal{G})$ to denote the set of edges, and not think about it as a symmetric subset of $X \times X$. The present downside to this is that defining the Borel structure and the edge measure can be done most naturally inside $X \times X$. For this reason let $\tilde{E}(\mathcal{G})$ denote the symmetric subset of $X \times X$ corresponding to the edges of $\mathcal{G}$:

$$\tilde{E}(\mathcal{G}) = \{(x, y) \in X \times X \mid xy \in E(\mathcal{G})\}.$$

We say $E(\mathcal{G})$ is a Borel edge set if $\tilde{E}(\mathcal{G}) \subseteq X \times X$ is Borel. Then $E(\mathcal{G})$ itself has a Borel σ-algebra corresponding to the sub-σ-algebra of symmetric Borel subsets of $\tilde{E}(\mathcal{G})$.

Moreover, the measure ν of a graphing $\mathcal{G}$ gives rise to a measure $\nu_{\tilde{E}}$ on $X \times X$ by defining

$$\nu_{\tilde{E}}(A \times B) = \frac{1}{2} \int_A \deg_B(x)\, d\nu(x)$$

for any measurable $A, B \subset X$. The measure $\nu_{\tilde{E}}$ is concentrated on $\tilde{E}(\mathcal{G})$.

We then define the *edge measure* ν_E on $E(\mathcal{G})$ by setting $\nu_E(F) = \nu_{\tilde{E}}(\tilde{F})$ for measurable subsets $F \subseteq E(\mathcal{G})$. ($\tilde{F}$ is defined analogously to $\tilde{E}(\mathcal{G})$.) Essentially we are restricting $\nu_{\tilde{E}}$ as defined above to the symmetric Borel subsets of $\tilde{E}(\mathcal{G})$. The factor $1/2$ is introduced so that the appropriate version of the usual edge double counting identity $2|E(G)| = \sum_{v \in V(G)} \deg(v)$ for finite graphs also holds for graphings:

$$2\nu_E\big(E(\mathcal{G})\big) = \int_X \deg(x)\, d\nu(x).$$

Example 13. Given a finitely generated group $\Gamma = \langle S \rangle$ and a p.m.p. action $\Gamma \curvearrowright (X, \nu)$, the Schreier graph $\mathrm{Sch}(\Gamma \curvearrowright X, S)$ is a graphing (after forgetting the orientation and S-labeling). The action being p.m.p. implies that the degree condition (2.1) holds.

2.7. Connection to Bernoulli graphings

We now introduce Bernoulli graphings, which are closely related to factors of iid.

For a unimodular quasi-transitive graph G, we define its Bernoulli graphing $\mathcal{G}$ as follows. The vertex set of $\mathcal{G}$ is Ω, the space of $[0, 1]$-decorated, rooted, connected graphs with degree bound D (up to rooted isomorphism). Elements of $V(\mathcal{G}) = \Omega$ are of the form (H, u, ω), where (H, u) is a connected, bounded-degree rooted graph, and $\omega : V(H) \to [0, 1]$ is a labeling. We connect (H, u, ω) with (H', u', ω') if and only if we can obtain (H', u', ω') from (H, u, ω) by moving the root u to one of its neighbors. We denote the resulting measurable edge set by $\mathcal{E}$.

It remains to define the probability measure on Ω. (Note that the vertex and edge sets are the same for every G, only the measure will be different.) G is quasi-transitive, so it has finitely many possible rooted versions, namely the (G, o_i) for $o_i \in \{o_1, \ldots, o_t\}$. Let us pick a random root o, choosing each o_i with probability $p(o_i)$. We also pick a random labeling $\omega \in [0, 1]^{V(G)}$ according to $\mathrm{u}^{V(G)}$. Recall that u stands for the uniform measure on $[0, 1]$. The triple (G, o, ω), considered up to rooted isomorphism, is a random element of Ω, let ν_G denote its distribution. The Bernoulli graphing of G is $\mathcal{G} = (\Omega, \mathcal{E}, \nu_G)$. $\mathcal{G}$ satisfies (2.1) because G is unimodular.

Given a unimodular quasi-transitive graph G, constructing an Aut(G)-factor of iid almost surely balanced orientation (Schreier decoration) of G is equivalent to constructing a measurable almost everywhere balanced orientation (Schreier decoration) of the Bernoulli graphing $\mathcal{G}$ built on G. Here, measurability means that the oriented edges (and the color classes) form $\nu_{\tilde{\mathcal{E}}}$-measurable subsets of $\Omega \times \Omega$.

Also note that a measurable Schreier decoration of any graphing $\mathcal{G}$ defines a p.m.p. action $F_d \curvearrowright V(\mathcal{G})$ that *generates* the graphing as in Example 13.

Therefore, an equivalent formulation of our main motivating question is the following: given a (quasi-transitive unimodular) $2d$-regular graph G, is the Bernoulli graphing $\mathcal{G}$ generated by a p.m.p. action of F_d? The answer is no for the bi-infinite line, and one can also construct $2d$-regular counterexamples for every d, see [3]. However, as far as the authors are aware, all such known counterexamples are 2-ended. Some cases when the answer is positive are also established in [3].

It would of course be even better to answer this question for all unimodular random rooted graphs.

2.8. Perfect matchings in expanding graphings

Finding measurable perfect matchings in non-hyperfinite graphings is usually achieved through expansion properties. (There are important results in the hyperfinite case as well, see e.g. [19], announced a few months after the first version of the current paper was made available online. Results in the hyperfinite world however tend to use a rather different set of tools.)

We will use the following two results, both based on the argument of Lyons and Nazarov in [7].

Theorem 14 (*Lyons-Nazarov, [7]*). *Let $\mathcal{G} = (X, E, \nu)$ be a graphing with no odd cycles. Assume it has vertex expansion at least $c > 1$. That is, for any $A \subset X$ such that $0 < \nu(A) \le 1/2$, we have*

$$\frac{\nu(N_{\mathcal{G}}(A))}{\nu(A)} \ge c.$$

Then $\mathcal{G}$ has a Borel matching that covers all vertices up to a nullset.

Note that $\mathcal{G}$ having no odd cycles in the previous theorem means that each connected component is bipartite, but $\mathcal{G}$ itself might not have a *measurable* bipartition of its vertex set. If such a measurable bipartition exists, we say $\mathcal{G}$ is *measurably bipartite*. In fact we will also need a variation of the above for measurably bipartite graphings, because in that case the expansion assumption in Theorem 14 cannot hold. (A measure 1/2 subset of the larger side of the bipartition violates the inequality.)

Theorem 15 (*Lyons-Nazarov, Theorem 9.1 in [20]*). *Let $\varepsilon > 0$. Let $\mathcal{G} = (X_1, X_2, E, \nu)$ be a measurably bipartite graphing with $\nu(X_1) = \nu(X_2)$. Assume it has bipartite vertex-expansion at least $1 + \varepsilon$. That is, for any $A \subseteq X_1$ and $B \subseteq X_2$, we have*

$$\nu(N_{\mathcal{G}}(A)) \ge \min\left\{(1+\varepsilon)\nu(A), \frac{1}{4} + \varepsilon\right\} \qquad \textit{and} \qquad \nu(N_{\mathcal{G}}(B)) \ge \min\left\{(1+\varepsilon)\nu(B), \frac{1}{4} + \varepsilon\right\}.$$

Then $\mathcal{G}$ has a Borel matching that covers all vertices up to a nullset.

3. Spectral gap for non-amenable quasi-transitive graphs

In this section we prove our spectral theoretic result, Theorem 2.

For a graphing $\mathcal{G}$ on (X, ν), we denote by ν_{st} the degree-biased version of ν, that is

$$\nu_{\mathrm{st}}(A) = \int_{x\in A} \deg(x)d\nu \Big/ \int_{x\in X} \deg(x)d\nu.$$

As the notation suggests, ν_{st} is stationary with respect to the Markov operator $\mathcal{M}$ of $\mathcal{G}$ that is defined by

$$(\mathcal{M}f)(G, o, \omega) = \frac{1}{\deg_G(o)} \sum_{ov\in E} f(G, v, \omega).$$

$\mathcal{M}$ is a self-adjoint operator on $L^2(\Omega, \nu_{\mathrm{st}})$. To get a bound on the spectral radius of $\mathcal{M}$ (on the appropriate subspace), we will use the following lemma.

Lemma 16 (*Lemma 2.4 of [10]*)**.** *For a bounded self-adjoint operator T on a Hilbert space $\mathcal{H}$ and for any spanning subset H of $\mathcal{H}$, we have*

$$\rho(T) = \|T\| = \sup_{v\in H} \left(\limsup_{k\to\infty} \left| \frac{\langle v, T^k v\rangle}{\langle v, v\rangle} \right|^{1/k} \right).$$

The following two theorems deal with the non-bipartite and bipartite case separately. Recall that we denote the Markov operator of G on $\ell^2(G, m_{\mathrm{st}})$ by M, where m_{st} denotes the degree-biased version of the counting measure on $V(G)$. As G is non-amenable we have $\rho = \|M\| < 1$. Recall also that $\rho_T = \max(\{0\} \cup \{|\lambda_i| \mid 1 < i \le t \text{ and } \lambda_i > -1\})$ is defined through the finite state Markov chain M_T, and $\rho_T < 1$.

Theorem 17. *Let G be as in Theorem 2, and assume also that M_T is not bipartite. Let $L^2_0(\Omega, \nu_{\mathrm{st}})$ denote the orthogonal complement of the subspace of constant functions. Then the spectral radius of $\mathcal{M}$ on $L^2_0(\Omega, \nu_{\mathrm{st}})$ is at most $\max\{\rho, \rho_T\} < 1$.*

Theorem 18. *Let G be as in Theorem 2, and assume that M_T is bipartite. Let $\rho < 1$ denote the spectral radius of G on $\ell^2(G, m_{\mathrm{st}})$. The Bernoulli graphing $\mathcal{G}$ is measurably bipartite, with bipartition $X_1 \cup X_2 = V(\mathcal{G})$. Let $L^2_{00}(\Omega, \nu_{\mathrm{st}})$ denote the orthogonal complement of the subspace generated by the functions $\mathbb{1}_X$ and $\mathbb{1}_{X_1} - \mathbb{1}_{X_2}$. Then the spectral radius of $\mathcal{M}$ on $L^2_{00}(\Omega, \nu_{\mathrm{st}})$ is at most $\max\{\rho, \rho_T\} < 1$.*

Proof of Theorem 2. The content of Theorem 2 is exactly Theorems 17 and 18. □

We first prove Theorem 17 and then use it to prove Theorem 18.

Proof of Theorem 17. As before, let $p_k(o, y)$ denote the probability that the random walk on G starting at o arrives at y after k steps. We have $\limsup_{k\to\infty} \left(p_k(o, y)\right)^{1/k} = \rho$, so for every $\varepsilon > 0$ there exists some $C_0(o, y, \varepsilon) \in \mathbb{R}$ such that $p_k(o, y) \le C_0(o, y, \varepsilon)(\rho + \varepsilon)^k$ for all k.

During this proof, we will write μ for $\mathrm{u}^{V(G)}$. We will use Lemma 16 in the following setting. Let $H \subseteq L^2_0(\Omega, \nu_{\mathrm{st}})$ be the set of functions f such that

- f has zero mean, i.e.

$$\int_{(G,o,\omega)\in\Omega} f(G, o, \omega)\, d\nu_{\mathrm{st}} = \sum_{i=1}^{t} \widetilde{p}(o_i) \int_{\omega\in[0,1]^{V(G)}} f(G, o_i, \omega)\, d\mu = 0;$$

- f has norm 1, i.e.

$$\int_{(G,o,\omega)\in\Omega} f^2(G, o, \omega)\, d\nu_{\mathrm{st}} = \sum_{i=1}^{t} \widetilde{p}(o_i) \int_{\omega\in[0,1]^{V(G)}} f^2(G, o_i, \omega)\, d\mu = 1;$$

- there exists some $r \geq 0$ such that if we change labels of vertices further than r from the root then the value of f does not change.

The set H is a spanning subset of $L_0^2(\Omega, \nu_{\mathrm{st}})$. (Note that a measurable function $f : \Omega \to \mathbb{R}$ defines an $\mathbb{R}$-valued factor of iid on any graph G. Indeed, for $\omega \in [0,1]^{V(G)}$ one defines $\big(\Phi(\omega)\big)(v) = f(G, v, \omega)$.) This is equivalent to saying that any factor of iid process is a limit of block factors; see [1].

Let us fix an element $f \in H$. Then

$$\begin{aligned}\langle \mathcal{M}^k f, f\rangle &= \sum_{i=1}^{t} \widetilde{p}(o_i) \int_{\omega\in[0,1]^{V(G)}} \sum_{y\in B_k(o_i)} p_k(o_i, y) f(G, o_i, \omega) f(G, y, \omega) d\mu \\ &= \sum_{i=1}^{t} \widetilde{p}(o_i) \sum_{y\in B_k(o_i)} p_k(o_i, y) \int_{\omega\in[0,1]^{V(G)}} f(G, o_i, \omega) f(G, y, \omega) d\mu.\end{aligned}$$

We split the sum depending on the distance between o_i and y:

$$\langle \mathcal{M}^k f, f\rangle = \sum_{i=1}^{t} \widetilde{p}(o_i) \sum_{y\notin B_{2r}(o_i)} p_k(o_i, y) \int_{\omega\in[0,1]^{V(G)}} f(G, o_i, \omega) f(G, y, \omega) d\mu \tag{3.1}$$

$$+ \sum_{i=1}^{t} \widetilde{p}(o_i) \sum_{y\in B_{2r}(o_i)} p_k(o_i, y) \int_{\omega\in[0,1]^{V(G)}} f(G, o_i, \omega) f(G, y, \omega) d\mu. \tag{3.2}$$

If the distance between o_i and y is bigger than $2r$ then (by the third property of f) the values $f(G, o_i, \omega)$ and $f(G, y, \omega)$ depend on labels at disjoint sets of vertices. Since those labels are independent, we have

$$\int_{\omega\in[0,1]^{V(G)}} f(G, o_i, \omega) f(G, y, \omega) d\mu = \int_{\omega\in[0,1]^{V(G)}} f(G, o_i, \omega) d\mu \int_{\omega\in[0,1]^{V(G)}} f(G, y, \omega) d\mu.$$

Therefore the first term, (3.1) is

$$\begin{aligned}&\sum_{i=1}^{t} \widetilde{p}(o_i) \sum_{y\notin B_{2r}(o_i)} p_k(o_i, y) \int_{\omega\in[0,1]^{V(G)}} f(G, o_i, \omega) f(G, y, \omega) d\mu \\ &\quad= \sum_{i=1}^{t} \widetilde{p}(o_i) \sum_{y\notin B_{2r}(o_i)} p_k(o_i, y) \int_{\omega\in[0,1]^{V(G)}} f(G, o_i, \omega) d\mu \int_{\omega\in[0,1]^{V(G)}} f(G, y, \omega) d\mu \\ &\quad= \sum_{i=1}^{t} \widetilde{p}(o_i) \int_{\omega\in[0,1]^{V(G)}} f(G, o_i, \omega) d\mu \sum_{y\notin B_{2r}(o_i)} p_k(o_i, y) \int_{\omega\in[0,1]^{V(G)}} f(G, y, \omega) d\mu \\ &\quad= \sum_{i=1}^{t} \widetilde{p}(o_i) \int_{\omega\in[0,1]^{V(G)}} f(G, o_i, \omega) d\mu \sum_{y\in V(G)} p_k(o_i, y) \int_{\omega\in[0,1]^{V(G)}} f(G, y, \omega) d\mu \\ &\qquad- \sum_{i=1}^{t} \widetilde{p}(o_i) \int_{\omega\in[0,1]^{V(G)}} f(G, o_i, \omega) d\mu \sum_{y\in B_{2r}(o_i)} p_k(o_i, y) \int_{\omega\in[0,1]^{V(G)}} f(G, y, \omega) d\mu\end{aligned}$$

$$= \sum_{i=1}^{t} \widetilde{p}(o_i) \int_{\omega\in[0,1]^{V(G)}} f(G, o_i, \omega) d\mu \sum_{j=1}^{t} p_k^{M_T}(o_i, o_j) \int_{\omega\in[0,1]^{V(G)}} f(G, o_j, \omega) d\mu \tag{3.3}$$

$$- \sum_{i=1}^{t} \widetilde{p}(o_i) \int_{\omega\in[0,1]^{V(G)}} f(G, o_i, \omega) d\mu \sum_{y\in B_{2r}(o_i)} p_k(o_i, y) \int_{\omega\in[0,1]^{V(G)}} f(G, y, \omega) d\mu. \tag{3.4}$$

Along the calculation, we used that $\int_{\omega\in[0,1]^{V(G)}} f(G,y,\omega)\,d\mu$ only depends on the orbit that y is in. Then we grouped all the $y \in \Gamma.o_j$ together, and used the fact that

$$\sum_{y\in\Gamma.o_j} p_k(o_i,y) = p_k^{M_T}(o_i,o_j).$$

Indeed, the probability of the random walk on G started from o_i ending up at some $y \in \Gamma.o_j$ after k steps is the same as the probability of the finite Markov chain M_T, starting from o_i ending up in o_j after k steps.

We now use that from any initial state, the finite Markov chain converges to the stationary distribution. That is, we use Lemma 11, with the vector

$$v : o_j \mapsto \int_{\omega\in[0,1]^{V(G)}} f(G,o_j,\omega)d\mu.$$

We get that there exists some $C_1 \in \mathbb{R}$ such that

$$\left|\sum_{j=1}^{t} p_k^{M_T}(o_i,o_j)\int_{\omega\in[0,1]^{V(G)}} f(G,o_j,\omega)d\mu - \sum_{j=1}^{t}\widetilde{p}(o_j)\int_{\omega\in[0,1]^{V(G)}} f(G,o_j,\omega)\,d\mu\right| \le C_1\rho_T^k.$$

Note that C_1 might depend on f, but not on k. The first property of f says the second term in the absolute value is 0, so we have

$$\left|\sum_{j=1}^{t} p_k^{M_T}(o_i,o_j)\int_{\omega\in[0,1]^{V(G)}} f(G,o_j,\omega)d\mu\right| \le C_1\rho_T^k.$$

We use this to bound the term (3.3):

$$\left|\sum_{i=1}^{t}\widetilde{p}(o_i)\int_{\omega\in[0,1]^{V(G)}} f(G,o_i,\omega)d\mu \sum_{j=1}^{t} p_k^{M_T}(o_i,o_j)\int_{\omega\in[0,1]^{V(G)}} f(G,o_j,\omega)d\mu\right|$$

$$\le \sum_{i=1}^{t}\widetilde{p}(o_i)\left|\int_{\omega\in[0,1]^{V(G)}} f(G,o_i,\omega)d\mu\right| \left|\sum_{j=1}^{t} p_k^{M_T}(o_i,o_j)\int_{\omega\in[0,1]^{V(G)}} f(G,o_j,\omega)d\mu\right|$$

$$\le \sum_{i=1}^{t}\widetilde{p}(o_i)\left|\int_{\omega\in[0,1]^{V(G)}} f(G,o_i,\omega)d\mu\right| C_1\rho_T^k = C_2\rho_T^k.$$

To recap, we had $\langle \mathcal{M}^k f, f\rangle = (3.1) + (3.2) = (3.3) - (3.4) + (3.2)$. We have already bounded the absolute value of (3.3), so we now bound the absolute values of (3.2) and (3.4). These terms, however, correspond to cases where the random walk on G arrives close to the starting point after k steps. As G is non-amenable, the probability of this happening decays exponentially in k.

Formally, let us recall that $p_k(o_i,y) \le C_0(o_i,y,\varepsilon)(\rho+\varepsilon)^k$. We write

$$|(3.2)| = \left|\sum_{i=1}^{t}\widetilde{p}(o_i)\sum_{y\in B_{2r}(o_i)} p_k(o_i,y)\int_{\omega\in[0,1]^{V(G)}} f(G,o_i,\omega)f(G,y,\omega)d\mu\right|$$

$$\le (\rho+\varepsilon)^k \underbrace{\sum_{i=1}^{t}\widetilde{p}(o_i)\sum_{y\in B_{2r}(o_i)} C_0(o_i,y,\varepsilon)\left|\int_{\omega\in[0,1]^{V(G)}} f(G,o_i,\omega)f(G,y,\omega)d\mu\right|}_{C_3},$$

$$\left|(3.4)\right| = \left|\sum_{i=1}^{t} \widetilde{p}(o_i) \sum_{y \in B_{2r}(o_i)} p_k(o_i, y) \int_{\omega \in [0,1]^{V(G)}} f(G, o_i, \omega) d\mu \int_{\omega \in [0,1]^{V(G)}} f(G, y, \omega) d\mu\right|$$

$$\leq (\rho + \varepsilon)^k \underbrace{\sum_{i=1}^{t} \widetilde{p}(o_i) \sum_{y \in B_{2r}(o_i)} C_0(o_i, y, \varepsilon) \left|\int_{\omega \in [0,1]^{V(G)}} f(G, o_i, \omega) d\mu \int_{\omega \in [0,1]^{V(G)}} f(G, y, \omega) d\mu\right|}_{C_4}.$$

Note that the constants C_3 and C_4 depend on f, but not on k. We now combine our bounds and get

$$\limsup_{k\to\infty} \left|\langle \mathcal{M}^k f, f\rangle\right|^{1/k} \leq \limsup_{k\to\infty} \left(C_2 \rho_T^k + (C_3 + C_4)(\rho + \varepsilon)^k\right)^{1/k}$$
$$= \lim_{k\to\infty} \left(C_2 \rho_T^k + (C_3 + C_4)(\rho + \varepsilon)^k\right)^{1/k} = \max(\rho_T, \rho + \varepsilon).$$

This holds for any $\varepsilon > 0$, so we have $\limsup_{k\to\infty} \left|\langle \mathcal{M}^k f, f\rangle\right|^{1/k} \leq \max(\rho_T, \rho)$. By Lemma 16, we now have $\|\mathcal{M}|_{L_0^2(\Omega,\nu_{\mathrm{st}})}\| \leq \max(\rho_T, \rho)$, which completes the proof. □

Proof of Theorem 18. Note that by bipartiteness, the subspaces $\mathbb{1}_{X_1}^{\perp}$ and $\mathbb{1}_{X_2}^{\perp}$ of $L^2(\Omega)$ are invariant under the action of $\mathcal{M}^2$, so $\mathcal{M}^2$ is well-defined as an operator on $S_1 = \{f \in L_{00}^2(\Omega, \nu_{\mathrm{st}}) : f|_{X_2} \equiv 0\}$ and on $S_2 = \{f \in L_{00}^2(\Omega, \nu_{\mathrm{st}}) : f|_{X_1} \equiv 0\}$. Moreover, $L_{00}^2(\Omega, \nu_{\mathrm{st}})$ can be written as the internal direct sum

$$L_{00}^2(\Omega, \nu_{\mathrm{st}}) = S_1 \oplus S_2,$$

so the spectrum of $\mathcal{M}^2|_{L_{00}^2(\Omega,\nu_{\mathrm{st}})}$ satisfies

$$\sigma\left(\mathcal{M}^2|_{L_{00}^2(\Omega,\nu_{\mathrm{st}})}\right) = \sigma\left(\mathcal{M}^2|_{S_1}\right) \cup \sigma\left(\mathcal{M}^2|_{S_2}\right).$$

M_T^2 restricted to T_1 or T_2 is not bipartite, so the proof of Theorem 17, applied to $\mathcal{M}^2$ on S_1 and on S_2, yields that the two spectral radii are both at most $\max(\rho^2, \rho_T^2)$.

Finally, we have $\sigma(\mathcal{M}^2|_{L_{00}^2}) = \{\lambda^2 \mid \lambda \in \sigma(\mathcal{M}_{L_{00}^2})\}$, which completes the proof. □

4. Perfect matchings in quasi-transitive graphs

In this section we prove Corollary 3. The hard work was done in Section 3 to establish our spectral-theoretic results, here we can mostly follow the proof of Lyons and Nazarov. In order to prove their main result, they obtain the necessary expansion properties from the spectral gap through [7, Lemma 2.3]. We recall this as Lemma 19 below. We also need to complement it with a measurably bipartite version, proved very similarly, which will be Lemma 20.

Lemma 19 ([7], Lemma 2.3)**.** *Let $\mathcal{G} = (X, E, \nu)$ be a graphing, and let $\rho_{\mathcal{G}} = \rho\left(\mathcal{M}|_{L_0^2(X,\nu_{st})}\right)$. Let $B \subseteq X$ be a measurable subset, and let $b = \nu_{st}(B)/\nu_{st}(X)$ denote the degree-biased density of B in X. Let $b' = \nu_{st}(N(B))/\nu_{st}(X)$ denote the degree-biased density of the neighbors of B in X. Then*

$$b' \geq \frac{1}{\rho_{\mathcal{G}}^2(1-b)+b} \cdot b.$$

Lemma 20. *Let $\mathcal{G} = (X_1, X_2, E, \nu)$ be a measurably bipartite graphing, and let $\rho_{\mathcal{G}} = \rho\left(\mathcal{M}|_{L_{00}^2(X,\nu_{st})}\right)$. Let $B \subseteq X_1$ be a measurable subset, and let $b = \nu_{st}(B)/\nu_{st}(X_1)$ denote the degree-biased density of B in X_1. Let $b' = \nu_{st}(N(B))/\nu_{st}(X_2)$ denote the degree-biased density of the neighbors of B in X_2. Then*

$$b' \geq \frac{1}{\rho_{\mathcal{G}}^2(1-b)+b} \cdot b.$$

The same holds for measurable subsets $B \subseteq X_2$.

Proof, following Lemma 2.3 in [7]. First we note that by the graphing condition (2.1), we must have $\nu_{\mathrm{st}}(X_1) = \nu_{\mathrm{st}}(X_2) = \frac{1}{2}$:

$$\nu_{\mathrm{st}}(X_1) = \frac{\int_{X_1} \deg(x) d\nu}{\int_X \deg(x) d\nu} = \frac{\int_{X_1} \deg_{X_2}(x) d\nu}{\int_X \deg(x) d\nu} = \frac{\int_{X_2} \deg_{X_1}(x) d\nu}{\int_X \deg(x) d\nu} = \frac{\int_{X_2} \deg(x) d\nu}{\int_X \deg(x) d\nu} = \nu_{\mathrm{st}}(X_2).$$

Since $\mathcal{M}\mathbb{1}_B$ is constant 0 on the complement of $B' = N(B)$, we have

$$\nu_{\mathrm{st}}(B) = \langle \mathbb{1}_B, \mathbb{1} \rangle = \langle \mathbb{1}_B, \mathcal{M}\mathbb{1} \rangle = \langle \mathcal{M}\mathbb{1}_B, \mathbb{1} \rangle = \langle \mathcal{M}\mathbb{1}_B, \mathbb{1}_{B'} \rangle.$$

Consequently,

$$\nu_{\mathrm{st}}(B)^2 = \langle \mathcal{M}\mathbb{1}_B, \mathbb{1}_{B'} \rangle^2 \leq \|\mathcal{M}\mathbb{1}_B\|^2 \cdot \|\mathbb{1}_{B'}\|^2 = \|\mathcal{M}\mathbb{1}_B\|^2 \cdot \nu_{\mathrm{st}}(B') = \|\mathcal{M}\mathbb{1}_B\|^2 \cdot \frac{b'}{2}. \tag{4.1}$$

We split $\mathbb{1}_B$ as follows: $\mathbb{1}_B = b\mathbb{1}_{X_1} + f_B$, where $f_B = \mathbb{1}_B - b\mathbb{1}_{X_1} = (1-b)\mathbb{1}_B + (-b)\mathbb{1}_{X_1 \setminus B}$. Notice that $f_B \perp \mathbb{1}$ and $f_B \perp \mathbb{1}_{X_1} - \mathbb{1}_{X_2}$, therefore $\|\mathcal{M}f_B\| \leq \rho_{\mathcal{G}} \cdot \|f_B\|$. Moreover,

$$\|f_B\|^2 = (1-b)^2 \cdot \nu_{\mathrm{st}}(B) + b^2 \cdot \nu_{\mathrm{st}}(X_1 \setminus B) = (1-b)^2 \cdot \frac{b}{2} + b^2 \cdot \frac{1-b}{2} = \frac{b(1-b)}{2}.$$

Now $\mathcal{M}\mathbb{1}_B = b \cdot \mathcal{M}\mathbb{1}_{X_1} + \mathcal{M}f_B = b\mathbb{1}_{X_2} + \mathcal{M}f_B$. Again, $\mathbb{1}_{X_2} \perp \mathcal{M}f_B$ because $\langle \mathbb{1}_{X_2}, \mathcal{M}f_B \rangle = \langle \mathcal{M}\mathbb{1}_{X_2}, f_B \rangle = \langle \mathbb{1}_{X_1}, f_B \rangle = 0$. Hence,

$$\|\mathcal{M}\mathbb{1}_B\|^2 = b^2 \|\mathbb{1}_{X_2}\|^2 + \|\mathcal{M}f_B\|^2 \leq b^2 \cdot \nu(X_2) + \rho_{\mathcal{G}}^2 \cdot \|f_B\|^2 = \frac{1}{2}\left(b^2 + \rho_{\mathcal{G}}^2 b(1-b)\right). \tag{4.2}$$

Putting (4.1) and (4.2) together, we get

$$b' \geq \frac{2\nu_{\mathrm{st}}(B)^2}{\|\mathcal{M}\mathbb{1}_B\|^2} = \frac{b^2}{2\|\mathcal{M}\mathbb{1}_B\|^2} \geq \frac{b^2}{b^2 + \rho_{\mathcal{G}}^2 b(1-b)} = \frac{1}{\rho_{\mathcal{G}}^2(1-b) + b} \cdot b. \quad \square$$

Proof of Corollary 3. Let M_T denote the finite state Markov chain defined by the quasi-transitive graph G described in Section 2.5. If M_T is not bipartite, we have spectral gap on $L_0^2(V(\mathcal{G}), \nu)$ by Theorem 17, which implies vertex expansion by Lemma 19, and Theorem 14 provides the perfect matching. If M_T is bipartite, the Bernoulli graphing $\mathcal{G}$ is measurably bipartite and has spectral gap by Theorem 18. This implies bipartite expansion by Lemma 20. The bipartite expansion implies the existence of a perfect matching by Theorem 15.

Note that we use the regularity of G, as it implies that the probability measures ν (used in Theorems 14 and 15) and ν_{st} (used in Lemmas 19 and 20) coincide. $\square$

Remark 21. Abért, Csóka, Lippner and Terpai show in [8] that any infinite transitive graph has a perfect matching. The following example shows that this is not true for quasi-transitive graphs. Therefore if we want to extend the result of Lyons and Nazarov on factor-of-iid perfect matchings beyond transitive graphs, assuming G to be bipartite is necessary.

Let G be any unimodular transitive non-amenable $2d$-regular graph, e.g. the tree T_{2d}. Let us now attach two pendant K_{2d+5}^- (the complete graph minus an edge) to every vertex of G so that the resulting graph $\widetilde{G}$ is $2d+4$-regular and has three orbits. $\widetilde{G}$ is quasi-isometric to G, and so it is non-amenable. To see that it is unimodular as well, we refer to [21], where it is shown that performing certain local changes preserves unimodularity. Every vertex v in $\widetilde{G}$ corresponding to an original vertex in G is now a cut vertex, and at least two of the components left in $\widetilde{G}$ when v is removed are finite and having odd order. $\widetilde{G}$ has therefore no perfect matching at all, let alone a factor-of-iid one.

5. Balanced orientations

In this subsection, we prove Theorem 1. We will use an auxiliary bipartite graph G^* whose perfect matchings correspond to balanced orientations of our graph G. This connection is implicit in Schrijver's paper about counting Eulerian orientations [5]. The first explicit constructions of pairs of graphs in which a balanced orientation of one is a perfect matching of the other were given by Mihail and Winkler [6]. The auxiliary graph G^* is constructed from G by local transformations, which makes sure that it is quasi-isometric to G.

Definition 22. Let G be a simple graph in which every vertex has an even degree. Then we define a simple graph G^* as follows (see also Fig. 1). G^* has a vertex for every edge $e \in E(G)$ and $\deg(v)/2$ vertices for every vertex $v \in V(G)$, i.e.

$$V(G^*) = \{x_e : e \in E(G)\} \cup \{v_i : v \in V(G), i \in [\deg(v)/2]\}.$$

Then every vertex corresponding to a former edge is joined to all copies of its former endpoints.

$$E(G^*) = \{x_{uv} v_i : uv \in E(G), i \in [\deg(v)/2]\}.$$

The vertices $x_e \in V(G^*)$ are called *edge-type* vertices of G^*, and $v_i \in V(G^*)$ are called *vertex-type* vertices. Any perfect matching M in G^* then defines a balanced orientation of G by orienting an edge $e \in E(G)$ towards its endpoint v if and only if x_e and v_i are matched by M for some $i \in [\deg(v)/2]$.

We now introduce the same construction starting from a graphing $\mathcal{G}$. Recall that for a graphing $(\mathcal{G}, \nu)$ we denote by ν_E the edge measure on $E(\mathcal{G})$.

Definition 23. Let $(\mathcal{G}, \nu)$ be a graphing with finite average degree $\overline{\deg} = 2\nu_E(E(\mathcal{G})) < \infty$ in which almost every vertex has even degree. Then we define the measurably bipartite auxiliary graphing $(\mathcal{G}^*, \nu^*)$ as follows.

- $V(\mathcal{G}^*) = X_1 \cup X_2$, where $X_1 = E(\mathcal{G})$ and $X_2 = \bigcup_{i=1}^{\infty} Y_i \times \{i\}$, where $Y_i = \{x \in V(\mathcal{G}) \mid \deg(x) \geq 2i\}$. Let us denote by $\pi : X_2 \to V(\mathcal{G})$ the projection onto the first coordinate.
- The measure ν^* is defined by

$$\nu^*|_{X_1} = \frac{1}{2\nu_E\big(E(\mathcal{G})\big)} \nu_E, \quad \nu^*|_{Y_i \times \{i\}} = \frac{1}{2\nu_E\big(E(\mathcal{G})\big)} \nu|_{Y_i}.$$

- For $e \in X_1, x \in X_2$ there is an edge $ex \in E(\mathcal{G}^*)$ connecting them if and only if $\pi(x) \in e$.

To check that $\mathcal{G}^*$ is indeed a graphing, we compute for any $A \subseteq X_1$ and $B \subseteq X_2$ that

$$\int_B \deg_A(v) d\nu^*(v) = \frac{\int_{V(\mathcal{G})} \left|\pi^{-1}(v) \cap B\right| \cdot \left|\{a \in A \mid v \text{ is incident to } a\}\right| d\nu(v)}{\int_{V(\mathcal{G})} \deg(u) d\nu(u)}$$
$$= \frac{\int_A |\pi^{-1}(u) \cap B| + |\pi^{-1}(v) \cap B| d\nu_E(uv)}{2\nu_E(E(\mathcal{G}))} = \int_A \deg_B(e) d\nu^*(e).$$

As in the discrete case, a measurable matching $M \subseteq E(\mathcal{G}^*)$ defines a measurable balanced orientation of $\mathcal{G}$ by orienting an edge $e \in E(\mathcal{G})$ towards its endpoint v if e and (v, i) are matched by M for some $i \in [\deg(v)/2]$.

We now go on to relate expansion properties of $\mathcal{G}$ to those of $\mathcal{G}^*$. Let us define the *Cheeger constant* of $\mathcal{G}$ as

$$\Phi_{\text{st}} = \inf \left\{ \frac{\int_S \deg_{N_{\mathcal{G}}(S) \setminus S}(u) d\nu(u)}{\nu_{\text{st}}(S)} \;\middle|\; 0 < \nu_{\text{st}}(S) \leq \frac{1}{2} \right\}.$$

Note that in this degree-biased version, we may have $\Phi_{\text{st}} > 0$ even when the set of isolated vertices has positive ν-measure.

Lemma 24. *Let $(\mathcal{G}, \nu)$ be a graphing with bounded average degree $\overline{\deg} < \infty$ and Cheeger constant $\Phi_{st}(\mathcal{G}) > 0$. Then $(\mathcal{G}^*, \nu^*)$ has bipartite expansion, that is, there is an $\varepsilon > 0$ such that for any $A \subseteq X_1$ and $B \subseteq X_2$, we have*

$$\nu^*\big(N_{\mathcal{G}^*}(A)\big) \geq \min\left\{(1 + \varepsilon)\, \nu^*(A), \frac{1}{4} + \varepsilon\right\} \quad \textit{and}$$

$$\nu^*\big(N_{\mathcal{G}^*}(B)\big) \geq \min\left\{(1+\varepsilon)\,\nu^*(B),\ \frac{1}{4}+\varepsilon\right\}.$$

In particular, $\varepsilon = \min\left\{\frac{3}{20}, \frac{\Phi_{st}(\mathcal{G})}{4\overline{\deg}}\right\}$ *satisfies this.*

Proof. First, we observe that for $B \subseteq X_2$ we have

$$\nu^*(B) \leq \frac{1}{2\nu_E\big(E(\mathcal{G})\big)} \int_{\pi(B)} \frac{\deg(v)}{2} d\nu(v) = \frac{1}{2}\nu_{\mathrm{st}}(\pi(B)).$$

For ease of notation we will write $B' = N_{\mathcal{G}^*}(B)$, $A' = N_{\mathcal{G}^*}(A)$, and $E = E(\mathcal{G})$. The set B' consists exactly of those edges of $\mathcal{G}$ that have at least one vertex in $\pi(B)$. Consequently,

$$\begin{aligned}
\nu^*(B') = \frac{1}{2\nu_E(E)}\nu_E(B') &= \frac{\frac{1}{2}\int_{\pi(B)} \deg(u)d\nu(u) + \frac{1}{2}\int_{\pi(B)'\setminus\pi(B)} \deg_{\pi(B)}(u)d\nu(u)}{2\nu_E(E)} \\
&\geq \frac{\frac{1}{2}\int_{\pi(B)} \deg(u)d\nu(u)}{\int_{V(\mathcal{G})} \deg(v)d\nu(v)} + \frac{\Phi_{\mathrm{st}} \min\{\nu_{\mathrm{st}}(\pi(B)), 1-\nu_{\mathrm{st}}(\pi(B))\}}{2\overline{\deg}} \\
&\begin{cases} \geq \nu^*(B) + \frac{\Phi_{\mathrm{st}}}{\overline{\deg}} \cdot \frac{\nu_{\mathrm{st}}(\pi(B))}{2} \geq \left(1+\frac{\Phi_{\mathrm{st}}}{\overline{\deg}}\right)\nu^*(B) & \text{if } \nu_{\mathrm{st}}(\pi(B)) \leq \frac{1}{2} \\
= \frac{1}{2}\nu_{\mathrm{st}}(\pi(B)) + \frac{\Phi_{\mathrm{st}}(1-\nu_{\mathrm{st}}(\pi(B)))}{2\overline{\deg}} \geq \frac{1}{4} + \frac{\Phi_{\mathrm{st}}}{4\overline{\deg}} & \text{if } \nu_{\mathrm{st}}(\pi(B)) \geq \frac{1}{2},
\end{cases}
\end{aligned}$$

where $\pi(B)' = N_{\mathcal{G}}\big(\pi(B)\big)$.

Second, let us consider $A \subseteq X_1$. In this case, A' is all possible lifts of the vertices induced by A in $\mathcal{G}$. That is, if $S \subseteq V(\mathcal{G})$ is the set of vertices that are incident to at least one edge from A, then $A' = \pi^{-1}(S)$. Thus

$$\begin{aligned}
\nu^*(A') &= \frac{1}{\int_{V(\mathcal{G})}\deg(v)d\nu(v)} \int_S |\pi^{-1}(u)|d\nu(u) = \frac{1}{2\nu_E(E)} \int_S \frac{1}{2}\deg(u)d\nu(u) \\
&\geq \frac{1}{2\nu_E(E)}\left(\frac{1}{2}\int_S \deg_S(u)d\nu(u) + \frac{1}{2}\Phi_{\mathrm{st}}\min\{\nu_{\mathrm{st}}(S), 1-\nu_{\mathrm{st}}(S)\}\right) \\
&\begin{cases} \geq \frac{\frac{1}{2}\int_S \deg_S(u)d\nu(u)}{2\nu_E(E)}\left(1+\frac{\Phi_{\mathrm{st}}}{\overline{\deg}}\right) \geq \nu^*(A)\left(1+\frac{\Phi_{\mathrm{st}}}{\overline{\deg}}\right) & \text{if } \nu_{\mathrm{st}}(S) \leq \frac{1}{2} \\
\geq \frac{\frac{1}{2}\int_S \deg_S(u)d\nu(u)}{2\nu_E(E)}\left(1+\frac{\Phi_{\mathrm{st}}(1-\nu_{\mathrm{st}}(S))}{\int_S \deg(u)d\nu(u)}\right) \geq \nu^*(A)\left(1+\frac{\Phi_{\mathrm{st}}}{\overline{\deg}}\cdot\frac{1-\nu_{\mathrm{st}}(S)}{\nu_{\mathrm{st}}(S)}\right) & \text{if } \nu_{\mathrm{st}}(S) \geq \frac{1}{2}.
\end{cases}
\end{aligned}$$

We hence have that $\nu^*(A') \geq \left(1+\frac{\Phi_{\mathrm{st}}}{4\overline{\deg}}\right)\nu^*(A)$ for all $A \subseteq X_1$ such that $\nu_{\mathrm{st}}(S) \leq \frac{4}{5}$. Moreover, $\nu^*(A') = \frac{1}{2}\nu_{\mathrm{st}}(S)$, which means that $\nu_{\mathrm{st}}(A') \geq \frac{1}{4} + \frac{3}{20}$ whenever $A \subseteq X_1$ is such that $\nu_{\mathrm{st}}(S) \geq \frac{4}{5}$. □

Proof of Theorem 1. We aim to find a factor-of-iid balanced orientation of the quasi-transitive graph G, that is we aim to find a measurable balanced orientation (up to nullsets) in its Bernoulli graphing $(\mathcal{G}, \nu)$.

The spectrum of the Markov operator $\mathcal{M}_{\mathcal{G}}$ restricted to $L_0^2\big(V(\mathcal{G}), \nu_{\mathrm{st}}\big)$ is bounded away from 1 (though not necessarily bounded away from -1). This is given for $\mathcal{M}_{\mathcal{G}}$ with non-bipartite M_T by Theorem 17. For $\mathcal{M}_{\mathcal{G}}$ with bipartite M_T, we deduce this by observing that $L_0^2\,(V(\mathcal{G}), \nu_{\mathrm{st}})$ can be written as the direct sum $L_{00}^2\,(V(\mathcal{G}), \nu_{\mathrm{st}}) \oplus \langle \mathbb{1}_{X_1} - \mathbb{1}_{X_2}\rangle$ and applying Theorem 18.

By a standard argument this spectral gap "at the top of the spectrum" implies that $\mathcal{G}$ has positive Cheeger constant. See e.g. [22, Proposition 3.3.6] for a formulation and proof for finite graphs that generalizes to graphings (with the appropriate vertex- and edge measures). Consequently by Lemma 24 the auxiliary graphing $\mathcal{G}^*$ has bipartite vertex expansion, which means it has a measurable perfect matching M by Theorem 15. Then M defines a measurable balanced orientation of $\mathcal{G}$ as described after Definition 23. □

6. Other decorations

6.1. Schreier decorations of T_{2d}

In this section we prove the four items of Proposition 4.

We start by pointing out that for T_{2d}, there are in fact unique $\mathrm{Aut}(T_{2d})$-invariant measures μ_{bo} and μ_{Sch} on the spaces $\mathtt{BalOr}(T_{2d})$ and $\mathtt{Sch}(T_{2d})$ respectively. The reason is that both the balanced orientation and Schreier decoration are essentially unique on T_{2d}, meaning that $\mathrm{Stab}_{\mathrm{Aut}(T_{2d})}(o)$ acts transitively on both $\mathtt{BalOr}(T_{2d})$ and $\mathtt{Sch}(T_{2d})$. Here o denotes an arbitrary root vertex in T_{2d}.

One can construct μ_{bo} and μ_{Sch} by starting at o and defining the balanced orientation or Schreier decoration on the incident edges uniformly at random. Then continue moving radially outwards through the vertices of T_{2d}, always extending the structure to the $2d-1$ outwards edges where it is not yet defined, doing so by choosing uniformly randomly among the possible extensions, independently at each vertex.

Note that μ_{bo} is a factor of μ_{Sch}, simply by forgetting the colors. In fact, there is an intermediate object, which we can obtain from μ_{Sch} by forgetting the order of the last two colors c_{d-1} and c_d. This gives μ_{Sch^*}, the unique invariant measure on $\mathtt{Sch}^*(T_d)$, the space of Schreier decorations of T_{2d} with the colors $\{c_{d-1}, c_d\}$ unordered. So the more detailed picture is that μ_{bo} is a factor of μ_{Sch^*}, which is itself a factor of μ_{Sch}.

Theorem 1 implies that μ_{bo} is a factor of iid. For $d > 1$, one could show that μ_{Sch} is a factor of iid if T_d had a factor of iid proper edge d-coloring. (However, the existence of such a coloring is an open question [1].)

Proof of (i). A balanced orientation of T_{2d} gives rise to a decomposition of the edges into infinitely many edge-disjoint d-regular subtrees, with each subtree having either only incoming or only outgoing edges at every vertex it covers. Each vertex is covered by exactly two such d-regular subtrees.

We construct a balanced orientation (and the resulting decomposition) as a factor of iid by Theorem 1. Furthermore, we can assume that each vertex v still has two independent uniform random labels $l_{\mathrm{in}}(v)$ and $l_{\mathrm{out}}(v)$ to be used in each of the two subtrees covering it. Then by using the assumed factor-of-iid proper edge d-coloring on each subtree, we obtain a Schreier decoration. □

In [1], Lyons presents a partial result towards constructing the unique invariant measure μ_{col} on proper edge colorings of T_d with d colors as a factor-of-iid. He obtains a factor-of-iid proper edge coloring, but with the last two colors being unordered. This allows us to prove part (ii), which states that even μ_{Sch^*} is a factor of iid.

Proof of (ii). We follow the construction of part (i), and use the factor-of-iid proper edge coloring with two colors unordered from [1] on the d-regular subtrees. To complete the construction, at every vertex of the tree, we have to match the colors of the $\{c_{d-1}, c_d\}$-colored incoming edges to the two outgoing $\{c_{d-1}, c_d\}$-colored edges. So each vertex chooses a random bijection between these incoming and outgoing edges, placing the paired edges in the same color class from $\{c_{d-1}, c_d\}$. □

Notice that the map forgetting the order of colors from $\mathtt{Sch}(T_{2d})$ to $\mathtt{Sch}^*(T_{2d})$ is a 2-to-1 cover. In a sense, we are only lacking a coin flip to find a Schreier decoration. However, this is exactly the kind of randomness that cannot be used when constructing factors of iid — vertices far away cannot generate a common random value because of correlation decay. In part (iii), we explicitly show that "finishing the construction" starting from a balanced orientation of T_4 is not possible.

Proof of (iii). Take two oriented edges $\vec{e}$ and $\vec{f}$ of $\vec{T}_4$ that are the first and the last edge on a path on which the orientation is alternating. Coloring one of them determines the color of the other in a Schreier decoration that respects the orientation. If the path consists of an odd number of edges then $\vec{e}$ and $\vec{f}$ have to have the same color.

On the other hand, the action of $\mathrm{Aut}(\vec{T}_4)$ is edge-transitive, which implies that if we pick $\vec{e}$ and $\vec{f}$ further and further apart, the correlation between their colors must decay. Hence, there can be no factor-of-iid Schreier decoration respecting the orientation. □

Finally we prove part (iv), namely that if μ_{Sch} is a factor of iid for T_{2d}, then it is a factor of iid also for T_{2d+2i} for all $i \in \mathbb{N}$.

Proof of (iv). Let us first construct two disjoint factor-of-iid perfect matchings on T_{2d+2} as in [7]. Then after disregarding the edges in these matchings, we are left with infinitely many copies of T_{2d}, in which we can find, by assumption, a factor-of-iid Schreier decoration with colors $c_1, \ldots, c_d$. Let us now in the tree T_{2d+2} disregard the edges colored with c_1, so that we again are left with infinitely many T_{2d}-s. We delete the $\{c_2, \ldots, c_d\}$, orientation, and matching decorations in these trees, and construct on them, anew, a Schreier decoration with colors $\{c_2, \ldots, c_{d+1}\}$. Together with the edges decorated with c_1, this gives a Schreier decoration of the tree T_{2d+2}. □

6.2. A connection to measured group theory

Part (ii) of Proposition 4 also has the following interpretation.

The $2d$-regular tree is the Cayley graph of F_d, the free group on d generators, but also of the group $(\mathbb{Z}/2\mathbb{Z})^{*2d}$, the $2d$-fold free product of $(\mathbb{Z}/2\mathbb{Z})$ with itself. A Schreier decoration corresponds to an action of F_d, while a proper edge coloring corresponds to an action $(\mathbb{Z}/2\mathbb{Z})^{*2d}$.

Let $\Gamma = (\mathbb{Z}/2\mathbb{Z})^{*2d}$. Consider the *Bernoulli shift* $F_d \curvearrowright \left([0,1]^{F_d}, \mathrm{u}^{F_d}\right)$, and similarly $\Gamma \curvearrowright \left([0,1]^{\Gamma}, \mathrm{u}^{\Gamma}\right)$. Let S and T denote the standard generating sets of F_2 and Γ respectively.

One can ask whether the two Bernoulli shifts are equivalent in the strong sense that there exists a measure-preserving bijection $\Phi : [0,1]^{F_d} \to [0,1]^{\Gamma}$ such that (on a subset of measure 1) whenever $s.\omega = \omega'$ for $\omega, \omega' \in [0,1]^{F_d}$, and $s \in S$, then there is some $t \in T$ such that $t.\Phi(\omega) = \Phi(\omega')$. (Note that this is much stronger than Orbit Equivalence, we require t to be from the finite generating set T. We require that the distances defined by the word length on the orbits are preserved.)

As far as the authors are aware, this question is open. The existence of such Φ would imply that F_d has a p.m.p. action on $[0,1]^{\Gamma}$ that defines the same distance on orbits as Γ and vice versa. So disproving the equivalence could be achieved by showing that one of these actions does not exist. This is a fruitful approach when considering the same problem for groups with Cayley graphs isomorphic to the square lattice.

The results of [1] and part (ii) of Proposition 4 respectively say that Γ acts on a 2-cover of $[0,1]^{F_d}$ defining the same distance on orbits, and F_d acts on a 2-cover of $[0,1]^{\Gamma}$ and defines the same distance on orbits.

6.3. Decorations of G^*

In this subsection, we further study the connection between balanced orientations of G and perfect matchings of the auxiliary graph G^*.

We will first finish proving the equivalence of a balanced orientation of G with a perfect matching on G^* started in Section 5, and then use the perfect matching to construct Schreier decorations and proper edge colorings of G^*.

Lemma 25. *Let G be a simple graph with all degrees even. There is a (finitary)* Aut(G)*-factor of iid balanced orientation of G if and only if there is a (finitary)* Aut(G^*)*-factor of iid perfect matching.*

Proof. Suppose G^* has a factor-of-iid perfect matching. Given random labels on $V(G)$, we can deterministically produce labels on $V(G^*)$ as we will describe below. We use the factor-of-iid perfect matching to deterministically compute matching M in G^*, which again deterministically defines a balanced orientation of G. As all the steps are Aut(G)-equivariant, their composition is a factor-of-iid balanced orientation of G.

By decomposing our original labels, we can assume that we have $\frac{3}{2}\deg(v)$ independent random labels at each $v \in V(G)$ at the beginning. We make each v give one of these labels to all the v_i as well as all x_e for edges e incident to v. Then each x_e takes the two labels it got from its endpoints and composes them to get a label. This way each vertex of $V(G^*)$ obtains a label. The joint distribution of these labels is uniform iid, which completes the construction in this direction.

On the other hand, suppose G has a factor-of-iid balanced orientation. Without loss of generality, we will assume that G is connected. If $G \neq P$ and $G \neq C_k$, then any $y \in V(G^*)$ can determine whether it is of edge-type or vertex-type. Also, if y is of vertex-type (say $y = v_i$), it can identify all other vertices of G^* that correspond to the same vertex of G as y (all vertices of the form v_j, $j \in [\deg(v)/2]$). If the v_j, $j \in [\deg(v)/2]$ compose their labels to get a label $l(v)$ for each $v \in V(G)$, then any $y \in V(G^*)$ of vertex-type can simulate the factor-of-iid balanced orientation on "its neighborhood in G". The balanced orientation determines which $\deg(v)/2$ vertices of the form x_{vu} get matched to the v_i. The v_i can together choose the matching between $\{v_i \mid i \in [\deg(v)/2]\}$ and $\{x_{vu} \mid uv \text{ is oriented towards } v \text{ in } G\}$ randomly, yielding a factor-of-iid perfect matching of G^*.

If $G = P$ is the bi-infinite path then $G^* = G$ and it has neither factor of iid perfect matching nor balanced orientation. If $G = C_k$ for some $k \in \mathbb{N}$ then $G^* = C_{2k}$, and so there is both a balanced orientation on G and a perfect matching on G^*. □

Remark 26. From a more algebraic point of view, in the proof above, we use the fact that $\mathrm{Aut}(G)$ acts on $[0,1]^{V(G^*)}$. There is a natural embedding $\varphi : \mathrm{Aut}(G) \to \mathrm{Aut}(G^*)$, which in turn defines the translation action of $\mathrm{Aut}(G)$ on $[0,1]^{V(G^*)}$. By decomposing and combining the labels as explained, we have in fact shown that $\mathrm{Aut}(G) \curvearrowright \left([0,1]^{V(G^*)}, \mathrm{u}^{V(G^*)}\right)$ is a factor of $\mathrm{Aut}(G) \curvearrowright \left([0,1]^{V(G)}, \mathrm{u}^{V(G)}\right)$. We can then utilize the existence of an $\mathrm{Aut}(G^*)$-factor of iid perfect matching of G^* and the correspondence with balanced orientations of G to finish the proof by composing the appropriate factor maps.

In the other direction, we aim to build a factor map from $\mathrm{Aut}(G^*) \curvearrowright \left([0,1]^{V(G^*)}, \mathrm{u}^{V(G^*)}\right)$ to $\mathrm{Aut}(G^*) \curvearrowright \left(\mathtt{PM}(G^*), \mu_{\mathrm{pm}}\right)$ through the factor from $\mathrm{Aut}(G) \curvearrowright \left([0,1]^{V(G)}, \mathrm{u}^{V(G)}\right)$ to $\mathrm{Aut}(G) \curvearrowright \left(\mathtt{BalOr}(G), \mu_{\mathrm{bo}}\right)$. But in order to do that we have to consider $\left([0,1]^{V(G)}, \mathrm{u}^{V(G)}\right)$ as an $\mathrm{Aut}(G^*)$-space. This is possible exactly when G has a vertex of degree at least 4, or equivalently when vertices of G^* can determine their type. In this case every element of $\mathrm{Aut}(G^*)$ is an element of $\mathrm{Aut}(G)$ up to permuting the sets $\{v_i \mid i \in [\deg(v)/2]\}$.

As an immediate corollary, we get factor-of-iid perfect matchings on G^*.

Corollary 27. *Let G be a unimodular, quasi-transitive, non-amenable graph with all degrees even. Then G^* is a unimodular, quasi-transitive, non-amenable graph that has a factor-of-iid perfect matching.*

Proof. Follows from Theorem 1 and Lemma 25. □

Note that even though we obtained the factor-of-iid balanced orientation in Theorem 1 through perfect matchings, there we used the auxiliary graphing $\mathcal{G}^*$ of the Bernoulli graphing $\mathcal{G}$ (of G). Whereas here we claim that the Bernoulli graphing of the graph G^* has a measurable balanced orientation.

We are now also ready to prove Proposition 5. For the reader's convenience we restate it here.

Proposition 5. *For every $2d$-regular graph G, the bipartite graph G^* is also $2d$-regular, and the following are equivalent.*

1. *G^* has got a factor-of-iid proper edge $2d$-coloring.*
2. *G^* has got a factor-of-iid perfect matching.*
3. *G^* has got a factor-of-iid Schreier decoration.*

Moreover, if any of these is a finitary factor, the others are too.

Even though proving three implications would be enough, we show five to emphasize the techniques that could be used more widely for other suitable bipartite graphs too.

Proof. Let v be a vertex of G whose neighbors are $u^1, \ldots, u^{2d}$. Then for every $i \in \left[\frac{\deg(v)}{2}\right] = [d]$, the neighbors of v_i in G^* are exactly $x_{vu^1}, \ldots, x_{vu^{2d}}$. Also for any edge uv in G, the neighbors of x_{uv} in

G^* are $u_1, \ldots, u_d, v_1, \ldots, v_d$, and so G^* is also $2d$-regular. Let us denote by A_V the set of vertices of G^* that are of vertex-type, and by A_E the set of vertices of edge-type.

$1 \implies 2$. Choose one of the $2d$ color classes to obtain a perfect matching.

$3 \implies 2$. Every finite bipartite $2d$-regular graph has a perfect matching and choosing one at random is a factor-of-iid process, so we can assume G is infinite. $P^* = P$ does not admit a factor-of-iid Schreier decoration, so let us suppose that $d \geq 2$. Then as in the proof of Lemma 25, every vertex can determine whether it belongs to $A_V \subset V(G^*)$ or $A_E \subset V(G^*)$. To obtain a perfect matching, let us fix a color c of the decoration and let each $x \in A_V$ pick the outgoing edge of color c and each $x \in A_E$ the incoming edge of color c.

$3 \implies 1$. Suppose the Schreier decoration uses colors $c_1, \ldots, c_d$ and that we want to produce proper coloring with colors $c'_1, \ldots, c'_{2d}$. Similarly as in the proof of $3 \implies 2$, let each edge of color c_i going from A_E to A_V get color c'_{2i} and each edge of color c_i going from A_V to A_E the color c'_{2i-1}.

$2 \implies 3$. For every $v \in V(G)$, the d copies of v in G^* together with the d vertices they are matched to induce a $K_{d,d}$. Let us note that the collection of these $K_{d,d}$-s is vertex-disjoint. Let us randomly pick a proper edge d-coloring on each of these $K_{d,d}$-s and orient all their edges from A_E to A_V. After removing the decorated edges, we are again left with a collection of vertex-disjoint $K_{d,d}$-s, this time in each of which one part is formed by $v_1, \ldots, v_d$ for some $v \in V(G)$ and the other by the neighbors of v_i, $i \in [d]$ that are matched towards some u_j where $uv \in E(G)$. Each of these $K_{d,d}$-s again picks a proper d-coloring at random, but this time we will orient each edge from A_V to A_E.

$1 \implies 3$. Suppose $E(G^*)$ is colored with $c_1, \ldots, c_{2d}$. Let every edge of color $c_1, c_3, \ldots$ retain it and become oriented from A_E to A_V. Then all edges of a color c_i, $i \in [d]$ will get recolored to c_i and get oriented from A_V to A_E. □

7. Open questions

Question 28. Is the unique $\mathrm{Aut}(T_{2d})$-invariant measure μ_{Sch} on $\mathtt{Sch}(T_{2d})$ a factor of iid?

We believe that this question, which has already been asked in [23] for the case $d = 2$, is the most natural and important one at this point. A very similar question asking for any Cayley diagram, not just of F_d, as a factor of iid on the regular tree was asked by Thornton [4, Problem 4.16]. Our positive examples of Schreier decorations in [3] so far seem fundamentally different from the tree in the sense that none of them even have infinite monochromatic paths, which would be automatic on T_{2d}. This leads us to the following question (also included in [3]).

Question 29. Is there a factor-of-iid Schreier decoration on a unimodular transitive graph that has infinite monochromatic paths with positive probability?

The Schreier decoration of T^*_{2d} obtained from a factor-of-iid perfect matching according to the $2 \implies 3$ part of Proposition 5 has infinite monochromatic paths, but T^*_{2d} is not transitive.

The following is the question discussed in Section 6.2. We encountered it during personal communication with Matthieu Joseph.

Question 30. Is there a measurable bijection Φ between the Bernoulli shifts of the free group F_d and the free product $(\mathbb{Z}/2\mathbb{Z})^{*2d}$ that preserves the distance defined by word length on almost all orbits?

Regarding our spectral result on quasi-transitive unimodular graphs, a natural question is to ask for an extension to unimodular random graphs.

Question 31. Let (G, o) be an invariantly non-amenable (a.k.a. non-hyperfinite) unimodular random rooted graph, and let $\mathcal{G}$ denote the Bernoulli graphing on (G, o). Does the Markov operator $\mathcal{M}$ on $\mathcal{G}$ have spectral gap? Maybe under some stronger assumption of non-amenability?

Addendum. After the first version of this paper was made available online, Abért, Fraczyk, and Hayes answered Question 31 negatively. They construct a unimodular random rooted graph that is non-amenable almost surely, but its Bernoulli graphing does not have spectral gap.

Acknowledgments

The authors would like to thank Miklós Abért, Jan Grebík, Matthieu Joseph, Gábor Kun, Gábor Pete, and Václav Rozhoň for inspiring discussions about various parts of this work. The first author was supported by the NKFIH (National Research, Development and Innovation Office, Hungary) grant KKP-133921, "Structure, chaos and simplification". The second and third authors were partially supported by the NKFIH grant KKP-139502, "Groups and graph limits".

References

[1] Russell Lyons, Factors of IID on trees, Combin. Probab. Comput. 26 (2) (2017) 285.
[2] László Márton Tóth, Invariant Schreier decorations of unimodular random networks, Ann. Henri Lebesgue 4 (2021) 1705–1726.
[3] Ferenc Bencs, Aranka Hrušková, Lászlo Márton Tóth, Factor of iid Schreier decoration of lattices in Euclidean spaces, 2021, Preprint arXiv:2101.12577.
[4] Riley Thornton, Factor maps for automorphism groups via Cayley diagrams, 2020, arXiv preprint arXiv:2011.14604.
[5] Alexander Schrijver, Bounds on the number of Eulerian orientations, Combinatorica 3 (3–4) (1983) 375–380.
[6] Milena Mihail, Peter Winkler, On the number of Eulerian orientations of a graph, Algorithmica (1996).
[7] Russell Lyons, Fedor Nazarov, Perfect matchings as IID factors on non-amenable groups, European J. Combin. 32 (7) (2011) 1115–1125.
[8] Endre Csóka, Gábor Lippner, Invariant random perfect matchings in Cayley graphs, Groups, Geom., Dyn. 11 (1) (2017) 211–244.
[9] Riley Thornton, Orienting Borel graphs, Proc. Amer. Math. Soc. 150 (04) (2022) 1779–1793.
[10] Ágnes Backhausz, Balázs Szegedy, Bálint Virág, Ramanujan graphings and correlation decay in local algorithms, Random Struct. Algorithms 47 (3) (2015) 424–435.
[11] Harry Kesten, Full Banach mean values on countable groups, Math. Scand. (1959) 146–156.
[12] Alkida Balliu, Brandt Sebastian, Yuval Efron, Juho Hirvonen, Yannic Maus, Dennis Olivetti, Jukka Suomela, Classification of distributed binary labeling problems, in: International Symposium on Distributed Computing, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020.
[13] Manuela Fischer, Mohsen Ghaffari, Sublogarithmic distributed algorithms for lovász local lemma, and the complexity hierarchy, in: 31st International Symposium on Distributed Computing (DISC 2017), Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.
[14] Gábor Kun, Expanders have a spanning Lipschitz subgraph with large girth, 2013, Preprint arXiv:1303.4982.
[15] Anton Bernshteyn, Measurable versions of the Lovász Local Lemma and measurable graph colorings, Adv. Math. 353 (2019) 153–223.
[16] Russell Lyons, Yuval Peres, Probability on Trees and Networks, Vol. 42, Cambridge University Press, 2017.
[17] David Aldous, Russell Lyons, Processes on unimodular random networks, Electron. J. Probab. 12 (2007) 1454–1508.
[18] László Lovász, Large Networks and Graph Limits, Vol. 60, American Mathematical Soc., 2012.
[19] Matthew Bowen, Gábor Kun, Marcin Sabok, Perfect matchings in hyperfinite graphings, 2021, arXiv preprint arXiv:2106.01988.
[20] Oleg Pikhurko, Borel combinatorics of locally finite graphs, 2020, Preprint arXiv:2009.09113.
[21] Dorottya Beringer, Gábor Pete, Ádám Timár, On percolation critical probabilities and unimodular random graphs, Electron. J. Probab. 22 (2017) 26.
[22] Emmanuel Kowalski, An Introduction To Expander Graphs, Société Mathématique de France, 2019.
[23] Karen Ball, Factors of independent and identically distributed processes with non-amenable group actions, Ergodic Theory Dynam. Systems (2005) 711–730.

Appendix B

Factor-of-iid Schreier decorations of lattices in Euclidean spaces

by Ferenc Bencs, Aranka Hrušková, László Márton Tóth

Discrete Mathematics 347 (2024) 114056

Contents lists available at ScienceDirect

Discrete Mathematics

journal homepage: www.elsevier.com/locate/disc

Factor-of-iid Schreier decorations of lattices in Euclidean spaces

Ferenc Bencs [a], Aranka Hrušková [b,c,*], László Márton Tóth [c]

[a] University of Amsterdam, Netherlands
[b] Department of Mathematics and its Applications, Central European University, Budapest, Hungary
[c] Alfréd Rényi Institute of Mathematics HUN-REN, Budapest, Hungary

ARTICLE INFO

Article history:
Received 1 February 2023
Received in revised form 31 March 2024
Accepted 18 April 2024
Available online xxxx

Keywords:
Transitive graph
Factor of iid
Schreier graph
Site percolation
Archimedean lattice
Planar lattice

ABSTRACT

A Schreier decoration is a combinatorial coding of an action of the free group F_d on the vertex set of a $2d$-regular graph. We investigate whether a Schreier decoration exists on various countably infinite transitive graphs as a factor of iid. We show that $\mathbb{Z}^d, d \geq 3$, the square lattice and also the three other Archimedean lattices of even degree have finitary-factor-of-iid Schreier decorations, and, answering a question of Thornton, exhibit examples of transitive graphs of arbitrary even degree in which obtaining such a decoration as a factor of iid is impossible.

We also prove that symmetrical planar lattices with all degrees even have a factor-of-iid balanced orientation, meaning the indegree of every vertex is equal to its outdegree, and demonstrate the existence of transitive graphs G whose classical chromatic number $\chi(G)$ is equal to their factor-of-iid chromatic number.

1. Introduction

Let G be a simple connected $2d$-regular graph. A *Schreier decoration* of G is a colouring of the edges with d colours together with an orientation such that at every vertex, there is exactly one incoming and one outgoing edge of each colour.

It is a folklore result in combinatorics that every finite $2d$-regular graph has a Schreier decoration, and it is easy to extend this to infinite $2d$-regular graphs as well. As a generalisation of the finite result, the third author proved in [29] that all $2d$-regular unimodular random rooted graphs admit an invariant random Schreier decoration. It is natural to ask whether such an invariant random Schreier decoration can be obtained as a factor of iid. In this article, we investigate this question for some infinite deterministic graphs, more specifically the d-dimensional Euclidean grids, symmetrical planar lattices, and graphs quasi-isometric to the bi-infinite path P of the form $H \times P$. We study non-amenable graphs in a separate paper [2].

Informally speaking, a Schreier decoration is a factor of iid if it is produced by a certain randomised local algorithm. To start with, each vertex of G gets a random label from $[0,1]$ independently and uniformly. Then it makes a deterministic measurable decision about the Schreier decoration of its incident edges, based on the labelled graph that it sees from itself as a root. Neighbouring vertices must make a consistent decision regarding the edge between them. The factor is *finitary* if the decision is based only on a random finite-radius neighbourhood of each vertex. The precise definition is given in Section 2. Obtaining combinatorial structures or certain models in statistical mechanics as factors of iid is a central topic in ergodic theory. See [19] and the references therein for a recent overview.

* Corresponding author at: Alfréd Rényi Institute of Mathematics HUN-REN, Budapest, Hungary.
E-mail addresses: ferenc.bencs@gmail.com (F. Bencs), umim.cist@gmail.com (A. Hrušková), toth.laszlo.marton@renyi.hu (L.M. Tóth).

https://doi.org/10.1016/j.disc.2024.114056

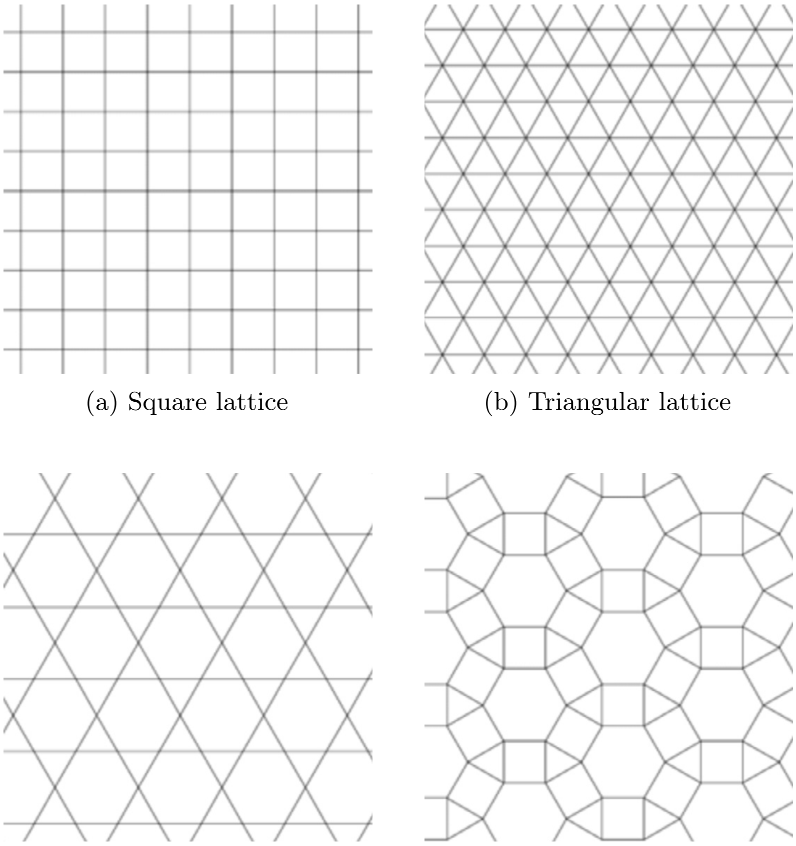

(a) Square lattice (b) Triangular lattice (c) Kagomé lattice (d) $(3,4,6,4)$ lattice

Fig. 1. Archimedean lattices of even degree.

A partial result towards a Schreier decoration on a $2d$-regular graph is a balanced orientation of the edges. An orientation of the edges of a graph with all degrees even is *balanced* if the indegree of any vertex is equal to its outdegree. For finite graphs, the term *Eulerian orientation* is often used [23], and when G is 4-regular in particular, a balanced orientation is known as an ice configuration in statistical physics [30,1]. Every Schreier decoration gives a balanced orientation by forgetting the colours.

The main results of the present paper are the following.

Theorem 1.1. *Let Λ be any of the four Archimedean lattices with even degrees (pictured in Fig. 1): the square lattice, the triangular lattice, the Kagomé lattice or the $(3,4,6,4)$ lattice. There is a finitary $\mathrm{Aut}(\Lambda)$-factor of $\left([0,1]^{V(\Lambda)}, \mu_\Lambda\right)$ which is a.s. a Schreier decoration of Λ. Moreover, it has almost surely no infinite monochromatic paths.*

Theorem 1.2. *Let $\Lambda_\square^d$ denote the d-dimensional Euclidean grid. For every $d \geq 3$, there is a finitary $\mathrm{Aut}(\Lambda_\square^d)$-factor of iid which is a.s. a Schreier decoration of $\Lambda_\square^d$. Moreover, it has almost surely no infinite monochromatic paths.*

As far as we are aware, this had not been known, not even in the case of the square lattice. Ray and Spinka show that balanced orientations (under the name “6-vertex-model”) exist on the square lattice as finitary factors of iid [21, Remark 5.3.]. Schreier decorations (that would be the “24-vertex model” in their terminology) are not investigated. In [25], Thornton studies factors of iid which are approximate and exact Cayley diagrams, but we do not allow for a small-probability local error here.

Parts of the proof have to be adapted to the individual lattices, but our approach remains the same throughout. First we break the lattices into a hierarchy of finite pieces (also called toasts elsewhere in the literature, e.g. in [17]) using percolation theory. Similar hierarchies or their weaker forms called cell processes were already employed in a measurable setting by Holroyd, Schramm, and Wilson [18] and by Spinka [24] and in a Borel setting by Gao, Jackson, Krohne, and Seward [15,16] and by Marks and Unger [20] who all build them from sequences of Voronoi cells with sparser and sparser centres. Then for each piece independently, we choose an edge-d-colouring scheme such that we can ensure that every monochromatic connected subgraph is a finite cycle. Each cycle will orient itself strongly. We also use a similar hierarchy argument to find balanced orientations on symmetrical planar lattices; see Theorem 3.8.

Next we show examples in which it is impossible to obtain a Schreier decoration as a factor of iid.

Theorem 1.3. *For every $d \geq 1$, there exists a $2d$-regular transitive graph that has no factor-of-iid balanced orientation. In particular, it has no factor of iid Schreier decoration.*

Theorem 1.3 answers Thornton's question [26, Problem 3.4]. However, the examples we construct are somewhat unsatisfactory because they are all quasi-isometric to the bi-infinite path P. We do not know of any examples of transitive graphs with no factor-of-iid Schreier decoration that are not quasi-isometric to P.

Furthermore, Theorems 1.1 and 1.2 can be phrased in terms of Bernoulli graphings over the respective graphs, yielding the following corollary.

Corollary 1.4. *The Bernoulli graphings of the Archimedean lattices as well as that of $\Lambda_{\square}^d$ admit a probability-measure-preserving action of the free group F_d which satisfies that two vertices x, y of the graphing are adjacent if and only if there is a generator γ of F_d such that $\gamma.x = y$.*

This connection is another important motivation for studying factor-of-iid Schreier decorations because we would like to understand which $2d$-regular graphings are so generated by actions of F_d. Our main question can be equivalently formulated as follows: given a (transitive unimodular) $2d$-regular graph G, is the Bernoulli graphing of G generated by a p.m.p. action of F_d? The connection is spelled out in our parallel paper [2], where we investigate non-amenable graphs and utilise it in the other direction. Our main result there is the following.

Theorem 1.5 *([2]). Every non-amenable quasi-transitive unimodular $2d$-regular graph has a factor of iid that is almost surely a balanced orientation.*

The existence of a factor-of-iid balanced orientation of non-amenable Cayley graphs, a prime example of unimodular graphs, was also showed by Thornton in [26].

The structure of the paper is as follows. In Section 2, we give definitions, present examples for Theorem 1.3, and treat invariance under quasi-isometry. Section 3 is concerned with general planar lattices, breaking them into clusters organised in a hierarchy and obtaining balanced orientations. We also give a proof, based on ideas of [7], of that the existence of cell-processes implies the existence of a hierarchy. Building on some of these results, Section 4 gives the proofs of Theorem 1.1 separately for each lattice and of Theorem 1.2. In Section 5, we explore what other combinatorial structures we can obtain using the results and ideas presented thus far. Open questions are collected in Section 6.

2. Notation and basics

We call a graph G *transitive* if and only if it is vertex-transitive, i.e., the automorphism group $\mathrm{Aut}(G)$ acts transitively on the vertex set $V(G)$. We call G *edge-transitive* whenever $\mathrm{Aut}(G)$ acts transitively on $E(G)$.

For the rest of the paper, for every $d \geq 2$, let $\Lambda_{\square}^d$ stand for the $2d$-regular d-dimensional grid, that is the Cayley graph of $\mathbb{Z}^d$ with the standard generators after forgetting the labelling. In the planar case, we also use simply $\Lambda_{\square}$ in place of $\Lambda_{\square}^2$.

2.1. Schreier graphs

Given a finitely generated group $\Gamma = \langle S \rangle$ and an action $\Gamma \curvearrowright X$ on some set X, the *Schreier graph* $\mathrm{Sch}(\Gamma \curvearrowright X, S)$ of the action is defined as follows. The set of vertices is X, and for every $x \in X$, $s \in S$, we introduce an oriented s-labelled edge from x to $s.x$.

Rooted connected Schreier graphs of Γ are in one-to-one correspondence with pointed transitive actions of Γ, which in turn are in one-to-one correspondence with subgroups of Γ. Trivially, a graph with a Schreier decoration is a Schreier graph of the free group F_d on d generators. A special case is the (left) *Cayley graph* of Γ, denoted $\mathrm{Cay}(\Gamma, S)$, which is the Schreier graph of the (left) translation action $\Gamma \curvearrowright \Gamma$.

2.2. Factors of iid

Let Γ be a group. A Γ-*space* is a measurable space X with an action $\Gamma \curvearrowright X$. A map $\Phi : X \to Y$ between two Γ-spaces is a Γ-*factor* if it is measurable and Γ-equivariant, that is $\gamma.\Phi(x) = \Phi(\gamma.x)$ for every $\gamma \in \Gamma$ and $x \in X$. (Contrary to the convention in dynamics but in line with the probabilistic literature on factors of iid, we do not require Φ to be surjective.)

A measure μ on a Γ-space X is *invariant* if $\mu(\gamma.A) = \mu(A)$ for all $\gamma \in \Gamma$ and all measurable $A \subseteq X$. We say an action $\Gamma \curvearrowright (X, \mu)$ is *probability-measure-preserving* (p.m.p.) if μ is a Γ-invariant probability measure.

Let G be a graph and $\Gamma \leq \mathrm{Aut}(G)$. Let u denote the Lebesgue measure on $[0, 1]$. We endow the space $[0, 1]^{V(G)}$ with the product measure $\mathrm{u}^{V(G)}$. The translation action $\Gamma \curvearrowright [0, 1]^{V(G)}$ is defined by

$$(\gamma.f)(v) = f(\gamma^{-1}.v),\ \forall \gamma \in \Gamma, v \in V(G).$$

The action $\Gamma \curvearrowright ([0, 1]^{V(G)}, \mathrm{u}^{V(G)})$ is p.m.p.

An orientation of G can be thought of as a function on $E(G)$ sending every edge to one of its endpoints. Viewed like this, orientations of G form a measurable function space $\mathtt{Or}(G)$ on which Γ acts. The set $\mathtt{BalOr}(G) \subseteq \mathtt{Or}(G)$ of balanced orientations is Γ-invariant and measurable, so it is a Γ-space in itself. Similarly, the set of all Schreier decorations of G forms the Γ-space $\mathtt{Sch}(G)$.

Definition 2.1. A Γ-factor of iid Schreier decoration (respectively, balanced orientation) of a graph G is a Γ-factor $\Phi : \big([0,1]^{V(G)}, \mathrm{u}^{V(G)}\big) \to \mathtt{Sch}(G)$ (respectively, to $\mathtt{BalOr}(G)$). If the group Γ is not specified, we mean an $\mathrm{Aut}(G)$-factor.

Remark 2.2. We allow Φ to not be defined on a $\mathrm{u}^{V(G)}$-null subset $X_0 \subseteq [0,1]^{V(G)}$.

Let us now recall some special classes of iid processes on graphs. For a fixed vertex $x \in V(G)$, let $\big(\Phi(\omega)\big)(x)$ denote the restriction of $\Phi(\omega)$ to the edges incident to x. We say Φ is a *finitary* factor of iid if for almost all $\omega \in [0,1]^{V(G)}$, there exists an $R \in \mathbb{N}$ such that $\big(\Phi(\omega)\big)(x)$ is already determined by $\omega|_{B_G(x,R)}$. That is, if we change ω outside $B_G(x,R)$, the decoration $\Phi(\omega)$ does not change around x. This radius R can depend on the particular ω. If it does not then we say Φ is a *block factor*.

When constructing factors of iid algorithmically, one often makes use of the fact that a uniform $[0,1]$ random variable can be decomposed into countably many independent uniform $[0,1]$ random variables. In practice, this means that we can assume that a vertex has multiple labels or that a new independent random label is always available after a previous one was used.

We will use a reverse operation as well: the composition of countably many uniform $[0,1]$ random variables is again a uniform $[0,1]$ random variable. When the number of variables combined is finite, one can do this in a permutation-invariant way. In particular, this allows finite graphs to make joint random decisions as factors of iid by composing the labels of their vertices. So any $\mathrm{Aut}(G)$-invariant random process on a finite graph G is a factor of iid. We frequently make use of this on finite subgraphs of our infinite graphs.

For $\mathrm{Aut}(G)$-factors of iid on an infinite transitive graph G, the fact that the factor map intertwines the actions of Γ implies that the further two vertices are from each other, the more independently the process behaves around them. We make use of the existence of this correlation decay both later in this section and in Section 6.

2.3. Graphs quasi-isometric to P

In this subsection, we present regular graphs of arbitrary even degree with the same large-scale geometry which in one case have factor-of-iid Schreier decorations but in the other not even a factor-of-iid balanced orientation. In both cases, our examples are quasi-isometric to the bi-infinite path P, that is the graph with $V(P) = \{v_i : i \in \mathbb{Z}\}$ in which v_i and v_j are adjacent if and only if $|i-j| = 1$. Proposition 2.3 below answers the question of Thornton about orienting regular graphs [26, Section 3].

For simple graphs G_1 and G_2, let the graph $G_1 \times G_2$ be defined by having $V(G_1 \times G_2) = V(G_1) \times V(G_2)$ with vertices (u,v) and (u',v') being adjacent if and only if $u = u'$ and $vv' \in E(G_2)$ or $v = v'$ and $uu' \in E(G_1)$.

Proposition 2.3. *Let H be a finite $(2d-2)$-regular graph with an odd number of vertices. The $2d$-regular graph $H \times P$ has no $\mathrm{Aut}(H \times P)$-factor of iid balanced orientation.*

Proof of Theorem 1.3. In Proposition 2.3, take H to be the $(2d-2)$-regular clique K_{2d-1}. As K_{2d-1} is transitive, so is $K_{2d-1} \times P$. □

Proof of Proposition 2.3. First we note that P has no $\mathrm{Aut}(P)$-factor of iid balanced orientation. This is because P has only two balanced orientations, and so orienting any single edge determines the whole balanced orientation. So two edges at an arbitrary distance get oriented in the same direction with probability 1. In a factor of iid orientation however, given two edges that are far enough from each other, the probability of them being oriented in the opposite direction is close to 1/2.

Second, suppose a balanced orientation of $H \times P$ is given. For adjacent vertices $v, v' \in V(P)$, define

$$n(v,v') = \Big|\{u \in V(H) \mid \big((u,v),(u,v')\big) \text{ is an oriented edge in } H \times P\}\Big|,$$

and note that $n(v_i, v_{i+1}) + n(v_{i+1}, v_i) = |V(H)|$.

We also claim that $n(v_i, v_{i+1}) = n(v_{i+1}, v_{i+2})$ for every $i \in \mathbb{Z}$. Indeed, as the orientation of $H \times P$ is balanced, we have

$$\sum_{u \in V(H)} \mathrm{indeg}\big((u, v_{i+1})\big) = \sum_{u \in V(H)} \mathrm{outdeg}\big((u, v_{i+1})\big). \tag{2.1}$$

On one hand, oriented edges of $H \times P$ of the form $\big((u, v_{i+1}), (u', v_{i+1})\big)$ contribute 1 to both sides of (2.1). On the other hand, $n(v_i, v_{i+1})$ counts the edges of the form $\big((u, v_i), (u, v_{i+1})\big)$ contributing to the left-hand side of (2.1), and $n(v_{i+1}, v_{i+2})$ counts the edges of the form $\big((u, v_{i+1}), (u, v_{i+2})\big)$ contributing to the right-hand side. Therefore $n(v_i, v_{i+1}) = n(v_{i+1}, v_{i+2})$ as claimed.

As $|V(H)|$ is odd, we either have $n(v_i, v_{i+1}) > n(v_{i+1}, v_i)$ or $n(v_i, v_{i+1}) < n(v_{i+1}, v_i)$. We now have an argument similar to the one we presented to prove there is no factor of iid balanced orientation of P.

Our claim shows that if we have $n(v_i, v_{i+1}) > n(v_{i+1}, v_i)$ for some i, then we have it for all i. That is, for $i, j \in \mathbb{Z}$, and *any* random balanced orientation, we have

$$\mathbb{P}\left[n(v_i, v_{i+1}) > n(v_{i+1}, v_i) \text{ and } n(v_j, v_{j+1}) < n(v_{j+1}, v_j)\right] = 0.$$

However, if the balanced orientation of $H \times P$ was a factor of iid, for any $\varepsilon > 0$ we could find i and j far enough such that

$$\left|\mathbb{P}\left[n(v_i, v_{i+1}) > n(v_{i+1}, v_i) \text{ and } n(v_j, v_{j+1}) < n(v_{j+1}, v_j)\right] - \frac{1}{2}\right| \leq \varepsilon.$$

Consequently, $H \times P$ has no factor of iid balanced orientation. □

These non-examples are all quasi-isometric to P, which renders them somewhat unsatisfactory (see Question 6.4). Being quasi-isometric to P, however, does not imply being a non-example.

Proposition 2.4. *Let H be a finite $(2d-2)$-regular graph which has a perfect matching. Then the $2d$-regular graph $H \times P$ has an $\mathrm{Aut}(H \times P)$-factor of iid Schreier decoration.*

Proof. Let S be a non-empty factor-of-iid 4-independent subset of P, that is a subset such that for every $v_i, v_j \in S$, the distance $d_P(v_i, v_j)$ is at least 5. For each $v \in S$, let v' be a neighbour of v in P, chosen uniformly at random. Then let both $H \times \{v\}$ and $H \times \{v'\}$ fix the same perfect matching M. Together with the $|V(H)|$ edges between $H \times \{v\}$ and $H \times \{v'\}$, these form a 2-factor of $H \times \{v, v'\}$ (whose all cycles are even).

Next we are going to complement these 2-factors to obtain a 2-factor of the entire $H \times P$. Let $S' = \{v' : v \in S\}$, and suppose that $i \in \mathbb{Z}, j \in \mathbb{N}$ are such that $v_{i-1}, v_{i+j+1} \in S \cup S'$ but $v_k \notin S \cup S'$ for any $k \in [i, i+j]$. We note that j is finite and greater than zero because S is a factor of iid and 4-independent, respectively. Now let $H \times \{v_i\}$ fix the same perfect matching as $H \times \{v_{i-1}\}$ does and $H \times \{v_{i+j}\}$ the same as $H \times \{v_{i+j+1}\}$. Together with the $j \cdot |V(H)|$ edges of the form $(u, v_{i+k-1})(u, v_{i+k}), k \in [j]$, these form a 2-factor of $H \times \{v_i, \ldots v_{i+j}\}$ (whose all cycles are even).

Let us give colour c_1 to every edge in this 2-factor of $H \times P$ and orient each of its finite cycles strongly. Then after removing all the decorated edges, we are left with a $(2d-2)$-regular graph whose all connected components are finite; in particular, they are all isomorphic to either H or a connected component of $(H \setminus M) \times \{v_i, v_{i+1}\}$. To complete the construction, let each connected component pick at random a Schreier decoration with colours $c_2, \ldots, c_d$. □

With a couple more technicalities in the proof, we could in fact tweak the Schreierisation so that all oriented cycles of colour c_1 would have length at most $3 \cdot |V(H)|$ and all oriented cycles of the other colours at most $2 \cdot |V(H)|$. This means that $H \times P$ is the Schreier graph of a factor-of-iid action of many more groups than just F_d.

Most $(2d-2)$-regular graphs with an even number of vertices do have a perfect matching. For example, if they are bipartite or contain a Hamiltonian path, a perfect matching must exist. However, there are instances for every $d \geq 3$ which do not have a perfect matching despite having an even number of vertices [8,4]. For these, we do not know whether their product with P admits a factor-of-iid Schreier decoration or a perfect matching; see also subsection 5.1 and Question 6.5.

Finally, taking $H = K_{2d-1}$ in Proposition 2.3 and $H = K_{2d-2,2d-2}$ in Proposition 2.4 shows that the property of having a factor-of-iid Schreier decoration is not invariant under quasi-isometry not even when we restrict to transitive graphs of a given even degree.

3. Finite clusters of arbitrary thickness and their hierarchy

In this section, we will develop the main tools to manipulate planar lattices. Our goal is to break a lattice into pieces of finite size such that each of the pieces is “wide” enough and surrounded by another one – such a partition of the vertex set will be called a hierarchy. (In the toast language our pieces in the hierarchy being wide translates to the pieces of toast having boundaries that are far apart.)

In the first subsection, we prove that given any hierarchy, one can coarsen it in a factor of iid way so that the pieces of the resulting hierarchy have arbitrarily large width. This thickening will become key both in subsection 3.2, in which we show how to obtain a factor-of-iid balanced orientation of any planar lattice with m-fold symmetry and all degrees even, and in Section 4. In these, we use percolation theory to obtain an initial hierarchy, but need non-adjacent vertex clusters to be far from one another to have enough space for the desired combinatorial constructions. In the last subsection, we show how a hierarchy can be derived from the cell processes of Spinka [24] and Timár [28].

We remark that the existence of these hierarchies (shown in Theorem 3.6 and in Corollary 3.11) has also been developed for toasts in a more general setting. The contents of subsection 3.3 follow from the fact that hyperfinite graphings admit a toast off a nullset (essentially due to Conley and Miller [10]), and that the Bernoulli graphings we consider are hyperfinite (by Kaimanovich's Theorem, see e.g. [14, Prop. 2.2]).

Throughout this and the following section, for any graph G, let μ_G denote the usual product measure on $[0,1]^{V(G)}$, with each coordinate getting the Lebesgue measure.

3.1. General cluster hierarchy and its thickness

To obtain Schreier decorations of three of the Archimedean lattices, we partition their vertex set V into finite parts, which we shall call clusters, such that for each cluster C, there is a unique cluster C^+ surrounding it. That is, a unique part C^+ such that C is in its convex hull and such that there are adjacent vertices $u \in C$, $v \in C^+$. Such a partition of V can be described by a one-ended infinite tree T whose vertices are the finite clusters; $B, C \in V(T)$ are adjacent if and only if one surrounds the other. Moreover, every vertex $C \in V(T)$ has a well-defined parent C^+ distinguishable from C's children, should there be any.

Definition 3.1 *(Hierarchy).* Let G be a graph and $\mathbf{H}$ a partition of $V(G)$. We say that two distinct parts $C, D \in \mathbf{H}$ are adjacent if and only if there is $u \in C$ and $v \in D$ which are adjacent in G. Then $\mathbf{H}$ is a *hierarchy on* G if the following holds for every $C \in \mathbf{H}$.

1) C is finite,
2) there is a *unique* $C^+ \in \mathbf{H}$ such that C and C^+ are adjacent and for all $v \in V(G)$ but finitely many, any path from C to v contains a vertex from C^+,
3) whenever $B \in \mathbf{H}$ is adjacent to C, either $B = C^+$ or $C = B^+$.

Note that it is not necessarily the case that the subgraph of G induced by a cluster C is connected.

A key feature of the hierarchies we use later is that any two non-adjacent clusters are far from one another. Let us describe how starting with any hierarchy, we can obtain one in which non-adjacent clusters are as far from one another as we wish.

Definition 3.2 *(k-spaced hierarchy).* Let G be a graph and k a natural number. A hierarchy $\mathbf{H}$ on G is *k-spaced* if for all non-adjacent $B, C \in \mathbf{H}$, the graph distance $d(B,C) = \min_{u\in B, v\in C} d_G(u,v)$ is at least k.

Proposition 3.3. *Let G be a graph and suppose there is a (finitary)* $\mathrm{Aut}(G)$*-factor of iid hierarchy* $\mathbf{H}$ *on* G*. Then* $\forall k \in \mathbb{N}$*, there is a (finitary)* $\mathrm{Aut}(G)$*-factor of iid k-spaced hierarchy* $\mathbf{H}_k$ *on* G*.*

Moreover, $\forall c, k \in \mathbb{N}$*, there is a (finitary)* $\mathrm{Aut}(G)$*-factor of iid pair* $(J_{c,k}, \eta : J_{c,k} \to [c])$ *where* $J_{c,k}$ *is a k-spaced hierarchy and* $\eta : J_{c,k} \to [c]$ *is a colouring of the parts with c colours such that* $\forall C \in J_{c,k}$*, if* C *has colour* i *then* C^+ *has colour* $i+1 \pmod{c}$*.*

Proof. Consider the infinite tree $T_{\mathbf{H}}$ whose vertices are the parts of $\mathbf{H}$, i.e., the clusters of the hierarchy. Let us now colour uniformly randomly the vertices of $T_{\mathbf{H}}$ green and yellow, independently of each other (in other words, we carry out a site percolation with $p = \frac{1}{2}$). For any vertex $z \in V(T_{\mathbf{H}})$, write z^+ for the parent of z and $z^{n+} := (z^{(n-1)+})^+$ whenever $n \geq 2$. Then with probability 1, $\forall z \in V(T_{\mathbf{H}})$ $\exists n$ such that z^{n+} has the opposite colour than z. Together with the fact that every vertex $x \in V(T_{\mathbf{H}})$ is an ancestor of only finitely many other vertices (i.e., $|\{y \in V(T_{\mathbf{H}}) : \exists n \text{ such that } x = y^{n+}\}| < \infty$), this means that with probability 1, all monochromatic connected components of the coloured $T_{\mathbf{H}}$ are finite.

For every $z \in V(T_{\mathbf{H}})$, let us now merge it with z^+ if and only if z and z^+ have the same colour or z is yellow and z^+ is green. The modification this makes on the tree $T_{\mathbf{H}}$ is that of contracting all edges whose two endpoints have the same colour and those edges zz^+ for which z is yellow and z^+ is green. This corresponds to coarsening $\mathbf{H}$ into a new hierarchy $\mathbf{H}^{\wr}$ which has the property that for each cluster C, the distance $d(C, C^{++}) = \min_{u\in C, v\in C^{++}} d_G(u,v)$ is at least 3. This is because every path from C to C^{++} must go through C^+, but $C^+ \in \mathbf{H}^{\wr}$ consists of at least two clusters of the original hierarchy $\mathbf{H}$. More precisely, C is the finite union of some former clusters $A_1, \dots, A_m \in \mathbf{H}$, C^+ of $B_1, \dots, B_n$ and C^{++} of $D_1, \dots, D_l$. C^+ surrounds C, so there must be unique $i \in [m]$, $j \in [n]$ such that $A_i^+ = B_j$. C and C^+ being distinct clusters in $\mathbf{H}^{\wr}$ implies that the vertex $A_i \in V(T_{\mathbf{H}})$ was coloured green and $B_j \in V(T_{\mathbf{H}})$ yellow. But then B_j cannot be the part of the union $\cup_r^n B_r$ through which C^+ neighbours C^{++}, that is $B_j^+ \in \{B_1, \dots, B_n\}$, and so $d(C, C^{++}) = d(A_i, C^{++}) \geq d(A_i, B_j^{++}) = d(A_i, A_i^{+++}) \geq 3$.

If we repeat such contraction of the hierarchy tree m times to obtain $\mathbf{H}^{m\wr}$, then $d(C, C^{++}) \geq 2^m + 1$ for every $C \in \mathbf{H}^{m\wr}$. However, for $B, C \in \mathbf{H}^{m\wr}$ such that $B^+ = C^+$, we can still have $d(B,C) = 2$. To fix this, let us divide every $C \in \mathbf{H}^{m\wr}$ to

$$C_{\text{outer}} = \{v \in C : d(v, C^+) = \min_{u\in C^+} d_G(v,u) < 2^{m-1}\}$$

$$\text{and } C_{\text{inner}} = \{v \in C : d(v, C^+) = \min_{u\in C^+} d_G(v,u) \geq 2^{m-1}\},$$

and define a hierarchy $\mathbf{H}_{2^{m-1}} = \{B_{\text{inner}} \cup \bigcup_{B=C^+} C_{\text{outer}} : B \in \mathbf{H}^{m\wr}\}$. See Fig. 2 for illustration. Then for every $B \in \mathbf{H}^{m\wr}$ such that $B_{\text{inner}} \neq \emptyset$,

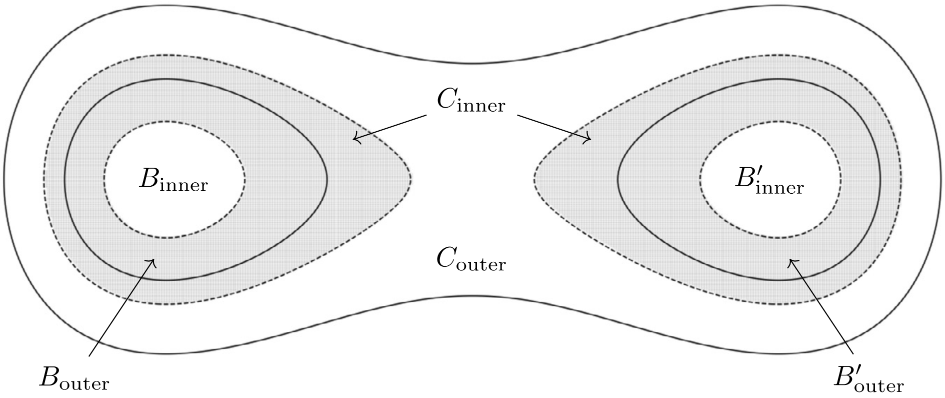

Fig. 2. Splitting clusters C, B, and B' into their inner and outer parts. The shaded region shows $C_{\text{inner}} \cup \bigcup_{B^+=C} B_{\text{outer}}$.

$$\begin{aligned}
& d\Big(B_{\text{inner}} \cup \bigcup_{B=C^+} C_{\text{outer}}, \big(B_{\text{inner}} \cup \bigcup_{B=C^+} C_{\text{outer}}\big)^{++}\Big) \\
&= d\Big(B_{\text{inner}} \cup \bigcup_{B=C^+} C_{\text{outer}}, B^{++}_{\text{inner}} \cup \bigcup_{B^{++}=C^+} C_{\text{outer}}\Big) \\
&= d\Big(B_{\text{inner}}, \bigcup_{B^{++}=C^+} C_{\text{outer}}\Big) = d(B_{\text{inner}}, B^+_{\text{outer}}) \geq d(B_{\text{inner}}, B^+) = 2^{m-1}.
\end{aligned}$$

Two clusters in $\mathbf{H}_{2^{m-1}}$ have the same parent if and only if they are of the form $B_{\text{inner}} \cup \bigcup_{B=C^+} C_{\text{outer}}$ and $D_{\text{inner}} \cup \bigcup_{D=C^+} C_{\text{outer}}$ for some $B, D \in \mathbf{H}^{m\wr}$ with $B^+ = D^+$. For such a pair, we observe that

$$\begin{aligned}
d\Big(B_{\text{inner}} \cup \bigcup_{B=C^+} C_{\text{outer}}, D_{\text{inner}} \cup \bigcup_{D=C^+} C_{\text{outer}}\Big) &= d(B_{\text{inner}}, D_{\text{inner}}) \\
&\geq d(B_{\text{inner}}, B^+) + d(D^+, D_{\text{inner}}) = 2 \cdot 2^{m-1}.
\end{aligned}$$

Having obtained $\mathbf{H}_k$ with $d(B, C) \geq k$ whenever B, C are non-adjacent for every $k \in \mathbb{N}$, we want to find coloured hierarchies $J_{c,k}$. For fixed $c, k \in \mathbb{N}$, start by considering a hierarchy $\mathbf{H}_{ck}$. Then divide every $C \in \mathbf{H}_{ck}$ to parts $C_1, \ldots, C_c$ where

$$\begin{aligned}
& C_i = \{v \in C : d(v, C^+) \in ((i-1)k, ik]\} \text{ for every } i \in [c-1] \\
\text{and } & C_c = \{v \in C : d(v, C^+) > (c-1)k\}.
\end{aligned}$$

Colouring C_i with colour $c + 1 - i$ yields the desired pair $(J_{c,k}, \eta : J_{c,k} \to [c])$. $\square$

3.2. Percolation clusters in planar lattices and their matching lattices

Let Λ be a *planar lattice*, that is a connected, locally finite plane graph, with $V(\Lambda)$ a discrete subset of $\mathbb{R}^2$, such that there are translations T_{v_1} and T_{v_2} through two independent vectors v_1 and v_2 both of which act on Λ as a graph isomorphism [5, p. 138]. Note that any planar lattice Λ is necessarily quasi-transitive. We wish to use site percolation in a way that would partition $V(\Lambda)$ into finite clusters with a hierarchy. For lattices satisfying $\theta^s\left(\frac{1}{2}\right) = 0$, i.e., on which site percolation does not occur when $p = \frac{1}{2}$, we could colour the vertices $V(\Lambda)$ uniformly at random yellow and green and consider the monochromatic connected components. However, though this produces, with probability 1, only finite clusters, there isn't necessarily a clear hierarchy associated to them; this inconvenience is best illustrated by the random vertex 2-colouring of the square lattice. Nevertheless, when the lattice possesses some mild symmetry properties, a two-step solution presents itself. The first step is a result on matching pairs with m-fold symmetry.

Definition 3.4 *(m-fold symmetry, [6]).* For $m \geq 2$, a plane lattice Λ has *m-fold symmetry* if the rotation about the origin through an angle of $2\pi/m$ maps the plane graph Λ into itself.

Based on ideas of Zhang, Bollobás and Riordan prove the following [5,6].

Theorem 3.5. *Let Λ be a plane lattice with m-fold symmetry for some $m \geq 2$ and $\Lambda^\times$ its matching lattice, i.e., the graph obtained from Λ by adding all diagonals to all faces of Λ. Then for every $p \in [0, 1]$, the percolation probabilities $\theta^s_\Lambda(p)$, $\theta^s_{\Lambda^\times}(1-p)$ satisfy that $\theta^s_\Lambda(p) = 0$ or $\theta^s_{\Lambda^\times}(1-p) = 0$. Furthermore, $p^s_H(\Lambda) + p^s_H(\Lambda^\times) = 1$, where p^s_H is the Hammersley critical probability.*

This result tells us that choosing to randomly 2-colour the vertices of our planar lattice Λ such that a site is yellow with probability $p_H^s(\Lambda)$ and green with probability $p_H^s(\Lambda^\times)$ is well defined. The definition of a cluster hierarchy on Λ then comes naturally, with yellow clusters being exactly the yellow connected components of the random 2-colouring and green clusters the would-be green connected components on $\Lambda^\times$ (that is in Λ, the green clusters are unions of connected components which have vertices appearing in a same face). The trouble is that the theorem does not guarantee that *neither* of the colours will have an infinite cluster, or in other words that percolation does not occur at criticality at *neither* the planar lattice nor its matching lattice. Benjamini and Schramm conjectured in 1996 that this indeed *is* the case [3]. However, for now, to obtain this crucial property, one needs to analyse lattices one by one (e.g. Russo showed that both critical percolations die for the square lattice $\Lambda_\square$ and its matching lattice $\Lambda_\boxtimes$ [22]).

The second step in our solution to the hierarchy problem is therefore to add a vertex to every non-triangular face of a lattice Λ with m-fold symmetry and connect it to all the vertices of that face. Let us call the resulting lattice $\Lambda^\bullet$ and observe that it also has m-fold symmetry. $\Lambda^\bullet$ is self-matching, and so Theorem 3.5 tells us that $p_H^s(\Lambda^\bullet) = \frac{1}{2}$ and percolation does not occur at criticality.

This finding gives us a factor-of-iid hierarchy on Λ as follows. Colour the vertices of Λ yellow or green uniformly at random. For each non-triangular face, decide uniformly at random whether either all of its yellow vertices will be treated as if they were connected through the face or all of its green vertices will be so treated. This results, with probability one, in a well-defined cluster hierarchy on Λ, which we can combine with Proposition 3.3 to conclude the following.

Theorem 3.6. *Let Λ be a planar lattice with m-fold symmetry, $m \geq 2$. Then for every $k \in \mathbb{N}$, there is a finitary* $\mathrm{Aut}(\Lambda)$*-factor of* $\left([0,1]^{V(\Lambda)}, \mu_\Lambda\right)$ *which is almost surely a k-spaced hierarchy on Λ.*

Theorem 3.6 now allows us to obtain a finitary factor of iid balanced orientation of any planar lattice with m-fold symmetry in which all degrees are even. To carry out the construction, we first need to clearly demarcate the clusters.

Definition 3.7 *(Boundary ∂B).* Let Λ be a planar lattice and $\mathbf{H}$ a hierarchy on Λ. The *boundary* ∂B of a cluster $B \in \mathbf{H}$ is a set of edges with both endpoints in B as follows. An edge $uv \in E(B)$ is in ∂B if and only if one of its two contiguous faces F_1, F_2 consists only of vertices of the cluster B and its offspring, while the other contains a vertex from B^+.

Note that in general, it may be the case that $F_1 = F_2$, but when all degrees of Λ are even, these two faces must be distinct.

Theorem 3.8. *Let Λ be a planar lattice with m-fold symmetry, $m \geq 2$, in which all degrees are even. There is a finitary* $\mathrm{Aut}(\Lambda)$*-factor of* $\left([0,1]^{V(\Lambda)}, \mu_\Lambda\right)$ *which is a balanced orientation of Λ almost surely.*

Proof. $V(\Lambda)$ is a discrete set in $\mathbb{R}^2$, i.e., it has no accumulation points, and so there is $\ell \in \mathbb{N}$ such that every face of Λ consists of at most ℓ vertices. Let us fix a finitary factor of iid $\frac{\ell+1}{2}$-spaced hierarchy $\mathbf{H}$ on Λ, so that no face contains vertices of non-adjacent clusters, and consider the boundaries ∂B, $B \in \mathbf{H}$.

We observe that for every vertex $v \in B$, the number of edges from ∂B that are incident to v is even. Indeed, let r be a positive number such that the disc D with radius r centred at v contains no other vertices than v, and let $x_1, \ldots, x_{\deg(v)}$ be the neighbours of v as we traverse them clockwise. Then the edges $vx_1, \ldots, vx_{\deg(v)}$ split D into $\deg(v)$ many areas $A_{i-1,i}$, where vx_{i-1} moves through $A_{i-1,i}$ clockwise towards vx_i (all indices are taken $\pmod{\deg(v)}$). Now call $F_{i-1,i}$ the faces spanned by $A_{i-1,i}$; note that we might have $F_{i-1,i} = F_{j-1,j}$ even when $i \neq j$. The faces $F_{i-1,i}$ are of two types – those with all vertices from the cluster B and its offspring and those with a vertex from B^+. As no face contains vertices of non-adjacent clusters, the edge vx_i is in ∂B if and only if the two consecutive faces $F_{i-1,i}, F_{i,i+1}$ are of different types. But when we start in $F_{1,2}$ and make one full circle clockwise back to $F_{1,2}$, the change of types must occur even number of times.

Even though the construction of a spaced hierarchy $\mathbf{H}$ on Λ based on Proposition 3.3 and Theorem 3.6 does not ensure that all clusters are connected, it does say that $\cup_{n=1}^\infty C^{n+}$ is connected (and infinite) for every $C \in \mathbf{H}$. Therefore, there is no vertex from $\cup_{n=1}^\infty C^{n+}$ inside ∂C (each cycle of ∂C splits the plane into a finite and an infinite region; by inside, we mean the union of the finite regions). On the other hand, all vertices of any offspring of C are inside ∂C, so any cluster which has a child must have non-empty boundary. Let us now reassign any vertices of a cluster C which are outside ∂C to C^+ – these are exactly the vertices such that all faces they are in contain a vertex from C^+.

Every cluster C will now randomly choose one of two allowed orientation patterns which will be given to all non-boundary edges with at least one endpoint in C and none in C^+. These are exactly the edges in the finite region enclosed by ∂C, but outside the finite regions enclosed by ∂B for all B with $B^+ = C$. The patterns are as follows: for a vertex v and its neighbours $x_1, \ldots, x_{\deg(v)}$, ordered as we traverse them clockwise, we want the orientations of $vx_1, \ldots, vx_{\deg(v)}$ to be alternating between into v and out of v. In other words, we choose randomly one of the two chessboard colourings of the faces in this region, and orient cycles bounding black faces clockwise and cycles bounding white faces counter-clockwise.

For every cluster C, if C and C^+ happen to choose the same pattern, we will propagate it to ∂C as well. If C and C^+ choose different patterns, then ∂C randomly picks one of its balanced orientations (recall that ∂C has all degrees even, so a balanced orientation exists).

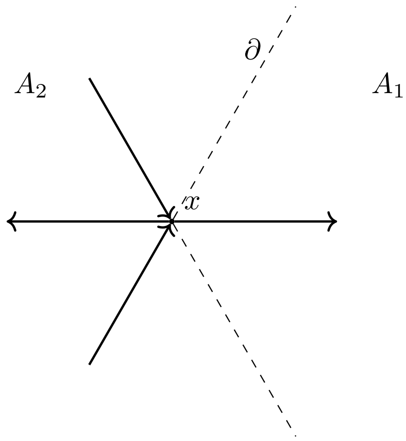

Fig. 3. A disagreement between patterns around a vertex of degree 6.

Finally, we claim that the resulting global orientation of Λ is balanced. If a vertex $v \in C$ is not in ∂C or if C and C^+ agree on their patterns, then all the edges incident to v follow the same pattern, i.e., they alternate between in and out and the orientation is balanced at v. Now suppose C and C^+ have different patterns and some of the edges incident to v are in ∂C. Then let D be a disc centred at v as before – the boundary edges split it into areas $A_0, \dots, A_{2n-1}$, ordered as we traverse them clockwise, for some integer n, where in any pair of adjacent areas, one is inside ∂C and one is outside. The patterns at adjacent areas disagree, i.e., they are not the restriction of the same alternating orientation at v (see Fig. 3). Let us note that vx and vy have the same orientation with respect to v if and only if there is an even number of edges incident to v between them and they are in regions following the same pattern, or there are an odd number of regions between them and they are in regions following different patterns.

Inside each area A_i, the number of non-boundary edges is either even or odd. If it is even, half of these edges are oriented towards v and half from v because they all follow the same pattern. If the number is odd, let j be the smallest positive integer such that A_{i+j} also contains odd number of non-boundary edges, where $i+j$ is understood mod $2n$. If j is even then A_i and A_{i+j} follow the same pattern. Also, the total number of edges separating them at v, counted from A_i to A_{i+j} in the clockwise direction (including those in ∂C) is even. These two observations imply that the orientation restricted to the edges (at v) in $A_i \cup A_{i+j}$ is alternating. Consequently, exactly half of these edges is oriented towards v and half from v. If j is odd then A_i and A_{i+j} follow different patterns, but the total number of edges separating them (including those in ∂C) is odd. These together also imply that the orientation restricted to the edges in $A_i \cup A_{i+j}$ is alternating (see Fig. 3), so again exactly half of these is oriented towards v and half from v. Combined with the fact that due to the balanced orientation of ∂C, half of the boundary edges go into v and half out of v, this means that the orientation at v is balanced. $\square$

Let us note that the first half of the proof uses neither that the lattice Λ is symmetric nor that all degrees are even, so we can also use it to deduce the following technical lemma that we will heavily use in Section 4.

Lemma 3.9. *Let Λ be a planar lattice and suppose there is a (finitary)* $\mathrm{Aut}(\Lambda)$*-factor of iid which is almost surely a hierarchy on Λ. Then for every $k \in \mathbb{N}$, there is a (finitary)* $\mathrm{Aut}(\Lambda)$*-factor of iid which is almost surely a k-spaced hierarchy with the additional property that the boundary ∂C of any cluster C is a union of edge-disjoint cycles and any path between any $u \in C$, $v \in \cup_{n=1}^{\infty} C^{n+}$ must cross ∂C. Subsequently, for any distinct clusters B, C, we have that $d(\partial B, \partial C) \geq k - \frac{\ell}{2}$, where ℓ is an upper bound on the number of vertices forming any face of Λ.*

Proof. Given $k \in \mathbb{N}$, let us fix a $\max\{k, \frac{\ell+1}{2}\}$-spaced hierarchy $\mathbf{H}_k$ (whose existence is ensured by Proposition 3.3), and observe that we can repeat verbatim the proof of that for every vertex $v \in B$, the number of edges from ∂B that are incident to v is even. It also remains true that for any cluster $C \in \mathbf{H}_k$, the vertices $v \in C$ which remained outside of ∂C are exactly those such that all faces they are in contain a vertex from C^+. Let us repeat the reassignment of vertices outside ∂C to C^+ so that for any cluster C, $C = \{v \in V(\Lambda) : v \text{ is inside } \partial C \text{ but outside } \partial B \text{ for every } B \text{ such that } B^+ = C\} \cup \partial C$ as required. In particular, this reassignment leaves the boundaries ∂C, $C \in \mathbf{H}_k$ intact and ensures the second property.

Finally, if B and C are distinct non-adjacent clusters then $d(\partial B, \partial C) \geq k$ simply because the hierarchy is k-spaced. On the other hand, for any $C \in \mathbf{H}_k$ and any vertex $u \in \partial C^+$, the definition of boundary tells us that $d(u, C^{++}) \leq \frac{\ell}{2}$. But $d(\partial C, C^{++}) \geq k$ as C and C^{++} do not neighbour, so by the triangle inequality,

$$k \leq d(\partial C, C^{++}) \leq d(\partial C, \partial C^+) + d(\partial C^+, C^{++}) \leq d(\partial C, \partial C^+) + \frac{\ell}{2}.$$

Thus $d(\partial C, \partial C^+) \geq k - \frac{\ell}{2}$ as required. $\square$

3.3. Deriving a hierarchy from a cell process

Let G be a graph. Following Spinka [24], we call a random sequence $A=(A_1,A_2,\dots)$ of subsets of $V(G)$ a *cell process* whenever almost surely $A_1\subset A_2\subset\dots$, for each $n\geq 1$, all connected components of A_n are finite, and $A_1\cup A_2\cup A_3\dots=V(G)$.

Given an invariant cell process, we would like to consider the partition $A_1\cup\{A_{n+1}\setminus A_n\mid n\geq 1\}$ of $V(G)$, but this is not necessarily a hierarchy because the definition does not guarantee that all neighbours of all vertices of A_n are in A_{n+1}. However, using some ideas from [7, Section 4], one can almost surely get to a factor-of-iid hierarchy starting from a factor-of-iid cell process.

Lemma 3.10. *Let G be a graph and A a (finitary) factor-of-iid cell process on G. Then there is a (finitary) factor-of-iid hierarchy on G.*

Proof. Let us define a sequence of indices recursively as follows: let $n_1=1$ and for $i>1$ let n_i be chosen as the smallest such an integer, such that:

$$\mathbb{P}[N(A_{n_{i-1}}(v))\subseteq A_{n_i}(v)\mid v\in A_{n_{i-1}}]\geq 1-2^{-i}$$

where $A_i(v)$ denote the connected component of v in A_i.

Now let us create, for every vertex $v\in V(G)$, a new sequence (where each member of the sequence can be obtained as a finitary factor-of-iid process) as follows: let $\bar{s}(v)_0=\min\{i\mid N(A_{n_i}(v))\subseteq A_{n_{i+1}}(v)\}$ and for $j>0$,

$$\bar{s}(v)_j=\min\{i>\bar{s}(v)_{j-1}\mid N(A_{n_i}(v))\subseteq A_{n_{i+1}}(v)\}.$$

First of all observe that with probability 1, there is a finite $i\in\mathbb{N}$ such that $v\in A_{n_i}$, since A_i was a cell-process.

Second we claim that $s(v)=\mathbb{N}\setminus\{\bar{s}(v)_j\}_{j\in\mathbb{N}}$, which we think of as the set of bad indices for v, is finite with probability 1. Indeed, if it was infinite with positive probability, then it would mean that with positive probability, infinitely many of the events

$$E_i=\{N(A_{n_i}(v))\not\subseteq A_{n_{i+1}}(v)\}$$

happened where i is sufficiently large. But by Borel-Cantelli lemma, we know that the probability that infinitely many E_i occur has probability 0 because $\sum_i 2^{-i}<\infty$.

Now we are ready to define the hierarchy as the "remained sets", i.e., for any $i\geq 1$

$$B_i=\{v\in V\mid i\notin s(v)\}\subseteq A_{n_i}.$$

To obtain the desired hierarchical partition, we take

$$\mathbf{H}=\mathcal{C}(B_1)\cup\left\{L\setminus\cup_{j<i+1}B_j\mid i\geq 1, L\in\mathcal{C}(B_{i+1})\right\},$$

where $\mathcal{C}(B)$ is the collection of connected subgraphs induced by B. □

Using the results of Spinka [24] and Timár [28], we can now conclude the following.

Corollary 3.11. *Let G be an amenable graph. If G is transitive, there is a finitary-factor-of-iid hierarchy on G. If G is unimodular, there is a factor-of-iid hierarchy on G.*

Proof. Spinka constructs a finitary-factor-of-iid cell process for any amenable transitive graph in [24], which together with Lemma 3.10 gives the first part of the statement.

In [28], Timár constructs a factor-of-iid sequence $\Gamma_1\subset\Gamma_2\subset\dots$ of *non-induced* subgraphs witnessing hyperfiniteness of any amenable unimodular G such that $\bigcup_n\Gamma_n=E(G)$ and for every n, all connected components of Γ_n are finite. Taking $A_n:=\{v\in V(G): uv\in\Gamma_n \text{ for all } u \text{ adjacent to } v\}$ gives a cell process, and we again apply Lemma 3.10. □

4. Schreier decorations of Archimedean lattices and $\Lambda_\square^d$, $d\geq 3$ as finitary factors of iid

An *Archimedean lattice* is a vertex-transitive tiling of the plane by regular polygons. It is known that there are ten of them (eleven if we count separately the two mirror images of the lattice $(3^4,6)$; see e.g. [5, Chapter 5]), out of which four have even regularity. These are the infinite grid $\Lambda_\square$ (that is the Cayley graph of $\mathbb{Z}^2$ with the standard generators), the triangular lattice T (the only 6-regular planar lattice), the Kagomé lattice (i.e., the line graph of the hexagonal lattice H), and the 4-regular lattice $(3,4,6,4)$ where the numbers denote the orders of the faces when we go round a vertex.

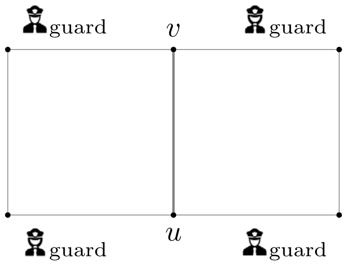

Fig. 4. Each edge has four guards.

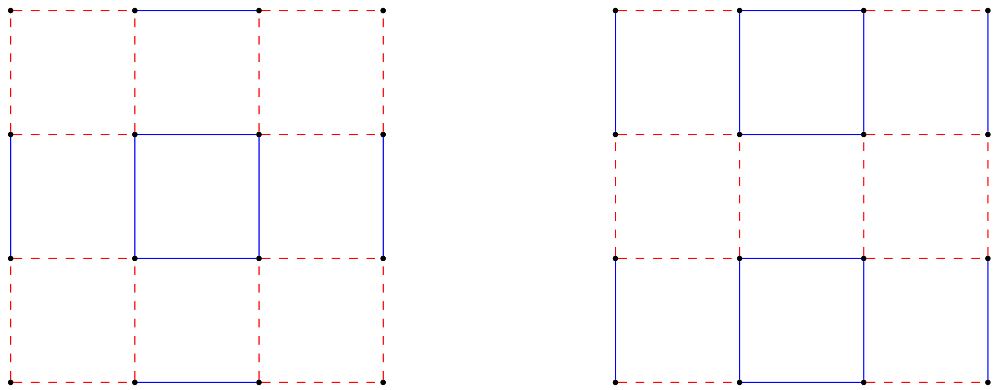

Fig. 5. Examples of what the inner pattern can look like on a fixed 16-vertex square. (Here, blue is identical to full lines, and red is identical to dashed lines.)

4.1. The square lattice

Proof of Theorem 1.1 for $\Lambda = \Lambda_\square$. We start with independent $[0,1]$-labels on $V(\Lambda)$.

Step 1 **Spaced hierarchy.** The square lattice has a 4-fold symmetry, so we can use Theorem 3.6, Proposition 3.3 and Lemma 3.9 to create a k-spaced hierarchy $\mathbf{H}_{2,k}$ based on percolation clusters where k is a sufficiently large integer. Moreover, each part $C \in \mathbf{H}_{2,k}$ is assigned a number from $\{1,2\}$ such that every child has the other number than its parent.

Step 2 **Definition of guards.** The *guards* of an edge $e = uv \in E(\Lambda_\square)$ are the four of its six incident edges that are perpendicular to it. That is, the four incident edges that appear with e in a C_4 (see Fig. 4).

Step 3 **Red-blue edge colourings inside clusters.** Each cluster numbered 1 wants to determine for itself an edge colouring consisting of monochromatic C_4-s. That is, each such cluster B finds a vertex $v \in B$ (e.g. the one with the largest label) and chooses one of the four C_4-s containing v to have all its edges coloured red. As we want our B-pattern to be the union of blue C_4-s and red C_4-s and such that every $u \in B$ has two incident edges blue and two red, fixing one red C_4 determines the rest of the pattern. We call this colouring the *inner pattern* of B (Fig. 5).
On the other hand, each cluster numbered 2 wants to determine for itself an edge colouring in which all parallel edges have the same colour (and so every C_4 has two edges red and two blue in an alternating manner, and every edge has different colour than its guards). We call such a colouring an *interface pattern*.

Step 4 **Amalgamating the patterns.** For every cluster C numbered 2 (i.e., a cluster with an interface pattern), we will apply its pattern to all edges with at least one endpoint in C. Let B be a cluster numbered 1 and e an edge with both endpoints in B. Let S be the set of the guards of e whose other endpoint (the one not shared with e) is not in B (note that S might be empty). Then we colour e as follows: If there is $g \in S$ on which the inner pattern of B (should it be extended there) disagrees with the interface pattern by which g is coloured, we colour e according to the interface pattern, i.e., with the opposite colour than g has. This may or may not coincide with the colour the inner pattern would give to e. If there is no such guard $g \in S$, e is coloured according to the inner pattern of B. Our hierarchy being k-spaced (where $k \geq 4$) ensures that this is well defined.

Step 5 **Claim that the colouring is balanced.** In the edge colouring defined above, every vertex $v \in V(\Lambda)$ has exactly two red and two blue incident edges a.s.

Proof of claim. If $v \in C$ where C is numbered 2, then all its four incident edges follow the same interface pattern. As any interface pattern is internally balanced, so are the colours at v.

So let us now assume that v is in a cluster B numbered 1, and consider the 9-vertex square centred on v. Regardless of what clusters the remaining 8 vertices belong to, the final colouring of the four edges incident to v can be described by say-

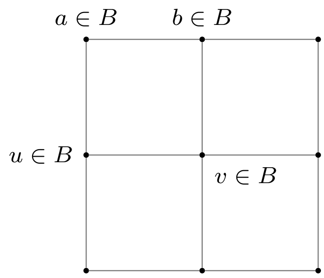

Fig. 6. Vertex v has two red and two blue incident edges.

ing that we first colour all four of them with the inner pattern (so that a perpendicular pair is red and the complementary one blue), and then recolour if necessary.

First observe that thanks to the hierarchy being k-spaced where $k \geq 5$, the 9-vertex square does not contain both a vertex from B^+ and a vertex from a B^-. Hence there is only one interface pattern according to which edges may wish to recolour themselves. Also note that regardless of the inner and the interface pattern, exactly one red and exactly one blue edge out of the four edges incident to v already follow the interface pattern at the beginning, so they will not get recoloured in any case. That is, at most two of the four edges will get recoloured.

If two edges get recoloured, then as explained above, one of them is initially blue and one is initially red. So recolouring both of them does not change the multiplicity of colours incident to v, and as this multiplicity was balanced to start with, it will remain balanced. If no edges get recoloured, then again, we invoke that the initial colouring was already balanced. Finally, we would like to argue that it cannot happen that exactly one of the four edges gets recoloured.

Let uv and bv be the two perpendicular edges on which the inner pattern and the interface pattern disagree, and suppose on the contrary that only uv gets recoloured. Note that uv is a guard of bv, so if $u \notin B$ then *both* uv and bv would get recoloured. Also bv is a guard of uv, so by the same logic we also must have $b \in B$. If a, the fourth corner of the square spanned by b, u and v, was not in B, then again *both* bv and uv would get recoloured. Finally, the inner pattern and the interface pattern agree on the two remaining guards of uv, so they will not be recoloured in any case, meaning uv also does not get recoloured. But that is in contradiction with the assumption that uv does get recoloured. Hence it cannot happen that exactly one of the four edges incident to v gets recoloured (Fig. 6). □

Step 6 Claim that there are no infinite paths. There are no infinite monochromatic paths in the balanced colouring defined above a.s. That is, the set of blue edges is a union of vertex-disjoint finite blue cycles, and similarly the set of red edges is a union of vertex-disjoint finite red cycles.

Proof of claim. Suppose there is an infinite monochromatic path P. Let A be a cluster whose intersection with P is non-empty. Then the path must share a vertex with every cluster in the sequence $(A^{n+})_{n=0}^{\infty}$, where $A^{0+} := A$ and $A^{n+} := \left(A^{(n-1)+}\right)^+$ for every positive integer n. Pick the $n \in \{1, 2\}$ such that A^{n+} is numbered 1, and so gets the inner pattern. There must be a section $u, v_1, \ldots, v_m$ of P such that $v_m \in \partial A^{n+}$, $v_i \in A^{n+} \setminus \partial A^{n+}$ for every $i \in [m-1]$, and $u \in \partial B$ for some cluster B such that $B^+ = A^{n+}$. We used Lemma 3.9 when building the k-spaced hierarchy, so we can now deduce that $m \geq k-2$. Moreover, for every $i \in [m-(k-5), m]$ and every cluster C such that $C^+ = A^{n+}$,

$$k-2 \leq d(C, \partial A^{n+}) \leq d(C, v_i) + d(v_i, \partial A^{n+}) \leq d(C, v_i) + k - 5,$$

and so $d(C, v_i) \geq 3$. This means that the edges $v_{m-(k-4)}v_{m-(k-5)}, \ldots, v_{m-1}v_m$ must have empty S, and so must all follow the inner pattern of A^{n+}. In particular, as $k \geq 8$, this is at least four edges. But no monochromatic walk of length four in an inner pattern can contain five distinct vertices (i.e., be a path) because inner patterns are made up of monochromatic C_4-s. This would be a contradiction, and so there is no infinite monochromatic path almost surely. □

Step 7 Orientation. Finally, let each monochromatic cycle choose one of the two strong orientations for itself (e.g. by finding the edge with the largest sum of labels and orienting it from the vertex with the larger label to the one with the smaller label). Then every vertex lies in one strongly oriented blue cycle and one strongly oriented red cycle, and so the coloured orientation is balanced. □

Remark 4.1. In fact, it is not necessary to rely on Proposition 3.3 to prove Theorem 1.1 for $\Lambda = \Lambda_{\square}$. It makes the proof neater, but one can also make do just with the basic percolation clusters for the matching pair $\{\Lambda_{\square}, \Lambda_{\boxtimes}\}$. The approach not using Proposition 3.3 (in which every yellow cluster is merged with its green parent into a blob and edges bridging different blobs get an interface pattern) heavily relies on the fact that critical percolation on either of the lattices in this matching pair has no infinite cluster a.s. (though expected sizes of the clusters will be infinite) [22]. Interestingly, the related

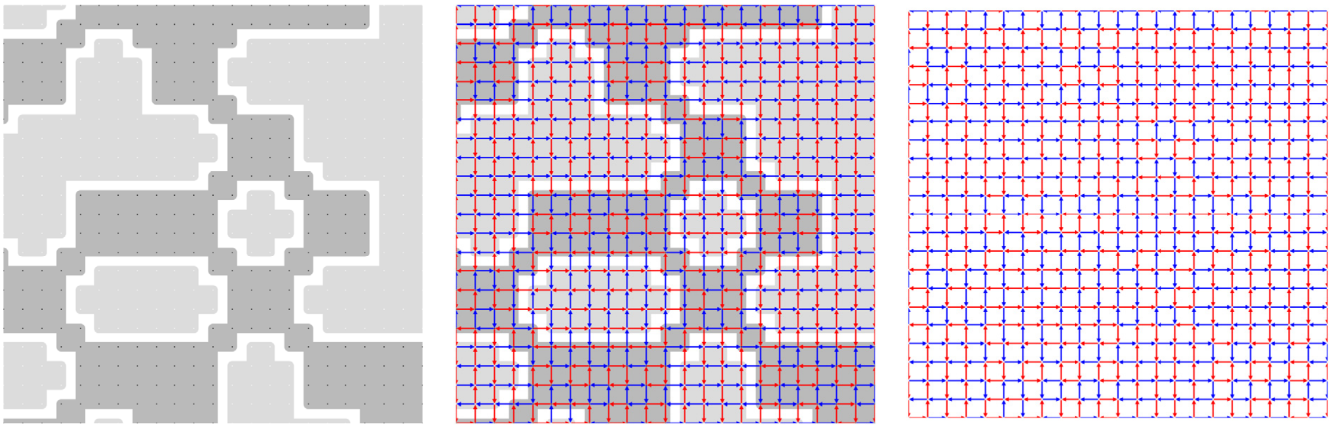

(a) Spaced hierarchy (b) Schreier decoration on hierarchy (c) Schreier decoration

Fig. 7. Constructing Schreier decoration of $\Lambda_\square$. (Subfigures (b) and (c) are coloured. To view them in their coloured version, the reader is referred to the web version of this article.)

conjecture of there being no percolation at criticality (site or bond) in $\mathbb{Z}^d$ has not been settled for $3 \le d \le 18$ (see [13] for a partial progress).

4.2. Grids in higher dimensions

Our construction on $\Lambda_\square$ made use of the hierarchical structure of clusters, as well as of their thickness preventing local recolourings of edges close to boundaries and other tweaks to overlap (Fig. 7). We achieved these desired properties by classical percolation-theoretic results, which are sadly not available in higher dimensions. However, we can utilise subsection 3.3 now and combine it with the patterns developed for $\Lambda_\square$ to obtain a factor-of-iid Schreierisation of all $\Lambda_\square^d, d \ge 2$.

While we were writing this paper up, Jan Grebík and Václav Rozhoň informed us that using techniques developed in [9], they can prove the existence of a (Borel) hierarchy in any Borel graph with connected components isomorphic to $\Lambda_\square^d$. This should allow one to construct Borel Schreier decorations on all such Borel graphs [17].

Proof of Theorem 1.2. We will first construct a balanced edge colouring with colours $c_1, \ldots, c_d$, and then give orientation separately to each monochromatic cycle.

Step 1 **Spaced hierarchy.** Let $\mathbf{H}$ be the finitary-factor-of-iid hierarchy on $\Lambda_\square^d$ given by Corollary 3.11, and let us use Proposition 3.3 to obtain from it $(\mathbf{H}_{2d-2,k}, \eta : \mathbf{H}_{2d-2,k} \to [2d-2])$ where k is a sufficiently large integer and $\mathbf{H}_{2d-2,k}$ is a k-spaced hierarchy.

Step 2 **Edge colouring of even-numbered clusters.** The edges of $\Lambda_\square^d$ travel in d different directions. Let us observe that if we remove all edges in $d-2$ of them, we are left with disjoint copies of the square lattice. Each cluster $C \in \mathbf{H}_{2d-2,k}$ such that $\eta(C)$ is odd will be assigned an *interface pattern* in direct analogy to subsection 4.1. By an *interface pattern*, we understand a balanced edge colouring in which the edge directions of $\Lambda_\square^d$ and the d colours are in bijection. That is, after removing all edges of any $d-2$ directions (= all edges of any $d-2$ colours), we would end up with an interface pattern on disjoint copies of the square lattice.

In particular, each cluster $C \in \mathbf{H}_{2d-2,k}$ numbered $2d-2$ identifies its ancestor $C^{(2d-2)+}$ (which is also numbered $2d-2$) and then all its $2d-2$-th generation offspring $C = C_1, \ldots, C_n$. The clusters $C_1, \ldots, C_n$ will make a common choice of the interface pattern.

Subsequently, each $\{2, \ldots, 2d-4\}$-numbered cluster $C \in \mathbf{H}_{2d-2,k}$ identifies its nearest $2d-2$-numbered ancestor $C^{\text{above}} := C^{(2d-2-\eta(C))+}$ and should there be any, also its nearest $2d-2$-numbered offspring and their interface patterns. The role of the $\{2, \ldots, 2d-4\}$-numbered clusters, as well as of the odd-numbered clusters, is to provide a balanced transition between the patterns of the $2d-2$-numbered clusters.

If a $\{2, \ldots, 2d-4\}$-numbered cluster C has no $2d-2$-numbered offspring then it will simply copy the pattern from C^{above}. Otherwise, we will assign interface patters so that for any cluster C with $\eta(C) = 2i \le 2d-4$, the directions of the colours $c_1, \ldots, c_i$ in C and C^{above} agree. In particular, we will first assign patterns to 2-numbered clusters, then 4-numbered etc. all the way up to $2d-4$. A $\{2, \ldots, 2d-4\}$-numbered cluster C shall identify the interface patterns of its grandchildren and of C^{above}. If these agree in the direction of the colour c_i then C will simply adopt the pattern of its grandchildren without a change. If they disagree then C will employ the interface pattern obtained from that of its grandchildren by swapping c_i and the colour that currently travels in the direction that c_i takes in C^{above}.

Step 3 **Edge colouring of odd-numbered clusters.** Each cluster $C \in \mathbf{H}_{2d-2,k}$ such that $\eta(C)$ is odd will be assigned an *inner pattern*. By an *inner pattern*, we understand a balanced edge d-colouring in which all the edges in $d-2$ chosen directions have the same colour, while the remaining copies of the square lattice get a pattern consisting of monochromatic C_4-s of the remaining two colours (that is an inner pattern in the sense of subsection 4.1).
In particular, each odd-numbered cluster $C \in \mathbf{H}_{2d-2,k}$ identifies its parent C^+ and its interface pattern and its children (should there be any) and their interface pattern. By construction, these two interface patterns are either the same or one can be obtained from the other by a single transposition, that is by swapping two colours. In the former case and also when C has no children, C will randomly choose two colours and put monochromatic C_4-s to the grids spanned by their two directions. In the latter case, the C_4-s shall be built from the two colours on which the two interface patterns disagree. Anyhow, the colours and directions of the colours not chosen to form C_4-s will propagate through C without change of direction.

Step 4 **Boundaries between clusters.** Whenever two clusters C and C^+ meet, there are $d-2$ colours which simply travel in a fixed direction, which is moreover the same in C and C^+. We will let these $d-2$ colours propagate in their directions also on all edges with one endpoint in C and one in C^+. After disregarding these, we are left with copies of the square grid, and on every such copy, we see an inner pattern in the sense of subsection 4.1 coming from one of the clusters and interface pattern from the other one. We amalgamate these into a balanced colouring as in subsection 4.1.
Finally, we observe that for every $C \in \mathbf{H}_{2d-2,k}$ and $i \in [d]$, there is, with probability 1, an $n \in \mathbb{N}$ such that C^{n+} has an inner pattern and one of the colours of its C_4-s is c_i. Therefore, there cannot be any infinite monochromatic path, and every colour class is almost surely a union of finite cycles. To finish the construction, let each monochromatic cycle randomly choose one of the two strong orientations for itself. □

4.3. The triangular lattice

Similarly as in the case of the higher-dimensional grids, once we have a spaced enough hierarchy, we can use the patterns developed for $\Lambda_\square$ to obtain a Schreier decoration of T.

Proof of Theorem 1.1 for $\Lambda = T$. We will first construct a red-blue-green edge colouring and then as in the previous case, give orientation separately to each monochromatic cycle.

Step 1 **Spaced hierarchy.** T has a 3-fold symmetry and is self-matching, so Theorem 3.5 tells us that its site-percolation critical probability is $\frac{1}{2}$ and percolation does not occur at criticality. Let $\mathbf{H}$ be the finitary factor of iid hierarchy on T given by the percolation clusters, and let us use Proposition 3.3 to obtain from it $(\mathbf{H}_{4,k}, \eta : \mathbf{H}_{4,k} \to [4])$ where k is a sufficiently large integer and $\mathbf{H}_{4,k}$ is a k-spaced hierarchy.

Step 2 **Edge colouring of odd-numbered clusters.** The edges of T travel in three different directions. Let us observe that if we remove all edges in one of them, we are left with the square lattice. Each cluster $C \in \mathbf{H}_{4,k}$ numbered 1 or 3 will be assigned an *inner pattern*. By an *inner pattern*, we understand a balanced edge 3-colouring in which all the edges in a chosen direction have the same colour, while the remaining square lattice gets a pattern consisting of monochromatic C_4-s of the remaining two colours (that is an inner pattern in the sense of subsection 4.1). In particular, each cluster $C \in \mathbf{H}_{4,k}$ numbered 1 identifies its grandparent C^{++} (which is numbered 3) and then all its grandchildren $C = C_1, \ldots, C_n$. The clusters $C_1, \ldots, C_n$ will make a common choice of the inner pattern. They randomly choose one of the three edge directions, colour all edges in the chosen direction green, and put blue and red C_4-s on the remaining square lattice.
Subsequently, each cluster $C \in \mathbf{H}_{4,k}$ numbered 3 identifies its grandchildren (should there be any) and their inner pattern and its grandparent C^{++} and its inner pattern. Then it chooses the direction not chosen by either of the patterns and colours all its edges blue. If there is a choice between two directions, siblings (i.e., clusters with the same parent C^+) make it together.
Every odd-numbered cluster C will apply its chosen inner pattern to all edges which are inside ∂C, but outside the boundaries of its children.

Step 3 **Edge colouring of even-numbered clusters.** Each cluster $C \in \mathbf{H}_{4,k}$ numbered 2 or 4 will be assigned an *interface pattern* in direct analogy to subsection 4.1. By an *interface pattern*, we understand a balanced edge colouring in which the edge directions of T and the three colours are in bijection. That is, after removing all edges of any of the directions (= all edges of any of the colours), we would end up with an interface pattern on the square lattice (Fig. 8).
In particular, each even-numbered cluster $C \in \mathbf{H}_{4,k}$ identifies its parent C^+ and its inner pattern and its children (should there be any) and their inner pattern. By construction, these two inner patterns disagree both in their chosen direction and their chosen colour. To the direction of C^+, C will assign the colour chosen by C^+ (i.e., green if C^+ is numbered 1 and blue if it is numbered 3). To the direction of the pattern of its children, C will assign the colour chosen by the children. This fully determines the interface pattern of C (if C has no children, then it shall randomly choose one of the two suitable patterns).

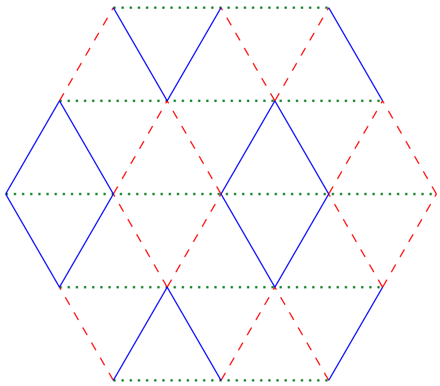

(a) An example of an inner pattern

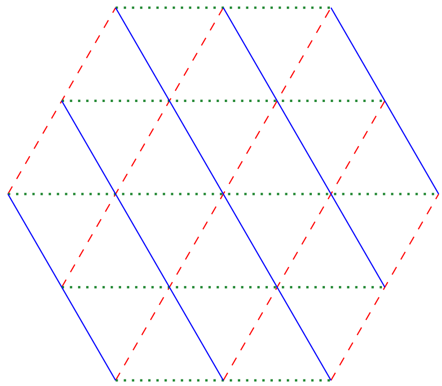

(b) An example of an interface pattern

Fig. 8. Examples of patterns used in odd- and even-numbered clusters.

Step 4 Boundaries. Let ∂C, $C \in \mathbf{H}_{4,k}$ be the cluster boundaries as per the definition in subsection 3.2. For every $C \in \mathbf{H}_{4,k}$, we want the boundary ∂C between C and C^+ to travel only in the two directions not chosen by the odd-numbered cluster. Suppose that $uv \in \partial C$ travels in the undesired direction and x, y are the common neighbours of u and v. Then by the definition of a boundary, exactly one of x, y is in C and one is in C^+ (without loss of generality, $x \in C^+$). We will replace every such uv in the boundary with the pair ux, xv. As $\mathbf{H}_{4,k}$ is k-spaced and $k \geq 3$, such changes cannot cause that boundaries of two different clusters would touch (share a vertex).
Now for a cluster C, there is exactly one colour c such that there is a direction in which all the edges of both C and C^+ have colour c. Moreover, the boundary ∂C never travels in this direction, so after forgetting the edges of colour c, we are left with two neighbouring clusters in the square lattice, one of which has the inner pattern and the other the interface pattern. We proved in subsection 4.1 that there is a 2-colouring of the boundary which amalgamates the two patterns into a balanced 2-colouring.
Finally, as in subsection 4.1, there are no infinite monochromatic paths a.s. To finish the construction, let each monochromatic cycle randomly choose one of the two strong orientations for itself. □

Let us also note that a strategy a little simpler but analogous to the above, of transforming a colour scheme to another one, one colour at a time can be used to obtain a vertex 4-colouring of T. A very similar idea has already been used by Holroyd, Schramm, and Wilson (who employ similar techniques as us to obtain a hierarchy, though they lack the function η from Proposition 3.3) to construct vertex 3-colourings of $\Lambda^d_\square$ [18].

Proposition 4.2. *There is a finitary* $\mathrm{Aut}(T)$*-factor of iid which is a.s. a proper vertex 4-colouring of* T*.*

Proof. There are six proper vertex colourings of T with the colours R, G, and B, which can be described as RGB, RBG, BGR, BRG, GBR, and GRB according to the colouring of one fixed triangle. Let us get a 2-spaced hierarchy J and a numbering $\eta : J \to [4]$ of its parts as given by Proposition 3.3. For each cluster numbered 1, pick one of the six colourings above uniformly at random. Then use the three remaining layers and the fourth colour c_4 to transition between the clusters numbered 1 (e.g., RGB→RGc_4 →BGc_4 →BRc_4 →BRG). □

Suppose G is a graph such that any proper vertex $\chi(G)$-colouring $c : V(G) \to [\chi(G)]$ has the property that knowing $c|_S$ for some finite subset $S \subset V(G)$ determines the entire c – e.g., in the case of T, any two adjacent vertices act as such a subset, and whenever G is bipartite and connected, any singleton would do. Then whenever G has a factor-of-iid hierarchy, we can always construct a factor-of-iid vertex $\chi(G) + 1$-colouring as in the proposition above, while correlation decay prevents the existence of a factor-of-iid $\chi(G)$-colouring. On the other hand, if no such finite subset S exists, then a factor-of-iid vertex $\chi(G)$-colouring may exist as demonstrated in subsection 5.2.

4.4. The Kagomé and the $(3, 4, 6, 4)$ *lattice*

The remaining two Archimedean lattices differ from the first two ones in that not all faces are bounded by cycles of the same length (the lattice $(3, 4, 6, 4)$ is not even edge-transitive). This slightly increased heterogeneity makes the proofs much easier, and in the case of the $(3, 4, 6, 4)$ lattice even the length of monochromatic cycles is bounded, meaning the resulting decoration is a Schreier graph of actions of many more groups than just F_2.

Proof of Theorem 1.1 for $\Lambda = K$. Split the lattice into finite clusters organised in a hierarchy $\mathbf{H}$ as in Lemma 3.9. Each finite cluster picks a pattern composed of monochromatic (red and blue) triangles. There are two ways to randomly place this

(a) Both triangles incident to x are in the same cluster.

(b) One triangle incident to x is in C and one in C^+.

Fig. 9. Kagomé lattice – two cases that can occur at a boundary vertex.

pattern on the Kagomé lattice. Consider the boundaries ∂C where $C \in \mathbf{H}$. If a boundary travels on two out of the three edges of a given triangle abc, replace the two edges in the boundary (say $ab, bc \in \partial C$) by the third one (ac) – as a side effect, each boundary becomes a union of *vertex*-disjoint cycles.

Let each cluster C apply its pattern to all non-boundary edges with at least one endpoint in C. If two neighbouring clusters C and C^+ agree in their pattern, simply propagate it to ∂C as well. If C and C^+ disagree, we will colour the boundary as follows. Suppose uv is in ∂C and w is the common neighbour of u and v. Then by construction, either all three edges uv, vw, wu are in the boundary (in which case we will propagate the pattern of C^+ to them) or neither vw nor wu is in the boundary. In the latter case, vw and wu follow the same pattern and are in the same triangle, so they are assigned the same colour. Let uv get the opposite colour.

As remarked earlier, every vertex $x \in C$ is incident to either two or none edges from ∂C. If it is incident to none and also if the patterns of C and C^+ agree, the colouring at x is an inner pattern, which is balanced. If x is incident to two boundary edges xy and xz and $yz \in E(K)$, the same reasoning applies. Finally, if the patterns disagree and y and z are not adjacent, let y' be the common neighbour of y and x and z' the common neighbour of z and x. Then xy has the opposite colour than xy' by the definition of the colouring on boundaries, and similarly xz has the opposite colour than xz'. Hence the amalgamated colouring is balanced at every vertex. (See Fig. 9.)

As in the previous cases, the colour classes are thus unions of finite cycles, and we let each monochromatic cycle randomly pick one of the two possible strong orientations to finish the decoration. □

Proof of Theorem 1.1 for $\Lambda = (3, 4, 6, 4)$. The $(3, 4, 6, 4)$ lattice is not edge-transitive which makes constructing a Schreier decoration easy in this case. Every edge is either part of a triangle or a hexagon, but not both.

Therefore, we can simply colour all triangles red and all hexagons blue. To complete the Schreier decoration, we choose a random strong orientation for each triangle and hexagon independently. □

Constructing a factor of iid perfect matching on the $(3, 4, 6, 4)$ lattice is similarly easy: each hexagon can choose one of its two perfect matchings independently at random, and these matchings together form a perfect matching of the whole lattice.

5. Ramifications for proper edge and vertex colourings and perfect matchings

Proper edge $2d$-colourings encode actions of the group $(\mathbb{Z}/2\mathbb{Z})^{*2d}$, the $2d$-fold free product of $(\mathbb{Z}/2\mathbb{Z})$ with itself, just like Schreier decorations encode actions of the free group F_d. Both groups have T_{2d} as their standard Cayley graph, and T_{2d} is the universal cover of any $2d$-regular graph. In this sense, looking for these structures on $2d$-regular graphs is equally natural, but we focus more on Schreier decorations because all $2d$-regular graphs have a (deterministic) Schreier decoration, while this is not the case for proper edge $2d$-colourings. When searching for these structures as factors of iid, however, we do not know of a single transitive example, where one exists and the other does not.

5.1. Proper edge $2d$-colourings and perfect matchings

Most of our results in Section 4 are ultimately based on the absence of infinite monochromatic paths, or in other words on decomposing the graphs into monochromatic cycles such that for every colour c and vertex x, there is a cycle of colour c going through x. This decomposition is useful more generally, for example to construct a proper $2d$-colouring whenever the graph in question is bipartite, so we isolate a statement about its existence here.

Corollary 5.1 *(of Theorems 1.1 and 1.2). For every $d \geq 2$, there is a finitary $Aut(\Lambda^d_\square)$-factor of iid that is a partition of the edges of $\Lambda^d_\square$ into d colour classes such that each colour class is a spanning 2-regular graph consisting of finite cycles.*

Corollary 5.2 *(of Lemma 5.1). For every $d \geq 2$, there is a finitary $Aut(\Lambda_{\square}^d)$-factor of iid which is a proper edge $2d$-colouring of $\Lambda_{\square}^d$ almost surely. Subsequently, there is a finitary $Aut(\Lambda_{\square}^d)$-factor of iid which is a perfect matching of $\Lambda_{\square}^d$ almost surely [27].*

Proof. By Lemma 5.1, we can decompose $\Lambda_{\square}^d$ into finite monochromatic cycles. $\Lambda_{\square}^d$ is bipartite, so all of these cycles are even. To obtain a proper $2d$-colouring, let each cycle of colour $c_i, i \in [d]$ choose randomly one of the two proper 2-colourings of its edges by c_i^{light} and c_i^{dark}.

By choosing one colour class (say c_1^{light}), we obtain the second statement. □

In fact, one can now obtain many more corollaries simply by noticing when graphs have a decomposition into several locally identifiable copies of $\Lambda_{\square}^d$. The square lattice with diagonals added is an example of a transitive graph that decomposes into three copies of the square lattice. This decomposition does not even need randomness, one can detect whether an edge is diagonal or not by looking at neighbourhoods of vertices. We use 2 colours for the horizontal and vertical edges (this is one copy of $\Lambda_{\square}$), and 2 other colours for the diagonals (these form two disjoint copies of $\Lambda_{\square}$, we decorate them independently).

Corollary 5.3. *There is an* $\text{Aut}(\Lambda_{\boxtimes})$*-factor of* $\left([0,1]^{V(\Lambda_{\boxtimes})}, \mu_{\boxtimes}\right)$ *which is a Schreier decoration a.s. Moreover, it has almost surely no infinite monochromatic paths. Subsequently, there is an* $\text{Aut}(\Lambda_{\boxtimes})$*-factor of* $\left([0,1]^{V(\Lambda_{\boxtimes})}, \mu_{\boxtimes}\right)$ *which is a perfect matching of* $\Lambda_{\boxtimes}$ *a.s.*

Unfortunately the same argument does not provide a proper edge 6-colouring on the triangular lattice, because there is no guarantee that the finite cycles have even length. Nevertheless in the next example, if care is taken while constructing a Schreier decoration, bipartiteness is not necessary.

Corollary 5.4 *(of Proposition 2.4). Let H be a finite $(2d-2)$-regular graph whose chromatic index is $\chi'(H) = 2d-2$. Then there is a finitary $Aut(H \times P)$-factor of iid which is a proper edge $2d$-colouring of $H \times P$ almost surely.*

Proof. Let us in the proof of Proposition 2.4 always choose a matching M which is given by a colour class of a proper edge colouring. We made sure in the proof to spell out that all cycles of colour c_1 have even length. On each of those, pick randomly one of the two proper edge colourings with colours c_1^{light} and c_1^{dark}.

Now on each connected component left after removing the decorated edges which is isomorphic to H, pick randomly a proper edge colouring with colours $c_3, \ldots, c_{2d}$. On each component isomorphic to $(H \setminus M) \times \{v_i, v_{i+1}\}$, let the $|V(H)|$ edges between $(H \setminus M) \times \{v_i\}$ and $(H \setminus M) \times \{v_{i+1}\}$ all have colour c_3, while we properly colour both $(H \setminus M) \times \{v_i\}$ and $(H \setminus M) \times \{v_{i+1}\}$ with $c_4, \ldots, c_{2d}$. □

We can again obtain a perfect matching of $H \times P$ by choosing one of the colour classes, but there is already the trivial construction of simply picking a perfect matching on each of the P copies of H.

5.2. Line graphs

Proposition 5.5. *Let G be a $2d$-regular graph. Then if G admits a factor-of-iid Schreier decoration, so does its line graph $L(G)$.*

Proof. If $d = 1$ then $G = L(G)$, so the statement is a tautology.

If $d \geq 2$ then there is a 1-to-1 correspondence between vertices of G and the cliques K_{2d} in $L(G)$. Also every vertex v in $L(G)$ is in exactly two cliques K_{2d} because it has two endpoints in G.

Suppose now that K is isomorphic to K_{2d} and its vertices are $c_1^{\text{in}}, c_1^{\text{out}}, \ldots, c_d^{\text{in}}, c_d^{\text{out}}$. Let us fix a proper edge $2d-1$-colouring of K with colours $c_1', \ldots, c_{2d-1}'$. Then put on top an orientation such that for all $i \in [d]$, $j \in [2d-1]$, the edge of colour c_j' incident to c_i^{in} is oriented towards c_i^{in} if and only if the edge of colour c_j' incident to c_i^{out} is oriented from c_i^{out}. This decorated K will serve as a template for decorating the cliques of $L(G)$.

Indeed, the Schreier decoration of G gives rise to a vertex d-colouring of $L(G)$ with colours $c_1, \ldots, c_d$ such that moreover, every vertex of $L(G)$ labels one of the cliques it is in as 'in' and the other as 'out'. Also every clique K_{2d} in $L(G)$ has vertices which are *relatively to this clique* labelled $c_1^{\text{in}}, c_1^{\text{out}}, \ldots, c_d^{\text{in}}, c_d^{\text{out}}$. Given this vertex decoration, let every K_{2d} in $L(G)$ decorate its edges as dictated by the template K. Since every c_i^{in} is c_i^{out} in its other clique, this gives a Schreier decoration of $L(G)$. □

Interestingly, we cannot use the same strategy to show that a proper edge $2d$-colouring of G implies a proper edge $4d-2$-colouring of $L(G)$. This is exactly because every vertex v of the line graph would get a decoration which is the same as viewed from either of the two cliques v is in. However, when d is even, we get the following.

Proposition 5.6. *Let G be a $2d$-regular graph where d is an even positive integer. Then if G admits a factor-of-iid balanced orientation, its line graph $L(G)$ has a factor-of-iid perfect matching.*

Proof. Every vertex v of G gives rise to K_{2d} in $L(G)$ with d vertices labelled 'in' and d labelled 'out'. As d is even, in each such clique, we can pick a (random) perfect matching on the vertices labelled 'in' – crucially, these are labelled 'out' with respect to the other cliques they are in. We claim that this is a perfect matching. Similarly as before, the balanced orientation of G gives a vertex decoration of $L(G)$ such that every vertex is labelled exactly once with 'in' and once with 'out', therefore at every vertex, there is exactly one matched edge to be chosen. □

Finally, the line graphs of $\Lambda_{\square}^{d}$, $d \geq 2$, are the first example in the literature that the authors are aware of infinite transitive graphs that admit proper vertex χ-colouring as a factor of iid, where χ is the classical chromatic number.

Proposition 5.7. *For every $d \geq 2$, there is a finitary $\mathrm{Aut}(L(\Lambda_{\square}^{d}))$-factor of iid which is a proper vertex $\chi(L(\Lambda_{\square}^{d}))$-colouring of $L(\Lambda_{\square}^{d})$ almost surely.*

Proof. Follows from Corollary 5.2. □

6. Open questions and remarks

Question 6.1. Is it true that for all $2d$-regular graphs G which are vertex- or edge-transitive, the following are equivalent?

1. There is a factor of iid which is a proper edge $2d$-colouring of G a.s.
2. There is a factor of iid which is a perfect matching of G a.s.
3. There is a factor of iid which is a Schreier decoration of G a.s.
4. There is a factor of iid which is a balanced orientation of G a.s.

We have demonstrated in this paper that all four of the structures exist as factors of iid for $\Lambda_{\square}^{d}, d \geq 2$, for $H \times P$ where H is a finite $2d$-regular graph with $\chi'(H) = 2d$, and more loosely speaking for graphs that are made up of locally identifiable copies of these. We have also constructed Schreier decorations and balanced orientations on more planar lattices than $\Lambda_{\square}$ and suspect that the hierarchies which are available on them should make it possible to obtain proper edge $2d$-colourings too. (We only achieved this for $\Lambda_{\square}$ and $\Lambda_{\boxtimes}$ in subsection 5.1.) Also pointing in this direction is the fact that all amenable Cayley graphs have invariant random perfect matchings [11,12] – the lattices T and K are indeed Cayley graphs of $G_T = \langle a, b, c \mid a^3 = b^3 = c^3 = abc = e \rangle$ and $G_K = \langle a, b \mid a^3 = b^3 = (ab)^3 = e \rangle$ respectively.

On the other hand, $H \times P$ where $|V(H)|$ is odd forms the prominent example of a class of graphs where none of the four factors exist. We have shown in Proposition 2.3 that no balanced orientation of $H \times P$ is a factor of iid, and the proof that there is no factor-of-iid perfect matching follows similar lines. Indeed, suppose a perfect matching M of $H \times P$ is given, and let $n(i)$ be the number of edges in M of the form $(u, v_i)(u, v_{i+1})$. Then for every $i \in \mathbb{Z}$, if $n(i)$ is odd then $n(i+1)$ must be even and if $n(i)$ is even then $n(i+1)$ must be odd. This is, however, in contradiction with the existence correlation decay in factor of iid processes.

In our parallel paper [2, Definition 22], we further show a class $\mathcal{C}_{2d}^{*} = \{G^{*} : G \text{ is } 2d\text{-regular}\}$ of $2d$-regular bipartite graphs for which a proper edge $2d$-colouring exists as a factor of iid if and only if perfect matching does if and only if Schreier decoration does. All the above leads us to tentatively conjecture that the answer to Question 6.1 is yes for amenable graphs.

It also seems that if the transitivity assumption is relaxed, the four structures are listed in a strictly decreasing order of difficulty. As we said, $H \times P$ where $|V(H)|$ is odd does not have a factor-of-iid balanced orientation, which by [2, Lemma 25] means $(H \times P)^{*} \in \mathcal{C}_{2d}^{*}$ has no factor-of-iid perfect matching. Thus $(H \times P)^{*}$ has no factor-of-iid edge $2d$-colouring either by [2, Proposition 5]. The non-transitivity of $(H \times P)^{*}$, however, allows for locally recognisable factorisation of the graph into copies of $K_{d,2d}$, which implies the existence of factor-of-iid balanced orientation whenever d is even.

To get graphs with a factor-of-iid Schreier decoration but without a factor-of-iid perfect matching, we can start with any $2d$-regular G satisfying Question 6.1 in its positive sense and by attaching two pendant K_{2d+5}^{-} to every vertex obtain a $2d+4$-regular graph which has no perfect matching at all. Alternatively, we can construct graphs like the one in Fig. 10 which have deterministic, but not factor-of-iid perfect matchings. Any perfect matching contains two or no edges of the big squares in an alternating fashion, which again violates correlation decay. However, Schreier decoration can be constructed by colouring the two edges to the left of a cut vertex red and to the right blue, or vice versa, chosen randomly and independently at each cut vertex. Each remaining finite component properly extends this 2-colouring and monochromatic cycles get a random strong orientation.

Finally, any finite $2d$-regular graph has a factor-of-iid Schreier decoration, so whenever its chromatic index equals $2d+1$, we again get an example in which a proper edge $2d$-colouring does not exist, but when a perfect matching does, like in the Meredith graph, we get an instance for the last type which has exactly three of the four structures as factors of iid.

Question 6.2. Is there a $2d$-regular quasi-transitive graph which has a factor-of-iid proper edge $2d$-colouring or perfect matching, but not a factor-of-iid Schreier decoration or balanced orientation?

We also point out that all the Schreier decorations constructed in this paper have no infinite monochromatic paths, i.e., no infinite orbits.

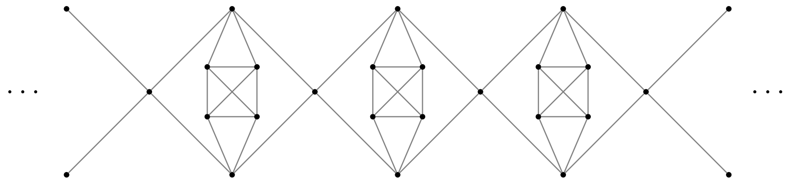

Fig. 10. An infinite graph which admits a factor-of-iid Schreier decoration but not a factor-of-iid perfect matching.

Question 6.3. Is there a factor of iid Schreier decoration on the square lattice that has infinite monochromatic paths with positive probability? Or on any of the other lattices we consider? Or on any transitive graph?

By ergodicity, this would imply that the factor of iid Schreier decoration has infinite monochromatic paths with probability 1. The next question aims to improve our counterexamples.

Question 6.4. Give an example of a transitive graph that is not quasi-isometric to $\mathbb{Z}$ and has no factor of iid balanced orientation. Or has no factor of iid Schreier decoration.

In fact, the best would be to have a precise structural description of when these decorations exist as factors of iid.

Question 6.5. Give a necessary and sufficient condition for a $2d$-regular transitive graph to have a factor of iid balanced orientation and/or Schreier decoration.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Acknowledgement

The authors would like to thank Miklós Abért, Jan Grebík, Matthieu Joseph, Gábor Kun, Gábor Pete, and Václav Rozhoň for inspiring discussions about various parts of this work.

The first author is supported by the NKFIH (National Research, Development and Innovation Office, Hungary) grant KKP-133921. The second and third author are partially supported by the NKFIH (National Research, Development and Innovation Office, Hungary) grant KKP-139502, "Groups and graph limits". The second author was also partially supported by ERC Synergy Grant 810115 DYNASNET, and the third author was partially supported by ERC Consolidator Grant 772466 NOISE.

References

[1] R.J. Baxter, Exactly Solved Models in Statistical Mechanics, Academic Press, London, 1982.
[2] Ferenc Bencs, Aranka Hrušková, László Márton Tóth, Factor-of-iid balanced orientation of non-amenable graphs, Eur. J. Comb. 115 (2024) 103784.
[3] Itai Benjamini, Oded Schramm, Percolation beyond $\mathbb{Z}^d$, many questions and a few answers, Electron. Commun. Probab. (1996) 71–82.
[4] Zoltán L. Blázsik, Jay Cummings, Willem H. Haemers, Cospectral regular graphs with and without a perfect matching, Discrete Math. 338 (3) (2015) 199–201.
[5] Béla Bollobás, Oliver Riordan, Percolation, Cambridge University Press, 2006.
[6] Béla Bollobás, Oliver Riordan, Percolation on dual lattices with k-fold symmetry, Random Struct. Algorithms 32 (4) (2008) 463–472.
[7] Matthew Bowen, Gábor Kun, Marcin Sabok, Perfect matchings in hyperfinite graphings, preprint arXiv:2106.01988, 2021.
[8] Carlos Carbonera, Jason Shepherd, On the Existence of a Perfect Matching for 4-Regular Graphs Derived from Quadrilateral Meshes, SCI Institute, 2006.
[9] Clinton T. Conley, Steve C. Jackson, Andrew S. Marks, Brandon M. Seward, Robin D. Tucker-Drob, Borel asymptotic dimension and hyperfinite equivalence relations, Duke Math. J. 172 (16) (2023) 3175–3226.
[10] Clinton T. Conley, Benjamin D. Miller, A bound on measurable chromatic numbers of locally finite Borel graphs, Math. Res. Lett. 23 (6) (2016) 1633–1644.
[11] T. Conley, A. Kechris, R. Tucker-Drob, Ultraproducts of measure preserving actions and graph combinatorics, Ergod. Theory Dyn. Syst. 33 (2013) 334–374.
[12] Endre Csóka, Gábor Lippner, Invariant random perfect matchings in Cayley graphs, Groups Geom. Dyn. 11 (1) (2017) 211–244.
[13] Michael Damron, Charles M. Newman, Vladas Sidoravicius, Absence of site percolation at criticality in $\mathbb{Z}^2 \times \{0, 1\}$, Random Struct. Algorithms 47 (2) (2015) 328–340.
[14] Gábor Elek, Finite graphs and amenability, J. Funct. Anal. 263 (9) (2012) 2593–2614.
[15] Su Gao, Steve Jackson, Edward Krohne, Brandon Seward, Forcing constructions and countable Borel equivalence relations, J. Symb. Log. 87 (3) (2022) 873–893.
[16] Su Gao, Steve Jackson, Edward Krohne, Brandon Seward, Continuous combinatorics of abelian group actions, preprint arXiv:1803.03872, 2018.

[17] Jan Grebík, Václav Rozhoň, Local problems on grids from the perspective of distributed algorithms, finitary factors, and descriptive combinatorics, Adv. Math. 431 (2023) 109241.
[18] Alexander E. Holroyd, Oded Schramm, David B. Wilson, Finitary coloring, Ann. Probab. 45 (5) (2017) 2867–2898.
[19] Russell Lyons, Factors of IID on trees, Comb. Probab. Comput. 26 (2) (2017) 285.
[20] Andrew S. Marks, Spencer T. Unger, Borel circle squaring, Ann. Math. (2017) 581–605.
[21] Gourab Ray, Yinon Spinka, Finitary codings for gradient models and a new graphical representation for the six-vertex model, Random Struct. Algorithms 61 (1) (2022) 193–232.
[22] Lucio Russo, On the critical percolation probabilities, Z. Wahrscheinlichkeitstheor. Verw. Geb. 56 (2) (1981) 229–237.
[23] Alexander Schrijver, Bounds on the number of Eulerian orientations, Combinatorica 3 (3–4) (1983) 375–380.
[24] Yinon Spinka, Finitely dependent processes are finitary, Ann. Probab. 48 (4) (2020) 2088–2117.
[25] Riley Thornton, Factor maps for automorphism groups via Cayley diagrams, preprint arXiv:2011.14604v2, 2021.
[26] Riley Thornton, Orienting Borel graphs, Proc. Am. Math. Soc. 150 (04) (2022) 1779–1793.
[27] Ádám Timár, Invariant matchings of exponential tail on coin flips in $\mathbb{Z}^d$, preprint arXiv:0909.1090, 2009.
[28] Ádám Timár, One-ended spanning trees in amenable unimodular graphs, Electron. Commun. Probab. 24 (2019) 1–12.
[29] László Márton Tóth, Invariant Schreier decorations of unimodular random networks, Ann. Henri Lebesgue 4 (2021) 1705–1726.
[30] D.J.A. Welsh, The computational complexity of some classical problems from statistical physics, in: Disorder in Physical Systems, 1990.

Appendix C

Limits of action convergent graph sequences with unbounded (p, q)-norms

by Aranka Hrušková

Limits of action convergent graph sequences with unbounded (p, q)-norms

Aranka Hrušková
Alfréd Rényi Institute of Mathematics, Reáltanoda utca 13-15, Budapest, Hungary

Abstract

The recently developed notion of action convergence by Backhausz and Szegedy [2] unifies and generalises the dense (graphon) and local-global (graphing) convergences of graph sequences. This is done through viewing graphs as operators and examining their dynamical properties. Suppose $(A_n)_n^\infty$ is a sequence of operators representing graphs, Cauchy with respect to the action metric. If $(A_n)_n^\infty$ has uniformly bounded (p, q)-norms where (p, q) is any pair in $[1, \infty) \times (1, \infty)$, then Backhausz and Szegedy prove that $(A_n)_n^\infty$ has a limit operator which, moreover, must be self-adjoint and positivity-preserving. In the present work, we construct a large class of graph sequences whose only uniformly bounded (p, q)-norm is the $(\infty, 1)$-norm, but which converge nonetheless. We show that the limit operators in this case are not unique, not self-adjoint, and need not be positivity-preserving. In particular, in the action convergence language, this means that the space of graphops is not compact. By identifying these multiple limits, we also demonstrate that c-regularity is not invariant under weak equivalence, where c is the eigenvalue of the identity function, when the identity function is an eigenfunction.

1 Introduction

The central object of the relatively young field of graph limit theory is a sequence $(G_n)_{n=1}^\infty$ of finite graphs, for which we seek to find a limit object. The two most thoroughly developed convergence notions, based on different methods of sampling small subgraphs from large graphs, are, however, only applicable when either the number of edges in G_n is asymptotically quadratic in terms of the number of vertices [6, 12, 13] or when the maximum degree $\Delta(G_n)$ is uniformly bounded above [12, 17]. We call the former sequences *dense* and have *graphons* for them as limit objects, while the latter are an extreme case of *sparse* sequences and their limits are *graphings* [9]. This leaves out the territory of sequences in which the number of edges grows subquadratically but superlinearly in terms of the number of vertices – that includes for example the hypercubes, the incidence graphs of finite projective planes, and many regimes of the Erdős-Rényi model $\mathcal{G}(n, p(n))$. A number of authors have in recent years defined various extensions of the classical notions mentioned above with the aim of reaching the world of sequences with intermediate densities [3, 4, 5, 7, 8, 11, 14, 15]. In this paper, we are interested in *action convergence* introduced by Backhausz and Szegedy [2], which unifies and generalises graphons and graphings in the common framework of *P-operators.*

Definition (P-operator)**.** *Let $(\Omega, \mathcal{A}, \mu)$ be a probability space. Then we call a linear operator $A\colon L^\infty(\Omega) \to$*

$L^1(\Omega)$ *a P-operator if the norm*

$$\|A\|_{\infty\to 1} = \sup_{v\in L^\infty(\Omega)} \frac{\|Av\|_1}{\|v\|_\infty}$$

is finite.

As the name and the definition of the limit object above suggest, action convergence lays emphasis on the dynamical properties of graphs. There is a number of natural operators associated to a finite graph, of which the most prominent are the adjacency matrix, the discrete Laplacian, and the transition matrix of the simple random walk; in this paper, we will always identify a finite graph with its adjacency matrix. The key to the definition of the convergence is the notion of *profiles* that allow us to compare operators even when they act on L-spaces of different probability spaces. A k-profile $\mathcal{S}_k(A)$ of a P-operator A is a collection of probability measures on $\mathbb{R}^{2k}$ encoding the actions of A and giving rise to the following metrisation.

Definition. *Let* $A\colon L^\infty(\Omega_A)\to L^1(\Omega_A)$ *and* $B\colon L^\infty(\Omega_B)\to L^1(\Omega_B)$ *be P-operators. Then their action distance is*

$$d_M(A,B) = \sum_{k=1}^{\infty} \frac{d_H(\mathcal{S}_k(A),\mathcal{S}_k(B))}{2^k},$$

where d_H *is the Hausdorff distance. A sequence* $(A_n)_n^\infty$ *of P-operators* action converges *to a P-operator* A *if and only if* $\lim_n d_M(A_n,A)=0$*. When* $d_M(A,B)=0$*, we say that* A *and* B *are* weakly equivalent.

Importantly, Backhausz and Szegedy show in [2] that a sequence $(G_n)_n^\infty$ of graphs with uniformly bounded maximum degree converges locally-globally to a graphing $\mathcal{G}$ if and only if their adjacency matrices $(A(G_n))_n^\infty$ action converge to $\mathcal{G}$, and that a sequence $(G_n)_n^\infty$ of graphs converges to a graphon W if and only if their scaled adjacency matrices $\left(\frac{A(G_n)}{|V(G_n)|}\right)_n^\infty$ action converge to W. In our setting, a graphon $W\colon [0,1]^2\to[0,1]$ becomes the P-operator that sends $f\in L^2([0,1],\lambda)$ to

$$(Wf)(x) = \int_0^1 W(x,y)f(y)\,d\lambda(y),$$

and a graphing $(\mathcal{G},\nu)$ becomes the P-operator that sends $v\in L^2(V(\mathcal{G}),\nu)$ to

$$(\mathcal{G}v)(x) = \sum_{xy\in E(\mathcal{G})} v(y).$$

In particular, both graphons and graphings are not just operators from L^∞ to L^1, but from L^2 to L^2, and satisfy that $\|W\|_{2\to 2}\le 1$ and $\|\mathcal{G}\|_{2\to 2}\le d$, where d is the maximum degree of the graphing $\mathcal{G}$. Both of them are also self-adjoint and positivity-preserving – any P-operator that has these two properties is called a *graphop*. Given a sequence of P-operators, we can in general deduce information about the existence of its limit and about the limit's properties if we can assume uniform boundedness of some (p,q)-norms like the $(2,2)$-norms above.

Lemma 1.1 (Sequential compactness, Lemma 2.6 in [2])**.** *Let* $(A_n)_{n=1}^\infty$ *be a sequence of P-operators with uniformly bounded* $\|\cdot\|_{\infty\to 1}$ *norms. Then* $(A_n)_{n=1}^\infty$ *has a Cauchy subsequence with respect to the distance* d_M*.*

Theorem 1.2 (Existence of limit object, Theorem 2.9 in [2])**.** *Let* $p \in [1, \infty)$ *and* $q \in [1, \infty]$*. Let* $(A_n)_{n=1}^{\infty}$ *be a sequence of* P*-operators, Cauchy with respect to the distance* d_M *and with uniformly bounded* $\|\cdot\|_{p\to q}$ *norms. Then there is a* P*-operator* A *such that* $\lim_n d_M(A_n, A) = 0$ *and* $\|A\|_{p\to q} \le \limsup_n \|A_n\|_{p\to q}$*.*

For $c \in \mathbb{R}$, a P-operator is called c-regular if the identity function $\mathbb{1}$ is an eigenfunction with eigenvalue c, i.e., $A\mathbb{1} = c\mathbb{1}$.

Proposition 1.3 (Section 3 in [2])**.** *Let* $p, q \in [1, \infty]$ *be fixed and let* $(A_n)_n^{\infty}$ *be a sequence of* P*-operators with uniformly bounded* (p, q)*-norms. Suppose that* $(A_n)_n^{\infty}$ *action converges to a* P*-operator* A*. Then*

(a) if $q \notin \{1, \infty\}$ *and* A_n *is self-adjoint for every* n*, then* A *is also self-adjoint,*

(b) if $p \neq \infty$ *and* A_n *is positivity-preserving for every* n*, then* A *is also positivity-preserving, and*

(c) if $p \neq \infty$*,* $c \in \mathbb{R}$ *and* A_n *is c-regular for every* n*, then* A *is also c-regular.*

In the present work, we are interested in optimality of the restrictions on p and q in Theorem 1.2 and Proposition 1.3. We quickly show that part (a) of Proposition 1.3 must hold for all $(p, q) \in [1, \infty]^2 \setminus \{(\infty, 1)\}$, but our main result is identifying a large class of sequences $(A_n)_n^{\infty}$ in which A_n is self-adjoint and positivity-preserving for every n and whose only uniformly bounded (p, q)-norms are the $(\infty, 1)$-norms, but which at the same time have multiple limit objects, none of which are self-adjoint and some of which are not positivity-preserving. In other words, we show that a sequence of graphops does not necessarily action converge to a graphop. The graphops in our sequences arise as the adjacency matrices of finite simple graphs G that contain a vertex of degree $|V(G)| - 1$, that is a vertex adjacent to every other vertex of G. We call G^+ the graph on $|V(G)| + 1$ vertices obtained from G by adding a vertex like that.

Theorem 1.4. *Let* $(G_n)_n^{\infty}$ *be a sequence of finite graphs with* $|V(G_n)| \to \infty$*, whose adjacency operators action converge to a* P*-operator* $A\colon L^{\infty}(\Omega, \nu) \to L^1(\Omega, \nu)$*, where* (Ω, ν) *is separable. Then there is a* ν*-filter* $\mathcal{F}$ *on* Ω *such that for any* ν*-ultrafilter* $\mathcal{U}$ *extending* $\mathcal{F}$*, both*

$$A^+\colon L^{\infty}(\Omega, \nu) \to L^1(\Omega, \nu) \quad \text{given by } \left(A^+ g\right)(\omega) = (Ag)(\omega) + \phi_{\mathcal{U}}(g)$$

and

$$A^-\colon L^{\infty}(\Omega, \nu) \to L^1(\Omega, \nu) \quad \text{given by } \left(A^- g\right)(\omega) = (Ag)(\omega) - \phi_{\mathcal{U}}(g)$$

are action limits of $(G_n^+)_n^{\infty}$*, where* $\phi_{\mathcal{U}}\colon L^{\infty}(\Omega, \nu) \to \mathbb{R}$ *is the functional sending* $\lim_{n\to\infty} \sum_{i=1}^{n} \alpha_{n,i}\chi_{E_{n,i}}$ *to* $\lim_{n\to\infty} \sum_{i=1}^{n} \alpha_{n,i}\mathbb{1}_{E_{n,i}\in\mathcal{U}}$*.*

In the construction of the limit objects, we utilise a functional $\phi_{\mathcal{U}}$ from the dual space $(L^{\infty})^*$ of L^{∞} which are not in the canonical embedding of L^1 in $(L^{\infty})^*$. In particular, $\phi_{\mathcal{U}}\colon L^{\infty} \to \mathbb{R}$ arises from a finitely additive measure given by an ultrafilter $\mathcal{U}$, where $\mathcal{U}$ is any extension of a suitable countably generated filter $\mathcal{F}$. In the course of proving that such a filter $\mathcal{F}$ exists, we establish Theorem 4.3 which we believe is of independent interest. It states that for *any* linear operator $A\colon L^{\infty}(\Omega) \to L^1(\Omega)$, where Ω is a finite measure space, there is an ultrafilter $\mathcal{U}$ such that for any function $f \in L^{\infty}$ and a number $a \in \mathbb{R}$, we can find functions f_a which are arbitrarily close to f in the 1-norm and also whose images under A are arbitrarily close to Af in the 1-norm, but such that $\phi_{\mathcal{U}}(f_a) = a$, independently of the value of $\phi_{\mathcal{U}}(f)$.

The rest of the paper is organised as follows. In Section 2, we define k-profiles and prove the extension of part (a) of Proposition 1.3. In Section 3, we warm up with the most basic case of our construction which is the star graphs $(S_n)_n^\infty$. We recall the structure of $(L^\infty)^*$, explain how the functional $\phi_\mathcal{U}$ plays the role of the high-degree vertex, prove that the $(\infty, 1)$-norm is not continuous with respect to d_M by showing that $\lim_n \|S_n\|_{\infty\to 1} > \|\lim_n S_n\|_{\infty\to 1}$, establish that no action limit of S_n can be self-adjoint, and prove a special case of Theorem 1.4. The main results are in Section 4 in which we first describe what the limiting k-profiles of $(G_n^+)_n^\infty$ must be. Then we prove Theorem 4.3 and relying on it, we go on to prove Theorem 1.4. We close off in Section 6 with two open questions.

2 Preliminaries

Let $A\colon L^\infty(\Omega) \to L^1(\Omega)$ be a P-operator. Then for every $f \in L^\infty(\Omega)$, the pair (f, Af) represents an observation of the dynamical properties of A. We want to compress this observation into a form which will not involve Ω and also neglect some inessential features of A, and do so by considering the law of the random variable

$$\begin{aligned}(f, Af)\colon \Omega &\to \mathbb{R}^2 \\ \omega &\mapsto (f(\omega), Af(\omega)).\end{aligned}$$

More generally, taking a k-tuple $(f_1, \dots, f_k)$ of functions in $L^\infty(\Omega)$ provides an even finer observation $(f_1, \dots, f_k, Af_1, \dots, Af_k)$, and by compressing and collecting all of these, we arrive at a set of probability measures which captures the various ways in which A acts on functions.

Definition (k-profile)**.** *Let $A\colon L^\infty(\Omega) \to L^1(\Omega)$ be a P-operator and k a positive integer. The k-profile $\mathcal{S}_k(A)$ of A is*

$$\mathcal{S}_k(A) := \left\{\mathcal{D}(f_1, \dots, f_k, Af_1, \dots, Af_k) : f_1, \dots, f_k \in B_1^{L^\infty}\right\} \subset \mathcal{P}\left(\mathbb{R}^{2k}\right),$$

where $B_1^{L^\infty}$ is the closed unit ball of $L^\infty(\Omega)$ and $\mathcal{D}(f_1, \dots, f_k, Af_1, \dots, Af_k)$ is the joint distribution of $f_1, \dots, f_k$, $Af_1, \dots, Af_k$, i.e., the image measure given by the map $\omega \mapsto (f_1(\omega), \dots, f_k(\omega), Af_1(\omega), \dots, Af_k(\omega))$.

We will also be using the shorthand

$$\mathcal{D}_A(f_1, \dots, f_k) := \mathcal{D}\left(f_1, \dots, f_k, Af_1, \dots, Af_k\right).$$

For a hands-on example, when A is an $n \times n$ matrix then its k-profile is the set of all discrete probability measures on $\mathbb{R}^{2k}$ of the form

$$\frac{1}{n}\sum_{j=1}^{n} \delta_{((v_1)_j, \dots, (v_k)_j, (v_1 A)_j, \dots, (v_k A)_j)}$$

where $v_1, \dots, v_k$ are real vectors with entries in $[-1, 1]$ and δ_x denotes the Dirac measure concentrated on $x \in \mathbb{R}^{2k}$. (We assumed the uniform distribution on $[n]$ here.)

As explained in the Introduction, we measure the similarity of k-profiles of different P-operators by the Hausdorff distance.

Definition (Hausdorff pseudometric)**.** *Let (M, d) be a metric space. Then the Hausdorff pseudometric d_H on the power set $\mathscr{P}(M)$ is given by*

$$d_H(X, Y) = \max\left\{\sup_{x\in X}\inf_{y\in Y} d(x, y), \sup_{y\in Y}\inf_{x\in X} d(x, y)\right\}$$

for all subsets X, Y of M.

Note that $d_H(X, Y) = 0$ if and only if $\overline{X} = \overline{Y}$.

It remains to determine what metric we consider on the set $\mathcal{P}\left(\mathbb{R}^{2k}\right)$ of probability measures on $\mathbb{R}^{2k}$. Since we want to metrise the weak convergence of measures, we choose the Lévy-Prokhorov metric.

Definition (Lévy-Prokhorov metric)**.** *Let (M, d) be a metric space, $\mathcal{B}(M)$ its associated Borel σ-algebra, and $\mathcal{P}(M)$ the set of all probability measures on $(M, \mathcal{B}(M))$. The Lévy-Prokhorov metric d_{LP} on $\mathcal{P}(M)$ is given by*

$$d_{LP}(\eta, \mu) = \inf\{\varepsilon > 0 : \eta(U) \leq \mu(U^\varepsilon) + \varepsilon \text{ and } \mu(U) \leq \eta(U^\varepsilon) + \varepsilon \text{ for every Borel set } U \subset M\},$$

where $U^\varepsilon = \{x \in M : d(x, U) < \varepsilon\}$.

Note that for any two probability measures $\eta, \mu \in \mathcal{P}(\mathbb{R}^n)$, we have $d_{LP}(\eta, \mu) \leq 1$, hence also for any two P-operators A, B, their distance satisfies $d_M(A, B) \leq \sum_{k=1}^{\infty} \frac{1}{2^k} = 1$.

Let now $(\cdot, \cdot)_A$ denote the bilinear form on functions from $L^\infty(\Omega)$ given by a P-operator A as follows:

$$(f, g)_A := \int_\Omega (Af)g \, d\mu = \mathbb{E}[(Af)g].$$

Definition. *A P-operator $A\colon L^\infty(\Omega) \to L^1(\Omega)$ is*

1. self-adjoint *if $(f, g)_A = (g, f)_A$ for all $g, f \in L^\infty(\Omega)$,*
2. positivity-preserving *if $(Af)(\omega) \geq 0$ holds for a.a. $\omega \in \Omega$ whenever $f(\omega) \geq 0$ holds for a.a. $\omega \in \Omega$,*
3. c-regular *if $A\mathbb{1} = c\mathbb{1}$, for $c \in \mathbb{R}$,*
4. *a* graphop *if it is self-adjoint and positivity-preserving.*

While the definition of a P-operator requires the $(\infty, 1)$-operator norm to be finite, the general definition of a (p, q)-norm, for $p, q \in [1, \infty]$, is

$$\|A\|_{p\to q} := \sup_{f\in L^\infty} \frac{\|Af\|_q}{\|f\|_p}.$$

Since the q-norms $\|\cdot\|_q$ are increasing with q, the operator norms $\|\cdot\|_{p\to q}$ are increasing with q and decreasing with p, meaning in particular that

$$\|A\|_{\infty\to 1} \leq \|A\|_{p\to q}$$

for any linear operator A and any $p, q \in [1, \infty]$.

Following the lines of the argument made in [10] for finite matrices, we quickly prove the following.

Lemma 2.1. *Let p, q be in $[1, \infty]$, and let p', q' be their Hölder conjugates. Let A and A^* be P-operators satisfying that*

$$(v, w)_A = (w, v)_{A^*} \text{ for all } v, w \in L^\infty.$$

Then $\|A\|_{p\to q} = \|A^*\|_{q'\to p'}$.

Proof.

$$\begin{aligned}
\|A\|_{p\to q} = \sup_{f\in L^\infty} \frac{\|Af\|_q}{\|f\|_p} &= \sup_{f\in L^\infty} \left\{\|Af\|_q : \|f\|_p \leq 1\right\} \\
&= \sup_{f\in L^\infty} \left\{ \sup_{g\in L^\infty} \left\{ \left|\int (Af)g\right| : \|g\|_{q'} \leq 1 \right\} : \|f\|_p \leq 1 \right\} \\
&= \sup_{g\in L^\infty} \left\{ \sup_{f\in L^\infty} \left\{ \left|\int f(A^*g)\right| : \|f\|_p \leq 1 \right\} : \|g\|_{q'} \leq 1 \right\} \\
&= \sup_{g\in L^\infty} \left\{\|A^*g\|_{p'} : \|g\|_{q'} \leq 1\right\} \\
&= \sup_{g\in L^\infty} \frac{\|A^*g\|_{p'}}{\|g\|_{q'}} \\
&= \|A^*\|_{q'\to p'}
\end{aligned}$$

□

Corollary 2.2. *The assumption in Proposition 1.3 (a) can be extended from $(p, q) \in [1, \infty] \times (1, \infty)$ to $(p, q) \in [1, \infty]^2 \setminus \{(\infty, 1)\}$.*

Proof. Let $p \in (1, \infty)$. If $\|A\|_{p\to 1}$ or $\|A\|_{p\to\infty}$ are uniformly bounded then so are $\|A\|_{\infty\to p'}$ or $\|A\|_{1\to p'}$ by Lemma 2.1, and so Proposition 1.3 (a) gives that A is self-adjoint.

If $\|A\|_{1\to 1}$, $\|A\|_{1\to\infty}$ or $\|A\|_{\infty\to\infty}$ are uniformly bounded then so are, say,

$$\|A\|_{2\to 1} = \|A\|_{\infty\to 2},$$

by monotonicity of the (p, q)-norms. Now we can apply Proposition 1.3 (a) to $\|A\|_{\infty\to 2}$ to conclude that A is self-adjoint. □

Finally, as remarked in [2, Section 2], the (p, q)-norms can be read out from the 1-profiles of P-operators (and are hence invariant under weak equivalence). For a measure μ on $\mathbb{R}^2$, let $\mu_y \in \mathcal{P}(\mathbb{R})$ denote its y-axis marginal. Then in the case of the $(\infty, 1)$-norm, we can get it as follows:

$$\begin{aligned}
\infty > \|B\|_{\infty\to 1} = \sup_{f\in L^\infty} \frac{\|Bf\|_1}{\|f\|_\infty} &= \sup\left\{\|fB\|_1 : \|f\|_\infty \leq 1\right\} \\
&= \sup\left\{ \int_\Omega |Bf|\, d\nu : f \in B_1^{L^\infty} \right\} \\
&= \sup\left\{ \int_{\mathbb{R}} |x|\, d\mathcal{D}(Bf) : f \in B_1^{L^\infty} \right\} \\
&= \sup\left\{ \int_{\mathbb{R}} |x|\, d\mu_y(x) : \mu \in \mathcal{S}_1(B) \right\}.
\end{aligned}$$

This enables us to prove that the $(\infty,1)$-norms are lower semicontinuous with respect to action convergence.

Lemma 2.3. *Let $(A_n)_n^\infty$ be a Cauchy sequence of P-operators with uniformly bounded $\|\cdot\|_{\infty\to 1}$ norms. Suppose further that $\lim_{n\to\infty} d_M(A_n, A) = 0$ for a P-operator A. Then $\|A\|_{\infty\to 1} \leq \liminf_{n\to\infty} \|A_n\|_{\infty\to 1}$.*

Proof. Suppose that a sequence $(\mu_n)_n^\infty$ of measures on $\mathbb{R}^2$ converges in d_{LP} to a measure μ. Then for every n, $d_{LP}((\mu_n)_y, \mu_y) \leq d_{LP}(\mu_n, \mu)$, and so the measures $(\mu_n)_y$ converge in d_{LP} to μ_y. The Lévy-Prokhorov distance is on $\mathbb{R}$ a metrisation of weak convergence of measures, and so by the portmanteau theorem,

$$\int_{\mathbb{R}} |x|\, d\mu_y(x) \leq \liminf_{n\to\infty} \int_{\mathbb{R}} |x|\, d(\mu_n)_y(x)$$

because the function $|x|$ is continuous and bounded below. This implies both that the equality

$$\|B\|_{\infty\to 1} = \sup\left\{ \int_{\mathbb{R}} |x|\, d\mu_y(x) : \mu \in \mathcal{S}_1(B) \right\}$$

from above can be extended to

$$\|B\|_{\infty\to 1} = \sup\left\{ \int_{\mathbb{R}} |x|\, d\mu_y(x) : \mu \in \mathcal{S}_1(B) \right\} = \sup\left\{ \int_{\mathbb{R}} |x|\, d\mu_y(x) : \mu \in \overline{\mathcal{S}_1(B)} \right\}$$

and that any measure $\mu \in \mathcal{S}_1(A)$ and any sequence $(\mu_n)_n^\infty$ with $\mu_n \in \mathcal{S}_1(A_n)$ and $\mu_n \xrightarrow{d_{LP}} \mu$ must satisfy

$$\int_{\mathbb{R}} |x|\, d\mu_y(x) \leq \liminf_{n\to\infty} \int_{\mathbb{R}} |x|\, d(\mu_n)_y(x) \leq \liminf_{n\to\infty} \|A_n\|_{\infty\to 1},$$

implying

$$\|A\|_{\infty\to 1} = \sup\left\{ \int_{\mathbb{R}} |x|\, d\mu_y(x) : \mu \in \mathcal{S}_1(B) \right\} \leq \liminf_{n\to\infty} \|A_n\|_{\infty\to 1}.$$

□

However, we will see in Section 3 that the $(\infty,1)$-norm is not continuous.

3 The case of stars

The question naturally arises whether Theorem 1.2 is as good as it gets, that is whether there are Cauchy sequences with uniformly bounded norms $\|\cdot\|_{\infty\to q}$ for some $q \in [1,\infty]$ which, however, do not action converge to any P-operator.

For positivity-preserving operators A,

$$\|A\|_{\infty\to q} = \|A\mathbb{1}\|_q,$$

where $\mathbb{1} = \chi_\Omega$ is the constant function with value 1. At the same time, we have seen that for self-adjoint operators A,

$$\|A\|_{\infty\to q} = \|A\|_{q'\to 1}.$$

Suppose then that we want to look for a candidate Cauchy sequence $(A_n)_n^\infty$ which would exemplify the impossibility of extending Theorem 1.2 beyond $p \neq \infty$ – if we want to find it among graphops then these two observations tell us that $(\|A_n \mathbb{1}\|_q)_n^\infty$ must be unbounded for every $q \in (1, \infty]$. Let us note that if A is the adjacency matrix of a graph then $A\mathbb{1} = \underline{d}$ is (an ordering of) its degree sequence, and so the sequence $(S_n)_n^\infty$ of stars, having as large a difference between minimum and maximum degree as possible in a simple graph, becomes in immediate candidate to investigate.

In particular, we denote by S_n the n-vertex tree with $n-1$ leaves. Then we can check that for $q \in (1, \infty)$, the q-norm of its degree sequence is unbounded:

$$\|\underline{d}\|_q = \left(\frac{1}{n}\left(1^q + \cdots + 1^q + (n-1)^q\right)\right)^{\frac{1}{q}} = \left(\frac{n-1}{n} + \frac{(n-1)^q}{n}\right)^{\frac{1}{q}} \geq \frac{n-1}{n^{\frac{1}{q}}} \to \infty \text{ as } n \to \infty.$$

By the monotonicities of (p,q)-norms, this means that the $(\infty, 1)$-norm

$$\|S_n\|_{\infty \to 1} = \frac{1}{n}(1 \times (n-1) + (n-1) \times 1) < 2$$

is the *only* uniformly bounded (p,q)-norm of the star sequence, where we took the liberty of identifying S_n with its adjacency matrix.

In the rest of the section, we will investigate what a putative action limit A of $(S_n)_n^\infty$ would have to satisfy, showing that necessarily, the continuity of the $(\infty, 1)$-norm cannot hold, i.e.,

$$\lim_{n\to\infty} \|S_n\|_{\infty \to 1} \neq \|A\|_{\infty \to 1},$$

and that A cannot be self-adjoint, so in particular, it cannot be a graphop. We will derive an action limit at the end of the section.

Lemma 3.1. *A P-operator A satisfies $\lim_{n\to\infty} d_M(S_n, A) = 0$ if and only if*

$$\overline{\mathcal{S}_k(A)} = \mathcal{P}\left([-1,1]^k\right) \times \left\{\delta_z : z \in [-1,1]^k\right\}$$

for every positive integer k.

We shall in fact deduce this lemma from the following more general result.

Lemma 3.2 (Uniform approximability of $\mathcal{P}(M)$)**.** *Let (M, d) be a totally bounded metric space. Then for all $\varepsilon > 0$, there is $N = N(\varepsilon)$ such that for every $\mu \in \mathcal{P}(M, \mathcal{B}(M))$, there is a sequence $\left(\mu_n = \frac{1}{n}\sum_{i=1}^n \delta_{x_i}\right)_n^\infty$ of discrete probability measures such that*

$$\forall n \geq N(\varepsilon) \quad d_{LP}(\mu, \mu_n) < \varepsilon.$$

In particular, N does not depend on μ.

Proof. For a positive integer k, let us fix a finite set $\{z_1, \ldots, z_{m(k)}\} \subseteq M$ whose $\frac{1}{k}$-balls cover M and a partition $M_1, \ldots, M_{m(k)}$ in $\mathcal{B}(M)$ of M such that $z_i \in M_i \subseteq B_{1/k}(z_i)$ for every $i \in [m] = [m(k)]$.

Given a probability measure μ on M and a partition $M_1, \ldots, M_m$ as described, let $\kappa_m \in \mathcal{P}(M)$ be $\kappa_m = \sum_i^m \mu(M_i)\delta_{z_i}$. Then for every measurable A,

$$\mu(A) = \sum_{i=1}^m \mu(A \cap M_i) = \sum_{\substack{i\in[m] \\ A\cap M_i \neq \emptyset}} \mu(A\cap M_i) \leq \sum_{\substack{i\in[m] \\ A\cap M_i \neq \emptyset}} \mu(M_i) = \sum_{\substack{i\in[m] \\ A\cap M_i \neq \emptyset}} \kappa_m(\{z_i\}) \leq \kappa_m\left(A^{1/k}\right)$$

and

$$\kappa_m(A) = \sum_{\substack{i\in[m] \\ z_i\in A}} \kappa_m(\{z_i\}) = \sum_{\substack{i\in[m] \\ z_i\in A}} \mu(M_i) \leq \mu(A^{1/k}),$$

so $d_{LP}(\mu, \kappa_m) \leq \frac{1}{k}$.

The next step is to approximate the discrete measure κ_m by $\mu_n = \frac{1}{n}\sum_i^n \delta_{x_i}$ where for every $i \in [n]$ we will put $x_i = z_j$ for some $j \in [m]$. Let μ_n be any probability measure on $\{z_1, \ldots, z_m\}$ such that $\mu_n(\{z_j\}) = \lfloor n\kappa_m(z_j)\rfloor/n$ or $\lceil n\kappa_m(z_j)\rceil/n$ for every $j \in [m]$. Then $|\kappa_m(z_j) - \mu_n(z_j)| < \frac{1}{n}$ for all $j \in [m]$, and so $d_{LP}(\kappa_m, \mu_n) < \frac{m}{n}$. Finally, the triangle inequality implies

$$d_{LP}(\mu, \mu_n) \leq d_{LP}(\mu, \kappa_m) + d_{LP}(\kappa_m, \mu_n) < \frac{1}{k} + \frac{m}{n},$$

and hence for every $\mu \in \mathcal{P}(M)$ and $n \geq m(k)\cdot k$, we have that $d_{LP}(\mu, \mu_n) < 2/k$ as desired. □

Proof of Lemma 3.1. Let us first note that $\lim_{n\to\infty} d_M(S_n, A) = 0$ if and only if

$$\lim_{n\to\infty} d_H(\mathcal{S}_k(S_n), \mathcal{S}_k(A)) = 0$$

for every positive integer k, so let us fix k for the rest of the proof and show that $\lim_{n\to\infty} d_H(\mathcal{S}_k(S_n), \mathcal{S}_k(A)) = 0$ if and only if

$$\overline{\mathcal{S}_k(A)} = \mathcal{P}\left([-1,1]^k\right) \times \left\{\delta_z : z \in [-1,1]^k\right\}.$$

Lemma 3.2 implies that for every $\varepsilon > 0$, there is an $N = N(\varepsilon)$ such that $\forall n \geq N$,

$$d_H\left(\left\{\frac{1}{n}\sum_{i=1}^n \delta_{x_i} : x_i \in [-1,1]^k\right\} \times \left\{\delta_z : z\in[-1,1]^k\right\}, \mathcal{P}\left([-1,1]^k\right) \times \left\{\delta_z : z \in [-1,1]^k\right\}\right) \leq \varepsilon.$$

On the other hand, S_n sends the vector $(a_1, a_2, \ldots, a_n) \in [-1,1]^n$ to $(\sum_{j=2}^n a_j, a_1, \ldots, a_1)$, which means that every choice of vectors $v_1, \ldots, v_k$ with entries in $[-1,1]$ gives an element $\frac{1}{n}\big(\delta_{((v_1)_1,\ldots,(v_k)_1,\sum_{j=2}^n (v_1)_j,\ldots,\sum_{j=2}^n (v_k)_j)} + \sum_{j=2}^n \delta_{((v_1)_j,\ldots,(v_k)_j,(v_1)_1,\ldots,(v_k)_1)}\big)$ of $\mathcal{S}_k(S_n)$ which satisfies

$$d_{LP}\Bigg(\frac{1}{n}\Big(\delta_{((v_1)_1,\ldots,(v_k)_1,\sum_{j=2}^n (v_1)_j,\ldots,\sum_{j=2}^n (v_k)_j)} + \sum_{j=2}^n \delta_{((v_1)_j,\ldots,(v_k)_j,(v_1)_1,\ldots,(v_k)_1)}\Big),$$
$$\frac{1}{n}\Big(\delta_x + \sum_{j=2}^n \delta_{((v_1)_j,\ldots,(v_k)_j)}\Big) \times \delta_{((v_1)_1,\ldots,(v_k)_1)}\Bigg) \leq \frac{1}{n}$$

for every $x \in [-1,1]^k$. This means that

$$d_H\left(\mathcal{S}_k(S_n), \left\{\frac{1}{n}\sum_{i=1}^{n}\delta_{x_i} : x_i \in [-1,1]^k\right\} \times \left\{\delta_z : z \in [-1,1]^k\right\}\right) \leq \frac{1}{n},$$

and so the triangle inequality tells us that $\forall n \geq N$,

$$\begin{aligned} & d_H\big(\mathcal{S}_k(S_n), \mathcal{P}\left([-1,1]^k\right) \times \{\delta_z : z \in [-1,1]^k\}\big) \\ \leq\ & d_H\Big(\mathcal{S}_k(S_n), \Big\{\frac{1}{n}\sum_{i=1}^{n}\delta_{x_i} : x_i \in [-1,1]^k\Big\} \times \{\delta_z : z \in [-1,1]^k\}\Big) \\ & \quad + d_H\Big(\Big\{\frac{1}{n}\sum_{i=1}^{n}\delta_{x_i} : x_i \in [-1,1]^k\Big\} \times \left\{\delta_z : z \in [-1,1]^k\right\}, \mathcal{P}\left([-1,1]^k\right) \times \left\{\delta_z : z \in [-1,1]^k\right\}\Big) \\ \leq\ & \quad \varepsilon + \frac{1}{n}. \end{aligned}$$

Now if $\overline{\mathcal{S}_k(A)} = \mathcal{P}\left([-1,1]^k\right) \times \{\delta_z : z \in [-1,1]^k\}$ then $\forall n \geq N$,

$$\begin{aligned} & d_H\big(\mathcal{S}_k(S_n), \mathcal{S}_k(A)\big) \\ & \leq d_H\left(\mathcal{S}_k(S_n), \mathcal{P}\left([-1,1]^k\right) \times \{\delta_z : z \in [-1,1]^k\}\right) + d_H\left(\mathcal{P}\left([-1,1]^k\right) \times \{\delta_z : z \in [-1,1]^k\}, \mathcal{S}_k(A)\right) \\ & = d_H\left(\mathcal{S}_k(S_n), \mathcal{P}\left([-1,1]^k\right) \times \{\delta_z : z \in [-1,1]^k\}\right) \\ & \leq \varepsilon + \frac{1}{n}, \end{aligned}$$

and so $\lim_{n\to\infty} d_H(\mathcal{S}_k(S_n), \mathcal{S}_k(A)) = 0$.

On the other hand, if $\lim_n d_H(\mathcal{S}_k(S_n), \mathcal{S}_k(A)) = 0$ and $\varepsilon > 0$ then there is $M = M(\varepsilon)$ such that $\forall n \geq M$ $d_H(\mathcal{S}_k(S_n), \mathcal{S}_k(A)) < \varepsilon$. Combining this with the computations above, we get that $\forall n \geq \max\{N(\varepsilon), M(\varepsilon)\}$,

$$\begin{aligned} & d_H\big(\mathcal{S}_k(A), \mathcal{P}([-1,1]^k) \times \{\delta_z : z \in [-1,1]^k\}\big) \\ & \leq d_H\left(\mathcal{S}_k(A), \mathcal{S}_k(S_n)\right) + d_H\big(\mathcal{S}_k(S_n), \mathcal{P}([-1,1]^k) \times \{\delta_z : z \in [-1,1]^k\}\big) \\ & < 2\varepsilon + \frac{1}{n}, \end{aligned}$$

and so in fact $d_H\left(\mathcal{S}_k(A), \mathcal{P}([-1,1]^k) \times \{\delta_z : z \in [-1,1]^k\}\right) = 0$. As $\mathcal{P}\left([-1,1]^k\right) \times \left\{\delta_z : z \in [-1,1]^k\right\}$ is closed, this is equivalent to $\overline{\mathcal{S}_k(A)} = \mathcal{P}\left([-1,1]^k\right) \times \left\{\delta_z : z \in [-1,1]^k\right\}$ as desired. □

Let us now assume that there is a P-operator $A\colon L^\infty(\Omega,\mu) \to L^1(\Omega,\mu)$ for some probability space $(\Omega, \mathcal{A}, \mu)$ such that $\lim_{n\to\infty} d_M(S_n, A) = 0$.

Every $f \in B_1^{L^\infty}$ gives an element $\mathcal{D}(f, Af)$ of $\mathcal{S}_1(A)$ which, by Lemma 3.2, is of the form $\nu \times \delta_{c_f}$ for some constant $c_f \in [-1,1]$. That is, A is in fact an element of $L^\infty(\Omega)^*$ and sends f to c_f.

Lemma 3.3. *For every $f \in B_1^{L^\infty}$, it must be the case that $|c_f| \leq \|f\|_\infty$.*

Proof.

$$\begin{aligned}\left\|\frac{f}{\|f\|_\infty}\right\|_\infty = 1 &\Rightarrow \frac{f}{\|f\|_\infty} \in B_1^{L^\infty}\\ &\Rightarrow \left|A\frac{f}{\|f\|_\infty}\right| \equiv |c_{\frac{f}{\|f\|}}| \le 1\\ &\Rightarrow \|f\|_\infty \left|A\frac{f}{\|f\|_\infty}\right| \le \|f\|_\infty\\ &\Rightarrow |c_f| \equiv |Af| \le \|f\|_\infty\end{aligned}$$

□

Next we note that for every $x \in [-1,1] \setminus \{0\}$, the Dirac measure $\delta_{(0,x)}$ is in $\overline{\mathcal{S}_1(A)}$, while the lemma above says it cannot be in the profile $\mathcal{S}_1(A)$ itself. Thus for every $x \in [-1,1] \setminus \{0\}$ and for every $\epsilon > 0$ there must be $f_{x,\epsilon} \in B_1^{L^\infty}$ such that $d_{LP}\big(\delta_{(0,x)}, \mathcal{D}_A(f_{x,\epsilon})\big) < \epsilon$. Now

$$\begin{aligned}&d_{LP}\big(\delta_{(0,x)}, \mathcal{D}_A(f_{x,\epsilon})\big) < \epsilon\\ \Leftrightarrow\ &\inf\{\delta > 0 : \delta_{(0,x)}(U) \le \mathcal{D}_A(f_{x,\epsilon})(U^\delta) + \delta \text{ and } \mathcal{D}_A(f_{x,\epsilon})(U) \le \delta_{(0,x)}(U^\delta) + \delta \text{ for every Borel set } U \subset \mathbb{R}^2\} < \epsilon\\ \Leftrightarrow\ &\delta_{(0,x)}(U) \le \mathcal{D}_A(f_{x,\epsilon})(U^\epsilon) + \epsilon \text{ and } \mathcal{D}_A(f_{x,\epsilon})(U) \le \delta_{(0,x)}(U^\epsilon) + \epsilon \text{ for every Borel set } U \subset \mathbb{R}^2\\ \Leftrightarrow\ &1 \le \mathcal{D}_A(f_{x,\epsilon})\big(B_\epsilon((0,x))\big) + \epsilon\\ \Rightarrow\ &c_{f_{x,\epsilon}} \in (x-\epsilon, x+\epsilon) \cap [-1,1]. \qquad (3.1)\end{aligned}$$

Also let $\Omega^x_{<\epsilon}$ be the subset $\{\omega : |f_{x,\epsilon}(\omega)| < \epsilon\}$ of Ω, and similarly let $\Omega^x_{\ge\epsilon} = \{\omega : |f_{x,\epsilon}(\omega)| \ge \epsilon\} \subset \Omega$. Then apart from the line (3.1) above, $1 \le \mathcal{D}_A(f_{x,\epsilon})\big(B_\epsilon((0,x))\big) + \epsilon$ also implies that $\mu\left(\Omega^x_{<\epsilon}\right) \ge 1-\epsilon$ and $\mu\left(\Omega^x_{\ge\epsilon}\right) \le \epsilon$.

Corollary 3.4. *If a P-operator A is an action limit of $(S_n)_n^\infty$ then $\|A\|_{\infty\to1} = 1$.*

Proof.

$$\|A\|_{\infty\to1} = \sup_{f\in B_1^{L^\infty}} \frac{\|Af\|_1}{\|f\|_\infty} = \sup_{f\in B_1^{L^\infty}} \frac{|c_f|}{\|f\|_\infty} \le \sup_{f\in B_1^{L^\infty}} \frac{\|f\|_\infty}{\|f\|_\infty} = 1.$$

On the other hand, for every $\varepsilon > 0$

$$\|A\|_{\infty\to1} \ge \frac{\|Af_{1,\varepsilon}\|_1}{\|f_{1,\varepsilon}\|_\infty} = \frac{|c_{f_{1,\varepsilon}}|}{\|f_{1,\varepsilon}\|_\infty} > \frac{1-\varepsilon}{\|f_{1,\varepsilon}\|_\infty} \ge 1-\varepsilon,$$

so $\|A\|_{\infty\to1} \ge 1$, completing the proof. □

Since we will later construct an action limit of $(S_n)_n^\infty$, Corollary 3.4 shows that, as claimed at the beginning of the section, the $(\infty,1)$-norm is not continuous with respect to action convergence. In particular, if $\lim_n d_M(S_n, A) = 0$ then

$$\lim_{n\to\infty} \|S_n\|_{\infty\to1} = \lim_{n\to\infty} \frac{2n-2}{n} = 2 \ne 1 = \|A\|_{\infty\to1}.$$

Using the notation set up before Corollary 3.4, we will now deliver on our second promise of proving that A cannot possibly be self-adjoint.

Proposition 3.5. *Suppose that a P-operator A is an action limit of $(S_n)_n^\infty$. Then A is not self-adjoint.*

Proof. Let $\varepsilon \in (0, 1/3)$ and $\mathbb{1}$ be the characteristic function χ_Ω. Then

$$|(\mathbb{1}, f_{1,\varepsilon})_A| = \left|\int (A\mathbb{1}) f_{1,\varepsilon}\, d\mu\right| = \left|c_{\mathbb{1}} \int f_{1,\varepsilon}\, d\mu\right| \leq \left|\int_{\Omega^1_{<\varepsilon}} f_{1,\varepsilon}\, d\mu + \int_{\Omega^1_{\geq\varepsilon}} f_{1,\varepsilon}\, d\mu\right| \leq \varepsilon\mu\left(\Omega^1_{<\varepsilon}\right) + \mu\left(\Omega^1_{\geq\varepsilon}\right) \leq 2\varepsilon$$

while

$$|(f_{1,\varepsilon}, \mathbb{1})_A| = \left|\int (A f_{1,\varepsilon}) \mathbb{1}\, d\mu\right| = |c_{f_{1,\varepsilon}}| > 1 - \varepsilon.$$

But then

$$|(f_{1,\varepsilon}, \mathbb{1})_A| > 1 - \varepsilon > 2\varepsilon \geq |(\mathbb{1}, f_{1,\varepsilon})_A|,$$

and so $(f_{1,\varepsilon}, \mathbb{1})_A \neq (\mathbb{1}, f_{1,\varepsilon})_A$ and A is not self-adjoint. $\square$

Again, since we will actually show the existence of an action limit A at the end of the section, Proposition 3.5 demonstrates a limitation of Proposition 1.3 (a) and together with Corollary 2.2 effectively answers the question of necessary self-adjointness of a limit of a sequence of self-adjoint P-operators.

Proposition 3.6. *There is no $g \in L^1(\Omega, \mu)$ such that $Af \equiv \int_\Omega gf\, d\mu$ for every $f \in L^\infty(\Omega, \mu)$.*

Proof. Suppose that on the contrary, $g \in L^1(\Omega, \mu)$ is as stated. Then

$$\|g\|_1 = \int_\Omega |g| \mathrm{d}\mu = \int_\Omega g\left(\mathbb{1}_{\{\omega: g(\omega)\geq 0\}} - \mathbb{1}_{\{\omega: g(\omega)<0\}}\right) \mathrm{d}\mu = A\left(\mathbb{1}_{\{\omega: g(\omega)\geq 0\}} - \mathbb{1}_{\{\omega: g(\omega)<0\}}\right) \leq 1.$$

Now let $x \in [-1, 1] \setminus \{0\}$ and $\epsilon > 0$. By the discussion between Lemma 3.3 and Corollary 3.4, we have that

$$x - \epsilon < |c_{f_{x,\epsilon}}| = \left|\int g f_{x,\epsilon} \mathrm{d}\mu\right| \leq \int |g f_{x,\epsilon}| = \int_{\Omega^x_{<\epsilon}} |g f_\epsilon| + \int_{\Omega^x_{\geq\epsilon}} |g f_\epsilon| < \epsilon \int_{\Omega^x_{<\epsilon}} |g| + \int_{\Omega^x_{\geq\epsilon}} |g| \leq \epsilon\Big(1 - \int_{\Omega^x_{\geq\epsilon}} |g|\Big) + \int_{\Omega^x_{\geq\epsilon}} |g|,$$

and so

$$\int_{\Omega^x_{\geq\epsilon}} |g| > \frac{x - 2\epsilon}{1 - \epsilon}.$$

Taking $x = 1$, we get

$$\int_{\Omega^1_{\geq\epsilon}} |g| > \frac{1 - 2\epsilon}{1 - \epsilon} \to 1 \text{ as } \epsilon \to 0, \tag{3.2}$$

and recalling $\int_{\Omega^1_{\geq\epsilon}} |g| \leq \int_\Omega |g| \leq 1$, we conclude that $\int_\Omega |g| = 1$.

Now for any integer $n \geq 2$, let $B_n = \bigcup_{i=n}^\infty \Omega^1_{\geq 2^{-i}} \subset \Omega$. Note that for every $m \geq n$, (3.2) tells us that

$$1 \geq \int_{B_n} |g| \geq \int_{\Omega^1_{\geq 2^{-m}}} |g| > 1 - 2^{1-m} \to 1 \text{ as } m \to \infty,$$

and so for every $n \geq 2$ we have $\int_{B_n} |g| = 1$ and thus $\int_{\Omega \setminus B_n} |g| = 0$. On the other hand, observe that

$$\mu(B_n) = \mu\left(\bigcup_{i=n}^{\infty} \Omega^1_{\geq 2^{-i}}\right) \leq \sum_{i=n}^{\infty} \mu\left(\Omega^1_{\geq 2^{-i}}\right) \leq \sum_{i=n}^{\infty} 2^{-i} = 2^{1-n},$$

which implies that the intersection of the chain $B_2 \supseteq B_3 \supseteq B_4 \supseteq B_5 \supseteq \dots$ has measure zero. But then

$$1 = \int_{\Omega} |g| = \int_{\cap_{n=2}^{\infty} B_n} |g| + \int_{\cup_{n=2}^{\infty} (\Omega \setminus B_n)} |g| = 0 + 0 = 0,$$

which is a contradiction. □

We have just shown that when $A\colon L^\infty(\Omega, \mathcal{A}) \to \mathbb{R}$ is an action limit of $(S_n)_n^\infty$ then

$$A \in L^\infty(\Omega)^* \setminus L^1(\Omega),$$

where we are abusing notation and identifying elements $g \in L^1(\Omega)$ with the functionals $f \mapsto \int fg\, d\mu$. Let us now recall the structure of the dual space $L^\infty(\Omega, \mathcal{A}, \mu)^*$. We mostly follow the terminology and notation of [16] and [1]. Let $ba(\mathcal{A})$ be the space of of bounded finitely additive signed measures, also known as signed charges. The total variation of a charge ν on the σ-algebra $\mathcal{A}$ is

$$|\nu|(\Omega) = \sup\left\{\sum_{i=1}^{n} |\nu(M_i)| : \{M_1, \dots, M_n\} \text{ is a measurable partition of } \Omega\right\},$$

and $\nu \in ba(\mathcal{A})$ if and only if $|\nu|(\Omega) < \infty$. The dual of $L^\infty(\mathcal{A}, \mu)$ is represented by the subset $ba(\mathcal{A}, \mu)$ of $ba(\mathcal{A})$ which consists of all the finitely additive signed measures ν that satisfy

$$\nu(N) = 0 \text{ whenever } \mu(N) = 0, \quad \text{for all } N \in \mathcal{A}.$$

Any such charge ν then gives a functional via

$$f \mapsto \int_{\Omega} f\, d\nu = \lim_{n\to\infty} \sum_{i=1}^{n} \alpha_{n,i} \nu\left(E_{n,i}\right),$$

where $\left(\sum_{i=1}^{n} \alpha_{n,i}\chi_{E_{n,i}}\right)_n^\infty$ is any sequence of simple functions which converge to f in the L^∞-norm, and every element of $L^\infty(\mathcal{A}, \mu)^*$ arises this way. Moreover, as shown for example in Section 6.2 of [1], the $(\infty, 1)$-norm of this functional is $|\nu|(\Omega)$. Having proved Corollary 3.4, we could have used this fact to streamline the first part of the proof of Proposition 3.6.

Every $\nu \in ba(\mathcal{A})$ can be uniquely written as $\nu = \kappa_\nu + \gamma_\nu$, where κ_ν is countably additive, i.e., κ_ν is a singed measure, and γ_ν is purely finitely additive. Having proved that if A is an action limit of stars, the purely finitely additive part of the finitely additive measure representing A must be non-trivial, we now zoom in on a particular subset of purely finitely additive measures. We start by introducing μ-ultrafilters, which turn out to be in one-to-one correspondence with the $\{0, 1\}$-valued elements of $ba(\mathcal{A}, \mu)$.

Definition (μ-filter)**.** *Let $(\Omega, \mathcal{A}, \mu)$ be a measure space. A non-empty collection $\mathcal{F}$ of measurable subsets of Ω is a* μ-filter *if and only if*

- $\mu(S) > 0$ *for all* $S \in \mathcal{F}$,
- *if* $S, T \in \mathcal{F}$ *then* $S \cap T \in \mathcal{F}$, *and*
- *if* $S \in \mathcal{F}$ *and* $S \subset T$ *then* $T \in \mathcal{F}$.

A maximal filter is called an ultrafilter.

Every μ-ultrafilter $\mathcal{U}$ now gives a finitely additive measure $\delta_{\mathcal{U}}$ by setting

$$\delta_{\mathcal{U}}(E) = \begin{cases} 1, & \text{if } E \in \mathcal{U} \\ 0, & \text{if } E \notin \mathcal{U}. \end{cases}$$

Moreover, every $\{0,1\}$-valued element of $ba(\mathcal{A}, \mu)$ arises this way, and if μ is atomless then each such charge is purely finitely additive.

Theorem 3.7. *Let $\mathcal{U}$ be a λ-ultrafilter on the probability space $(\Omega = [0,1), \mathcal{B}([0,1)), \lambda)$. Then the P-operator A given by*

$$\begin{aligned} A \colon L^\infty(\lambda) &\to L^1(\lambda) \\ f &\mapsto \left(\int f \, d\delta_{\mathcal{U}} \right) \cdot \chi_\Omega \end{aligned}$$

satisfies $\lim_{n\to\infty} d_M(S_n, A) = 0$.

Proof. By Lemma 3.1, this is equivalent to showing that for every positive integer k,

$$\overline{\mathcal{S}_k(A)} = \mathcal{P}\left([-1,1]^k\right) \times \{\delta_z : z \in [-1,1]^k\}.$$

Let k be fixed. The operator A sends every element $f \in L^\infty(\Omega)$ to a constant function where the absolute value of the constant is at most $\|f\|_\infty$, so in particular elements of $B_1^{L^\infty}$ are sent to $[-1,1]$. Thus $\mathcal{S}_k(A) \subset \mathcal{P}\left([-1,1]^k\right) \times \{\delta_z : z \in [-1,1]^k\}$. On the other hand, for every $\mu \in \mathcal{P}\left([-1,1]^k\right)$ and $z \in [-1,1]^k$, we will now construct a sequence $(\mathcal{D}_A(f_{n,1}, \ldots, f_{n,k}))_n^\infty$ in $\mathcal{S}_k(A)$ such that $\lim_n d_{LP}\left(\mu \times \delta_z, \mathcal{D}_A(f_{n,1}, \ldots, f_{n,k})\right) = 0$, implying that $\mathcal{P}\left([-1,1]^k\right) \times \{\delta_z : z \in [-1,1]^k\} \subseteq \overline{\mathcal{S}_k(A)}$.

We start by fixing a sequence $(E_n)_n^\infty$ in $\mathcal{U}$ such that for every $n \geq 1$, $E_n = \left[\frac{i-1}{n}, \frac{i}{n}\right)$ for some $i \in [n]$. Given $\mu \in \mathcal{P}\left([-1,1]^k\right)$ and $z \in [-1,1]^k$, let $(\mu_n = \frac{1}{n}\sum_i^n \delta_{x_i})_n^\infty$ be a sequence given by Lemma 3.2. Now for every $n \geq 1$ and $j \in [k]$ let

$$f_{n,j}(\omega) = \begin{cases} (x_i)_j, & \text{if } \omega \in \left[\frac{i-1}{n}, \frac{i}{n}\right) \setminus E_n \\ z_j, & \text{if } \omega \in E_n. \end{cases}$$

Then

$$\begin{aligned} d_{LP}\big(\mu_n \times \delta_z, \mathcal{D}_A(f_{n,1}, \ldots, f_{n,k})\big) &= d_{LP}\Bigg(\mu_n \times \delta_z, \frac{1}{n}\Big(\delta_z + \sum_{\substack{i \in [n] \\ [\frac{i-1}{n}, \frac{i}{n}) \neq E_n}} \delta_{x_i}\Big) \times \delta_z\Bigg) \\ &= d_{LP}\Bigg(\frac{1}{n}\sum_{i=1}^n \delta_{x_i}, \frac{1}{n}\Big(\delta_z + \sum_{\substack{i \in [n] \\ [\frac{i-1}{n}, \frac{i}{n}) \neq E_n}} \delta_{x_i}\Big)\Bigg) \leq \frac{1}{n}. \end{aligned}$$

Hence, by the triangle inequality,

$$\begin{aligned} d_{LP}\big(\mu \times \delta_z, \mathcal{D}_A(f_{n,1}, \dots, f_{n,k})\big) &\leq d_{LP}\left(\mu \times \delta_z, \mu_n \times \delta_z\right) + d_{LP}\big(\mu_n \times \delta_z, \mathcal{D}_A(f_{n,1}, \dots, f_{n,k})\big) \\ &\leq d_{LP}(\mu, \mu_n) + \frac{1}{n} \to 0 \text{ as } n \to \infty. \end{aligned}$$

□

Remark 3.8. *Note that the limit object in Theorem 3.7 is definitely not unique. Not only do we have a choice of the ultrafilter, but we could also take any* $(M, \mathcal{B}(M), \kappa)$ *instead of* $(\Omega = [0,1), \mathcal{B}([0,1)), \lambda)$*, where* (M, d) *is a totally bounded metric space such that for infinitely many positive integers* n*, it is possible to partition* M *into* $M_1, \dots, M_n \in \mathcal{B}(M)$ *with* $\kappa(M_i) = \frac{1}{n}$ *for all* $i \in [n]$*.*

4 Graphs G^+ with a vertex that neighbours everything

Stars are in fact just a special case of the following construction.

Definition. *Let* G *be a graph. Then* G^+ *is the graph on* $|V(G)| + 1$ *vertices formed from* G *by adding a single vertex* v *and connecting it to all the other vertices.*

Seen like this, a star with n leaves is E_n^+ where E_n is the empty graph on n vertices. This viewpoint allows us to extend the result from the previous section to more sequences of the form $(G_n^+)_n^\infty$ than just $(S_n)_n^\infty = (E_n^+)_n^\infty$. Before stating the full theorem, we set up the scene with a couple of lemmas and propositions.

Lemma 4.1. *Let* $(G_n)_n^\infty$ *be Cauchy in* d_M*. If* $\limsup_{n\to\infty} |V(G_n)| = \infty$ *then* $\liminf_{n\to\infty} |V(G_n)| = \infty$*, and for any action limit* $A\colon L^\infty(\Omega, \nu) \to L^1(\Omega, \nu)$ *of* $(G_n)_n^\infty$*,* (Ω, ν) *is atomless.*

Proof. Suppose on the contrary that there is a natural number k such that for infinitely many n, the graph G_n has k vertices. Then there is a subsequence $(G_{m_j})_j$ such that $|V(G_{m_i})| = k$ for all $i \geq 1$. On the other hand, let $(G_{n_i})_i$ be a subsequence with strictly increasing number of vertices. For a measure $\mu \in \mathcal{P}\left(\mathbb{R}^2\right)$, let $\mu_x \in \mathcal{P}(\mathbb{R})$ be its x-axis marginal. Then $d_{LP}\left(\mu_x, \kappa_x\right) \leq d_{LP}(\mu, \kappa)$ for all $\mu, \kappa \in \mathcal{P}(\mathbb{R}^2)$, and thus

$$d_H\left((\mathcal{S}_1)_x(G_n), (\mathcal{S}_1)_x(G_m)\right) \leq d_H\left(\mathcal{S}_1(G_n), \mathcal{S}_1(G_m)\right) \leq 2 d_M(G_n, G_m)$$

for any $n, m \in \mathbb{N}$, where

$$(\mathcal{S}_1)_x(G) := \{\mu_x : \mu \in \mathcal{S}_1(G)\} = \left\{\mathcal{D}(f) : f \in B_1^{L^\infty}\right\}.$$

Now $\lim_i (\mathcal{S}_1)_x(G_{n_i})$ contains the uniform measure on $[-1,1]$, which implies that both $\lim_j (\mathcal{S}_1)_x(G_{m_j})$ and $\overline{(\mathcal{S}_1)_x(A)}$ must contain it too. But $\lim_j (\mathcal{S}_1)_x(G_{m_j})$ only contains measures which are approximable arbitrarily well by k atoms, so we get a contradiction with $\liminf_n |V(G_n)| < \infty$. Similarly, $\overline{(\mathcal{S}_1)_x(A)}$ containing the uniform measure on $[-1,1]$ implies that (Ω, ν) is atomless. □

Lemma 4.1 implies that if $(G_n^+)_n^\infty$ is Cauchy in d_M and $\limsup_{n\to\infty} |V(G_n)| = \infty$ then the special added vertices v_n will have smaller and smaller weight, yet the fact that they are adjacent to all the other vertices in their graph means that the value at v_n has a great impact on the outcome after applying the adjacency operator. In particular, if two functions $f, g : V(G_n^+) \to \mathbb{R}$ only differ in the value they assign to the added vertex v_n, then as n grows larger, the resulting measures $\mathcal{D}_{G_n}(f)$ and $\mathcal{D}_{G_n}(g)$ will more and more look like they are equivalent up to a shift along the y-axis by $f(v_n) - g(v_n)$. To formalise the intuition forming upon this observation, we introduce the following definition.

Definition. *Let μ be a measure on $\mathbb{R}^n$ and $v \in \mathbb{R}^n$ a vector. Then $\mu \oplus v$ is the measure on $\mathbb{R}^n$ such that*

$$(\mu \oplus v)(T) = \mu(T \oplus \{-v\}) = \mu(\{t - v : t \in T\})$$

for any measurable subset T. For any set $\mathcal{S}$ of measures on $\mathbb{R}^n$ and set $V \subseteq \mathbb{R}^n$ of vectors,

$$\mathcal{S} \oplus V := \{\mu \oplus v : \mu \in \mathcal{S}, v \in V\}.$$

So as not to clutter the exposition with technicalities, we will only focus on sequences $(G_n)_n^\infty$ with $|V(G_n)| \to \infty$.

Proposition 4.2. *If $(G_n)_n^\infty$ is Cauchy in d_M and $\limsup_n |V(G_n)| = \infty$ then $(G_n^+)_n^\infty$ is Cauchy in d_M too. Moreover, if the limiting closures of the k-profiles of $(G_n)_n^\infty$ are $X_k := \lim_{n\to\infty} \overline{\mathcal{S}_k(G_n)}$ then the limiting closures of $\mathcal{S}_k(G_n^+)$ are $X_k \oplus V_k$ where $V_k = \{0\}^k \times [-1, 1]^k$.*

Proof. Let us fix $\mu \in X_k$ and $v = \{0\}^k \times (v_1, \dots, v_k) \in V_k$. By the definition of X_k, there is a sequence $(\mu_n)_n^\infty = (\mathcal{D}_{G_n}(f_1, \dots, f_k))_n^\infty$ of measures in $\mathcal{S}_k(G_n)$ such that $\lim_{n\to\infty} d_{LP}(\mu_n, \mu) = 0$. Let μ_n^{+v} be the measure in $\mathcal{S}_k(G_n^+)$ obtained by assigning the same values to the vertices of the subgraph G_n as $f_1, \dots, f_k$ do and giving the values $(v_1, \dots, v_k)$ to the additional vertex. Formally,

$$\mu_n^{+v} := \mathcal{D}_{G_n^+}\left(f_1^{+v_1}, \dots, f_k^{+v_k}\right) \quad \text{where} \quad f_i^{+v_i}(u) = \begin{cases} f_i(u) & \text{if } u \in V(G_n) \\ v_i & \text{if } \{u\} = V(G_n^+) \setminus V(G_n). \end{cases}$$

Then

$$d_{LP}(\mu_n^{+v}, \mu_n \oplus v) \leq \frac{1}{|V(G_n)| + 1}$$

because $(\mu \oplus v)(A \oplus \{v\}) = \mu(A)$ for any measurable A, and $\left((f_1^{+v_1}, \dots, f_k^{+v_k})G_n^+\right)(u) = ((f_1, \dots, f_k)G_n)(u) + (v_1, \dots, v_k)$. By the triangle inequality,

$$d_{LP}(\mu_n^{+v}, \mu \oplus v) \leq d_{LP}(\mu_n^{+v}, \mu_n \oplus v) + d_{LP}(\mu_n \oplus v, \mu \oplus v) \leq \frac{1}{|V(G)| + 1} + d_{LP}(\mu_n, \mu),$$

where the right-hand side tends to 0 by Lemma 4.1. Hence for every measure in $X_k \oplus V_k$, there is a sequence in $\mathcal{S}_k(G_n^+)$ converging to it.

Vice versa, suppose that $(\mu_n^+ = \mathcal{D}_{G_n^+}(f_1^+, \dots, f_k^+))_n^\infty$ is a sequence with $\mu_n^+ \in \mathcal{S}_k(G_n^+)$ which is convergent in d_{LP} to a measure μ^+. We want to show that $\mu^+ \in X_k \oplus V_k$. Let $v_n \in [-1, 1]^k$ be the values assigned by $f_1^+, \dots, f_k^+$ to the added vertex in G_n^+. The set $[-1, 1]^k$ is compact, so there is a subsequence $\left(\mu_{n_i}^+\right)_i^\infty$ on which v_n converges to some $v \in [-1, 1]^k$.

Employing the triangle inequality again gives

$$d_{LP}(\mu_{n_i} \oplus v, \mu^+) \leq d_{LP}(\mu_{n_i} \oplus v, \mu_{n_i} \oplus v_{n_i}) + d_{LP}(\mu_{n_i} \oplus v_{n_i}, \mu_{n_i}^+) + d_{LP}(\mu_{n_i}^+, \mu^+) \tag{4.1}$$

for every $i \in \mathbb{N}$. The second summand on the right-hand side is bounded above by $\frac{1}{|V(G_{n_i})|+1}$ like before and the third summand goes to 0 by the assumption of convergence of $(\mu_n^+)_n^\infty$. To bound the first summand, we observe two facts: that $d_{LP}(\eta \oplus w, \nu) = d_{LP}(\eta, \nu \oplus (-w))$ and that $d_{LP}(\eta, \eta \oplus w) \leq \|w\|$ for any measures η and ν and vector w. The first is true because for any $\varepsilon > 0$,

$$\begin{aligned}
&(\eta \oplus w)(U) \leq \nu\left(U^\varepsilon\right) + \varepsilon \text{ for all measurable } U \\
\Leftrightarrow\ &(\eta \oplus w)(U \oplus \{w\}) \leq \nu\left((U \oplus \{w\})^\varepsilon\right) + \varepsilon \text{ for all measurable } U \\
\Leftrightarrow\ &\eta(U) \leq \nu\left(U^\varepsilon \oplus \{w\}\right) + \varepsilon \text{ for all measurable } U \\
\Leftrightarrow\ &\eta(U) \leq \left(\nu \oplus (-w)\right)\left(U^\varepsilon\right) + \varepsilon \text{ for all measurable } U
\end{aligned}$$

and similarly the other way round. The second is true because $U \oplus \{w\} \subseteq U^{\|w\|+\varepsilon}$ for any $\varepsilon > 0$, implying that

$$\begin{aligned}
\eta(U) = (\eta \oplus w)(U \oplus \{w\}) &\leq (\eta \oplus w)\left(U^{\|w\|+\varepsilon}\right) \\
\text{and } (\eta \oplus w)(U) = \eta(U \oplus \{-w\}) &\leq \eta\left(U^{\|w\|+\varepsilon}\right).
\end{aligned}$$

Returning to inequality 4.1, we now get

$$d_{LP}\left(\mu_{n_i}, \mu^+ \oplus (-v)\right) = d_{LP}\left(\mu_{n_i} \oplus v, \mu^+\right) \leq \|v_{n_i} - v\| + \frac{1}{|V(G_{n_i})|+1} + d_{LP}(\mu_{n_i}^+, \mu^+) \to 0 \text{ as } i \to \infty.$$

Since $\mu_{n_i} \in \mathcal{S}_k(G_{n_i})$, we conclude that $\mu := \mu^+ \oplus (-v) \in X_k$, and so $\mu^+ = \mu \oplus v \in X_k \oplus V_k$ as claimed. □

Having established the form of $\lim_n \mathcal{S}_k\left(G_n^+\right)$, we could now show along more technical, but in essence similar lines as in the proof of Proposition 3.5 that no action limit of $(G_n^+)_n^\infty$, where $\lim |V(G_n)| = \infty$, can possibly be self-adjoint.

Notation. *Given a measure space $(\Omega, \mathcal{A}, \mu)$, we write χ_F for the characteristic function of a measurable set $F \in \mathcal{A}$, and $\mathbb{1}_E$ for the indicator functions of events E from any other measure space. We also write $B_r^{L^p}(g)$ for the closed L^p-ball of radius r around the function $g \in L^p$, and we denote by 0 the function which sends (almost every) element of Ω to 0. Finally, $B_1^{L^p}$ is a shorthand for the unit ball around 0.*

We are now preparing to prove that when $G_n \to A$, there indeed is a P-operator A^+ whose k-profiles are like those prescribed by Proposition 4.2 for the prospective limit of $(G_n^+)_n^\infty$. We would like to keep A in A^+ in some form because it encodes the possibly complicated structure of $(G_n)_n^\infty$ which is of course also present in $(G_n^+)_n^\infty$. But at the same time we must introduce the shifts arising from the presence of the special vertex in G_n^+ which has an outsized influence with respect to its increasingly negligible measure. Like we did in Section 3, we will use an ultrafilter-based functional

$$\begin{aligned}
\phi_{\mathcal{U}} \colon L^\infty &\to \mathbb{R} \\
\lim_{n\to\infty} \sum_{i=1}^n \alpha_{i,n} \chi_{E_{i,n}} &\mapsto \lim_{n\to\infty} \sum_{i=1}^n \alpha_{i,n} \mathbb{1}_{E_{i,n} \in \mathcal{U}}
\end{aligned}$$

for this – decreasing nested sequences of smaller and smaller sets in $\mathcal{U}$ will play the role that the special vertex played in the proof of Proposition 4.2. However, we must be careful that the ultrafilter $\mathcal{U}$ does not interfere with the key properties of the original limit A. The following theorem, which is also of independent interest, tells us that $\mathcal{U}$ can be chosen to satisfy our needs.

Theorem 4.3. *Let $A\colon L^\infty(\Omega,\mu) \to L^1(\Omega,\mu)$ be a linear operator, where (Ω,μ) is an atomless separable finite measure space. Then there is a μ-filter $\mathcal{F}$ such that any ultrafilter $\mathcal{U}$ containing $\mathcal{F}$ has the following property:*
for all $f \in B_1^{L^\infty}$, $\varepsilon > 0$, $a \in [-1,1]$, there is $f_{a,\varepsilon} \in B_1^{L^\infty}$ such that

I. $\|f - f_{a,\varepsilon}\|_1 < \varepsilon$

II. $\|Af - Af_{a,\varepsilon}\|_1 < \varepsilon$

III. $\phi_{\mathcal{U}}(f_{a,\varepsilon}) \in (a-\varepsilon, a+\varepsilon)$.

Proof. We will generate $\mathcal{F}$ in countably many steps; using the separability of $L^1(\Omega)$, we will generate countably many functions $f \in L^\infty(\Omega)$ that represent the action of A to an arbitrary precision (see Figure 1), and build a filter $\mathcal{F}$ that meshes well with the properties of these functions.

Let us first consider the image $A\left(B_1^{L^\infty}\right) \subseteq L^1(\Omega,\mu)$ of the unit ball of $L^\infty(\Omega)$ under the action of the operator A. Since $L^1(\Omega,\mu)$ is separable, we can pick, for any positive integer n, a countable collection $\mathcal{C}_n$ of functions $g \in L^1(\Omega,\mu)$ such that the union of their $\frac{1}{n}$-balls covers $A\left(B_1^{L^\infty}\right)$, that is,

$$A\left(B_1^{L^\infty}\right) \subseteq \bigcup_{g\in\mathcal{C}_n} B_{1/n}^{L^1}(g).$$

Next, we look at the preimages $A^{-1}\left(B_{1/n}^{L^1}(g)\right) \subseteq L^\infty(\Omega,\mu)$ of these covering $\frac{1}{n}$-balls. Given one such fixed preimage $A^{-1}\left(B_{1/n}^{L^1}(g)\right)$, we view it as a subset of $L^1(\Omega,\mu)$ and pick a countable collection $\mathcal{D}_{n,g}$ of functions in $A^{-1}\left(B_{1/n}^{L^1}(g)\right) \cap B_1^{L^\infty}$ whose $\frac{1}{n}$-balls cover $A^{-1}\left(B_{1/n}^{L^1}(g)\right) \cap B_1^{L^\infty}$. Altogether, this produces the countable family $\mathcal{B}_n = \bigcup_{g\in\mathcal{C}_n} \mathcal{D}_{n,g}$ of functions from the unit ball of $L^\infty(\Omega,\mu)$ which by construction satisfies that for every $f \in B_1^{L^\infty}$, there is some $h_f \in \mathcal{B}_n$ such that both $\|h_f - f\|_1 \leq \frac{1}{n}$ and $\|Ah_f - Af\|_1 \leq \frac{2}{n}$.

We will now construct a filter $\mathcal{F}$ based on the countable collection $\mathcal{B} = \bigcup_{n=1}^\infty \mathcal{B}_n$. Let us start by fixing an ordering $f_1, f_2, \ldots$ of the elements of $\mathcal{B} \subset B_1^{L^\infty}$.

For each $m \geq 1$, we will find a sequence $E_{m,1} \supset E_{m,2} \supset E_{m,3} \ldots$ of measurable sets such that

1. $\mu(E_{m,k}) > 0$ for all k, but $\mu(E_{m,k}) \to 0$ as $k \to \infty$,
2. $\left\|A\chi_{E_{m,k}}\right\|_1 \to 0$ as $k \to \infty$, and
3. the range of $f_m|_{E_{m,k}}$ shrinks, i.e. $\operatorname{ess\,sup}(f_m|_{E_{m,k}}) - \operatorname{ess\,inf}(f_m|_{E_{m,k}}) \to 0$ as $k \to \infty$.

The conditions 1.–3. will help us to prove the corresponding requirements I.–III. from the statement of the theorem. In particular, condition 3. helps us track the value of $\phi_{\mathcal{U}}(f_m)$, which we will be able to change by adding functions which are non-zero only on $E_{m,k}$. Moreover, the sets $E_{m,k}$ will also satisfy

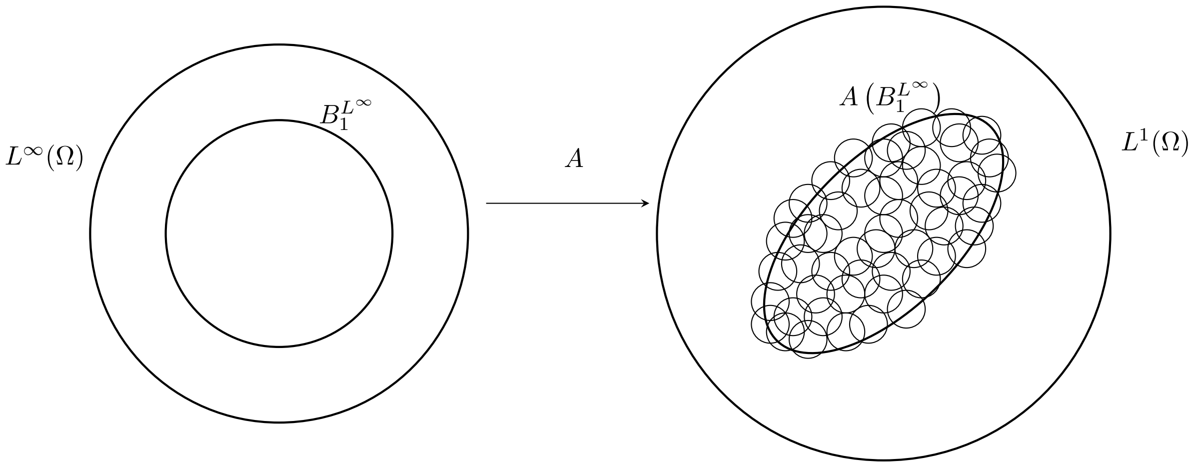

(a) We first cover $A\left(B_1^{L^\infty}\right)$ with $\frac{1}{n}$-balls, i.e., $A\left(B_1^{L^\infty}\right)$ is represented by the countable collection $\mathcal{C}_n$ up to an error of $1/n$ in the 1-norm.

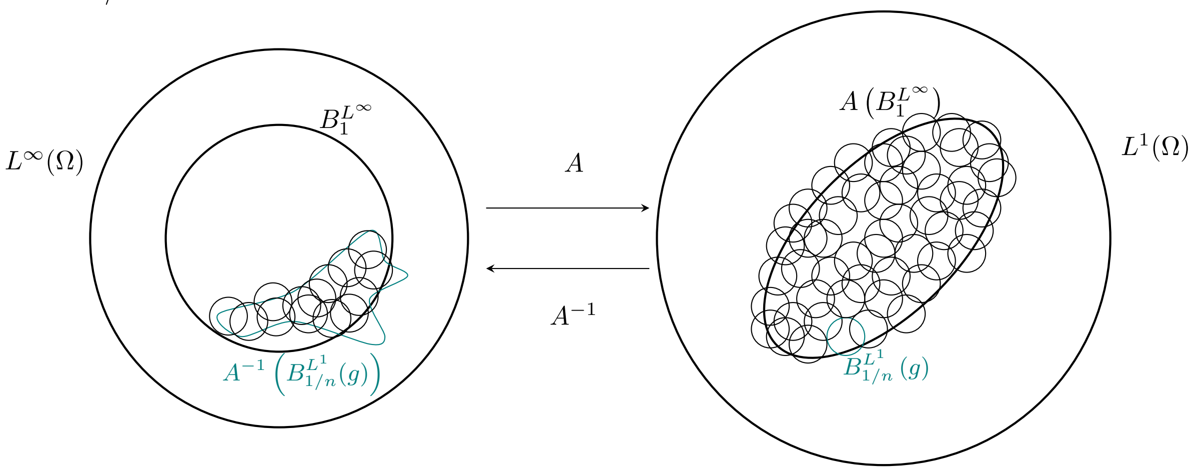

(b) Each function $f \in B_1^{L^\infty}$ belongs to a preimage $A^{-1}\left(B_{1/n}^1(g)\right)$, which can itself be represented up to a $\frac{1}{n}$-error in the 1-norm by a countable collection $\mathcal{D}_{n,g}$.

Figure 1: Constructing the countable family $\mathcal{B}_n = \bigcup_{g\in\mathcal{C}_n} \mathcal{D}_{n,g}$ of functions in the unit ball $B_1^{L^\infty}$ whose $\frac{1}{n}$-balls cover it

$$\begin{array}{ccccccc} E_{1,1} & \supset & E_{1,2} & \supset & E_{1,3} & \supset & \dots \\ \cup & & \cup & & \cup & & \\ E_{2,1} & \supset & E_{2,2} & \supset & E_{2,3} & \supset & \dots \\ \cup & & \cup & & \cup & & \\ E_{3,1} & \supset & E_{3,2} & \supset & E_{3,3} & \supset & \dots \\ \vdots & & \vdots & & \vdots & & \end{array}$$

Figure 2: The measurable sets $E_{m,k}$ satisfy that for any fixed $m \geq 1$, the real numbers $\mu(E_{m,k})$, $\left\|A\chi_{E_{m,k}}\right\|_1$, and $\sup(f_m|_{E_{m,k}}) - \inf(f_m|_{E_{m,k}})$ all tend to zero as k tends to infinity.

the inclusions $E_{1,k} \supset E_{2,k} \supset E_{3,k} \dots$ for every positive integer k (see Figure 2). Together with that $\mu(E_{m,k}) > 0$ for all $m, k \in \mathbb{N}$, this implies that

$$\mu(E_{m_1,k_1} \cap \dots \cap E_{m_n,k_n}) > 0 \tag{4.2}$$

for any $n \in \mathbb{N}$ and $m_1, \dots, m_n, k_1, \dots, k_n \in \mathbb{N}$ because $E_{m_1,k_1} \cap \dots \cap E_{m_n,k_n} \supseteq E_{\max_{i\in[n]}\{m_i\},\max_{i\in[n]}\{k_i\}}$. Finally, inequality (4.2) tells us that defining

$$\mathcal{F} := \{(E_{m_1,k_1} \cap \dots \cap E_{m_n,k_n}) \cup B : n \in \mathbb{N}, B \subseteq \Omega \text{ is measurable}\}$$

gives a μ-filter.

Let us now construct the sequences $(E_{m,k})_{k=1}^{\infty}$ by induction as follows. For $m = 1$, put

$$E_{1,1} = \begin{cases} f_1^{-1}([0,1]), & \text{if } \mu\left(f_1^{-1}([0,1])\right) > 0 \\ f_1^{-1}([-1,0)), & \text{if } \mu\left(f_1^{-1}([0,1])\right) = 0. \end{cases}$$

Note that $\mu\left(f_1^{-1}([-1,0)) \cup f_1^{-1}([0,1])\right) = \mu(\Omega)$ because $f_1 \in B_1^{L^\infty}$, so we necessarily have $\mu(E_{1,1}) > 0$. Having obtained $E_{1,k}$ as a subset of $f_1^{-1}\left([x_{1k}, x_{1k} + 2^{1-k}]\right)$ where x_{1k} is some number in $\left[-1, 1-2^{1-k}\right]$, we obtain $E_{1,k+1}$ by first restricting to R, where R is given by

$$R = \begin{cases} E_{1,k} \cap f_1^{-1}\left([x_{1k}, x_{1k} + 2^{-k}]\right), & \text{if } \mu\left(E_{1,k} \cap f_1^{-1}([x_{1k}, x_{1k} + 2^{-k}])\right) > 0 \\ E_{1,k} \cap f_1^{-1}\left([x_{1k} + 2^{-k}, x_{1k} + 2^{1-k}]\right), & \text{if } \mu\left(E_{1,k} \cap f_1^{-1}([x_{1k}, x_{1k} + 2^{-k}])\right) = 0, \end{cases}$$

and then to R', where R' is any measurable subset of R satisfying that $0 < \mu(R') \leq \frac{\mu(R)}{2}$ (see Figure 3). We then set $x_{1,k+1}$ to be equal to x_{1k} if R was chosen to be a subset of $f_1^{-1}\left([x_{1k}, x_{1k} + 2^{-k}]\right)$ and to $x_{1k} + 2^{-k}$ otherwise. Restricting to R and then R' ensures, respectively, that

$$\sup\left(f_1|_{E_{1,k+1}}\right) - \inf\left(f_1|_{E_{1,k+1}}\right) \leq 2^{-k} \quad \text{and} \quad \mu\left(E_{1,k+1}\right) \leq \frac{\mu(\Omega)}{2^k}.$$

Since (Ω, μ) is atomless, there is an uncountable collection $\mathcal{R}$ of measurable subsets of R' such that for any $S_1 \neq S_2 \in \mathcal{R}$, we have $\mu(S_1 \triangle S_2) > 0$ and either $S_1 \subset S_2$ or $S_2 \subset S_1$. We now consider the uncountable collection of functions $A\chi_S, S \in \mathcal{R}$ in $L^1(\Omega, \mu)$, and conclude that by separability of

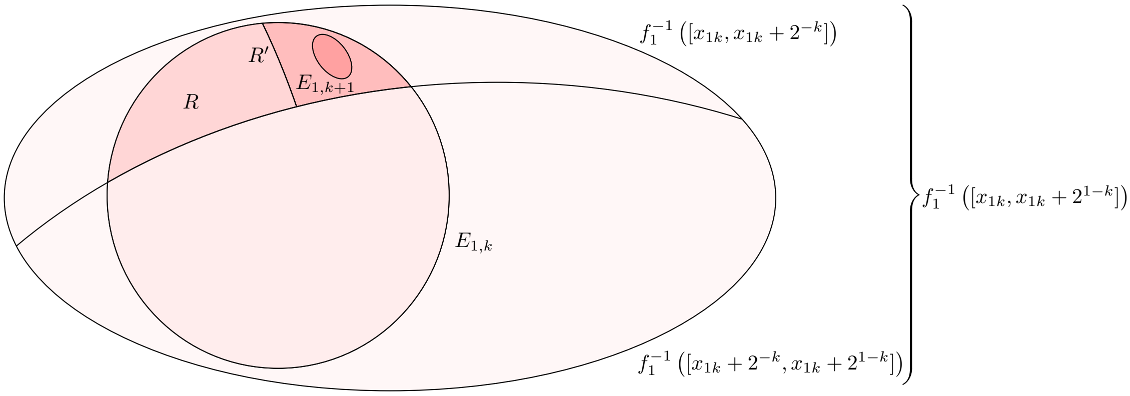

Figure 3: Constructing $E_{1,k+1}$ as a subset of $E_{1,k}$

$L^1(\Omega,\mu)$, there must be some $S \subset T \in \mathcal{R}$ such that $\|A\chi_T - A\chi_S\|_1 < \frac{1}{k}$. Linearity of A then implies that $\left\|A\chi_{T\setminus S}\right\|_1 < \frac{1}{k}$, and so we set $E_{1,k+1} := T \setminus S$.

For $m > 1$, suppose we already have $E_{m-1,1} \supset E_{m-1,2} \supset \dots$ as desired. Now if there is some $j \geq 1$ such that $\mu\left(E_{m-1,j} \cap f_m^{-1}([0,1])\right) = 0$ then we set $E_{m,1} := E_{m-1,1} \cap f_m^{-1}([-1,0))$, and we necessarily have that $\mu(E_{m,1} \cap E_{m-1,j}) > 0$ for all $j \in \mathbb{N}$. Otherwise $\mu\left(E_{m-1,j} \cap f_m^{-1}([0,1])\right) > 0$ for all $j \in \mathbb{N}$, and we set $E_{m,1} := E_{m-1,1} \cap f_m^{-1}([0,1])$. Suppose now that we have obtained $E_{m,k}$ as a subset of $f_m^{-1}\left([x_{mk}, x_{mk}+2^{1-k}]\right)$ for some $x_{mk} \in \left[-1, 1-2^{1-k}\right]$ and that $\mu(E_{m,k} \cap E_{m-1,j}) > 0$ for all $j \geq k$. We obtain $E_{m,k+1}$ by first restricting to

$$R = \begin{cases} E_{m-1,k+1} \cap E_{m,k} \cap f_m^{-1}\left([x_{mk}, x_{mk}+2^{-k}]\right), \\ \qquad\qquad \text{if } \mu\left(E_{m-1,j} \cap E_{m,k} \cap f_m^{-1}([x_{mk}, x_{mk}+2^{-k}])\right) > 0 \text{ for all } j > k \\ E_{m-1,k+1} \cap E_{m,k} \cap f_m^{-1}\left([x_{mk}+2^{-k}, x_{mk}+2^{1-k}]\right), \\ \qquad\qquad \text{if } \mu\left(E_{m-1,j} \cap E_{m,k} \cap f_m^{-1}([x_{mk}, x_{mk}+2^{-k}])\right) = 0 \text{ for some } j > k. \end{cases}$$

As before, we set $x_{m,k+1}$ to be x_{mk} if R is chosen to be a subset of $f_m^{-1}\left([x_{mk}, x_{mk}+2^{-k}]\right)$ and to be $x_{mk}+2^{-k}$ otherwise. Let us now partition R into $\bigcup_{j=k+1}^{\infty} R_j$, where

$$R_j := (E_{m-1,j} \setminus E_{m-1,j+1}) \cap R.$$

Importantly, we observe that we must have $\mu(R_j) > 0$ for infinitely many j: if there was some $J > k$ such that $\mu(R_j) = 0$ for all $j \geq J$ then we would conclude that

$$0 = \sum_{j=J}^{\infty} \mu(R_j) = \mu\left(\bigcup_{j=J}^{\infty} R_j\right) = \mu\left(E_{m-1,J} \cap R\right),$$

which contradicts the definition of R. Next, by atomlessness of Ω, for every $j > k$, there is a family $\{R_j(t) : t \in (0,1]\}$ of measurable subsets of R_j, which satisfies that

$$\mu(R_j(t)) = t\mu(R_j)$$

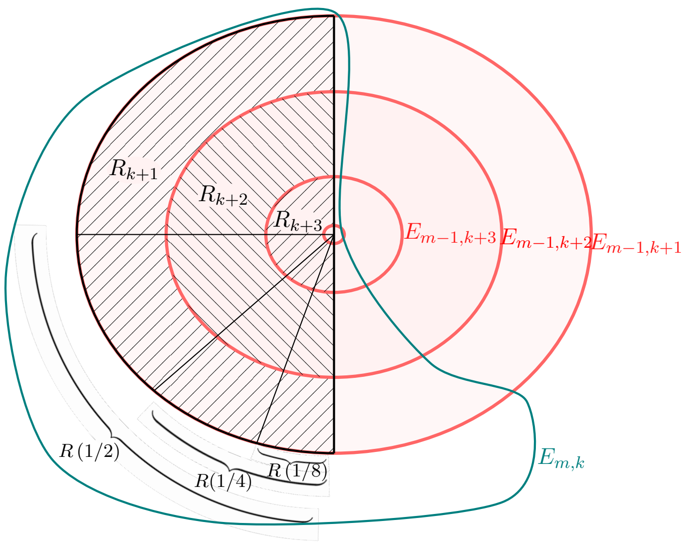

Figure 4: Constructing $E_{m,k+1}$ as a subset of $E_{m-1,k+1}$

and $R_j(t) \subseteq R_j(s)$ whenever $t \leq s$. Let us define $R(t) \subseteq R$ to be

$$R(t) := \bigcup_{j=k+1}^{\infty} R_j(t)$$

for any $t \in (0,1]$, and note that $R(t) \subsetneq R(s)$ whenever $t < s$ (see Figure 4). Crucially,

$$\mu\left((R(s) \setminus R(t)) \cap E_{m-1,j}\right) > 0$$

for all $t < s$ and $j > k$. This is because $(R(s) \setminus R(t)) \cap E_{m-1,j}$ contains

$$\left(R_{j'}(s) \setminus R_{j'}(t)\right) \cap E_{m-1,j'} = R_{j'}(s) \setminus R_{j'}(t)$$

where $j' \geq j$ is an integer such that $\mu(R_{j'}) > 0$.

We are now ready to consider the uncountable family $A\chi_{R(t)}$, $t \in (0,1]$ of functions in $L^1(\Omega)$. As before, by separability of $L^1(\Omega)$, there must be some $t < s$ such that $\left\|A\chi_{R(s)} - A\chi_{R(t)}\right\|_1 < \frac{1}{k}$. Linearity of A then implies that $\left\|A\chi_{R(s)\setminus R(t)}\right\|_1 < \frac{1}{k}$, and we set $E_{m,k+1} := R(s) \setminus R(t)$.

We have now constructed the sets $E_{m,k}$ as promised and, as already said, we set the μ-filter $\mathcal{F}$ to be

$$\mathcal{F} := \left\{(E_{m_1,k_1} \cap \cdots \cap E_{m_n,k_n}) \cup B : n \in \mathbb{N}, B \subseteq \Omega \text{ is measurable}\right\}.$$

Let now $\mathcal{U}$ be any ultrafilter containing $\mathcal{F}$. It remains to show that for a given $f \in B_1^{L^\infty}$, $\varepsilon > 0$ and $a \in [-1,1]$, there is $f_{a,\varepsilon} \in B_1^{L^\infty}$ satisfying the conditions 1.–3. We first note that for every $m \geq 1$, the sequence $(x_{mk})_{k=1}^{\infty}$ of real numbers converges. Let x_m be the limit, and let us observe that by the careful construction of $\mathcal{F}$, we can conclude that $\phi_{\mathcal{U}}(f_m) = x_m$ for every m. This follows from the general fact that if some measurable set E is in $\mathcal{U}$ then for any function $g \in L^\infty$, the number $\phi_\mathcal{U}(g)$ must be in $[\operatorname{ess\,inf}(g|_E), \operatorname{ess\,sup}(g|_E)]$. Next, let n be a positive integer such that $\frac{2}{n} < \varepsilon$ and let us recall what we noted at the beginning of the proof, that is that by the construction of $\mathcal{B}_n$, there is some $h_f \in \mathcal{B}_n$ such that $\|h_f - f\|_1 \leq \frac{1}{n}$ and $\|Ah_f - Af\|_1 \leq \frac{2}{n}$. Suppose that h_f appears as f_m in the ordering of $\mathcal{B} = \bigcup_n \mathcal{B}_n$. Then we would like to set $f_{a,\varepsilon}$ to be $f_m + (a - x_m)\chi_{E_{m,k}}$ for some large enough k because linearity of $\phi_\mathcal{U}$ and $E_{m,k}$ being in $\mathcal{F}$ gives us that

$$\phi_\mathcal{U}\left(f_m + (a - x_m)\chi_{E_{m,k}}\right) = \phi_\mathcal{U}(f_m) + (a - x_m)\phi_\mathcal{U}(\chi_{E_{m,k}}) = x_m + (a - x_m) = a. \tag{4.3}$$

However, $f_m + (a - x_m)\chi_{E_{m,k}}$ may not be in the unit ball $B_1^{L^\infty}$, so we will need to do some technical tinkering. On $\Omega \setminus E_{m,k}$, the functions f_m and $f_{a,\varepsilon}$ are equal, so we do not get out of $B_1^{L^\infty}$ there because $f_m = h_f$ is in $B_1^{L^\infty}$ to start with. On $E_{m,k}$, f_m takes values in $\left[x_{mk}, x_{mk} + 2^{1-k}\right]$, and so $f_m + (a - x_m)\chi_{E_{m,k}}$ takes values in $V = \left[x_{mk} + a - x_m, x_{mk} + 2^{1-k} + a - x_m\right]$. But x_m itself is in $\left[x_{mk}, x_{mk} + 2^{1-k}\right]$, and so $V \subset \left[a - 2^{1-k}, a + 2^{1-k}\right]$. Setting $f_{a,\varepsilon}$ to be

$$f_{a,\varepsilon} := \begin{cases} f_m + (a - x_m)\chi_{E_{m,k}}, & \text{if } a \in \left[-1 + 2^{1-k}, 1 - 2^{1-k}\right] \\ f_m + (1 - 2^{1-k} - x_m)\chi_{E_{m,k}}, & \text{if } a \in \left(1 - 2^{1-k}, 1\right] \\ f_m + (-1 + 2^{1-k} - x_m)\chi_{E_{m,k}}, & \text{if } a \in \left[-1, -1 + 2^{1-k}\right) \end{cases}$$

therefore ensures that $f_{a,\varepsilon}$ is in the unit ball $B_1^{L^\infty}$, and we fix k to be a positive integer large enough so that

$$\max\left\{\frac{\mu(\Omega)}{2^{k-2}}, \frac{2}{k-1}, 2^{1-k}\right\} < \varepsilon - \frac{2}{n}.$$

Finally, we check that $f_{a,\varepsilon}$ satisfies the conditions I.–III. Firstly,

$$\|f - f_{a,\varepsilon}\|_1 \leq \|f - f_m\|_1 + \|f_m - f_{a,\varepsilon}\|_1 < \frac{1}{n} + 2\left\|\chi_{E_{m,k}}\right\|_1 = \frac{1}{n} + 2\mu(E_{m,k}) \leq \frac{1}{n} + 2\mu(E_{1,k}) \leq \frac{1}{n} + \frac{\mu(\Omega)}{2^{k-2}} < \varepsilon.$$

Secondly,

$$\|Af - Af_{a,\varepsilon}\|_1 \leq \|Af - Af_m\|_1 + \|Af_m - Af_{a,\varepsilon}\|_1 \leq \frac{2}{n} + 2\left\|A\chi_{E_{m,k}}\right\|_1 \leq \frac{2}{n} + \frac{2}{k-1} < \varepsilon.$$

Thirdly, if $a \in \left[-1 + 2^{1-k}, 1 - 2^{1-k}\right]$ then $\phi_\mathcal{U}(f_{a,\varepsilon}) = a$ as per equation (4.3). If $a \in (1 - 2^{1-k}, 1]$ then

$$\phi_\mathcal{U}(f_{a,\varepsilon}) = 1 - 2^{1-k} = a + \left(1 - 2^{1-k} - a\right),$$

where $1 - 2^{1-k} - a \in [-2^{1-k}, 0)$, and similarly for $a \in [-1, -1 + 2^{1-k})$, so

$$|\phi_\mathcal{U}(f_{a,\varepsilon}) - a| \leq 2^{1-k} < \varepsilon$$

as desired. □

Remark 4.4 (on assumptions in Theorem 4.3)**.** *In the theorem above, finiteness of $\mu(\Omega)$ is implicitly used in viewing subsets of $L^\infty(\Omega)$ as subsets of $L^1(\Omega)$. Note, however, that we do* not *require $\|A\|_{\infty\to 1}$ to be bounded! One may hope to eliminate the condition of separability of $L^1(\Omega)$, and especially if we would instead insist that $\|A\|_{\infty\to 1}$ be finite, it does not seem unreasonable to believe that this might indeed be possible.*

We now restate and prove our main theorem, of which Theorem 3.7 is a special case.

Theorem 1.4. *Let $(G_n)_n^\infty$ be a sequence of finite graphs with $|V(G_n)| \to \infty$, whose adjacency operators action converge to a P-operator $A\colon L^\infty(\Omega,\nu) \to L^1(\Omega,\nu)$, where (Ω,ν) is separable. Then there is a ν-filter $\mathcal{F}$ on Ω such that for any ν-ultrafilter $\mathcal{U}$ extending $\mathcal{F}$, both*

$$A^+\colon L^\infty(\Omega,\nu) \to L^1(\Omega,\nu) \quad \text{given by } \left(A^+g\right)(\omega) = (Ag)(\omega) + \phi_{\mathcal{U}}(g) \qquad \text{and} \qquad A^-\colon L^\infty(\Omega,\nu) \to L^1(\Omega,\nu) \quad \text{given by } \left(A^-g\right)(\omega) = (Ag)(\omega) - \phi_{\mathcal{U}}(g)$$

are action limits of $(G_n^+)_n^\infty$, where $\phi_{\mathcal{U}}\colon L^\infty(\Omega,\nu) \to \mathbb{R}$ is the functional sending $\lim_{n\to\infty}\sum_{i=1}^n \alpha_{n,i}\chi_{E_{n,i}}$ to $\lim_{n\to\infty}\sum_{i=1}^n \alpha_{n,i}\mathbb{1}_{E_{n,i}\in\mathcal{U}}$.

As already mentioned, our limits A^+ and A^- cannot be self-adjoint, showing that Proposition 1.3 (a), which holds under the assumption of uniform boundedness of the (p,q)-norms of a Cauchy sequence, cannot be extended to include $(p,q) = (\infty,1)$. However, the situation is interestingly more subtle with being positivity-preserving. If the original limit A is positivity-preserving then so is A^+ because the value $\phi_{\mathcal{U}}(f)$ is always in the essential range of f. But that is to say that the functional $-\phi_{\mathcal{U}}$ is positivity-*reversing*, and so when we return to the case of the star sequence $(S_n)_n^\infty$, where $A \equiv 0$ is the trivial P-operator on an atomless space (Ω,ν), then the limit A^+ is positivity-preserving while A^- is not. This shows that Proposition 1.3 (b) cannot be extended to $(p,q) = (\infty,1)$, but raises the question whether a Cauchy sequence of graphops always has a positivity-preserving limit (see Section 6). In any case, we conclude that the property of being positivity-preserving is not invariant under weak equivalence.

Furthermore, if A is c-regular for some $c \in \mathbb{R}$ then A^+ is $c+1$-regular while A^- is $c-1$-regular. In other words, Theorem 1.4 shows that also c-regularity is not invariant under weak equivalence, unless we first restrict our consideration from the set of all P-operators to the set of P-operators with the (p,q)-norm bounded above by b, for some $p,q \in [1,\infty) \times [1,\infty]$ and fixed $b \in \mathbb{R}_{\geq 0}$. This in particular means that Proposition 1.3 (c) cannot be extended to include $(p,q) = (1,\infty)$.

Proof. If $\limsup_{n\to\infty}|V(G_n)| < \infty$ then by Lemma 4.1, $(G_n)_n^\infty$ is eventually constant, and so is $(G_n^+)_n^\infty$.

Suppose now that $\limsup_{n\to\infty}|V(G_n)| = \infty$. To prove that A^+ and A^- are action limits of $(G_n^+)_n^\infty$, we need to show that the closures of their k-profiles equal the limits of the closures of the k-profiles of G_n^+. In other words, that

$$\overline{\mathcal{S}_k(A^\pm)} = \lim_{n\to\infty}\overline{\mathcal{S}_k(G_n^+)}$$

for every k. By Proposition 4.2,

$$\lim_{n\to\infty}\overline{\mathcal{S}_k(G_n^+)} = X_k \oplus V_k$$

where $X_k = \lim_n \overline{\mathcal{S}_k(G_n)}$. We will now show first that $\overline{\mathcal{S}_k(A^\pm)} \subseteq X_k \oplus V_k$ and then that $X_k \oplus V_k \subseteq \overline{\mathcal{S}_k(A^\pm)}$.

For a given $k \in \mathbb{N}$, let $\mu = \mathcal{D}_{A^\pm}(f_1,\ldots,f_k)$ be any measure in $\mathcal{S}_k(A^\pm)$, given by some k functions in $B_1^{L^\infty}$. Then

$$\begin{aligned}\mu = \mathcal{D}_{A^\pm}(f_1,\ldots,f_k) &= \mathcal{D}\left(f_1,\ldots,f_k,A^\pm f_1,\ldots,A^\pm f_k\right)\\ &= \mathcal{D}\left(f_1,\ldots,f_k,Af_1 \pm \phi_{\mathcal{U}}(f_1)\cdot\mathbb{1},\ldots,Af_k \pm \phi_{\mathcal{U}}(f_k)\cdot\mathbb{1}\right)\\ &= \mathcal{D}\left(f_1,\ldots,f_k,Af_1,\ldots,Af_k\right) \oplus (0,\ldots,0,\pm\phi_{\mathcal{U}}(f_1),\ldots,\pm\phi_{\mathcal{U}}(f_k)) \in \mathcal{S}_k(A) \oplus V_k.\end{aligned}$$

This proves that both $\mathcal{S}_k(A^+)$ and $\mathcal{S}_k(A^-)$ are subsets of $\mathcal{S}_k(A) \oplus V_k$, and hence their closures are subsets of $\overline{\mathcal{S}_k(A) \oplus V_k} = X_k \oplus V_k$. In other words, $\overline{\mathcal{S}_k(A^\pm)} \subseteq X_k \oplus V_k$ as required.

On the other hand, let μ be a measure in X_k. Then there is a sequence $\left(\mu_n = \mathcal{D}_A(f^1, \ldots, f^k)\right)_n^\infty$ of measures that converge to μ in d_{LP}. The functions $f^1, \ldots, f^k \in B_1^{L^\infty}$ naturally depend on n, but we again drop this further index in the interest of readability. We also stress that the upper indices are just indices and do not stand for taking powers. Let also $v = \{0\}^k \times (v_1, \ldots, v_k)$ be an element of V_k. We will now show that we can approximate the measure $\mu \oplus v$ with elements from $\mathcal{S}_k(A^+)$ as well as from $\mathcal{S}_k(A^-)$.

Let $\mathcal{F}$ be a ν-filter given by Theorem 4.3 and let $\mathcal{U}$ be any ultrafilter extending it. For every $i \in [k]$, let $f^i_{v_i, \frac{1}{n}}$ and $f^i_{-v_i, \frac{1}{n}}$ be the functions also given by Theorem 4.3 upon applying it to f^i. Then we set

$$\mu_n^{+v} := \mathcal{D}_{A^+}\left(f^1_{v_1, \frac{1}{n}}, \ldots, f^k_{v_k, \frac{1}{n}}\right) \in \mathcal{S}_k(A^+)$$

and similarly

$$\mu_n^{-v} := \mathcal{D}_{A^-}\left(f^1_{-v_1, \frac{1}{n}}, \ldots, f^k_{-v_k, \frac{1}{n}}\right) \in \mathcal{S}_k(A^-).$$

By the triangle inequality,

$$\begin{aligned} d_{LP}\left(\mu_n^{\pm v}, \mu \oplus v\right) &\le d_{LP}\left(\mu_n^{\pm v}, \mu_n \oplus v\right) + d_{LP}\left(\mu_n \oplus v, \mu \oplus v\right) \\ &= d_{LP}\left(\mathcal{D}_{A^\pm}\left(f^1_{\pm v_1, \frac{1}{n}}, \ldots, f^k_{\pm v_k, \frac{1}{n}}\right), \mathcal{D}_A\left(f^1, \ldots, f^k\right) \oplus v\right) + d_{LP}\left(\mu_n, \mu\right) \end{aligned} \tag{4.4}$$

where the second summand goes to 0 by the assumption that $\mu_n \to \mu$. Shifting our attention to the first summand, we see that for A^+ it is equal to

$$\begin{aligned} &d_{LP}\left(\mathcal{D}\left(f^1_{v_1, \frac{1}{n}}, \ldots, f^k_{v_k, \frac{1}{n}}, A^+ f^1_{v_1, \frac{1}{n}}, \ldots, A^+ f^k_{v_k, \frac{1}{n}}\right), \mathcal{D}\left(f^1, \ldots, f^k, Af^1, \ldots, Af^k\right) \oplus v\right) \\ &= d_{LP}\left(\mathcal{D}\left(f^1_{v_1, \frac{1}{n}}, \ldots, f^k_{v_k, \frac{1}{n}}, Af^1_{v_1, \frac{1}{n}} + \phi_{\mathcal{U}}(f^1_{v_1, \frac{1}{n}})\mathbb{1}, \ldots, Af^k_{v_k, \frac{1}{n}} + \phi_{\mathcal{U}}(f^k_{v_k, \frac{1}{n}})\mathbb{1}\right), \mathcal{D}\left(f^1, \ldots, f^k, Af^1 + v_1\mathbb{1}, \ldots, Af^k + v_k\mathbb{1}\right)\right) \end{aligned}$$

and similarly when we start with $\mathcal{D}_{A^-}\left(f^1_{-v_1, \frac{1}{n}}, \ldots, f^k_{-v_k, \frac{1}{n}}\right)$ instead of $\mathcal{D}_{A^+}\left(f^1_{v_1, \frac{1}{n}}, \ldots, f^k_{v_k, \frac{1}{n}}\right)$, we arrive to the first summand being

$$\begin{aligned} &d_{LP}\left(\mathcal{D}\left(f^1_{-v_1, \frac{1}{n}}, \ldots, f^k_{-v_k, \frac{1}{n}}, A^- f^1_{-v_1, \frac{1}{n}}, \ldots, A^- f^k_{-v_k, \frac{1}{n}}\right), \mathcal{D}\left(f^1, \ldots, f^k, Af^1, \ldots, Af^k\right) \oplus v\right) \\ &\quad = d_{LP}\Big(\mathcal{D}\left(f^1_{-v_1, \frac{1}{n}}, \ldots, f^k_{-v_k, \frac{1}{n}}, Af^1_{-v_1, \frac{1}{n}} - \phi_{\mathcal{U}}(f^1_{-v_1, \frac{1}{n}})\mathbb{1}, \ldots, Af^k_{-v_k, \frac{1}{n}} - \phi_{\mathcal{U}}(f^k_{-v_k, \frac{1}{n}})\mathbb{1}\right), \\ &\qquad \mathcal{D}\left(f^1, \ldots, f^k, Af^1 + v_1\mathbb{1}, \ldots, Af^k + v_k\mathbb{1}\right)\Big). \end{aligned}$$

To bound this first summand, both for A^+ and A^-, we use that by Theorem 4.3, the functions $f^i_{v_i, \frac{1}{n}} - f^i$ and $f^i_{-v_i, \frac{1}{n}} - f^i$ as well as $Af^i_{v_i, \frac{1}{n}} - Af^i$ and $Af^i_{-v_i, \frac{1}{n}} - Af^i$ all have their 1-norms bounded above by $1/n$. Markov's inequality then gives that the $2k$ sets

$$N_i^+ = \left\{\omega \in \Omega : \left|\left(f^i_{v_i, \frac{1}{n}} - f^i\right)(\omega)\right| \ge \delta\right\}$$
$$M_i^+ = \left\{\omega \in \Omega : \left|\left(Af^i_{v_i, \frac{1}{n}} - Af^i\right)(\omega)\right| \ge \delta\right\}$$

as well as the $2k$ sets

$$N_i^- = \left\{\omega \in \Omega : \left|\left(f^i_{-v_i,\frac{1}{n}} - f^i\right)(\omega)\right| \geq \delta\right\}$$
$$M_i^- = \left\{\omega \in \Omega : \left|\left(Af^i_{-v_i,\frac{1}{n}} - Af^i\right)(\omega)\right| \geq \delta\right\}$$

all satisfy $\nu\left(N_i^{\pm}\right) \leq \frac{1}{\delta n}$ and $\nu\left(M_i^{\pm}\right) \leq \frac{1}{\delta n}$, where $\delta = \delta(n)$ is a positive number to be chosen later. Put together also with that

$$\phi_{\mathcal{U}}\left(f^i_{v_i,\frac{1}{n}}\right) \in \left(v_i - \frac{1}{n}, v_i + \frac{1}{n}\right),$$

this says that the set

$$M^+ := \left\{\omega \in \Omega : \exists i \in [k] \text{ such that } \left|\left(f^i_{v_i,\frac{1}{n}} - f^i\right)(\omega)\right| \geq \delta \text{ or } \left|\left(A(f^i_{v_i,\frac{1}{n}} - f^i) + (\phi_{\mathcal{U}}(f^i_{v_i,\frac{1}{n}}) - v_i)\mathbb{1}\right)(\omega)\right| \geq \delta + \frac{1}{n}\right\}$$

satisfies $M^+ \subseteq N_1^+ \cup \cdots \cup N_k^+ \cup M_1^+ \cup \cdots \cup M_k^+$. Analogously, with

$$M^- := \left\{\omega \in \Omega : \exists i \in [k] \text{ such that } \left|\left(f^i_{-v_i,\frac{1}{n}} - f^i\right)(\omega)\right| \geq \delta \text{ or } \left|\left(A(f^i_{-v_i,\frac{1}{n}} - f^i) + (-\phi_{\mathcal{U}}(f^i_{-v_i,\frac{1}{n}}) - v_i)\mathbb{1}\right)(\omega)\right| \geq \delta + \frac{1}{n}\right\}$$

we have that $M^- \subseteq N_1^- \cup \cdots \cup N_k^- \cup M_1^- \cup \cdots \cup M_k^-$, and so by the union bound,

$$\nu\left(M^{\pm}\right) \leq \sum_{i=1}^{k} \nu\left(N_i^{\pm}\right) + \sum_{i=1}^{k} \nu\left(M_i^{\pm}\right) \leq \frac{2k}{\delta n}.$$

The inequality above expresses that on most of Ω, the functions

$$\left(f^1_{\pm v_1,\frac{1}{n}}, \ldots, f^k_{\pm v_k,\frac{1}{n}}, Af^1_{\pm v_1,\frac{1}{n}} \pm \phi_{\mathcal{U}}(f^1_{\pm v_1,\frac{1}{n}})\mathbb{1}, \ldots, Af^k_{\pm v_k,\frac{1}{n}} \pm \phi_{\mathcal{U}}(f^k_{\pm v_k,\frac{1}{n}})\mathbb{1}\right) : \Omega \to \mathbb{R}^{2k}$$
$$\text{and} \quad \left(f^1, \ldots, f^k, Af^1 + v_1\mathbb{1}, \ldots, Af^k + v_k\mathbb{1}\right) : \Omega \to \mathbb{R}^{2k}$$

output real vectors that are close to one another. In particular, for every $\omega \in \Omega \setminus M^{\pm}$, we have that all the $2k$ coordinates of $\left(f^1_{\pm v_1,\frac{1}{n}}, \ldots, Af^k_{\pm v_k,\frac{1}{n}} \pm \phi_{\mathcal{U}}(f^k_{\pm v_k,\frac{1}{n}})\mathbb{1}\right)(\omega) - (f^1, \ldots, Af^k + v_k\mathbb{1})(\omega)$ are in $(-\delta - 1/n, \delta + 1/n)$. This implies that

$$\left(f^1_{\pm v_1,\frac{1}{n}}, \ldots, Af^k_{\pm v_k,\frac{1}{n}} \pm \phi_{\mathcal{U}}(f^k_{\pm v_k,\frac{1}{n}})\mathbb{1}\right)^{-1}(U) \subseteq \left(f^1, \ldots, Af^k + v_k\mathbb{1}\right)^{-1}\left(U^{\|(\delta+\frac{1}{n},\ldots,\delta+\frac{1}{n})\|}\right) \cup M^{\pm}$$
$$\text{and} \quad (f^1, \ldots, Af^k + v_k\mathbb{1})^{-1}(U) \subseteq \left(f^1_{\pm v_1,\frac{1}{n}}, \ldots, Af^k_{\pm v_k,\frac{1}{n}} \pm \phi_{\mathcal{U}}(f^k_{v_k,\frac{1}{n}})\mathbb{1}\right)^{-1}\left(U^{\|(\delta+\frac{1}{n},\ldots,\delta+\frac{1}{n})\|}\right) \cup M^{\pm}$$

for any measurable subset U of $\mathbb{R}^{2k}$. Taking the measure ν of both sides of these inclusions gives

$$\begin{aligned}
&\mathcal{D}\left(f^1_{\pm v_1,\frac{1}{n}}, \ldots, Af^k_{\pm v_k,\frac{1}{n}} \pm \phi_{\mathcal{U}}(f^k)\mathbb{1}\right)(U) \\
&= \nu\left(\left(f^1_{\pm}, \ldots, Af^k_{\pm v_k,\frac{1}{n}} \pm \phi_{\mathcal{U}}(f^k_{v_k,\frac{1}{n}})\mathbb{1}\right)^{-1}(U)\right) \leq \nu\left((f^1, \ldots, Af^k + v_k\mathbb{1})^{-1}\left(U^{\|(\delta+\frac{1}{n},\ldots,\delta+\frac{1}{n})\|}\right) \cup M^{\pm}\right) \\
&\leq \mathcal{D}\left(f^1, \ldots, Af^k + v_k\mathbb{1}\right)\left(U^{(\delta+1/n)\sqrt{2k}}\right) + \frac{2k}{\delta n}
\end{aligned}$$

and

$$\mathcal{D}\left(f^1,\dots,Af^k+v_k\mathbb{1}\right)(U)=\nu\left((f^1,\dots,Af^k+v_k\mathbb{1})^{-1}(U)\right)$$
$$\le\nu\left(\left(f^1_{\pm v_1,\frac{1}{n}},\dots,Af^k_{\pm v_k,\frac{1}{n}}\pm\phi_{\mathcal{U}}(f^k_{v_k,\frac{1}{n}})\mathbb{1}\right)^{-1}\left(U^{\|(\delta+\frac{1}{n},\dots,\delta+\frac{1}{n})\|}\right)\cup M^{\pm}\right)$$
$$\le\mathcal{D}\left(f^1_{\pm v_1,\frac{1}{n}},\dots,Af^k_{\pm v_k,\frac{1}{n}}\pm\phi_{\mathcal{U}}(f^k_{v_k,\frac{1}{n}})\mathbb{1}\right)\left(U^{(\delta+1/n)\sqrt{2k}}\right)+\frac{2k}{\delta n},$$

which means that

$$d_{LP}\left(\mathcal{D}\left(f^1_{\pm v_1,\frac{1}{n}},\dots,Af^k_{\pm v_k,\frac{1}{n}}\pm\phi_{\mathcal{U}}(f^k_{\pm v_k,\frac{1}{n}})\mathbb{1}\right),\mathcal{D}\left(f^1,\dots,Af^k+v_k\mathbb{1}\right)\right)\le\max\left\{\left(\delta+\frac{1}{n}\right)\sqrt{2k},\frac{2k}{\delta n}\right\}.$$

Finally, we set $\delta=\delta(n)$ to be $1/\sqrt{n}$, to get that

$$d_{LP}\left(\mu_n^{\pm v},\mu_n\oplus v\right)=d_{LP}\left(\mathcal{D}\left(f^1_{\pm v_1,\frac{1}{n}},\dots,Af^k_{\pm v_k,\frac{1}{n}}\pm\phi_{\mathcal{U}}(f^k_{\pm v_k,\frac{1}{n}})\mathbb{1}\right),\mathcal{D}(f^1,\dots,Af^k+v_k\mathbb{1})\right)$$
$$\le\max\left\{\frac{(1+\sqrt{n})\sqrt{2k}}{n},\frac{2k}{\sqrt{n}}\right\}\to 0.$$

This finishes the argument started at inequality (4.4) that $d_{LP}\left(\mu_n^{\pm v},\mu\oplus v\right)\to 0$, and so $X_k\oplus V_k\subseteq\overline{\mathcal{S}_k(A^{\pm})}$. ☐

Remark 4.5 (on assumptions in Theorem 1.4)**.** *The mild requirement that (Ω,ν) be separable could be replaced by that A be positivity-preserving or at least by the existence of a measurable set E with $\nu(E)>0$ such that for every $f\ge 0$ which is only non-zero on E, $Af\ge 0$. In that case, the construction of an appropriate ν-filter $\mathcal{F}$ becomes much easier than in the proof of Theorem 4.3.*

5 Subdivisions of the complete graphs K_n

Let us denote by $K_n^\bullet$ the graph on $\binom{n}{2}+n$ vertices obtained from the n-vertex clique by dividing each of its edges into two.

Theorem 5.1. *Let (Ω,μ) be an atomless probability space. There exists a μ^2-ultrafilter-valued random variable $\mathcal{U}$ on (Ω,μ) such that the P-operator $A\colon L^\infty(\Omega\times\Omega,\mu\times\mu)\to L^1(\Omega\times\Omega,\mu\times\mu)$ given by*

$$(Af)(x,y)=\phi_{\mathcal{U}(x)}(f)+\phi_{\mathcal{U}(y)}(f)$$

is an action limit of the subdivisions of K_n.

Lemma 5.2. *For every $k\in\mathbb{N}$, $\lim_{n\to\infty}\overline{\mathcal{S}_k(K_n^\bullet)}$ is equal to*

$$\left\{\kappa\in\mathcal{P}\left([-1,1]^k\times[-2,2]^k\right):\exists\mu_1,\dots,\mu_k\in\mathcal{P}([-1,1])\text{ such that }\kappa_{k+i}=law(X_i+Y_i)\text{ for }X_i,Y_i\overset{iid}{\sim}\mu_i\right\},$$

where κ_j is the j-th marginal of κ.

Similarly to Section 3, we will prove this lemma using a general result about approximability of divisible laws. For that purpose, we introduce the following notation.

$$\mathcal{R} := \left\{\kappa \in \mathcal{P}\left([-1,1]\times[-2,2]\right) \ : \ \exists \mu \in \mathcal{P}\left([-1,1]\right) \text{ such that } \kappa_y = \text{law}(X+Y) \text{ for } X, Y \overset{iid}{\sim} \mu\right\}$$

$$\mathcal{R}_n := \Bigg\{\kappa = \frac{1}{n^2}\sum_{i=1}^{n^2} \delta_{(a_i,b_i)} \ : \ (a_i,b_i) \in [-1,1]\times[-2,2]$$

$$\text{and } \exists x_1,\ldots,x_n \in [-1,1] \text{ such that } \kappa_y = \text{law}(X+Y) \text{ for } X,Y \overset{iid}{\sim} \frac{1}{n}\sum_{i=1}^{n}\delta_{x_i}\Bigg\}$$

and more generally for every positive integer k,

$$\mathcal{R}^k = \left\{\kappa \in \mathcal{P}\left([-1,1]^k \times [-2,2]^k\right) : \exists \mu_1,\ldots,\mu_k \in \mathcal{P}([-1,1]) \text{ such that } \kappa_{k+i} = \text{law}\left(X_i+Y_i\right) \text{ for } X_i, Y_i \overset{iid}{\sim} \mu_i\right\},$$

$$\mathcal{R}_n^k := \left\{\kappa = \frac{1}{n^2}\sum_{i,j}\delta_{\left(a^{(1)}_{\{i,j\}},\ldots,a^{(k)}_{\{i,j\}},b^{(1)}_i+b^{(1)}_j,\ldots,b^{(k)}_i+b^{(k)}_j\right)} \ : \ a^{(m)}_{\{i,j\}}, b_i \in [-1,1] \text{ for all } m\in[k], i,j\in[n]\right\}$$

Proof. Let θ be an element of $\mathcal{S}_k\left(K_n^\bullet\right)$, given by some k functions $f_1,\ldots,f_k \in L^\infty_{[-1,1]}\left([n+\binom{n}{2}]\right)$. Then $\theta = \frac{n}{n+\binom{n}{2}}\theta_v + \frac{\binom{n}{2}}{n+\binom{n}{2}}\theta_e$, where θ_v is the probability measure on $\mathbb{R}^{2k}$ given by sampling one of the vertices $v_1,\ldots,v_n$ of degree $n-1$ and getting the corresponding value $(f_1(v),\ldots,f_k(v),\sum_{e\sim v} f_1(e),\ldots,\sum_{e\sim v} f_k(e))$, while θ_e is the probability measure given by sampling one of the vertices $e_1,\ldots,e_{\binom{n}{2}}$ of degree 2 and getting the value $(f_1(e),\ldots,f_k(e),f_1(u)+f_1(v),\ldots,f_k(u)+f_k(v))$, where u and v are the two neighbours of e.

Similarly, any measure $\kappa \in \mathcal{R}_n^k$ can be written as $\kappa = \frac{n-1}{n}\kappa_d + \frac{1}{n}\kappa_s$, where κ_d is given by sampling two different indices among the available n, that is, $\kappa_d = \frac{1}{\binom{n}{2}}\sum_{i<j}\delta_{\left(a^{(1)}_{\{i,j\}},\ldots,a^{(k)}_{\{i,j\}},b^{(1)}_i+b^{(1)}_j,\ldots,b^{(k)}_i+b^{(k)}_j\right)}$, and κ_s comes from sampling two identical indices, i.e., $\kappa_s = \frac{1}{n}\sum_i \delta_{\left(a^{(1)}_{\{i\}},\ldots,a^{(k)}_{\{i\}},2b^{(1)}_i,\ldots,2b^{(k)}_i\right)}$. But we note that θ_e is of the form κ_d for any measure $\kappa \in \mathcal{R}_n^k$ whose κ_s is equal to $\frac{1}{n}\sum_i \delta_{\left(a^{(1)}_i,\ldots,a^{(k)}_i,2f_1(v_i),\ldots,2f_k(v_i)\right)}$, where $a^{(m)}_i$, $i\in[n], m\in[k]$ are any numbers in $[-1,1]$. Let, for concreteness, κ_θ be the measure in $\mathcal{R}_n^k$ whose κ_d is θ_e and $\kappa_s = \frac{1}{n}\sum_i \delta_{(0,\ldots,0,2f_1(v_i),\ldots,2f_k(v_i))}$. Then by the triangle inequality, $d_{LP}\left(\theta,\kappa_\theta\right) \le d_{LP}\left(\theta,\theta_\kappa\right) + d_{LP}\left(\kappa_d,\kappa_\theta\right)$. But for any $t\in[0,1]$ and probability measures α,β on $\mathbb{R}^{2k}$,

$$d_{LP}(\alpha, t\alpha + (1-t)\beta) \le 1-t, \tag{5.1}$$

and so

$$d_{LP}\left(\theta,\theta_\kappa\right) + d_{LP}\left(\kappa_d,\kappa_\theta\right) \le \frac{n}{n+\binom{n}{2}} + \frac{1}{n} < \frac{3}{n}.$$

In particular,

$$\sup_{\theta\in\mathcal{S}_k}\inf_{\kappa\in\mathcal{R}^k} d_{LP}(\theta,\kappa)\leq \sup_{\theta\in\mathcal{S}_k}\inf_{\kappa\in\mathcal{R}^k_n} d_{LP}(\theta,\kappa)\leq \frac{3}{n}. \tag{5.2}$$

On the other hand, let $\kappa=\frac{1}{n^2}\sum_{i,j}\delta_{\left(a^{(1)}_{\{i,j\}},\dots,a^{(k)}_{\{i,j\}},b^{(1)}_i+b^{(1)}_j,\dots,b^{(k)}_i+b^{(k)}_j\right)}$ be an element of $\mathcal{R}^k_n$. Then $\kappa_d=\frac{1}{\binom{n}{2}}\sum_{i<j}\delta_{\left(a^{(1)}_{\{i,j\}},\dots,a^{(k)}_{\{i,j\}},b^{(1)}_i+b^{(1)}_j,\dots,b^{(k)}_i+b^{(k)}_j\right)}$ is exactly of the form θ_e for $\theta_\kappa\in\mathcal{S}_k\left(K^\bullet_n\right)$ given by $f_k\left(v_i\right)=b^{(k)}_i$ and $f_k\left(v_iv_j\right)=a^{(k)}_{\{i,j\}}$, and so again by the triangle inequality,

$$d_{LP}\left(\kappa,\theta_\kappa\right)\leq d_{LP}\left(\kappa,\kappa_d\right)+d_{LP}\left(\theta_e,\theta_\kappa\right).$$

Reusing inequality (5.1), we obtain that

$$d_{LP}\left(\kappa,\kappa_d\right)+d_{LP}\left(\theta_e,\theta_\kappa\right)\leq \frac{n}{n+\binom{n}{2}}+\frac{1}{n}<\frac{3}{n},$$

and hence

$$\sup_{\kappa\in\mathcal{R}^k_n}\inf_{\theta\in\mathcal{S}_k} d_{LP}(\theta,\kappa)\leq \frac{3}{n}. \tag{5.3}$$

Inequalities (5.2) and (5.3) now combine to give $d_H\left(\mathcal{S}_k\left(K^\bullet_n\right),\mathcal{R}^k_n\right)\leq\frac{3}{n}$, and so we obtain that

$$d_H\left(\mathcal{S}_k\left(K^\bullet_n\right),\mathcal{R}^k\right)\leq d_H\left(\mathcal{S}_k\left(K^\bullet_n\right),\mathcal{R}^k_n\right)+d_H\left(\mathcal{R}^k_n,\mathcal{R}^k\right)\leq\frac{3}{n}+d_H\left(\mathcal{R}^k_n,\mathcal{R}^k\right).$$

But since $\mathcal{R}^k_n\subset\mathcal{R}^k$, Lemma 5.3 tells us that $d_H\left(\mathcal{R}^k_n,\mathcal{R}^k\right)$ tends to 0 as $n\to\infty$, and hence so does $d_H\left(\mathcal{S}_k\left(K^\bullet_n\right),\mathcal{R}^k\right)$. □

Lemma 5.3 (Uniform approximability of $\mathcal{R}$)**.** *For all $\varepsilon>0$, there is $N=N(\varepsilon)$ such that for every measure $\kappa\in\mathcal{R}$, there is a sequence $\left(\kappa_n=\frac{1}{n^2}\sum_{i=1}^{n^2}\delta_{(a_i,b_i)}\right)_n^\infty$ of discrete probability measures in $\mathcal{R}_n$ such that*

$$\forall n\geq N(\varepsilon)\quad d_{LP}(\kappa,\kappa_n)<\varepsilon.$$

In particular, N does not depend on κ, so since $\mathcal{R}_n\subset\mathcal{R}$, the lemma implies that $d_H\left(\mathcal{R},\mathcal{R}_n\right)\leq\varepsilon$ for all $n\geq N(\varepsilon)$.

Proof. We prove the lemma for the case $k=1$ only because it greatly eases the notation while already containing all the relevant ideas, which is, in essence, due to the fact that we are dealing with nearly transitive graphs.

Let $\varepsilon>0$ be fixed. We will show that the lemma holds with $N(\varepsilon)=\frac{5^4}{\varepsilon^4}$. Throughout, we will work with the Wasserstein distance instead of the Lévy-Prokhorov distance because the former one lends itself to easily obtaining upper bounds on it by constructing concrete transport plans. In particular, given $\kappa\in\mathcal{R}$ and $n\geq N(\varepsilon)$, we will produce a measure $\kappa_n\in\mathcal{R}_n$ whose construction is geared towards quickly getting ε^2 as an upper bound on $d_{\mathcal{W}_1}\left(\kappa,\kappa_n\right)$. Then since $d_{LP}(\mu,\nu)^2\leq d_{\mathcal{W}_1}(\mu,\nu)$ for all measures $\mu,\nu\in\mathcal{P}\left([-1,1]\times[-2,2]\right)$, the statement of the lemma follows.

Let κ and μ be fixed, where $\kappa \in \mathcal{R}$ and the y-marginal κ_y of κ is the law of $X+Y$ for $X, Y \overset{iid}{\sim} \mu$. Let us suppose that the cumulative density function $F_\kappa \colon (-\infty, 1] \times (-\infty, 2] \to [0,1]$ is continuous and further also that the values

$$\begin{aligned} h_i :=& \sup\left\{a \in [-1,1] : \mu\left((-\infty, a]\right) < \frac{i}{n}\right\} \\ =& \inf\left\{b \in [-1,1] : \mu\left((-\infty, b]\right) \geq \frac{i}{n}\right\} \\ =& \sup\left\{a \in [-1,1] : \kappa\left([-1,1] \times (-\infty, a]\right) < \frac{i}{n}\right\}, \; i \in [n] \end{aligned}$$

are such that $h_i + h_j = h_k + h_\ell$ only if $\{i,j\} = \{k,\ell\}$. We will denote by $v_1 < v_2 < \cdots < v_{\binom{n}{2}+n} \leq 2$ the values in $\{h_i + h_j : i, j \in [n]\}$, so that the $\kappa_n = \frac{1}{n^2}\sum_{i,j} \delta_{(w_{i,j}, h_i + h_j)}$ we are trying to obtain is of the form $\kappa_n = \frac{1}{n^2}\sum_{k \in O} \delta_{(w_k, v_k)} + \frac{2}{n^2}\sum_{k \in T} \delta_{(w_k, v_k)}$, where O, T is the partition

$$\begin{aligned} O &:= \{k : v_k = h_i + h_i \text{ for some } i \in [n]\} \\ T &:= \{k : v_k = h_i + h_j \text{ for some } i \neq j \in [n]\} \end{aligned}$$

of $\left[\binom{n}{2} + n\right]$. Let now μ_n be the discrete measure $\frac{1}{n}\sum_{i=1}^n \delta_{h_i}$ and let $X_n, Y_n \overset{iid}{\sim} \mu_n$.

For all $v \in \mathbb{R}$, we have that

$$\begin{aligned} \mathbb{P}\left(X_n + Y_n \leq v\right) &= \sum_{\substack{i,j \in [n] \\ h_i + h_j \leq v}} \mathbb{P}\left(X_n \in (h_{i-1}, h_i]\right) \cdot \mathbb{P}\left(Y_n \in (h_{j-1}, h_j]\right) \\ &= \sum_{\substack{i,j \in [n] \\ h_i + h_j \leq v}} \mathbb{P}\left(X \in (h_{i-1}, h_i]\right) \cdot \mathbb{P}\left(Y \in (h_{j-1}, h_j]\right) \leq \mathbb{P}\left(X + Y \leq v\right) \end{aligned}$$

and

$$\mathbb{P}\left(X + Y \leq v\right) - \frac{2n-1}{n^2} \leq \mathbb{P}\left(X_n + Y_n \leq v\right).$$

In particular, for all $u < v \in \mathbb{R}$,

$$\begin{aligned} \mathbb{P}\left(X_n + Y_n \in (u,v]\right) &= \mathbb{P}\left(X_n + Y_n \leq v\right) - \mathbb{P}\left(X_n + Y_n \leq u\right) \\ &\geq \mathbb{P}\left(X + Y \leq v\right) - \frac{2n-1}{n^2} - \mathbb{P}\left(X + Y \leq u\right) \\ &= \mathbb{P}\left(X + Y \in (u,v]\right) - \frac{2}{n} + \frac{1}{n^2}. \end{aligned}$$

Let us observe that for all $i \in \left[\binom{n}{2} + n\right]$, $\mathbb{P}\left(X_n + Y_n \in (v_{i-1}, v_i]\right)$ is equal to $\frac{1}{n^2}$ or $\frac{2}{n^2}$, and the first value occurs for n indices i, while the second for $\binom{n}{2}$ indices i.

Now we want to set up roughly $\sqrt{n}$ intervals J_ℓ and then roughly $\sqrt{n}$ intervals $I_{\ell,k}$ for every ℓ, dividing $[-1,1] \times [-2,2]$ into roughly n rectangles $I_{\ell,k} \times J_\ell$ such that $\kappa\left(I_{\ell,k} \times J_\ell\right) = \frac{1}{n}$ for most ℓ, k.

In particular, we let

$$\begin{aligned} J_1 &:= [-2, u_1], \\ J_\ell &:= (u_{\ell-1}, u_\ell], \text{ for } \ell = 2, \ldots, \left\lfloor \sqrt{n} \right\rfloor, \end{aligned}$$

where $u_1 < \cdots < u_{\lfloor\sqrt{n}\rfloor} = 2$ are chosen so that

$$\kappa\left([-1,1]\times J_\ell\right) = \mathbb{P}\left(X+Y\in J_\ell\right) = \frac{1}{\sqrt{n}} \quad \text{for all } \ell < \left\lfloor\sqrt{n}\right\rfloor.$$

Now that we spliced $[-2,2]$ up into $\lfloor\sqrt{n}\rfloor$ intervals J_ℓ, we want to further splice the rectangles $[-1,1]\times J_\ell$ into $I_{\ell,k}\times J_\ell$ so that also these rectangles have roughly the same measure.

For each $\ell\in[\lfloor\sqrt{n}\rfloor]$, let

$$\begin{aligned} w_{\ell,k} &:= \sup\left\{a\in[-1,1] : \kappa\left([-1,a]\times J_\ell\right) < \frac{k}{n}\right\}, \quad k\in\left[\left\lfloor\sqrt{n}\right\rfloor - 1\right], \\ w_{\ell,\lfloor\sqrt{n}\rfloor} &:= 1, \end{aligned}$$

and let us set

$$\begin{aligned} I_{\ell,1} &:= [-1, w_{\ell,1}] \\ I_{\ell,k} &:= (w_{\ell,k-1}, w_{\ell,k}], \text{ for } k = 2,\dots,\left\lfloor\sqrt{n}\right\rfloor. \end{aligned}$$

Then $\kappa\left(I_{\ell,k}\times J_\ell\right) = \frac{1}{n}$ for every $\ell\in[\lfloor\sqrt{n}\rfloor]$ and $k\in[\lfloor\sqrt{n}\rfloor-1]$.

Having obtained the partition $I_{\ell,k}\times J_\ell, \ell,k\in[\lfloor\sqrt{n}\rfloor]$ of $[-1,1]\times[-2,2]$ into $\lfloor\sqrt{n}\rfloor^2$ rectangles most of which have κ-measure $\frac{1}{n}$, we define κ_n to be

$$\kappa_n = \frac{1}{n^2}\sum_{i,j}\delta_{\left(w_{\ell(i,j),k(i,j)},h_i+h_j\right)},$$

where $\ell(i,j)$ is the least integer ℓ such that $h_i+h_j\le u_\ell$ and $k(i,j)$ is the greatest integer $k\in[\lfloor\sqrt{n}\rfloor]$ such that

$$\frac{k-1}{n} < \mathbb{P}\left(X_n+Y_n\in\left(u_{\ell(i,j)-1}, h_i+h_j\right]\right),$$

where $u_0 := -2$.

The measure κ_n then satisfies that

$$\kappa_n\left(I_{\ell,k}\times J_\ell\right)\in\left\{\frac{1}{n}-\frac{1}{n^2},\frac{1}{n},\frac{1}{n}+\frac{1}{n^2}\right\} \text{ for all } \ell\in\left[\left\lfloor\sqrt{n}\right\rfloor\right], k\in\left[\left\lfloor\sqrt{n}\right\rfloor-2\right].$$

This is because for all ℓ,k,

$$\begin{aligned} \kappa_n\left([-1,w_{\ell,k}]\times J_\ell\right) &= \kappa_n\left([-1,w_{\ell,k}]\times\left(u_{\ell-1},v_{m(\ell,k)}\right]\right) \\ &= \kappa_n\left([-1,1]\times\left(u_{\ell-1},v_{m(\ell,k)}\right]\right) \\ &= \mathbb{P}\left(X_n+Y_n\in\left(u_{\ell-1},v_{m(\ell,k)}\right]\right), \end{aligned}$$

where $v_{m(\ell,k)}$ is the largest h_i+h_j with $\ell(i,j)=\ell$, $k(i,j)\le k$. But we also have that for all $\ell\in[\lfloor\sqrt{n}\rfloor]$,

$$\mathbb{P}\left(X_n+Y_n\in J_\ell\right)\ge\mathbb{P}\left(X+Y\in J_\ell\right)-\frac{2}{n}+\frac{1}{n^2}\ge\frac{1}{\sqrt{n}}-\frac{2}{n}+\frac{1}{n^2}>\frac{\lfloor\sqrt{n}\rfloor-2}{n},$$

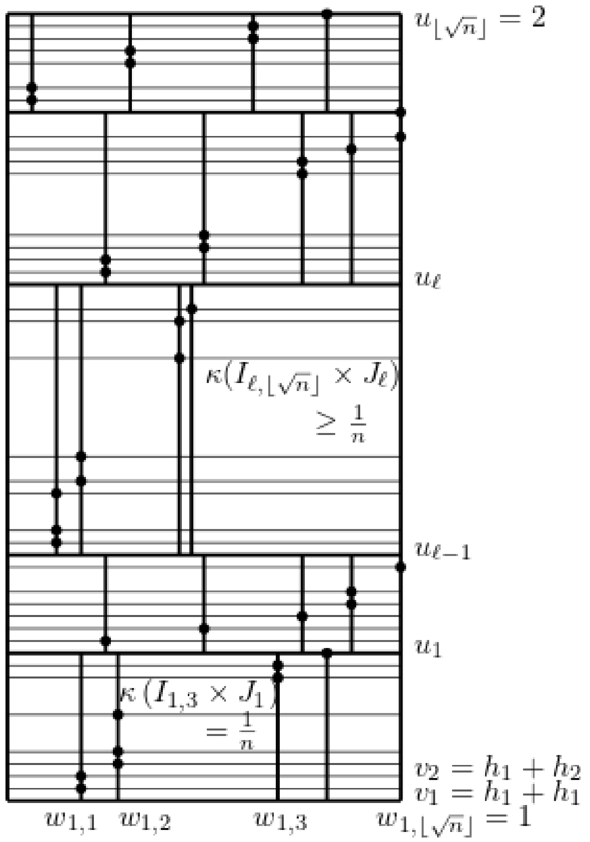

(a) To construct κ_n, we first partition $[-1,1] \times [-2,2]$ to $\lfloor\sqrt{n}\rfloor$ rectangles $I_{\ell,k} \times J_\ell$ satisfying $\kappa(I_{\ell,k} \times J_\ell) \geq \frac{1}{n}$, then we place Dirac masses at some intersections of the lines $y = v_j$ and $x = w_{\ell,k}$.

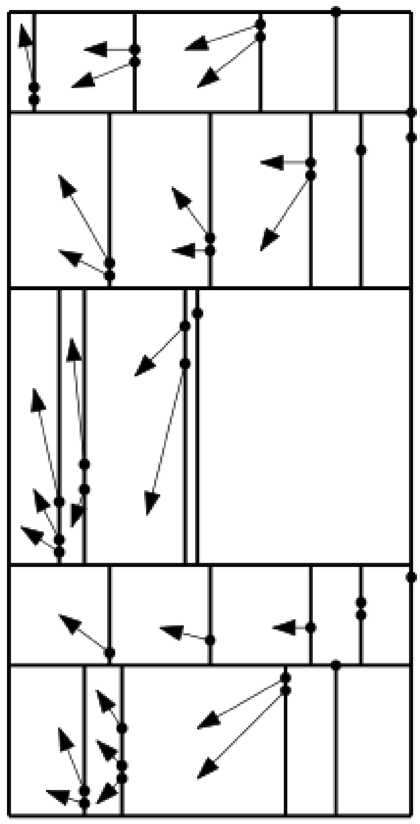

(b) After constructing κ_n, we consider transport plans from κ_n to κ which transport $\frac{1}{n}$ or $\frac{1}{n} - \frac{1}{n^2}$ of mass within the rectangles $I_{\ell,k} \times J_\ell$, $k \in [\lfloor\sqrt{n}\rfloor - 2]$.

Figure 5: Constructing $\kappa_n \in \mathcal{R}_n$ which is close to a given $\kappa \in \mathcal{R}$

meaning that for all $k \in [\lfloor\sqrt{n}\rfloor - 1]$, there exist $i, j \in [n]$ such that $\ell(i,j) = \ell$ and $k(i,j) = k$. In particular, for any $k \in [\lfloor\sqrt{n}\rfloor - 2]$, $v_{m(\ell,k)+1}$ is still in J_ℓ and hence satisfies

$$\mathbb{P}\left(X_n + Y_n \in \left(u_{\ell-1}, v_{m(\ell,k)}\right]\right) \leq \frac{k}{n} < \mathbb{P}\left(X_n + Y_n \in \left(u_{\ell-1}, v_{m(\ell,k)+1}\right]\right).$$

Since the increment $\left(v_{m(\ell,k)}, v_{m(\ell,k)+1}\right]$ has probability $\frac{1}{n^2}$ or $\frac{2}{n^2}$ under the law of $X_n + Y_n$ and because $\frac{k}{n} = \frac{kn}{n^2}$, we conclude that $\mathbb{P}\left(X_n + Y_n \in \left(u_{\ell-1}, v_{m(\ell,k)}\right]\right) = \frac{k}{n}$ or $\frac{k}{n} - \frac{1}{n^2}$. Subsequently,

$$\begin{aligned}\kappa_n\left(I_{\ell,k} \times J_\ell\right) &= \kappa_n\left([-1, w_{\ell,k}] \times J_\ell\right) - \kappa_n\left([-1, w_{\ell,k-1}] \times J_\ell\right) \\ &= \mathbb{P}\left(X_n + Y_n \in \left(u_{\ell-1}, v_{m(\ell,k)}\right]\right) - \mathbb{P}\left(X_n + Y_n \in \left(u_{\ell-1}, v_{m(\ell,k-1)}\right]\right)\end{aligned}$$

is equal to one of $\frac{1}{n} - \frac{1}{n^2}, \frac{1}{n}, \frac{1}{n} + \frac{1}{n^2}$.

In the last step, we construct a transport plan from κ_n to κ which is good enough to witness our desired inequality.

We first observe that

$$\kappa_n\left([-1,1]\times[-2,2]\setminus\bigcup_{\ell=1}^{\lfloor\sqrt{n}\rfloor}\bigcup_{k=1}^{\lfloor\sqrt{n}\rfloor-2} I_{\ell,k}\times J_\ell\right) = 1-\sum_{\ell=1}^{\lfloor\sqrt{n}\rfloor}\sum_{k=1}^{\lfloor\sqrt{n}\rfloor-2}\kappa_n\left(I_{\ell,k}\times J_\ell\right)$$
$$\leq 1-\lfloor\sqrt{n}\rfloor\left(\lfloor\sqrt{n}\rfloor-2\right)\left(\frac{1}{n}-\frac{1}{n^2}\right)$$
$$< 1-\left(\sqrt{n}-1\right)\left(\sqrt{n}-3\right)\left(\frac{1}{n}-\frac{1}{n^2}\right)$$
$$= \frac{4}{\sqrt{n}}-\frac{2n+4\sqrt{n}-3}{n^2} < \frac{4}{\sqrt{n}}$$

and hence the cost of any plan which transports as much mass within $I_{\ell,k}\times J_\ell$ for each $\ell\in[\lfloor\sqrt{n}\rfloor], k\in[\lfloor\sqrt{n}\rfloor-2]$ as possible is bounded above by

$$\sum_{\ell,k}\text{cost of transport within } I_{\ell,k}\times J_\ell+\sum_{\ell,k}\text{cost of transport from } I_{\ell,k}\times J_\ell \text{ outside}$$
$$< \sum_{\ell=1}^{\lfloor\sqrt{n}\rfloor}\sum_{k=1}^{\lfloor\sqrt{n}\rfloor-2}\frac{1}{n}\sqrt{|I_{\ell,k}|^2+|J_\ell|^2}+\sum_{\ell=1}^{\lfloor\sqrt{n}\rfloor}\sum_{k=1}^{\lfloor\sqrt{n}\rfloor-2}\frac{1}{n^2}\sqrt{2^2+4^2}+\frac{4}{\sqrt{n}}\sqrt{2^2+4^2}$$
$$\leq \frac{1}{n}\sum_{\ell=1}^{\lfloor\sqrt{n}\rfloor}\sum_{k=1}^{\lfloor\sqrt{n}\rfloor-2}\left(|I_{\ell,k}|+|J_\ell|\right)+\lfloor\sqrt{n}\rfloor\left(\lfloor\sqrt{n}\rfloor-2\right)\frac{\sqrt{20}}{n^2}+\frac{4\sqrt{20}}{\sqrt{n}}$$
$$\leq \frac{1}{n}\sum_{\ell=1}^{\lfloor\sqrt{n}\rfloor}\left(2+\left(\lfloor\sqrt{n}\rfloor-2\right)|J_\ell|\right)+\sqrt{n}\left(\sqrt{n}-2\right)\frac{\sqrt{20}}{n^2}+\frac{4\sqrt{20}}{\sqrt{n}}$$
$$\leq \frac{2\lfloor\sqrt{n}\rfloor}{n}+\frac{\lfloor\sqrt{n}\rfloor-2}{n}\cdot 4+\frac{\left(\sqrt{n}-2\right)\sqrt{20}}{n\sqrt{n}}+\frac{8\sqrt{5}}{\sqrt{n}}$$
$$\leq \frac{6+8\sqrt{5}}{\sqrt{n}}+\frac{\left(2\sqrt{5}-8\right)\sqrt{n}-4\sqrt{5}}{n\sqrt{n}} < \frac{6+8\sqrt{5}}{\sqrt{n}}.$$

For $\kappa\in\mathcal{R}$ with continuous cdf F_κ and such that the values $h_i+h_j, i,j\in[n]$ are distinct, we just constructed $\kappa_n\in\mathcal{R}_n$ satisfying $d_{\mathcal{W}_1}\left(\kappa,\kappa_n\right)<\frac{6+8\sqrt{5}}{\sqrt{n}}$. But for any $\kappa\in\mathcal{R}$, there are elements in $\mathcal{R}$ arbitrarily close to κ possessing these desired properties. In particular, there exists $\kappa'\in\mathcal{R}$ as above and with $d_{\mathcal{W}_1}\left(\kappa,\kappa'\right)\leq\frac{1}{\sqrt{n}}$. Then the triangle inequality gives us that

$$d_{\mathcal{W}_1}\left(\kappa,\kappa_n'\right)\leq d_{\mathcal{W}_1}\left(\kappa,\kappa'\right)+d_{\mathcal{W}_1}\left(\kappa',\kappa_n'\right)<\frac{1}{\sqrt{n}}+\frac{6+8\sqrt{5}}{\sqrt{n}}<\frac{25}{\sqrt{n}}.$$

Then
$$d_{LP}\left(\kappa,\kappa_n'\right)\leq\sqrt{d_{\mathcal{W}_1}\left(\kappa,\kappa_n'\right)}<\frac{5}{n^{1/4}},$$

which is bounded above by ε whenever $n\geq\frac{5^4}{\varepsilon^4}$. □

Proof of Theorem 5.1. For any probability space (Ω, μ) and any μ^2-ultrafilter-valued random variable $\mathcal{U}$ on (Ω, μ), the k-profiles of A satisfy $\mathcal{S}_k(A) \subseteq \mathcal{R}^k$ by virtue of $(\Omega \times \Omega, \mu \times \mu)$ being a product space and hence the two summands $\phi_{\mathcal{U}}(f)$ being distributed identically and independently of one another for every $f \in L^\infty(\Omega \times \Omega, \mu \times \mu)$.

Now we will construct a λ^2-ultrafilter-valued random variable $\mathcal{U}$ which also ensures $\mathcal{R}^k \subseteq \overline{\mathcal{S}_k(A)}$. To start with, let, for every $a \in [0, 1)$, $\mathcal{F}_a$ be the λ-filter

$$\mathcal{F}_a := \{[a, a+\varepsilon] \cup B : \varepsilon \in (0, 1-a), B \in \mathcal{B}([0,1])\}\,.$$

We note that for any $a, b \in [0, 1)$, the λ-filters $\mathcal{F}_a, \mathcal{F}_b$ are mutual shifts. That is, the bijection

$$\begin{aligned}\sigma_{b-a} \colon \mathcal{B}([0,1)) &\longrightarrow \mathcal{B}([0,1))\\ B &\mapsto \{x + b - a \pmod 1 : x \in B\}\end{aligned}$$

sends $\mathcal{F}_a$ to $\mathcal{F}_b$.

Let now $\mathcal{V}_0$ be a fixed λ-ultrafilter extending $\mathcal{F}_0$. Then for every $a \in [0, 1)$, the shift $\mathcal{V}_a := \sigma_a(\mathcal{V}_0)$ is a λ-ultrafilter extending $\mathcal{F}_a$.

Similarly, we can now construct λ^2-filters

$$\mathcal{G}_a := \left\{(V \times [0, \varepsilon]) \cup B : V \in \mathcal{V}_a, \varepsilon \in (0,1), B \in \mathcal{B}\left([0,1)^2\right)\right\}$$

and fix a λ^2-ultrafilter $\mathcal{U}_0$ extending $\mathcal{G}_0$. As before, $\mathcal{G}_a$, $\mathcal{G}_b$ are related by the shift σ_{b-a}, where now this is understood to be the map

$$\begin{aligned}\sigma_{b-a} \colon \mathcal{B}([0,1)^2) &\longrightarrow \mathcal{B}([0,1)^2)\\ B &\mapsto \{(x + b - a \pmod 1, y) : (x,y) \in B\}\,,\end{aligned}$$

and we define $\mathcal{U}_a$ to be $\sigma_a(\mathcal{U}_0)$. Then we claim that setting

$$\begin{aligned}\mathcal{U} \colon ([0,1), \lambda) &\longrightarrow \lambda^2\text{-ultrafilters}\\ x &\mapsto \mathcal{U}_x\end{aligned}$$

meets our requirements.

Let $\kappa \in \mathcal{R}^k$ be given. Then we want to find $[-1,1]$-valued functions $f_n^i \in L^\infty([0,1), \lambda)$ such that $\mathcal{D}_A\left(f_n^1, \ldots, f_n^k\right) \xrightarrow{d_{LP}} \kappa$ as $n \to \infty$. By definition of $\mathcal{R}^k$, there are measurable functions $f_1, \ldots, f_k \colon [0,1) \times [0,1) \longrightarrow [-1,1]$ and $g_1, \ldots, g_k \colon [0,1) \longrightarrow [-1,1]$ such that the pushforward $F_*(\lambda \times \lambda)$ of the measure $\lambda \times \lambda$ under the map

$$\begin{aligned}F \colon \Omega \times \Omega &\longrightarrow [-1,1]^k \times [-2,2]^k\\ (x,y) &\mapsto (f_1(x,y), \ldots, f_k(x,y), g_1(x) + g_1(y), \ldots, g_k(x) + g_k(y))\end{aligned}$$

is equal to κ. Now for $\varepsilon \in (0,1)$, for every $i \in [k]$, we define $f_i^\varepsilon \colon \Omega \times \Omega \longrightarrow [-1,1]$ by

$$\begin{aligned}f_i^\varepsilon|_{\Omega \times [\varepsilon, 1)} &= f_i|_{\Omega \times [\varepsilon,1)}\\ \text{and } f_i^\varepsilon(x,y) &= g_i^\varepsilon(x) \text{ whenever } (x,y) \in \Omega \times [0, \varepsilon),\end{aligned}$$

where g_i^ε is a continuous function that will be specified later. Then by construction of $\mathcal{U}$ and continuity of g_i^ε,

$$\left(Af_i^\varepsilon\right)(x,y) = \phi_{\mathcal{U}_x}\left(f_i^\varepsilon\right) + \phi_{\mathcal{U}_y}\left(f_i^\varepsilon\right) = \phi_{\mathcal{V}_x}\left(g_i^\varepsilon\right) + \phi_{\mathcal{V}_y}\left(g_i^\varepsilon\right) = g_i^\varepsilon(x) + g_i^\varepsilon(y).$$

What we would really like to see is this sum being equal to $g_i(x) + g_i(y)$ because that would exactly mean that the joint distribution $\mathcal{D}\left(Af_1^\varepsilon, \ldots, Af_k^\varepsilon\right)$ is equal to $\kappa_{>k}$ as desired. Unfortunately, for $k > 1$, we cannot ensure that for every $i \in [k]$, g_i is discontinuous only on a set of measure 0, which means that on a set of positive measure, it becomes difficult to control the values $\phi_{\mathcal{V}_x}\left(g_i\right)$. This is why we turn to Lusin's theorem, by which there are continuous functions $g_i^\varepsilon \colon \Omega \longrightarrow [-1,1]$ such that $\lambda\left(\{\omega : g_i^\varepsilon(\omega) \neq g_i(\omega)\}\right) < \varepsilon$. Then the distributions $\mathcal{D}_A\left(f_1^\varepsilon, \ldots, f_k^\varepsilon\right)$ weakly* converge to $F_*(\lambda \times \lambda) = \kappa$ as $\varepsilon \to 0$ because the function F differs from

$$\begin{aligned} F_\varepsilon \colon \Omega \times \Omega &\longrightarrow [-1,1]^k \times [-2,2]^k \\ (x,y) &\mapsto \left(f_1^\varepsilon(x,y), \ldots, f_k^\varepsilon(x,y), g_1^\varepsilon(x) + g_1^\varepsilon(y), \ldots, g_k^\varepsilon(x) + g_k^\varepsilon(y)\right) \end{aligned}$$

on a set of measure at most $\varepsilon + k \cdot (1 - (1-\varepsilon)^2) = \varepsilon + k \cdot (2\varepsilon - \varepsilon^2)$. But since $\mathbb{R}^{2k}$ is separable, we have that d_{LP} metrizes weak* convergence, and so $\mathcal{D}_A\left(f_1^\varepsilon, \ldots, f_k^\varepsilon\right) \xrightarrow{d_{LP}} \kappa$ as we were aiming to show. □

6 Two questions

We tried to push the boundaries of Theorem 3.7 and Proposition 1.3, and while we concluded that all three parts of the latter fail unless uniform boundedness of (p,q)-norms is assumed, we did not find a Cauchy sequence that would exemplify that the norm requirement could not be dropped from the statement of Theorem 3.7.

Question 6.1. *Let $(A_n)_n^\infty$ be a Cauchy sequence of graphops. Is there always a P-operator A with $\lim_n d_M(A_n, A) = 0$?*

We have also seen limits of graphop sequences which were not positivity-preserving, yet we were always able to find a positivity-preserving limit too.

Question 6.2. *Let $(A_n)_n^\infty$ be a graphop sequence which has an action limit A. Is there necessarily an action limit B of $(A_n)_n^\infty$ which is positivity-preserving?*

Acknowledgement

I would like to thank Jan Grebík and Konrad Królicki for discussions that set me on the right track.

References

[1] Charalambos Aliprantis and Kim Border. *Infinite Dimensional Analysis.* Springer Berlin, Heidelberg, 3rd edition, 2006.

[2] Ágnes Backhausz and Balázs Szegedy. Action convergence of operators and graphs. *Canadian Journal of Mathematics*, page 1–50, 2020.

[3] Christian Borgs, Jennifer T. Chayes, Henry Cohn, and Nina Holden. Sparse exchangeable graphs and their limits via graphon processes. *Journal of Machine Learning Research*, 18(1):7740–7810, jan 2017.

[4] Christian Borgs, Jennifer T. Chayes, Henry Cohn, and Yufei Zhao. An L^p theory of sparse graph convergence II: LD convergence, quotients and right convergence. *The Annals of Probability*, 46(1):337 – 396, 2018.

[5] Christian Borgs, Jennifer T. Chayes, Henry Cohn, and Yufei Zhao. An L^p theory of sparse graph convergence I: Limits, sparse random graph models, and power law distributions. *Transactions of the American Mathematical Society*, 372(5):3019 – 3062, 2019.

[6] Christian Borgs, Jennifer T. Chayes, László Lovász, Vera T. Sós, and Katalin Vesztergombi. Convergent sequences of dense graphs I: Subgraph frequencies, metric properties and testing. *Advances in Mathematics*, 219(6):1801–1851, 2008.

[7] Martin Doležal. Graph limits: An alternative approach to s-graphons. *Journal of Graph Theory*, 99(1):90–106, 2022.

[8] Péter Frenkel. Convergence of graphs with intermediate density. *Transactions of the American Mathematical Society*, 370, 02 2016.

[9] Hamed Hatami, László Lovász, and Balázs Szegedy. Limits of locally–globally convergent graph sequences. *Geometric and Functional Analysis*, 24:269–296, 2014.

[10] Anne-Louise Klaus and Chi-Kwong Li. Isometries for the vector (p,q) norm and the induced (p,q) norm. *Linear and Multilinear Algebra*, 38(4):315–332, 1995.

[11] Dávid Kunszenti-Kovács, László Miklós Lovász, and Balázs Szegedy. Measures on the square as sparse graph limits. *J. Comb. Theory, Ser. B*, 138:1–40, 2019.

[12] Lászlo Lovász. *Large Networks and Graph Limits.* Providence, R.I.: American Mathematical Society., 2012.

[13] László Lovász and Balázs Szegedy. Limits of dense graph sequences. *Journal of Combinatorial Theory, Series B*, 96(6):933–957, 2006.

[14] Jaroslav Nešetřil and Patrice Ossona de Mendez. *Sparsity.* Springer Berlin, Heidelberg, 2012.

[15] Balázs Szegedy. Sparse graph limits, entropy maximization and transitive graphs. *arXiv: Combinatorics 1504.00858*, 2015.

[16] John Toland. *The Dual of* $L_\infty(X, \mathcal{L}, \lambda)$*, Finitely Additive Measures and Weak Convergence.* Springer, 2020.

[17] Remco van der Hofstad. *Random Graphs and Complex Networks: Volume 2.* to appear, 2023.

Appendix D

Logarithmic convergence of projective planes

by Márton Borbényi, Aranka Hrušková, Ander Lamaison

Logarithmic convergence of projective planes

Márton Borbényi, Aranka Hrušková, Ander Lamaison

June 2022

Abstract

In this paper, we study the so-called log-convergence of graphs defined by Balázs Szegedy [4]. We answer positively his Question 4 whether the sequence of the incidence graphs of projective planes over finite fields log-converges and whether the limit coincides with that of a particular random graph model.

1 Introduction

Since its introduction in the 1980s, the theory of graph limits has been applied to numerous problems in extremal graph theory, and it has been generalized and extended to many different discrete structures. A particular focus has been placed on dense graph limits: limits of graph sequences in which the edge density is bounded away from 0. While there has been some success in defining a limit theory in other particular classes, such as graphs with bounded maximum degree, there is no universally agreed upon approach toward limits and convergence in the entire sparse regime.

One such approach, called logarithmic convergence, was proposed in 2015 by Balázs Szegedy [4]. This notion was motivated by Sidorenko's conjecture [3] in which the central inequality becomes linear after taking logarithms. For technical reasons, we will restrict ourselves to considering only bipartite graphs, although Szegedy's definition applies to general graphs.

We denote by $\mathcal{B}$ the set of finite bipartite graphs whose vertex classes are labelled V_1 and V_2. That is,

$$\mathcal{B} = \{G = (V_1(G), V_2(G), E(G)) : E(G) \subseteq V_1(G) \times V_2(G)\}.$$

A homomorphism between two elements of $\mathcal{B}$ is a graph homomorphism which is label-preserving. That is, we will only count maps $\varphi : H \to G$ which are graph homomorphisms in the classical sense *and* for which $\varphi(V_1(H)) \subseteq V_1(G)$ and $\varphi(V_2(H)) \subseteq V_2(G)$. Subsequently, this gives rise to the definition of the bipartite density of $H \in \mathcal{B}$ in $G \in \mathcal{B}$ as

$$t_B(H,G) := \frac{\mathsf{hom}(H,G)}{|V_1(G)|^{|V_1(H)|}|V_2(G)|^{|V_2(H)|}}.$$

where $\mathsf{hom}(H,G)$ is the number of homomorphisms from H to G.

The quantity $t_B(H,G)$ can be interpreted also as the probability that a randomly chosen label-preserving map $V(H) \to V(G)$ is a homomorphism, in which view

$$d(H,G) := -\log t_B(H,G)$$

denotes the Kullback-Leibler divergence of the uniform distribution on homomorphisms with respect to the uniform measure on all label-preserving maps.

Definition (log-convergence)**.** *Let $\mathcal{B}_0$ be the class of graphs of $\mathcal{B}$ with at least one edge. We say that a sequence $(G_n)_n^\infty$ of graphs in $\mathcal{B}_0$ is log-convergent if, for every pair of graphs $H_1, H_2 \in \mathcal{B}_0$, the sequence*

$$\frac{-\log t_B(H_1, G_n)}{-\log t_B(H_2, G_n)} \tag{1}$$

has a limit.

This definition can, however, be greatly simplified. Let $h(H, G)$ be the particular ratio

$$h(H,G) = \frac{d(H,G)}{d(K_2,G)} = \frac{-\log t_B(H,G)}{-\log t_B(K_2,G)},$$

where $K_2 \in \mathcal{B}_0$ is the edge with partition. Then the following lemma tells us that, in a sense, log-convergence is in fact only concerned with the quantities $\log t_B(H,G)$ modulo edge density.

Lemma 1 (Lemma 4.2 in [4])**.** *A graph sequence $(G_n)_n^\infty$ in $\mathcal{B}_0$ is log-convergent if and only if*

$$\lim_{n\to\infty} h(H, G_n) = \lim_{n\to\infty} \frac{-\log t_B(H, G_n)}{-\log t_B(K_2, G_n)}$$

exists for every $H \in \mathcal{B}_0$. Every graph sequence in $\mathcal{B}_0$ has a log-convergent subsequence.

The ratios of the form (1) are not defined when G_n is a complete bipartite graph. In this case, following the same criterion as Szegedy, we define $h(H, G) = |E(H)|$ for any complete bipartite graph G.

Lemma 1 tells us that the trivial limit object for a log-convergent sequence $(G_n)_n^\infty$ is the vector $(\lim_n h(H_1, G_n), \lim_n h(H_2, G_n), \dots) \in \mathbb{R}_{\geq 0}^{\mathcal{B}_0}$, where $H_1, H_2, \dots$ is some fixed enumeration of $\mathcal{B}_0$. As of now, the theory of log-convergence lacks a good analytical or algebraic limit object akin to graphon or graphing, and so whenever we will be talking about a limit, we will simply mean an element of $\mathbb{R}_{\geq 0}^{\mathcal{B}_0}$.

Hand in hand with any theory of graph convergence goes a random model reflecting its salient features. When analytical limit objects like graphons are at our disposal, we can typically use them to generate random graph sequences that almost surely converge to the limit object we started with. While as pointed out above, we currently do not have such a non-trivial limit, Szegedy introduced a random bipartite model that satisfies the almost sure convergence and captures a wide array of structural scenarios. As opposed to the Erdős-Rényi model $\mathcal{G}(n, p)$ which has only one parameter determining the asymptotic behaviour, namely the edge density p, Szegedy's model requires two parameters: β accounts for the edge density, and α accounts for the relative sizes of the two vertex classes in the bipartition. This random model $R(n, \beta, \alpha)$ will be defined in Section 2. Szegedy proved that, for any fixed choice of the two parameters, the sequence $(R(n, \beta, \alpha))_n^\infty$ of random bipartite graphs log-converges with probability 1, and gave an explicit description of the limit $R(\beta, \alpha) \in \mathbb{R}_{\geq 0}^{\mathcal{B}_0}$.

While studying the random model $R(n, \beta, \alpha)$, Szegedy asked whether the incidence graphs of finite projective planes are pseudorandom in the log-convergence framework. In particular, let $PG(2, q)$ be the projective plane over the finite field $\mathbb{F}_q$. Then its incidence graph G_q is the bipartite graph in which the class V_1 is the set of points of $PG(2, q)$, the class V_2 is the set of lines of $PG(2, q)$, and an edge is drawn between a point p and a line ℓ if $p \in \ell$. Szegedy asked whether the sequence of incidence graphs of projective planes is log-convergent, when the parameter q ranges over the set of primes. He further asked, in case the sequence is convergent, whether the limit is $R(3/4, 1/2)$, the same as the limit of a specific sequence of random bipartite graphs. We answer both questions in the affirmative, even when q is allowed to be a prime power.

Theorem 2. *The sequence $(G_q)_q^\infty$ of incidence graphs of projective planes $PG(2,q)$, where q ranges over the prime powers, is log-convergent. Moreover, its limit is $R(3/4,1/2)$.*

This statement is remarkable because if we change the definition of t_B to count only injective homomorphisms rather than all homomorphisms, the incidence graph of the projective plane behaves differently than the random bipartite graph. We will see some examples of this in the next section.

2 Preliminaries

In this paper, all graphs mentioned belong to $\mathcal{B}_0$, and in particular are finite and bipartite. Through the rest of the paper, $\mathsf{inj}(H,G)$ and $\mathsf{surj}(H,G)$ denote, respectively, the number of injective homomorphisms and surjective homomorphisms from H to G (remember that we only count those homomorphisms sending $V_i(H)$ to $V_i(G)$).

Earlier, we defined log-convergence in terms of the fraction $\frac{-\log t_B(H_1,G_n)}{-\log t_B(H_2,G_n)}$. Lemma 1 says that the only parameters that we are interested in are the *exponents* to which we need to raise the edge density in order to obtain the other densities $t_B(H,G)$. This is what allows us to distinguish graphs which, in the classical definition of convergence, would have as limit the zero graphon. This control over the exponents of densities is reflected in the following definition of a random bipartite graph model.

Definition (random graph model)**.** *Let $\beta \in (0,1]$ and $\alpha \in (0,1)$ be fixed. We denote by $G(n,\beta,\alpha)$ the distribution on bipartite graphs with $|V_1| = \lceil n^\alpha \rceil$, $|V_2| = \lceil n^{1-\alpha} \rceil$ given by including each of the $\lceil n^\alpha \rceil \lceil n^{1-\alpha} \rceil$ possible edges with probability $n^{\beta-1}$, independently of each other.*

Szegedy proves in [4] that for every fixed $\beta, \alpha > 0$, $G(n,\beta,\alpha)$ log-converges as n tends to ∞ with probability 1, and denotes by $R(\beta,\alpha)$ the collection of the limits $\lim_n h\,(H, G(n,\beta,\alpha))$. íThe limits attained for different values of β and α are distinct. Are they?

We now shift our attention to the graph sequence whose log-convergence we wish to prove. Let q be a (power of a) prime number. As we mentioned earlier, G_q is the incidence graph of the projective plane $PG(2,q)$. In this graph, both $V_1(G_q)$ and $V_2(G_q)$ contain q^2+q+1 vertices, and the graph is $q+1$-regular. By looking at the largest order terms, we can see that the graph $G(q^4,3/4,1/2)$ would have similar part sizes (q^2 each) and expected edge density ($\mathbb{E}\,[t_B(K_2,G)] = q^{-1}$) as G_q. Szegedy asks whether the sequence $(G_q)_q^\infty$ is pseudorandom, in particular, whether it converges to $R(3/4,1/2)$. In the rest of the paper, we will prove that $(G_q)_q^\infty$ converges to $R(3/4,1/2)$ by showing increasing similarity between G_q and $R_q := G(q^4,3/4,1/2)$.

The following definition plays a crucial role in the description of the limit in our main result, as well as in its proof:

Definition (collapse)**.** *A graph H' is a* collapse[1] *of a graph H if there is a graph homomorphism φ from H to H' which is both vertex- and edge-surjective. Let $\mathcal{C}(H)$ denote the set of collapses of H.*

Since we only care about the ratio of logarithms, we are free to choose the base of such logarithm. Given the parametrization of our graphs, it seems natural to take q as the base of the logarithm, because the edge density of our graph is $(1+o(1))q^{-1}$. That means that $\lim_{q\to\infty} -\log_q t_B(K_2,G_q) = 1$, cancelling the denominator in Lemma 1. In addition, observe that $\log_q t_B(H,G_q) = \log_q \mathsf{hom}(H,G_q) - |V(H)|\log_q(q^2 + q + 1)$, and so by Lemma 1, to guarantee the log-convergence of $(G_q)_q^\infty$ it is enough to verify that $\log_q \mathsf{hom}(H,G_q)$ has a limit for every $H \in \mathcal{B}_0$.

[1]Szegedy uses the term *homomorphic image* for collapses in [4].

First we are going to show that, by restricting ourselves to an adequate subsequence of prime powers q, we can assume that $\lim_{q\to\infty}\log_q \mathsf{inj}(H,G_q)$ and $\lim_{q\to\infty}\log_q \mathsf{hom}(H,G_q)$ exist for all graphs H, where $\log_q 0 = -\infty$.

Lemma 3. *There exists a sequence of prime powers $(q_i)_{i=1}^{\infty}$ such that $\log_q \mathsf{inj}(H,G_q)$ and $\log_q \mathsf{hom}(H,G_q)$ converge on this subsequence for every $H\in\mathcal{B}_0$. Moreover, if Theorem 2 is false, then such a subsequence can be taken in a way that $\lim_{q\to\infty}\log_q \mathsf{hom}(H,G_q)\neq\lim_{q\to\infty}\log_q \mathsf{hom}(H,R_q)$ for at least one $H\in\mathcal{B}_0$.*

Proof. For every H, we have that $\mathsf{inj}(H,G_q)$ and $\mathsf{hom}(H,G_q)$ are integers in the interval $\left[0,(q^2+q+1)^{|V(H)|}\right]$. Since $q^2+q+1<q^3$ for $q\geq 2$, we have that $\log_q \mathsf{inj}(H,G_q)$ and $\log_q \mathsf{hom}(H,G_q)$ both lie on $\{-\infty\}\cup[0,3|V(H)|]$. Therefore, by associating each G_q with its corresponding values of $\log_q \mathsf{inj}(H,G_q)$ and $\log_q \mathsf{hom}(H,G_q)$, we obtain a sequence $(x_q)_q^{\infty}$ in the space

$$X=\prod_{H\in\mathcal{B}_0}\left(\{-\infty\}\cup[0,3v(H)]\right)^2.$$

X is the product of countably many sequentially compact spaces, and is therefore sequentially compact. Hence there exists a subsequence of $(x_q)_q^{\infty}$ in which every coordinate converges.

If Theorem 2 is false, then for some $H\in\mathcal{B}_0$ the sequence $\left(\log_q \mathsf{hom}(H,G_q)\right)_q^{\infty}$ does not converge to $\lim_{q\to\infty}\log_q \mathsf{hom}(H,R_q)$. This means that this sequence must have a different accumulation point. Select a sequence of prime powers q that tends to this accumulation point, then use the sequential compactness of X to find a subsequence for which the corresponding points in X converge. □

From this point on, when we write $(G_q)_q^{\infty}$, we assume that the values of q are restricted to the sequence obtained in Lemma 3. This restriction is necessary: for example G_2, which is the incidence graph of the Fano plane, is a subgraph of G_q if and only if q is a power of 2, meaning that the sequence $\left(\log_q \mathsf{inj}(G_2,G_q)\right)_q^{\infty}$ does not converge when q ranges over all prime powers.

We are going to use the following notation:

$$\hat{\imath}(H)=\lim_{q\to\infty}\log_q \mathsf{inj}(H,G_q)$$
$$\hat{h}(H)=\lim_{q\to\infty}\log_q \mathsf{hom}(H,G_q)$$
$$\hat{\imath}_R(H)=2|V(H)|-|E(H)|$$
$$\hat{h}_R(H)=\max_{H'\in\mathcal{C}(H)}\hat{\imath}_R(H')$$

Given the notation that we are using, one might expect to find an analogy between $\hat{\imath}$ and $\hat{\imath}_R$, and between $\hat{h}$ and $\hat{h}_R$. Such a relation is not evident at first sight, but it can be found when considering the graph R_q. A simple first moment calculation reveals that

$$\lim_{q\to\infty}\log_q \mathbb{E}\left[\mathsf{inj}(H,R_q)\right]=2|V(H)|-|E(H)|=\hat{\imath}_R(H).$$

From there, classifying the homomorphisms from H to R_q by their image produces

$$\lim_{q\to\infty}\log_q \mathbb{E}\left[\mathsf{hom}(H,R_q)\right]=\lim_{q\to\infty}\log_q\sum_{H'\in\mathcal{C}(H)}\mathsf{surj}(H,H')\mathbb{E}\left[\mathsf{inj}(H',R_q)\right]=\max_{H'\in\mathcal{C}(H)}\hat{\imath}_R(H')=\hat{h}_R(H).$$

Szegedy ([4], see Theorem 3) proved a concentration result for the number of homomorphisms in the random bipartite graph $R(n, \beta, \alpha)$. His result implies that with probability 1, for every graph H the number $\log_q \mathsf{hom}(H, R_q)$ converges to $\hat{h}(H)$.

The next proposition states that it is possible to compute the $\hat{h}$ of a graph from the $\hat{\imath}$ of its collapses, in the same way as one can compute $\hat{h}_R$ from $\hat{\imath}_R$.

Proposition 4. *For every graph $H \in \mathcal{B}_0$, we have*

$$\hat{h}(H) = \max_{H' \in \mathcal{C}(H)} \hat{\imath}(H').$$

Proof. Let ϕ be a homomorphism from H to G_q. Its image is a collapse H' of H. In fact, ϕ can be expressed in a unique way as the composition of a surjective homomorphism from H to H' and an injective homomorphism from H' to G_q. Recall that $\mathsf{surj}(H, H')$ denotes the number of label-preserving surjective homomorphisms from H to H'. Then

$$\mathsf{hom}(H, G_q) = \sum_{H' \in \mathcal{C}(H)} \mathsf{surj}(H, H')\mathsf{inj}(H', G_q).$$

Note that

$$\max_{H' \in \mathcal{C}(H)} \hat{\imath}(H') = \max_{H' \in \mathcal{C}(H)} \lim_{q \to \infty} \log_q \mathsf{inj}(H', G_q) = \lim_{q \to \infty} \log_q \max_{H' \in \mathcal{C}(H)} \mathsf{inj}(H', G_q)$$

by finiteness of $\mathcal{C}(H)$, and thus

$$\begin{aligned}
\log_q \mathsf{hom}(H, G_q) &= \log_q \left(\sum_{H' \in \mathcal{C}(H)} \mathsf{surj}(H, H')\mathsf{inj}(H', G_q) \right) \\
&\leq \log_q \left(\left(\sum_{H' \in \mathcal{C}(H)} \mathsf{surj}(H, H') \right) \max_{H' \in \mathcal{C}(H)} \mathsf{inj}(H', G_q) \right) \\
&= \log_q \left(\sum_{H' \in \mathcal{C}(H)} \mathsf{surj}(H, H') \right) + \log_q \max_{H' \in \mathcal{C}(H)} \mathsf{inj}(H', G_q) \\
&\to \max_{H' \in \mathcal{C}(H)} \hat{\imath}(H') \quad \text{as } q \to \infty,
\end{aligned}$$

and

$$\begin{aligned}
\log_q \mathsf{hom}(H, G_q) &= \log_q \left(\sum_{H' \in \mathcal{C}(H)} \mathsf{surj}(H, H')\mathsf{inj}(H', G_q) \right) \\
&\geq \log_q \max_{H' \in \mathcal{C}(H)} \mathsf{inj}(H', G_q) \\
&\to \max_{H' \in \mathcal{C}(H)} \hat{\imath}(H') \quad \text{as } q \to \infty,
\end{aligned}$$

hence we obtained the desired identity. □

Throughout the paper we will use the following notations for graph operations.

Definition (graph operations)**.** *Let $G = (V_1(G), V_2(G), E(G)) \in \mathcal{B}$ be a labelled bipartite graph. For any $W \subseteq V_1(G) \cup V_2(G)$ let $G - W$ denote the bipartite graph that we obtain by deleting the vertices of W and all edges incident to these vertices. Let $v, w \in V_i(G)$ be vertices of G for some $i \in \{1, 2\}$. Let $G_{v,w}$ denote the bipartite graph obtained from G by first deleting the vertices v, w (obtaining $G - \{v, w\}$) and then adding a vertex u to $V_i(G)$ which is connected to a vertex x if and only if x was a neighbour of v or w in G. Do we ever use this?*

The wedge sum of two graphs H_1 and H_2, at vertices $v_1 \in V_i(H_1)$ and $v_2 \in V_i(H_2)$, is obtained by taking disjoint copies of H_1 and H_2, and identifying the vertices v_1 and v_2. The disjoint union of H_1 and H_2 connected by an edge at vertices $v_1 \in V_i(H_1)$ and $v_2 \in V_{3-i}(H_2)$ is, as its name indicates, the disjoint union of H_1 and H_2 together with an edge between v_1 and v_2.

Lemma 5. *Let k be a nonnegative integer and $H \in \mathcal{B}_0$. Let us set*

$$\mathsf{inj}_k(H, G_q) := \min_{v_1,\dots,v_k \in V(G_q)} \mathsf{inj}(H, G_q - \{v_1, \dots, v_k\}).$$

Then for every $H \in \mathcal{B}_0$ and any nonnegative k,

$$\hat{i}(H) = \lim_{q\to\infty} \log_q \mathsf{inj}_k(H, G_q).$$

Proof. We can assume that the left-hand side is not $-\infty$, because otherwise the lemma is trivial.

Let $w \in V(H)$, $v \in V(G_q)$. Then the probability that a uniformly randomly chosen injection f takes w to v is 0 or $1/(q^2 + q + 1)$ by the point/line transitivity of the projective planes. Thus by the union bound for any $A = \{v_1, \dots, v_k\} \subset V(G_q)$

$$\mathbb{P}(f(H) \cap A \neq \emptyset) \leq \frac{|V(H)|k}{q^2 + q + 1} = O(q^{-2}).$$

Hence

$$\log_q \mathsf{inj}_k(H, G_q) - \log_q \mathsf{inj}(H, G_q) = \log_q \frac{\mathsf{inj}_k(H, G_q)}{\mathsf{inj}(H, G_q)} = \log_q \left[\min_{|A|=k} \mathbb{P}(f(H) \cap A = \emptyset)\right] \to 0$$

as q tends to infinity, so we are done. □

Lemma 6. *Let H_1 and H_2 be two labelled bipartite graphs. Let H be obtained from H_1 and H_2 by taking either their disjoint union ($H_1 \dot\cup H_2$), wedge sum ($H_1 \vee H_2$) or the disjoint union connected by an edge. Then $\hat{i}(H)$, $\hat{i}_R(H)$, $\hat{h}(H)$ and $\hat{h}_R(H)$ are as in Table 1, independently of which two vertices are identified by the wedge or connected by the edge.*

Proof. The result is trivial for the quantitites $\hat{i}_R$ and $\hat{h}_R$. For $\hat{i}$ and $\hat{h}$, the first column comes from Lemma 5 because $\mathsf{hom}(H_1 \dot\cup H_2, G_q) = \mathsf{hom}(H_1, G_q)\mathsf{hom}(H_2, G_q)$ and

$$\mathsf{inj}(H_1, G_q)\mathsf{inj}_{|V(H_1)|}(H_2, q) \leq \mathsf{inj}(H_1 \dot\cup H_2, G_q) \leq \mathsf{inj}(H_1, G_q)\mathsf{inj}(H_2, G_q).$$

For the second column, we have $\mathsf{hom}(H_1 \vee H_2, G_q) = \frac{\mathsf{hom}(H_1,G_q)\mathsf{hom}(H_2,G_q)}{q^2+q+1}$ and $\mathsf{inj}(H_1 \vee H_2, G_q) \leq \frac{\mathsf{inj}(H_1,G_q)\mathsf{inj}(H_2,G_q)}{q^2+q+1}$. For the lower bound we need an argument similar to Lemma 5.

	$H = H_1 \dot\cup H_2$	$H = H_1 \vee H_2$	disjoint union connected by an edge
$\hat{\imath}(H)$	$\hat{\imath}(H_1) + \hat{\imath}(H_2)$	$\hat{\imath}(H_1) + \hat{\imath}(H_2) - 2$	$\hat{\imath}(H_1) + \hat{\imath}(H_2) - 1$
$\hat{\imath}_R(H)$	$\hat{\imath}_R(H_1) + \hat{\imath}_R(H_2)$	$\hat{\imath}_R(H_1) + \hat{\imath}_R(H_2) - 2$	$\hat{\imath}_R(H_1) + \hat{\imath}_R(H_2) - 1$
$\hat{h}(H)$	$\hat{h}(H_1) + \hat{h}(H_2)$	$\hat{h}(H_1) + \hat{h}(H_2) - 2$	$\hat{h}(H_1) + \hat{h}(H_2) - 1$
$\hat{h}_R(H)$	$\hat{h}_R(H_1) + \hat{h}_R(H_2)$	$\hat{h}_R(H_1) + \hat{h}_R(H_2) - 2$	$\hat{h}_R(H_1) + \hat{h}_R(H_2) - 1$

Table 1: Results of Lemma 6

We can assume that $\hat{\imath}(H_1), \hat{\imath}(H_2) \neq -\infty$, as otherwise the statement is trivial. Up to a change in the indices, we can suppose that $v_1 \in V_1(H_1)$. Fix a copy of H_1 in G_q, and consider the point p which is the image of v_1. Consider the subgroup of automorphisms of $PG(2,q)$ which fixes p. This group is transitive on the set of points distinct from p, transitive on the set of lines containing p, and transitive on the set of lines not containing p. Next consider a copy of H_2 in G_q for which the image of v_2 is p. By transitivity of G_q, there are exactly $\frac{\mathsf{inj}(H_2,G_q)}{q^2+q+1}$ such copies. Now consider the image of H_2 under the automorphism group described above. Given a vertex $w_1 \in V(H_1) \setminus \{v_1\}$ and a vertex $w_2 \in V(H_2) \setminus \{v_2\}$, the proportion of homomorphisms that send w_1 to w_2 is at most $\frac{1}{q+1}$. This means that the proportion of automorphisms of the copy of H_2 that intersect non-trivially the fixed copy of H_1 is at most $O(q^{-1})$, and so $\mathsf{inj}(H_1 \vee H_2, G_q) \geq (1 - O(q^{-1}))\frac{\mathsf{inj}(H_1,G_q)\mathsf{inj}(H_2,G_q)}{q^2+q+1}$.

Finally, the third column can be obtained from the second, because the disjoint union of H_1 and H_2 joined by an edge can be written as $(H_1 \vee K_2) \vee H_2$. □

Corollary 7. *Let H be a labelled bipartite graph, and let $v \in V(H)$.*

- *If $d(v) = 0$, then $\hat{\imath}(H - v) = \hat{\imath}(H) - 2$.*
- *If $d(v) = 1$, then $\hat{\imath}(H - v) = \hat{\imath}(H) - 1$.*
- *If $d(v) \geq 2$, then $\hat{\imath}(H - v) \geq \hat{\imath}(H)$.*

Proof. An isolated vertex has $\hat{\imath}(K_1) = 2$. If $d(v) = 0$, then H is the disjoint union of $H - v$ and K_1. If $d(v) = 1$, then H is the union of $H - v$ and K_1 connected by an edge. If $d(v) \geq 2$, then the fact that any pair of distinct vertices in G_q has at most one common neighbour implies that any injective homomorphism from $H - v$ to G_q extends in at most one way to H, so

$$\hat{\imath}(H) = \lim_{q\to\infty} \log_q \mathsf{inj}(H, G_q) \leq \lim_{q\to\infty} \log_q \mathsf{inj}(H - v, G_q) = \hat{\imath}(H - v).$$ □

Finally, we will need the following lemma. It seems out of place right now, but it will be used in the final steps of the proof of Theorem 2, in which 2-connectivity is important.

Lemma 8. *Let G be a 2-connected graph with $\delta(G) \geq 3$. Then there is an edge $e \in E(G)$ such that the graph $G - \{e\}$ obtained by removing e from G is still 2-connected.*

Proof. By the main theorem of [1], there is a vertex $v \in V(G)$ such that $G-v$ is still 2-connected. Suppose $u_1, \ldots, u_{\deg(v)} \in V(G)$ were the neighbours of v. Then adding a vertex w to $G-v$ and connecting it to $u_2, \ldots, u_{\deg(v)}$ is adding a vertex of degree at least 2 to a 2-connected graph, so the resulting graph $G-v+w$ is still 2-connected. But $G-v+w$ is isomorphic to $G-\{u_1v\}$, so u_1v is an edge whose removal preserves 2-connectivity. □

Before starting with the proof of our main theorem, let us rephrase it using the notation that we have introduced. The following statement is, using Lemma 3, equivalent to Theorem 2.

Theorem 9. *For every bipartite graph H we have $\hat{h}(H) = \hat{h}_R(H)$.*

If instead of homomorphisms we only count injective homomorphisms, we can see that the graphs G_q and R_q behave differently. An obvious example is C_4: in a projective plane, every pair of lines intersects in exactly one point, so $\hat{\imath}(C_4) = -\infty$, while we have $\hat{\imath}_R(C_4) = 4$. On the other hand, there are examples in which $\hat{\imath}(H) > \hat{\imath}_R(H)$, i.e., the graph appears more often in G_q than in R_q. Such examples include the incidence graphs of the point-line arrangements corresponding to Pappus's theorem (denoted by $\bar{P}$) and Desargues's theorem (denoted by $\bar{D}$). These graphs satisfy $\hat{\imath}_R(\bar{P}) = 9$, $\hat{\imath}(\bar{P}) = 10$, $\hat{\imath}_R(\bar{D}) = 10$ and $\hat{\imath}(\bar{D}) = 11$.

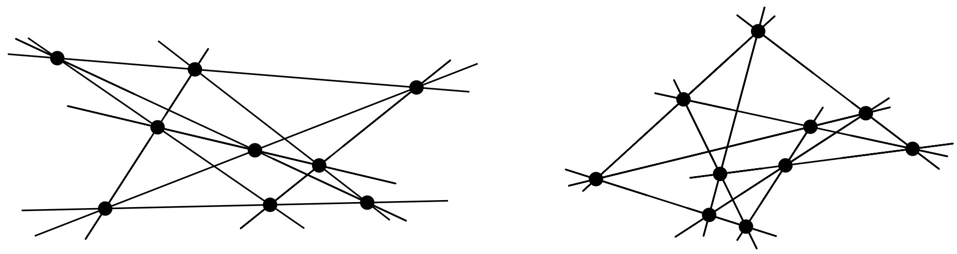

Figure 1: Point-line arrangements corresponding to Pappus's theorem (left) and Desargues's theorem (right).

The reason why these differences between G_q and R_q do not immediately disprove Theorem 2 is that a significant proportion of the homomorphisms from the graphs described above to G_q and R_q are not injective. By Proposition 4 and the definition of $\hat{h}_R$, in order to determine the values of $\hat{h}$ and $\hat{h}_R$ of a graph H one needs to consider the values of $\hat{\imath}$ and $\hat{\imath}_R$ not just for H itself, but also for its collapses. In these cases, there are collapses for which the values of $\hat{\imath}$ and $\hat{\imath}_R$ are greater than or equal to those of H. Indeed, consider the star $K_{1,t}$ which has $\hat{\imath}(K_{1,t}) = \hat{\imath}_R(K_{1,t}) = t+2$. Observe that $K_{1,2}$ is a collapse of C_4, $K_{1,9}$ is a collapse of $\bar{P}$, and $K_{1,10}$ is a collapse of $\bar{D}$, each obtained by mapping all vertices of one side of the bipartition to the center of the star.

3 Proof of Theorem 9

We mentioned in the previous paragraph how, for some graphs H, a constant proportion of homomorphisms from H to G_q and R_q are not injective. We will establish a definition for the opposite case, that is, when almost all homomorphisms are injective. This definition will play a crucial role in the proof of our main theorem:

Definition. *A graph H is* critical *in $(G_q)_q^\infty$ if $\hat{\imath}(H) > \hat{\imath}(H')$ for all proper collapses H' of H. Analogously, H is* critical *in $(R_q)_q^\infty$ if $\hat{\imath}_R(H) > \hat{\imath}_R(H')$ for all proper collapses H' of H.*

The following proposition, which will be the main part of our proof, implies Theorem 9:

Proposition 10. *Suppose that a graph $H \in \mathcal{B}_0$ satisfies $\hat{\imath}(H) \neq \hat{\imath}_R(H)$. Then H is not critical in $(G_q)_q^\infty$ nor in $(R_q)_q^\infty$.*

Proof of Theorem 9. Suppose that the theorem is not true. Let H be a graph with $\hat{h}(H) \neq \hat{h}_R(H)$. Suppose that $\hat{h}(H) > \hat{h}_R(H)$ (the other case is analogous). By Proposition 4, there exists at least one collapse H' of H which satisfies $\hat{\imath}(H') = \hat{h}(H)$. Out of those, select one that minimizes $|V(H')|$. Since the collapse relation is transitive, every collapse H'' of H' satisfies $\hat{\imath}(H'') < \hat{\imath}(H')$, as H'' is a collapse of H smaller than H'. This means that H' is critical in $(G_q)_q^\infty$, and so by Proposition 10 we have $\hat{\imath}_R(H') = \hat{\imath}(H')$. But again by Proposition 4 we have $\hat{h}_R(H) \geq \hat{\imath}_R(H') = \hat{\imath}(H') = \hat{h}(H)$, yielding the desired contradiction. □

This is not a big improvement because proving Proposition 10 is essentially as hard as proving Theorem 9. The main difference between the two is that Proposition 10 is written in a form that lends itself to be proved using induction. In general, it will be easier to relate the values of $\hat{\imath}(H)$ and $\hat{\imath}_R(H)$ to those of the subgraphs of H than the values of $\hat{h}(H)$ and $\hat{h}_R(H)$.

Proof of Proposition 10. Suppose the statement is not true and let H be a minimal counterexample (minimal with respect to number of vertices). We will do a thorough case analysis to show that H cannot exist.

We start by observing that H must be connected. Indeed, if H is not connected then it is the disjoint union of two non-empty graphs H_1, H_2. Then by Lemma 6, we know that $\hat{\imath}(H_1) + \hat{\imath}(H_2) = \hat{\imath}(H)$ and $\hat{\imath}_R(H_1) + \hat{\imath}_R(H_2) = \hat{\imath}_R(H)$. Thus we can assume that $\hat{\imath}(H_1) \neq \hat{\imath}_R(H_1)$. By minimality of H, we have that H_1 is not critical in $(G_q)_q^\infty$ nor in $(R_q)_q^\infty$. There exist proper collapses H_1' and H_1'' of H_1 such that $\hat{\imath}(H_1') \geq \hat{\imath}(H_1)$ and $\hat{\imath}_R(H_1'') \geq \hat{\imath}_R(H_1)$. But then again by Lemma 6, we obtain $\hat{\imath}(H_1' \cup H_2) \geq \hat{\imath}(H)$ and $\hat{\imath}_R(H_1'' \cup H_2) \geq \hat{\imath}_R(H)$. Since $H_1' \cup H_2$ and $H_1'' \cup H_2$ are proper collapses of H, we conclude that H is not critical in $(R_q)_q^\infty$ or $(G_q)_q^\infty$, proving that H is not a counterexample in the first place.

Next we consider different cases depending on the minimum degree of H:

Case 0: $\delta(H) = 0$. H is in $\mathcal{B}_0$, and so it cannot be an isolated vertex. But H must be connected, so we must conclude that $\delta(H) = 0$ cannot happen.

Case 1: $\delta(H) = 1$. Let v be a vertex in H of degree 1. Then Lemma 6 tells us that

$$\hat{\imath}(H) = \hat{\imath}(H - v) + 1.$$

On the other hand, the definition of $\hat{\imath}_R$ says that

$$\hat{\imath}_R(H-v) = 2|V(H-v)| - |E(H-v)| = 2(|V(H)|-1) - (|E(H)-1) = 2|V(H)| - |E(H)| - 1 = \hat{\imath}_R(H) - 1.$$

Since H is a counterexample, we must have $\hat{\imath}(H) \neq \hat{\imath}_R(H)$, and in particular by the two equations above,

$$\hat{\imath}(H - v) = \hat{\imath}(H) - 1 \neq \hat{\imath}_R(H) - 1 = \hat{\imath}_R(H - v).$$

But H was a *minimal* counterexample, so $H - v$ is not a counterexample, and hence it cannot be critical in $(G_q)_q^\infty$ nor in $(R_q)_q^\infty$. In particular, $H - v$ not being critical in $(G_q)_q^\infty$ tells us that there is a proper collapse $(H-v)'$ of $H-v$ satisfying $\hat{\imath}(H-v) \le \hat{\imath}((H-v)')$. Let $+v$ denote the operation of attaching the previously deleted vertex back to (the node representing) the original neighbour of v. With this notation, $H-v+v=H$, and $(H-v)'+v$ is a proper collapse of H. Then

$$\hat{\imath}(H-v) \le \hat{\imath}((H-v)')$$
$$\hat{\imath}(H) = \hat{\imath}(H-v)+1 \le \hat{\imath}((H-v)')+1 = \hat{\imath}((H-v)'+v)$$

meaning that H is not critical in $(G_q)_q^\infty$.

$H-v$ is also not critical in $(R_q)_q^\infty$, and so there is some proper collapse $(H-v)''$ of $H-v$ satisfying $\hat{\imath}_R(H-v) \le \hat{\imath}_R((H-v)'')$, which implies that

$$\begin{aligned}\hat{\imath}_R(H) = \hat{\imath}_R(H-v)+1 \le \hat{\imath}_R((H-v)'')+1 &= 2|V((H-v)'')| - |E((H-v)'')| + 1 \\ &= 2(|V((H-v)''+v)|-1) - (|E((H-v)''+v)|-1)+1 \\ &= 2|V((H-v)''+v)| - |E((H-v)''+v)| = \hat{\imath}_R\left((H-v)''+v\right).\end{aligned}$$

But this means that H is not critical in $(R_q)_q^\infty$ either, so in particular it is, after all, not a counterexample to the proposition.

Case 2: $\delta(H)=2$. Let v be a vertex in H of degree 2.

Case 2.1: H is critical in $(G_q)_q^\infty$. Then

$$\hat{h}(H) = \hat{\imath}(H) \le \hat{\imath}(H-v) \tag{2}$$

by Proposition 4 and Corollary 7. Now we consider two cases depending on whether $H-v$ is critical in $(G_q)_q^\infty$ or not.

Case 2.1.1: $H-v$ is not critical in $(G_q)_q^\infty$. Then there exists a proper collapse H' of $H-v$ such that $\hat{\imath}(H') \ge \hat{\imath}(H-v)$ and moreover, we can choose it in such a way that the collapse itself is critical in $(G_q)_q^\infty$.

By definition of collapse, H' is the image of $H-v$ after a homomorphism. Let x and y be the neighbours of v in H, and let x' and y' be their images in H' (potentially we have $x'=y'$). Construct a graph $H'+w$ by adding a new vertex w, and attaching it to x' and y'. Then $H'+w$ is a collapse of H, obtained by extending the homomorphism $H-v \to H'$ to map v to w.

Then

$$\hat{h}(H'+w) \ge \hat{h}(H') \ge \hat{\imath}(H') \ge \hat{\imath}(H-v) \ge \hat{h}(H),$$

contradicting the criticality of H in $(G_q)_q^\infty$.

Case 2.1.2: $H-v$ is critical in $(G_q)_q^\infty$. Since H is a minimal counterexample to the proposition, we know that $H-v$ being critical in $(G_q)_q^\infty$ must mean that

$$\hat{\imath}(H-v) = \hat{\imath}_R(H-v).$$

$$\hat{h}_R(H) \ge \hat{h}_R(H-v) \ge \hat{\imath}_R(H-v) = \hat{\imath}(H-v) \ge \hat{h}(H).$$

Consider $H' \in \mathcal{C}(H)$ such that H' is critical in $(R_q)_q^\infty$ and

$$\hat{h}_R(H) = \hat{h}_R(H') = \hat{\imath}_R(H').$$

If H is not critical in $(R_q)_q^\infty$ (thus $H' \neq H$), by minimality of H

$$\hat{\imath}_R(H') = \hat{\imath}(H') \leq \hat{h}(H') < \hat{h}(H),$$

and hence

$$\hat{h}_R(H) < \hat{h}(H),$$

which is a contradiction. If H is critical in $(R_q)_q^\infty$ (thus $H = H'$)

$$\hat{h}_R(H) = \hat{\imath}_R(H) = \hat{\imath}_R(H - v)$$

$$\hat{\imath}_R(H - v) = \hat{\imath}_R(H) \neq \hat{\imath}(H) = \hat{\imath}(H - v).$$

which produces a contradiction.

Case 2.2: ***H* is critical in $(R_q)_q^\infty$, but not in $(G_q)_q^\infty$.** Then

$$\hat{h}_R(H) = \hat{\imath}_R(H) = \hat{\imath}_R(H - v).$$

Consider $H' \in \mathcal{C}(H)$, such that H' is critical in $(G_q)_q^\infty$. Then by the minimality of H

$$\hat{h}(H) = \hat{h}(H') = \hat{\imath}(H') = \hat{\imath}_R(H') < \hat{\imath}_R(H).$$

Now we consider two cases depending on whether $\hat{\imath}(H - v) = \hat{\imath}_R(H - v)$ or not.

Case 2.2.1: $\hat{\imath}(H - v) \neq \hat{\imath}_R(H - v)$**.** In this case $H - v$ is not critical in $(R_q)_q^\infty$. Thus there exists a collapse $H'' \in \mathcal{C}(H - v)$ such that

$$\hat{\imath}_R(H'') \geq \hat{\imath}_R(H - v).$$

It is then possible to add a vertex w with $\deg(w) \leq 2$, such that $H'' + w \in \mathcal{C}(H)$, i.e. is a collapse of H (see Case 2.1.1). Then

$$\hat{\imath}_R(H'' + w) \geq \hat{\imath}_R(H'') \geq \hat{\imath}_R(H - v) = \hat{\imath}_R(H),$$

contradicting the criticality of H.

Case 2.2.2: $\hat{\imath}(H - v) = \hat{\imath}_R(H - v)$**.** First observe that

$$\hat{\imath}(H - v) = \max_{u \in V_v(H)} \left\{ \hat{\imath}(H), \hat{\imath}(H_{v,u}) \right\},$$

where $V_v(H)$ is the bipartition class of H containing v, and we get $H_{v,u}$ from H by collapsing vertices v and u together. Since $\hat{\imath}(H - v) = \hat{\imath}_R(H - v) = \hat{\imath}_R(H) \neq \hat{\imath}(H)$, the maximum is not attained at $\hat{\imath}(H)$. Select a vertex u for which the maximum holds.

Note that

$$\hat{\imath}(H_{v,u}) = \hat{\imath}(H - v) = \hat{\imath}_R(H - v) = \hat{\imath}_R(H) \neq \hat{\imath}_R(H_{v,u}),$$

hence $H_{v,u}$ is not critical. Thus there exists $H'_{v,u} \in \mathcal{C}(H_{v,u})$ such that $H'_{v,u}$ is critical in $(G_q)_q^\infty$.

$$\hat{\imath}_R(H'_{v,u}) = \hat{\imath}(H'_{v,u}) \geq \hat{\imath}(H_{v,u}) = \hat{\imath}_R(H),$$

contradicting the criticality of H in $(R_q)_q^\infty$.

Case 3: $\delta(H) \geq 3$.

Case 3.1: H **is not 2-connected.** In this case H has either at least two connected components, or a cut-vertex. In either case, the graph H can be expressed as the disjoint union or the wedge sum of two graphs H_1 and H_2, each with fewer vertices than H. By Lemma 6, if $\hat{\imath}(H) \neq \hat{\imath}_R(H)$ then either $\hat{\imath}(H_1) \neq \hat{\imath}_R(H_1)$ or $\hat{\imath}(H_2) \neq \hat{\imath}_R(H_2)$. Without loss of generality assume the former.

By the induction hypothesis, H_1 is not critical in $(G_q)_q^\infty$ nor in $(R_q)_q^\infty$. This means that there exist graphs H_1' and H_2'' which are proper collapses of H_1 and which satisfy $\hat{\imath}(H_1) = \hat{\imath}(H_1')$ and $\hat{\imath}_R(H_1) = \hat{\imath}_R(H_1'')$. If H is the disjoint union of H_1 and H_2, let H' be the disjoint union of H_1' and H_2. Otherwise, let v_1 and v_2 be the vertex at which H_1 and H_2 are joined to form H. Let H' be formed by taking the wedge sum of H_1' at the homomorphic image of v_1 in H_1', and of H_2 at v_2. Define H'' analogously. These two graphs are proper collapses of H and, by Lemma 6, we have $\hat{\imath}(H) = \hat{\imath}(H')$ and $\hat{\imath}_R(H) = \hat{\imath}_R(H'')$, thus H is not critical in $(G_q)_q^\infty$ nor in $(R_q)_q^\infty$, so H is not a counterexample to our statement.

Case 3.2: H **is** 2**-connected.** In this case, by Lemma 13(a), we have $\hat{\imath}(H) \leq \frac{|V(H)|}{2} + 1$. Let $V(H) = X \cup Y$ be the bipartition of H. Assume without loss of generality that $|X| \geq |Y|$. Then $\hat{\imath}(H) \leq |X| + 1$.

The star $K_{1,|X|}$ is a proper collapse of H, obtained by mapping all vertices of Y into the center of the star and all vertices of X into the leaves. We also have that $\hat{\imath}(K_{1,|X|}) = |X| + 2 > \hat{\imath}(H)$, since it can be obtained from a single vertex by repeatedly attaching leaves, each of which increases $\hat{\imath}$ by 1 by Proposition 6. We deduce that H is not critical in $(G_q)_q^\infty$. It is also not critical in $(R_q)_q^\infty$, since $\hat{\imath}_R(H) = 2|V(H)| - |E(H)| \leq \frac{|V(H)|}{2} \leq |X| < \hat{\imath}(K_{1,|X|})$. Hence H cannot be a counterexample to our statement. □

We have now reduced our problem to proving that $\hat{\imath}(H) \leq \frac{|V(H)|}{2} + 1$ for 2-connected graphs with $\delta(H) \geq 3$. This is a substantial step down from our main theorem. Unfortunately this statement by itself is not easy to prove by induction, so just like in Proposition 10, we will use induction on a more general statement. In order to state it properly, we will need to define some concepts related to connectivity:

Definition. *Let H be a graph. A* bridge *is an edge $e = uv$ satisfying that, in the graph $H - e$, the vertices u and v lie in different components.*

Definition. *Let H be a graph. A* block *of H is an induced subgraph $B \subseteq H$ which is maximally 2-connected (that is, it is not a proper superset of another induced 2-connected subgraph of H) and it is not a bridge. This definition differs slightly from the more traditional definition, in which bridges are not excluded.*

This redefinition of block means that several classical results about them need to be rephrased. For example, every edge which is not a bridge is contained in exactly one block. Another result that can be rephrased is the existence of the block-cut tree:

Lemma 11 ([2], Chapter 4)**.** *Let H be a connected graph, let S_1, S_2 and S_3 be the sets of its blocks, bridges and cut vertices, resectively. Construct a graph G by taking $S_1 \cup S_2 \cup S_3$ as the vertex set, and connecting each cut vertex to every block and every bridge that contain that vertex. Then G is a tree.*

Definition. *Let H be a graph. A block B of H is a* pseudoleaf *if B there is at most one vertex v in B which has edges to vertices outside of S. If such a vertex exists, we call it the* linking vertex *of B.*

Corollary 12. *A graph H with $\delta(H) \geq 3$ and which contains a bridge has at least two pseudoleaves.*

Proof. We can assume that H is connected, otherwise restrict to the component of H that contains the bridge. In the block-cut tree T of H, consider a leaf v. This leaf does not correspond to a cut-vertex of F, since a cut-vertex is incident to at least two blocks or bridges (since it separates them). v cannot be a bridge either, since both endpoints of a bridge are cut vertices. v must therefore be a block, which is incident to at most one cut vertex. In other words, v correspond to a pseudoleaf. Use the fact that any tree that is not a single vertex contains at least two leaves to complete the proof. □

Lemma 13. *(a) Let H be a 2-connected graph with $\delta(H) \geq 2$, and let $v_2(H)$ be the number of vertices with degree 2. Then $\hat{i}(H) \leq \frac{|V(H)|+\max\{v_2(H),2\}}{2}$.*

(b) Let H be a connected graph with $\delta(H) \geq 3$, and let $\ell(H)$ be the number of pseudoleaves. Then $\hat{i}(H) \leq \frac{|V(H)|+\max\{\ell(H),2\}}{2}$

The statement that we wanted to prove, namely $\hat{i}(H) \leq \frac{|V(H)|}{2} + 1$ for 2-connected graphs with $\delta(H) \geq 3$, is implied by each of a and b.

Before we start the proof, let us briefly discuss the role of $\max\{v_2(H), 2\}$ and $\max\{\ell(H), 2\}$ in Lemma 13. For the former, let H be a 2-connected subgraph. Suppose that H is an induced subgraph of a strictly larger graph H' with $\delta(H') \geq 3$, and that H is not a pseudoleaf in H'. What is the least possible number of edges in H' that go from H to the rest of the graph? There must be at least two, since H is not a pseudoleaf. And since $\delta(H') \geq 3$, every vertex of degree 2 in H must be incident to an edge connecting it to $H' \setminus H$, giving a lower bound of $\max\{v_2(H), 2\}$.

Next let H be a connected graph with $\delta(H) \geq 3$, and suppose that H is an induced subgraph of a strictly larger 2-connected graph H'. What is the least possible number of edges in H' that go from H to the rest of the graph? There must be at least two, since H' is 2-connected. In addition, if H is not 2-connected and S is a pseudoleaf of H, then there must exist an edge of H' that connects S to $H' \setminus H$ that does not hit the linking vertex of S. Since every vertex that is not a linking vertex is contained in at most one pseudoleaf, there are at least $\ell(H)$ edges from H to $H' \setminus H$, giving a lower bound of $\max\{\ell(H), 2\}$.

On the other hand, both maxima are necessary: for example, the incidence graph P of the Pappus configuration is a 2-connected cubic graph with $|V(P)| = 18$ and $\hat{i}(P) = 10$.

Proof of Lemma 13. Suppose the statement is not true and let H be a minimal counterexample (minimal with respect to number of edges this time) to either (a) or (b). We will do a thorough case analysis to show that H cannot exist. Observe that if H is 2-connected and $\delta(H) \geq 3$ then both parts of the statement give the same bound, since $v_2(H) = 0$ and $\ell(H) = 1$.

Case 1: $\delta(H) = 2$. We produce a graph sequence $H_0, H_1, H_2, \ldots$ as follows: we start with $H_0 = H$. Given H_i, if it contains at least one vertex with degree at most 2, choose one of them as v_i and set $H_{i+1} = H_i - v_i$. Otherwise we stop the sequence. The last term of the sequence H_k is either empty or satisfies $\delta(H_k) \geq 3$. Note that since $\delta(H) = 2$ we have $k > 0$.

Let $W_1, W_2, \ldots, W_r$ be the components of H_k (where we potentially have $r = 0$). By Lemma 6, and by induction, we have $\hat{i}(H_k) = \sum_{i=1}^{r} \hat{i}(W_i) \leq \sum_{i=1}^{r} \frac{|V(W_i)|+\max\{\ell(W_i),2\}}{2}$. Since v_i has degree at most 2 in H_i, by Corollary 7 we have $\hat{i}(H_i) \leq \hat{i}(H_{i+1}) + 2 - d(v_i)$. This means that $\hat{i}(H_i) + 2|V(H_i)| - |E(H_i)|$ is non-decreasing on i. Thus, comparing this value at $i = 0$ and $i = k$:

$$\hat{i}(H) \leq \hat{i}(H_k) + 2k - |E(H) \setminus E(H_k)|.$$

Let us estimate the value of $|E(H)\setminus E(H_k)|$. Consider the graph W_0 formed by the edges in $E(H)\setminus E(H_k)$. As pointed out before the start of the proof of Lemma 13, the vertices on each W_i are incident to at least $\max\{\ell(W_i), 2\}$ edges in W_0. On the other hand, since all but $v_2(H)$ vertices in $\{v_0, v_1, \dots, v_{k-1}\}$ have degree at least 3 in H (and the remaining vertices have degree 2), the sum of their degrees in W_0 is at least $3k - v_2(H)$. Adding the degrees of all vertices, we have that $|E(W_0)| \geq \frac{3k-v_2(H)}{2} + \sum_{i=1}^{r} \frac{\max\{\ell(W_i),2\}}{2}$. We conclude that

$$\begin{aligned}\hat{\imath}(H) \leq & \hat{\imath}(H_k) + 2k - |E(H) \setminus E(H_k)| \\ \leq & \sum_{i=1}^{r} \frac{|V(W_i)| + \max\{\ell(W_i), 2\}}{2} + 2k - \frac{3k - v_2(H)}{2} - \sum_{i=1}^{r} \frac{\max\{\ell(W_i), 2\}}{2} \\ = & \frac{k + v_2(H)}{2} + \sum_{i=1}^{r} \frac{|V(W_i)|}{2} \leq \frac{|V(H)| + \max\{v_2(H), 2\}}{2},\end{aligned}$$

as we wanted to prove. Hence H is not a counterexample to the lemma.

Case 2: $\delta(H) \geq 3$, H is 2-connected. By Lemma 8 we can remove an e edge 2-connected graph. The number injections could not decrease as we romve e, the number of vertices does not change, and $v_2(H) = 0$, $v_2(H - e) \leq 2$ thus $\max\{v_2(\cdot), 2\}$ does not change. Thus if H is a counterexample, then so is $H - e$, which has less number of edges.

Case 3: $\delta(H) \geq 3$, H is not 2-connected. In this case H has at least one cut-vertex, meaning that all the blocks of H satisfy part (a). The next lemma in our proof states precisely that in these circumstances H satisfies (b):

Proposition 14. *Let H be a connected graph with $\delta(H) \geq 3$ which is not 2-connected. If every block W_i of H satisfies $\hat{\imath}(W_i) \leq \frac{|V(W_i)|+\max\{v_2(H),2\}}{2}$, then H satisfies $\hat{\imath}(H) \leq \frac{|V(H)|+\max\{\ell(H),2\}}{2}$.*

This means that H is not a counterexample to the lemma, completing the proof. □

The reason we state Lemma 14 as a standalone proposition is because the induction step here is incompatible with the one on Lemma 13: on some occasions we will actually want to increase the number of edges (this will happen when we add bridges between blocks). In order to properly define our induction we will need the following proposition:

Proposition 15. *There exists a function f such that every connected graph H with k blocks and $\delta(G) \geq 3$ has at most $f(k)$ bridges.*

Proof. Consider the block-cut tree T of H (the tree obtained in Lemma 11), and let S_1 be the set of vertices corresponding to blocks of H. Because $\delta(H) \geq 3$, every leaf of T is a block, and thus T contains at most k leaves. As a consequence, T contains at most $k - 2$ vertices of degree at most 3. Denote the set of these vertices as W.

Consider a bridge uv, and its corresponding vertex z in T. Because $\delta(H) \geq 3$, each of the vertices u and v (which are themselves cut vertices) is either contained in a block of H, or it is incident to at least three bridges, meaining that they are elements of W. In either case, the bridge uv is the unique bridge separating two elements of $S_1 \cup W$. From this we deduce that H contains at most $\binom{2k-2}{2}$ bridges. □

Proof of Proposition 14. Suppose that the statement is not true. Consider a counterexample H that minimizes the number of blocks, and among those counterexamples, select one that *maximizes* the number of bridges. This is well-defined because by Proposition 15. We will do a thorough case analysis to show that H cannot exist.

Case 1: H **contains a vertex** v **contained in at least two blocks.** We can express the graph v as the wedge sum of two graphs H_1 and H_2, each containing at least one of the blocks containing H. Then $\hat{\imath}(H) = \hat{\imath}(H_1) + \hat{\imath}(H_2) - 2$, by Lemma 6. Let v_1 and v_2 be the vertices of H_1 and H_2 corresponding to v, and let H' be the union of H_1 and H_2 connected by the edge v_1v_2.

By Lemma 6 we have $\hat{\imath}(H') = \hat{\imath}(H_1) + \hat{\imath}(H_2) - 1 = \hat{\imath}(H) + 1$. H' contains as many blocks as H, as many pseudoleaves as H, one more vertex and one more bridge. In addition we have $\delta(H') \geq 3$, since v_1 and v_2 have degree at least 2 in the respective graphs because they lie in a block. Thus $\hat{\imath}(H') \leq \frac{|V(H')|+\max\{\ell(H'),2\}}{2}$, and $\hat{\imath}(H) \leq \frac{|V(H')|+\max\{\ell(H'),2\}}{2} - 1 < \frac{|V(H)|+\max\{\ell(H),2\}}{2}$. Thus H cannot be a counterexample.

Case 2: H **contains a vertex** v **not contained in any block.** In this case, every edge incident to v is a bridge. Let $W_1, W_2, \ldots, W_k$ be the components of $H - v$, and observe that $k = d(v) \geq 3$. Applying Lemma 6 repeatedly, we obtained that $\hat{\imath}(H) = -(k-2) + \sum_{i=1}^{k} \hat{\imath}(W_i)$.

For any $1 \leq i, j \leq k$, let W_{ij} be the union of W_i and W_j with an edge between the corresponding neighbours of v. Every pseudoleaf of H is contained in some W_i, and it is a leaf of W_{ij} for all $j \neq i$, which means that $\ell(H) = \sum_{i=1}^{k} \ell(W_{i(i+1)})/2$, where the subindex $k+1$ is identified with 1. In addition, by Lemma 6, and using Corollary 12 we have $\hat{\imath}(W_i) + \hat{\imath}(W_j) - 1 = \hat{\imath}(W_{ij}) \leq \frac{|V(W_i)|+|V(W_j)|+\ell(W_{ij})}{2}$, so

$$\begin{aligned}
\hat{\imath}(H) =& \sum_{i=1}^{k} \hat{\imath}(W_i) - k + 2 \leq \sum_{i=1}^{k} \frac{\hat{\imath}(W_i) + \hat{\imath}(W_{i+1})}{2} - k + 2 = \sum_{i=1}^{k} \frac{\hat{\imath}(W_i) + \hat{\imath}(W_{i+1}) - 1}{2} - \frac{k-4}{2} \\
\leq& \sum_{i=1}^{k} \frac{|V(W_i)| + |V(W_{i+1})| + \ell(W_{i(i+1)})}{4} - \frac{k-4}{2} = \frac{|V(H)| - 1 + \ell(H)}{2} - \frac{k-4}{2} \\
\leq& \frac{|V(H)| + \ell(H)}{2},
\end{aligned}$$

since $k \geq 3$. Thus H cannot be a counterexample.

Case 3: Every vertex in H **is contained in exactly one block.** Let H' be the graph obtained from H by contracting every block of H. Because every vertex is contained in exactly one block of H, the vertices of H' are in bijection with the blocks of H. H' is connected and acyclic, so it is a tree.

Let $\mathcal{B}$ be the set of blocks of H and let $B \in \mathcal{B}$. Notice that $\delta(B) \geq 2$, and any vertex in B with degree 2 must be incident to a bridge in H, so $d(v_B) \geq v_2(B)$. In addition, if B is not a pseudoleaf then it must be incident to at least two bridges, so $d(v_B) \geq \max\{v_2(B), 2\}$. On the other hand, if B is a pseudoleaf then there is exactly one vertex of B which is incident to a bridge (and thus might have degree 2 in B). This means that $d(v_B) \geq \max\{v_2(B), 2\} - 1$. All together, that makes

$$2|\mathcal{B}| - 2 = 2e(H') = \sum_{B \in \mathcal{B}} d(v_B) \geq -\ell(H) + \sum_{B \in \mathcal{B}} \max\{v_2(B), 2\}.$$

By applying Lemma 6 repeatedly we see that $\hat{\imath}(H) = \sum_{B \in \mathcal{B}} \hat{\imath}(B) - |\mathcal{B}| + 1$, as every bridge decreases $\hat{\imath}$ by one. Since each block satisfies $\hat{\imath}(B) \leq \frac{|V(B)|+\max\{v_2(B),2\}}{2}$, we can put this together:

$$\hat{\imath}(H) \leq \sum_{B \in \mathcal{B}} \frac{|V(B)| + \max\{v_2(B), 2\}}{2} - |\mathcal{B}| + 1 \leq \frac{|V(H)| + \ell(H)}{2}.$$

Thus H cannot be a counterexample, and the proof is complete.

□

4 Concluding remarks

In reality we don't need the fact that these projective planes are field planes. We only need the following conditions (here P_q is the projective plane, G_q is its incidence graph):

- P_q is isomorphic to its dual.
- P_q is transitive. That is, the automorphism group of P_q is transitive.
- The size of P_q tends to infinity.

This shows how the statement does not care about the prime divisor of the prime power, the proof only uses the basic geometric properties of the planes.

In the classical notion of graph convergence, the Chung-Graham-Wilson theorem establishes the equivalence of several conditions on graph sequences, which are collectively referred to as "quasirandomness". One of those conditions is convergence to the constant graphon, i.e., having the same limit as $G(n,p)$. We would be interested in finding analogous of some of these conditions. In particular, we would be interested in two types of conditions:

Question 16. *Does there exist a finite family $\mathcal{F} \subseteq \mathcal{B}_0$, such that for every infinite sequence $\{G_n\}_{n=1}^{\infty}$ of graphs in $\mathcal{B}_0$, if*

$$\lim_{n\to\infty} \frac{\log \mathsf{hom}(H, G_n)}{-\log t_B(K_2, G_n)} = \hat{h}_R(H)$$

holds for every $H \in \mathcal{F}$, then it also holds for every $H \in \mathcal{B}_0$?

One of the implications of the Chung-Graham-Wilson theorem is that $t(K_2, G_n)$ and $t(C_4, G_n)$ are enough to determine quasirandomness. In particular, if $t(K_2, G_n)$ converges and $\lim\limits_{n\to\infty} \frac{-\log t(C_4,G_n)}{-\log t(K_2,G_n)} = 4$, then $\{G_n\}_{n=1}^{\infty}$ is quasirandom. Szegedy asked whether a similar limit, with t_B replacing t, is enough to guarantee that the sequence $\{G_n\}_{n=1}^{\infty}$ converges to a linear combination of certain random graphs.

If this was the case, we could say that the sequence of all incidence graphs of projective planes (not necessarily standard) is quasirandom. This is because, if G is the incidence graph of a projective plane of order q (where q might not necessarily be restricted to prime powers), then $\mathsf{hom}(C_4, G) = \Theta(q^4)$, with the image of each homomorphism being either a single edge or a cherry.

Another, more restrictive condition, involves the spectrum of the adjacency matrix of the graphs. The eigenvalues of the incidence graph of a projective plane depend only on its order q, and can be computed explicitly. If the log-convergence of a sequence of graphs to the quasirandom limit depends only on its spectrum, as is the case in the classical notion of convergence, then that would also imply the log-convergence of all incidence graphs of projective planes.

Question 17. *Given an infinite sequence $\{G_n\}_{n=1}^{\infty}$ of graphs in $\mathcal{B}_0$, is it possible to determine whether it converges to the same limit as $\{R_n\}_{n=1}^{\infty}$ just by knowing the spectrum of the graphs G_n?*

Finally, we return to considering only standard projective planes. As we saw in Section 2, the sequence $\log_q \mathsf{inj}(G_2, G_q)$ does not converge, as it is equal to $-\infty$ when q is an odd prime power and non-negative when q is a power of 2. Could something similar happen if q is restricted to only take prime values? Remember that this restriction was included in Szegedy's question about log-convergence of projective planes, even though it turned out to be unnecessary. Could this restriction be enough to guarantee convergence on the number of injective homomorphisms?

Question 18. *If q ranges over the prime numbers, does $\log_q \mathsf{inj}(H, G_q)$ converge for all $H \in \mathcal{B}_0$?*

References

[1] Gary Chartrand, Agnis Kaugars, and Don R Lick. Critically n-connected graphs. *Proceedings of the American Mathematical Society*, 32(1):63–68, 1972.

[2] Frank Harary. *Graph Theory*. Addison-Wesley Publishing Company, 1969.

[3] Alexander Sidorenko. A correlation inequality for bipartite graphs. *Graphs and Combinatorics*, 9:201–204, 06 1993.

[4] Balázs Szegedy. Sparse graph limits, entropy maximization and transitive graphs. *arXiv: Combinatorics 1504.00858*, 2015.

Appendix E

Isometric rigidity of Wasserstein spaces over Euclidean spheres

by György Pál Gehér, Aranka Hrušková, Tamás Titkos, Dániel Virosztek

ISOMETRIC RIGIDITY OF WASSERSTEIN SPACES OVER EUCLIDEAN SPHERES

GYÖRGY PÁL GEHÉR, ARANKA HRUŠKOVÁ, TAMÁS TITKOS, AND DÁNIEL VIROSZTEK

Abstract. We study the structure of isometries of the quadratic Wasserstein space $\mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ over the sphere endowed with the distance inherited from the norm of $\mathbb{R}^{n+1}$. We prove that $\mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ is isometrically rigid, meaning that its isometry group is isomorphic to that of $\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$. This is in striking contrast to the non-rigidity of its ambient space $\mathcal{W}_2\left(\mathbb{R}^{n+1}, \varrho_{\|\cdot\|}\right)$ but in line with the rigidity of the geodesic space $\mathcal{W}_2\left(\mathbb{S}^n, \sphericalangle\right)$. One of the key steps of the proof is the use of mean squared error functions to mimic displacement interpolation in $\mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$. A major difficulty in proving rigidity for quadratic Wasserstein spaces is that one cannot use the Wasserstein potential technique. To illustrate its general power, we use it to prove isometric rigidity of $\mathcal{W}_p\left(\mathbb{S}^1, \varrho_{\|\cdot\|}\right)$ for $1 \le p < 2$.

Contents

1. Motivation and main result

In recent years, there has been considerable activity in characterising isometries of various metric spaces of measures. See e.g. [2–4,6,7,9–21,24,27] for results about the total variation, Lévy, Kuiper, Lévy-Prokhorov, Kolmogorov-Smirnov, and Wasserstein metrics. Among these, an interesting result is due to Kloeckner. In [18, Theorem 1.1 and Theorem 1.2], he shows that the quadratic Wasserstein space $\mathcal{W}_2\left(\mathbb{R}^{n+1}, \varrho_{\|\cdot\|}\right)$, where $\varrho_{\|\cdot\|}(x,y) = \|x-y\|$ is the metric induced by the norm, exhibits the rare phenomenon of not being isometrically rigid, meaning that not all isometries of $\mathcal{W}_2\left(\mathbb{R}^{n+1}, \varrho_{\|\cdot\|}\right)$ are induced by an isometry of $\left(\mathbb{R}^{n+1}, \varrho_{\|\cdot\|}\right)$. In this paper, we consider the metric subspace $\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ of the base space $\left(\mathbb{R}^{n+1}, \varrho_{\|\cdot\|}\right)$ and prove that the non-rigidity does not carry over: the exotic isometries

2020 *Mathematics Subject Classification.* Primary: 54E40; 46E27 Secondary: 60A10; 60B05.
Key words and phrases. Wasserstein space, isometric rigidity.
Gehér was supported by the Leverhulme Trust Early Career Fellowship (ECF-2018-125), and also by the Hungarian National Research, Development and Innovation Office (Grant no. K115383); Hrušková was supported by the DYNASNET ERC Synergy grant, agreement ID 810115 and by the KKP 139502 project; Titkos was supported by the Hungarian National Research, Development and Innovation Office - NKFIH (grant no. K115383) and by the Momentum Program of the Hungarian Academy of Sciences (grant no. LP2021-15/2021); Virosztek was supported by the Momentum Program of the Hungarian Academy of Sciences (grant no. LP2021-15/2021) and partially supported by the ERC Consolidator Grant no. 772466.

of $\mathcal{W}_2\left(\mathbb{R}^{n+1}, \varrho_{\|\cdot\|}\right)$ send measures supported on $\mathbb{S}^n$ to measures supported also outside of $\mathbb{S}^n$, while we gain no new exotic isometries by restricting to this smaller metric space. In general, when H is an arbitrary Borel subset of $\mathbb{R}^{n+1}$ then $\mathcal{W}_p(H, \varrho_{\|\cdot\|})$ embeds isometrically into $\mathcal{W}_p\left(\mathbb{R}^{n+1}, \varrho_{\|\cdot\|}\right)$, but this does not necessarily imply that there exists such a natural embedding for their isometry groups. To see an example, we mention the case of the real line $(\mathbb{R}, |\cdot|)$ with the subset $H = [0,1]$ (see [12, Theorem 2.5 and Theorem 3.7] for details): the isometry group of $\mathcal{W}_1([0,1], |\cdot|)$ is the Klein group, which cannot be embedded by a group homomorphism into the isometry group of $\mathcal{W}_1(\mathbb{R}, |\cdot|)$, which is isomorphic to the isometry group of the real line.

Finally, we draw attention to Santos-Rodríguez's paper [24] and our recent work [15]. In [24], the author considers (among others) Wasserstein spaces with $p > 1$ whose underlying metric space is a rank-one symmetric space, which class contains the sphere $\mathbb{S}^n$ with the spherical distance $\sphericalangle$, while in [15], we considered finite-dimensional tori and spheres with their geodesic distances for all parameters $p \geq 1$. Together, these two papers show that $\mathcal{W}_p\left(\mathbb{S}^n, \sphericalangle\right)$ is isometrically rigid for all $p \geq 1$. As explained above, in this paper, we replace the angular distance $\sphericalangle$ with another natural metric: the distance $\varrho_{\|\cdot\|}$ inherited from the norm of $\mathbb{R}^{n+1}$. We focus on the case of $p = 2$ because this is the only parameter value for which the ambient space $\mathcal{W}_p\left(\mathbb{R}^{n+1}, \varrho_{\|\cdot\|}\right)$ is *not* rigid. We expect that for $p \neq 2$, techniques similar to the ones used in [14] would lead to a proof of isometric rigidity. The situation is analogous to the case of the real line and the unit interval: the quadratic Wasserstein space is not rigid over $\mathbb{R}$ but it is rigid over the compact subset $[0,1]$, see [18, Theorem 1.1] and [12, Theorem 2.6]. Our main result reads as follows.

Theorem 1.1. *For all $n \in \mathbb{N}$, the quadratic Wasserstein space $\mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ is isometrically rigid. That is, for any isometry $\Psi\colon \mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right) \to \mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$, there exists an isometry $\psi\colon \mathbb{S}^n \to \mathbb{S}^n$ such that $\Psi = \psi_\#$.*

In our recent works [14, 15], recovering measures from their Wasserstein potentials — see (2.3) for precise definition — turned out to be a powerful method to prove isometric rigidity. However, this method cannot be used in the case of $\mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$, as shown by the following simple example. Let δ_x denote the Dirac measure concentrated at $x \in \mathbb{S}^n$, let $\mu_z := \frac{1}{2}(\delta_z + \delta_{-z})$ for $z \in \mathbb{S}^n$, and note that for any $x \in \mathbb{S}^n$ we have $d^2_{\mathcal{W}_2}(\mu_z, \delta_x) = \frac{1}{2}(\|x-z\|^2 + \|x+z\|^2) = 2$ independently of x and z — see (2.1) for the precise definition of the p-Wasserstein distance $d_{\mathcal{W}_p}$. This means that every element of the set $\{\mu_z \,|\, z \in \mathbb{S}^n\}$ has the same Wasserstein potential function, and hence potentials do not determine measures uniquely in general.

Our complimentary result Theorem 3.1 demonstrates sensitivity of the Wasserstein potential method to the parameter value p. Namely, we show that, at least in the special case of $\mathbb{S}^1$, measures are uniquely determined by their potentials if $1 \leq p < 2$, and hence $\mathcal{W}_p\left(\mathbb{S}^1, \varrho_{\|\cdot\|}\right)$ is isometrically rigid.

2. The Wasserstein space $\mathcal{W}_p\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ and the Wasserstein potential

In this section, we recall all the necessary notions and notations. Let (Y, m) be a complete and separable metric space, $p \geq 1$ a fixed real number, and $\mathcal{P}(Y)$ the set of all Borel probability measures on Y. The p-Wasserstein space $\mathcal{W}_p(Y, m)$, where $p \in [1, \infty)$, is then defined as the set

$$\left\{\mu \in \mathcal{P}(Y) \,\middle|\, \exists \hat{y} \in Y : \int_Y m(y, \hat{y})^p \, \mathrm{d}\mu(y) < \infty \right\}$$

of probability measures endowed with the *p-Wasserstein metric*

$$d_{\mathcal{W}_p}(\mu, \nu) := \left(\inf_{\pi \in \Pi(\mu,\nu)} \iint_{Y \times Y} m(x, y)^p \, \mathrm{d}\pi(x, y) \right)^{1/p}, \tag{2.1}$$

where the infimum is taken over the set $\Pi(\mu, \nu)$ of all couplings of μ and ν. A Borel probability measure π on $Y \times Y$ is called a *coupling* of μ and ν if $\pi(A \times Y) = \mu(A)$ and $\pi(Y \times B) = \nu(B)$ for all Borel sets $A, B \subseteq Y$. For more details about Wasserstein spaces, we refer the reader to the comprehensive textbooks [1,8,23,26]. Now we only mention that optimal couplings always exist, and the infimum in (2.1) becomes minimum [1, Theorem 1.5]. Furthermore, finitely supported measures are dense in Wasserstein spaces, see, e.g., [26, Theorem 6.18].

An *isometric embedding* between metric spaces (X, d) and (Y, m) is a map $\phi\colon (X, d) \to (Y, m)$ which preserves distances, i.e., a map such that $d(x, x') = m(\phi(x), \phi(x'))$ for all $x, x' \in X$. We shall use the term *isometry* for a surjective isometric embedding from a metric space onto itself. It is important to note that if (X, d) is a compact metric space then every isometric embedding $\phi\colon (X, d) \to (X, d)$ is surjective and hence an isometry [5, Theorem 1.6.14].

For a Borel-measurable map $\psi\colon Y \to Y$, its push-forward $\psi_\#\colon \mathcal{W}_p(Y, m) \to \mathcal{W}_p(Y, m)$ is defined by $\big(\psi_\#(\mu)\big)(A) := \mu(\psi^{-1}[A])$, where $A \subseteq Y$ is a Borel set and $\psi^{-1}[A] = \{x \in X \,|\, \psi(x) \in A\}$. In particular, when $\psi\colon Y \to Y$ is an isometry then so is $\psi_\#$ by the very definition of the Wasserstein distance, giving rise to a canonical embedding of the isometries of (Y, m) to the isometries of $\mathcal{W}_p(Y, m)$.

In this paper, we consider the compact metric space $\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$, where $\mathbb{S}^n$ is the unit sphere of $\mathbb{R}^{n+1}$

$$\mathbb{S}^n := \{x \in \mathbb{R}^{n+1} \colon \|x\| = 1\}$$

and $\varrho_{\|\cdot\|}\colon \mathbb{S}^n \times \mathbb{S}^n \to [0, 2]$ is the metric inherited from the euclidean norm of $\mathbb{R}^{n+1}$, i.e., $\varrho_{\|\cdot\|}(x, y) = \|x - y\|$ for all $x, y \in \mathbb{S}^n$. The point $-x$ is called the antipodal of x. Since $\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ is bounded, the Wasserstein space $\mathcal{W}_p\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ is the entire set $\mathcal{P}(\mathbb{S}^n)$ endowed with the distance

$$d_{\mathcal{W}_p}(\mu, \nu) := \left(\inf_{\pi \in \Pi(\mu,\nu)} \iint_{\mathbb{S}^n \times \mathbb{S}^n} \|x - y\|^p \, \mathrm{d}\pi(x, y) \right)^{1/p}. \tag{2.2}$$

We write $\mathcal{W}_p\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ instead of the usual $\mathcal{W}_p(\mathbb{S}^n)$ notation to avoid any confusion with the results in [15,24]. As the Wasserstein distance metrizes the weak convergence of probability measures over bounded metric spaces (see, e.g., [25, Theorem 7.12]), by Prokhorov's theorem, $\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ being compact tells us that $\mathcal{W}_p(\mathbb{S}^n, \varrho_{\|\cdot\|})$ is compact too — see also Remark 6.19 in [26]. This implies that every isometric embedding of $\mathcal{W}_p\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ into itself is an isometry.

For a measure $\mu \in \mathcal{P}(\mathbb{S}^n)$, its support $\mathrm{supp}(\mu)$ is the set of all points $x \in \mathbb{S}^n$ for which every open neighbourhood of x has positive measure. As usual, δ_x denotes the Dirac measure supported on the single point $x \in \mathbb{S}^n$.

The question arises whether it is possible to identify a measure if we know its distance from all Dirac measures. (Recall that $d_{\mathcal{W}_p}(\delta_x, \delta_y) = \|x-y\|$ for all $x, y \in \mathbb{S}^n$ and thus the set of all Dirac measures is an isometric copy of the underlying metric space.) To answer this question, we first introduce the notion of *Wasserstein potential* $\mathcal{T}_\mu^{(p)}$. For a given $\mu \in \mathcal{W}_p\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$, the Wasserstein potential is the function

$$\mathcal{T}_\mu^{(p)} \colon \mathbb{S}^n \to \mathbb{R}; \qquad \mathcal{T}_\mu^{(p)}(x) := d_{\mathcal{W}_p}^p(\delta_x, \mu) = \int_{\mathbb{S}^n} \|x-y\|^p \, \mathrm{d}\mu(y). \tag{2.3}$$

Now, the question above can be rephrased as follows: *does the Wasserstein potential determine the measure uniquely?*

3. Does the Wasserstein potential determine the measure uniquely?

The answer to this question is no, in general. A prominent example is $\mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ where measures supported on antipodal points with both weights equal to $\frac{1}{2}$ have the same (constant) potential function — see the example in Section 1, after Theorem 1.1. Beyond this, exotic isometries of $\mathcal{W}_2(\mathbb{R}, |\cdot|)$ (see [18, Section 5.1 and Section 5.2]) are also counterexamples. For the reader's convenience, we briefly recall some elements of Kloeckner's surprising result. Let us introduce the subsets

$$\Delta_1(\mathbb{R}) := \{\delta_x \,|\, x \in \mathbb{R}\}, \qquad \Delta_2(\mathbb{R}) := \{\lambda\delta_x + (1-\lambda)\delta_y \,|\, x, y \in \mathbb{R},\ \lambda \in [0,1]\} \tag{3.1}$$

of $\mathcal{P}(\mathbb{R})$. One can show that if Φ is an isometry of $\mathcal{W}_2(\mathbb{R}, |\cdot|)$, then Φ maps the set $\Delta_2(\mathbb{R})$ onto itself. Kloeckner made use of the fact that each element of $\Delta_2'(\mathbb{R}) := \Delta_2(\mathbb{R}) \setminus \Delta_1(\mathbb{R})$ can be written uniquely as

$$\mu(m, \sigma, r) := \frac{e^{-r}}{e^{-r}+e^r}\delta_{m-\sigma e^r} + \frac{e^r}{e^{-r}+e^r}\delta_{m+\sigma e^{-r}}, \tag{3.2}$$

for $m, r \in \mathbb{R}$, $\sigma \in \mathbb{R}\setminus\{0\}$ and showed that $d_{\mathcal{W}_2}^2\Big(\mu(m_1,\sigma_1,r_1), \mu(m_2,\sigma_2,r_2)\Big) = |m_1-m_2|^2 + \sigma_1^2 + \sigma_2^2 + 2\sigma_1\sigma_2 e^{-|r_1-r_2|}$. This identity implies that the map $\Phi_\psi \colon \mu(m,\sigma,r) \mapsto \mu(m,\sigma,\psi(r))$ is a distance-preserving bijection on $\Delta_2'(\mathbb{R})$ for any isometry $\psi \colon \mathbb{R} \to \mathbb{R}$. Kloeckner proved that Φ_ψ can be uniquely extended into an isometry of the whole space $\mathcal{W}_2(\mathbb{R}, |\cdot|)$, and this extension, denoted by $\widetilde{\Phi}_\psi$, leaves all Dirac measures fixed. If ψ is not the identity, then $\widetilde{\Phi}_\psi \neq \mathrm{id}_{\mathcal{W}_2(\mathbb{R},|\cdot|)}$, and therefore we can find a measure μ, such that $\mu \neq \widetilde{\Phi}_\psi(\mu) =: \nu$. But for these two measures, we have

$$\mathcal{T}_\mu^{(2)}(s) = d_{\mathcal{W}_2}^2(\delta_s, \mu) = d_{\mathcal{W}_2}^2\left(\widetilde{\Phi}_\psi(\delta_s), \widetilde{\Phi}_\psi(\mu)\right) = d_{\mathcal{W}_2}^2(\delta_s, \nu) = \mathcal{T}_\nu^{(2)}(s) \qquad \text{for all} \quad s \in \mathbb{R}.$$

This shows that indeed, the Wasserstein potential does not always determine the measure uniquely.

However, we will now prove that it does in the case of $\mathbb{S}^1 \simeq \mathbb{T} = \{z \in \mathbb{C} : |z| = 1\}$ equipped with the distance function $r(z,\omega) = \left|\frac{1}{2}(z-\omega)\right|$ for $1 \le p < 2$. This normalization of the distance is consistent with the one used in [27]. In this section, we assume that $1 \leq p < 2$, and we recall that the p-Wasserstein distance of $\mu, \nu \in \mathcal{W}_p(\mathbb{T})$ in this case is

$$d_{\mathcal{W}_p}(\mu, \nu) = \left(\inf_{\pi \in \Pi(\mu,\nu)} \iint\limits_{\mathbb{T}\times\mathbb{T}} \left| \frac{1}{2}(z - \omega) \right|^p \, \mathrm{d}\pi(z, \omega) \right)^{1/p},$$

and therefore for any $z \in \mathbb{T}$, the Wasserstein potential is of the form

$$\mathcal{T}_\mu^{(p)}(z) = d^p_{\mathcal{W}_p}(\delta_z, \mu) = \int\limits_{\mathbb{T}} \left| \frac{1}{2}(z - \omega) \right|^p \, \mathrm{d}\mu(\omega). \tag{3.3}$$

We showed in [15] that Fourier analytic methods can sometimes solve the problem of rigidity in a very elegant way. For example, we showed that isometric rigidity of $\mathcal{W}_2(\mathbb{T}, \sphericalangle)$ can be proved by using the Fourier transform of the Wasserstein potential, however, the same method fails in the case $\mathcal{W}_1(\mathbb{T}, \sphericalangle)$. As we will see, if we endow $\mathbb{T}$ with the distance $r(z, \omega) = \left|\frac{1}{2}(z - \omega)\right|$ then the situation changes: the same method works to prove isometric rigidity of $\mathcal{W}_1(\mathbb{T}, r)$, but fails in the case $\mathcal{W}_2(\mathbb{T}, r)$.

Now we recall the very basics of Fourier analysis on the abelian group $\mathbb{T}$. The main reason for doing so is to fix the notation. The continuous *characters* of $\mathbb{T}$ are exactly the power functions with an integer exponent. That is, if $\varphi_k(z) = z^k$ for all $k \in \mathbb{Z}$ and Γ is the *dual group* (i.e., the group of all continuous characters), then $\Gamma = \{\varphi_k \colon \mathbb{T} \to \mathbb{C} \,|\, k \in \mathbb{Z}\}$, and $\Gamma \cong \mathbb{Z}$. The group $\mathbb{T}$ is compact, hence it admits a unique *Haar* probability measure λ, which can be expressed explicitly as

$$\mathrm{d}\lambda(z) = \frac{\mathrm{d}z}{2\pi i z}.$$

The *Fourier transform* of a (complex-valued) function $f \in L^1(\mathbb{T}, \lambda)$ is defined by

$$\hat{f}(k) = \int\limits_{\mathbb{T}} f \overline{\varphi_k} \mathrm{d}\lambda = \frac{1}{2\pi i} \int\limits_{\mathbb{T}} f(z) z^{-(k+1)} \mathrm{d}z \qquad (k \in \mathbb{Z}).$$

Let us denote the set of all (complex-valued) measures of finite total variation by $\mathbb{M}(\mathbb{T})$. The Fourier transform of $\mu \in \mathbb{M}(\mathbb{T})$ is defined by

$$\hat{\mu}(k) = \int\limits_{\mathbb{T}} \overline{\varphi_n} \mathrm{d}\mu = \int\limits_{\mathbb{T}} z^{-k} \, \mathrm{d}\mu(z) \qquad (k \in \mathbb{Z}).$$

We note that L^1 functions can be naturally identified with absolutely continuous measures (with respect to the Haar measure), see [22, Subsection 1.3.4.]. The convolution of L^1 functions f and g is defined by

$$(f * g)(z) = \int\limits_{\mathbb{T}} f\left(z\omega^{-1}\right) g(\omega) \, \mathrm{d}\lambda(\omega) \tag{3.4}$$

and the convolution of $f \in L^1(\mathbb{T}, \lambda)$ and $\mu \in \mathbb{M}(\mathbb{T})$ is defined by

$$(f * \mu)(z) = \int\limits_{\mathbb{T}} f\left(z\omega^{-1}\right) \, \mathrm{d}\mu(\omega). \tag{3.5}$$

It is a key identity that the Fourier transform factorizes the convolution, that is,

$$\widehat{f * \nu} = \hat{f} \cdot \hat{\nu}. \tag{3.6}$$

Now we are ready to state and prove the main result of this section. It says that if $1 \leq p < 2$ then the Wasserstein space $\mathcal{W}_p(\mathbb{T}, r)$ is isometrically rigid.

Theorem 3.1. *Let $p \in [1,2)$ be a real number and let $\Psi\colon \mathcal{W}_p(\mathbb{T}, r) \to \mathcal{W}_p(\mathbb{T}, r)$ be an isometry. Then there exists an isometry $\psi\colon (\mathbb{T}, r) \to (\mathbb{T}, r)$ such that $\Psi = \psi_\#$.*

Proof. First observe that the diameter of $\mathcal{W}_p(\mathbb{T}, r)$ is 1, and $d_{\mathcal{W}_p}(\mu,\nu) = 1$ if and only if $\mu = \delta_x$ and $\nu = \delta_{-x}$ for some $x \in \mathbb{T}$. Since Ψ is an isometry, we have

$$1 = d_{\mathcal{W}_p}(\delta_x, \delta_{-x}) = d_{\mathcal{W}_p}\left(\Psi(\delta_x), \Psi(\delta_{-x})\right) \tag{3.7}$$

for all $x \in \mathbb{T}$, which implies that $\Psi(\delta_x)$ is a Dirac measure as well.

Let us define the map $\psi\colon \mathbb{T} \to \mathbb{T}$ via the identity $\Psi(\delta_x) = \delta_{\psi(x)}$ – this means that Ψ coincides with $\psi_\#$ on the set of Dirac measures. The map $\psi\colon (\mathbb{T}, r) \to (\mathbb{T}, r)$ is in fact an isometry:

$$r\left(\psi(x), \psi(y)\right) = d_{\mathcal{W}_p}\left(\delta_{\psi(x)}, \delta_{\psi(y)}\right) = d_{\mathcal{W}_p}\left(\Psi(\delta_x), \Psi(\delta_y)\right) = d_{\mathcal{W}_p}\left(\delta_x, \delta_y\right) = r(x,y)$$

for all $x, y \in \mathbb{T}$, and $(\mathbb{T}, r)$ is compact. These together combine into that $\left(\psi^{-1}\right)_\# \circ \Psi$ is an isometry which fixes all Dirac measures. If we now prove that any isometry of $\mathcal{W}_p(\mathbb{T}, r)$ which fixes all Dirac measures must be the identity, we are done: in that case, $\left(\psi^{-1}\right)_\# \circ \Psi = \mathrm{id}_{\mathcal{W}_p(\mathbb{T},r)}$, i.e., $\Psi = \psi_\#$ as claimed.

From now on, let us assume that $\Phi\colon \mathcal{W}_p(\mathbb{T}, r) \to \mathcal{W}_p(\mathbb{T}, r)$ is an isometry such that $\Phi(\delta_z) = \delta_z$ for all $z \in \mathbb{T}$. Then we have

$$\mathcal{T}_\mu^{(p)}(z) = d_{\mathcal{W}_p}^p\left(\delta_z, \mu\right) = d_{\mathcal{W}_p}^p\left(\Phi(\delta_z), \Phi(\mu)\right) = d_{\mathcal{W}_p}^p\left(\delta_z, \Phi\left(\mu\right)\right) = \mathcal{T}_{\Phi(\mu)}^p(z)$$

for all $z \in \mathbb{T}$ and $\mu \in \mathcal{W}_p(\mathbb{T}, r)$. The question is whether this implies $\mu = \Phi(\mu)$. The proof will be done once we prove that a measure $\mu \in \mathcal{W}_p(\mathbb{T}, r)$ is uniquely determined by its Wasserstein potential. To this end, assume that μ and ν are two measures such that

$$\mathcal{T}_\mu^{(p)}(z) = \mathcal{T}_\nu^p(z) \quad \text{for all} \quad z \in \mathbb{T}. \tag{3.8}$$

We need to show that (3.8) implies $\mu = \nu$. Let us introduce the map

$$f_p(z) := \left|\frac{1}{2}(z-1)\right|^p. \tag{3.9}$$

Then by (3.3) and (3.5) one can observe that $\mathcal{T}_\mu^{(p)}(z) = (f_p * \mu)(z)$ holds for all $z \in \mathbb{T}$ and $\mu \in \mathcal{W}_p(\mathbb{T})$. Indeed, we have

$$\begin{aligned}\mathcal{T}_\mu^{(p)}(z) = d_{\mathcal{W}_p}^p\left(\delta_z, \mu\right) &= \int_{\mathbb{T}} \left|\frac{1}{2}(z-\omega)\right|^p \,\mathrm{d}\mu\left(\omega\right) \\ &= \int_{\mathbb{T}} \left|\frac{1}{2}\left(z\omega^{-1}-1\right)\right|^p \,\mathrm{d}\mu\left(\omega\right) = \int_{\mathbb{T}} f_p\left(z\omega^{-1}\right) \,\mathrm{d}\mu\left(\omega\right) = \left(f_p * \mu\right)(z).\end{aligned} \tag{3.10}$$

The key observation is that the Fourier transform of f_p does not vanish anywhere, that is, $\hat{f}_p(n) \neq 0$ for all $n \in \mathbb{Z}$. For $n = 0$, we have

$$\hat{f}_p(0) = \int_{\mathbb{T}} \left|\frac{1}{2}(z-1)\right|^p \,\mathrm{d}\lambda(z) > 0,$$

while for $n \neq 0$, we use that

$$f_p(z) = \left(\left|\frac{1}{2}(z-1)\right|^2\right)^{\frac{p}{2}} = \left(\frac{1}{4}\left(2 - z - z^{-1}\right)\right)^{\frac{p}{2}} = \left(1 - \frac{1}{4}\left(2 + z + z^{-1}\right)\right)^{\frac{p}{2}},$$

and by the binomial series expansion we get that

$$f_p(z) = \sum_{k=0}^{\infty} \binom{\frac{p}{2}}{k} \left(\frac{-1}{4}\left(2 + z + z^{-1}\right)\right)^k, \tag{3.11}$$

where $\binom{\frac{p}{2}}{0} = 1$ and $\binom{\frac{p}{2}}{k} = \frac{\prod_{j=0}^{k-1}\left(\frac{p}{2}-j\right)}{k!}$. Using that the sign of $\binom{\frac{p}{2}}{k}(-1)^k$ is negative for all $k \geq 1$, equality (3.11) can be written as

$$f_p(z) = 1 - \left\{\frac{p}{2} \cdot \frac{2 + z + z^{-1}}{4} + \sum_{k=2}^{\infty} \left(\frac{\frac{p}{2}\prod_{j=1}^{k-1}\left(j - \frac{p}{2}\right)}{k!}\left(\frac{2 + z + z^{-1}}{4}\right)^k\right)\right\}. \tag{3.12}$$

It is a useful feature of the group $\mathbb{T}$ that the Fourier series of a function coincides with its power series. Therefore, the above binomial expansion gives us useful information about $\hat{f}_p$, namely, $\hat{f}_p(k)$ coincides with the coefficient of z^k in the expansion (3.11).

Let us note that for $n \neq 0$, the coefficient of z^n must be strictly negative because the expressions $\frac{p}{2}, 1 - \frac{p}{2}, 2 - \frac{p}{2}, \ldots$ are all positive – here we use the assumption that $p < 2$. So we obtained that $\hat{f}_p(0) > 0$ and $\hat{f}_p(n) < 0$ for $n \neq 0$ which means that $\hat{f}_p(n) \neq 0$ for all $n \in \mathbb{Z}$.

By (3.10), the assumption that $\mathcal{T}_\mu^{(p)}(z) = \mathcal{T}_\nu^{p}(z)$ for all $z \in \mathbb{T}$ implies that $f_p * \mu = f_p * \nu$. By (3.6), this means that $\hat{f}_p \cdot \hat{\mu} = \hat{f}_p \cdot \hat{\nu}$. Since $\hat{f}_p(n) \neq 0$ for every n, we can deduce that $\hat{\mu} = \hat{\nu}$, but the Fourier transform completely determines the measure [22, Chapter 1], hence $\mu = \nu$, and the proof is done.

□

4. Isometric rigidity of $\mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ — the proof of Theorem 1.1

The assumption $p < 2$ was crucial in the previous section, and therefore the quadratic case cannot be handled with the same Fourier analytic technique. In this section, we use a method that allows us to prove isometric rigidity in the quadratic case not only over the circle but over higher-dimensional spheres too. We start this section with three propositions which will be utilized later in the proof of Theorem 1.1.

The first proposition, which can be found also in the Appendix of [14] (see the proof of Lemma 3.13 there), helps us understand how a translation affects the Wasserstein distance. For $\mu \in \mathcal{P}\left(\mathbb{R}^{n+1}\right)$ and $v \in \mathbb{R}^{n+1}$, the translation of μ by v is the measure $(t_v)_{\#}\mu \in \mathcal{P}\left(\mathbb{R}^{n+1}\right)$ where $t_v\colon \mathbb{R}^{n+1} \to \mathbb{R}^{n+1}$, $x \mapsto x + v$ is the translation by v. Recall that elements of $\mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ can be considered as elements of $\mathcal{W}_2(\mathbb{R}^{n+1}, \varrho_{\|\cdot\|})$ because $\mathcal{W}_2(\mathbb{S}^n, \varrho_{\|\cdot\|})$ naturally embeds into $\mathcal{W}_2(\mathbb{R}^{n+1}, \varrho_{\|\cdot\|})$. The barycenter of $\mu \in \mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ is defined to be the point $m(\mu) = \int_{\mathbb{S}^n} x \, \mathrm{d}\mu(x) \in \mathbb{R}^{n+1}$.

Proposition 4.1. *Let $\mu, \nu \in \mathcal{W}_2\left(\mathbb{R}^{n+1}\right)$ and $v \in \mathbb{R}^{n+1}$. Then we have*

$$d^2_{\mathcal{W}_2}\left((t_v)_{\#}\mu, \nu\right) = d^2_{\mathcal{W}_2}(\mu, \nu) + \langle v, v + 2m(\mu) - 2m(\nu)\rangle. \tag{4.1}$$

In particular, substituting $v = m(\nu) - m(\mu)$ gives

$$d^2_{\mathcal{W}_2}(\mu, \nu) = d^2_{\mathcal{W}_2}\left((t_{-m(\mu)})_{\#}\mu, (t_{-m(\nu)})_{\#}\nu\right) + \|m(\nu) - m(\mu)\|^2. \tag{4.2}$$

Subsequently, ν is a translated version of μ if and only if $d_{\mathcal{W}_2}(\mu,\nu) = \|m(\nu)-m(\mu)\|$.

Proof. For any $\pi \in \Pi(\mu,\nu)$ and $v \in \mathbb{R}^{n+1}$, we have $(t_{(v,0)})_{\#}\pi \in \Pi\left((t_v)_{\#}\mu,\nu\right)$, and vice versa. (Here, 0 stands for $0 \in \mathbb{R}^{n+1}$.) Hence

$$\begin{aligned} d^2_{\mathcal{W}_2}\left((t_v)_{\#}\mu,\nu\right) &= \inf_{\pi\in\Pi(\mu,\nu)} \iint_{\mathbb{R}^{n+1}\times\mathbb{R}^{n+1}} \|x-y\|^2 \,\mathrm{d}\left((t_{(v,0)})_{\#}\pi\right)(x,y) \\ &= \inf_{\pi\in\Pi(\mu,\nu)} \iint_{\mathbb{R}^{n+1}\times\mathbb{R}^{n+1}} \|x+v-y\|^2 \,\mathrm{d}\pi(x,y) \\ &= \inf_{\pi\in\Pi(\mu,\nu)} \iint_{\mathbb{R}^{n+1}\times\mathbb{R}^{n+1}} \left(\|x-y\|^2 + \|v\|^2 + 2\langle x,v\rangle - 2\langle y,v\rangle\right) \,\mathrm{d}\pi(x,y) \\ &= d^2_{\mathcal{W}_2}(\mu,\nu) + \|v\|^2 + 2\int_{\mathbb{R}^{n+1}} \langle x,v\rangle \,\mathrm{d}\mu(x) - 2\int_{\mathbb{R}^{n+1}} \langle y,v\rangle \,\mathrm{d}\nu(y), \end{aligned}$$

which gives (4.1). The identity (4.2) follows if we translate both arguments in the left-hand side by the vector $m(\nu)$. □

In quadratic Wasserstein spaces over uniquely geodesic spaces, the α-weighted mean squared-error function

$$\rho \mapsto (1-\alpha)d^2_{\mathcal{W}_2}(\mu,\rho) + \alpha d^2_{\mathcal{W}_2}(\nu,\rho)$$

defined by μ and ν has a unique minimizer — provided that the optimal coupling of μ and ν is unique — which is the *displacement convex combination* or *displacement interpolation* of μ and ν with weights $(1-\alpha)$ and α [25,26]. Intuitively, this is the measure that we obtain if we start moving μ to ν according to the optimal transport plan, but stop at proportion α of the journey. A great challenge concerning $\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ is that it has no geodesics at all, and hence the quadratic Wasserstein space $\mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ has no geodesics either. Still, mean squared-error functions make perfect sense on $\mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$, they are invariant under isometries in an appropriate sense, and hence if the measures μ and ν defining them are fixed by an isometry Φ, then so are the unique minimizers — if they exist. We will prove in Proposition 4.2 that on $\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$, the minimizer of the α-weighted squared-error function is the projection of the displacement interpolation onto the sphere. This is similar to how for a measure $\mu \in \mathcal{P}(\mathbb{S}^n)$, its closest Dirac measure supported on a point in $\mathbb{R}^{n+1}$ is $\delta_{m(\mu)}$, while among those supported on $\mathbb{S}^n$, it is the projection of $\delta_{m(\mu)}$. We are going to exploit this characterisation in Step 6 of the proof of Theorem 1.1.

We will use the following projection of the α-weighted mean of two points x,y onto $\mathbb{S}^n$ frequently:

$$p_\alpha(x,y) := \frac{(1-\alpha)x+\alpha y}{\|(1-\alpha)x+\alpha y\|} \qquad \left(\alpha\in[0,1],\ x,y\in\mathbb{S}^n\right).$$

Note that $p_\alpha(x,y)$ is not defined when $\alpha = \frac{1}{2}$ and $x=-y$.

Let us define the cost $c_\alpha\colon \mathbb{S}^n\times\mathbb{S}^n \to [0,2]$ by

$$c_\alpha(x,y) := \min_{z\in\mathbb{S}^n}\left\{(1-\alpha)\,||x-z||^2 + \alpha\,||z-y||^2\right\} = 2(1-||(1-\alpha)x+\alpha y||). \tag{4.3}$$

If $\pi\in\mathcal{P}(\mathbb{S}^n\times\mathbb{S}^n)$ is a coupling of μ and ν, then let $\rho^{(\alpha)}_\pi \in \mathcal{W}_2\left(\mathbb{S}^n,\varrho_{\|\cdot\|}\right)$ be defined to be the displacement interpolation in $\mathbb{R}^{n+1}$ at time α between μ and ν according to the plan π, projected to $\mathbb{S}^n$. Formally,

$$\rho^{(\alpha)}_\pi := \left(\{(x,y)\mapsto p_\alpha(x,y)\}\right)_{\#}\pi. \tag{4.4}$$

Now let $\mu,\nu \in \mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ and consider the α-weighted mean squared error

$$Q_\alpha^{\mu,\nu}\colon \mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right) \to [0,\infty);\ \rho \mapsto Q_\alpha^{\mu,\nu}(\rho) := (1-\alpha) d^2_{\mathcal{W}_2}(\mu,\rho) + \alpha d^2_{\mathcal{W}_2}(\nu,\rho). \tag{4.5}$$

Proposition 4.2. *Let $\mu,\nu \in \mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ be such that there is a unique optimal transport plan π_* for them with respect to the cost c_α defined in (4.3). Suppose that $\alpha \neq \frac{1}{2}$ or $\alpha = \frac{1}{2}$ and $\pi_*\left(\{(z,-z) : z \in \mathbb{S}^n\}\right) = 0$. Then the mean squared error $Q_\alpha^{\mu,\nu}$ defined in (4.5) has a unique minimizer which is equal to $\rho^{(\alpha)}_{\pi_*}$, the push-forward of π_* by p_α — see (4.4) for the precise definition. If $\pi_*\left(\{(z,-z) : z \in \mathbb{S}^n\}\right) > 0$ and $\alpha = \frac{1}{2}$, then $\rho^{(\alpha)}_{\pi_*} = \rho^{(\frac{1}{2})}_{\pi_*}$ is not well-defined and $Q_\alpha^{\mu,\nu} = Q_{\frac{1}{2}}^{\mu,\nu}$ has infinitely many minimizers.*

Proof. We proceed by establishing a lower bound for (4.5) and taking care of the case of equality. Let $\rho \in \mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ be arbitrary, and let $\pi_{\mu,\rho}$ and $\pi_{\rho,\nu}$ be optimal transport plans (w.r.t. the quadratic distance) between μ and ρ, and ρ and ν, respectively. Let $\pi_{\mu,\rho,\nu} \in \mathcal{P}\left(\mathbb{S}^n \times \mathbb{S}^n \times \mathbb{S}^n\right)$ be the *gluing* of $\pi_{\mu,\rho}$ and $\pi_{\rho,\nu}$ — see [25, Lemma 7.6] for the precise definition. Then $\pi_{\mu,\nu} := (\pi_{\mu,\rho,\nu})_{1,3} \in \mathcal{P}\left(\mathbb{S}^n \times \mathbb{S}^n\right)$ is a coupling of μ and ν. Now

$$\begin{aligned} Q_\alpha^{\mu,\nu}(\rho) &= (1-\alpha) d^2_{\mathcal{W}_2}(\mu,\rho) + \alpha d^2_{\mathcal{W}_2}(\nu,\rho) \\ &= (1-\alpha) \iint_{\mathbb{S}^n\times\mathbb{S}^n} \|x-z\|^2 \,\mathrm{d}\pi_{\mu,\rho}(x,z) + \alpha \iint_{\mathbb{S}^n\times\mathbb{S}^n} \|z-y\|^2 \mathrm{d}\pi_{\rho,\nu}(z,y) \\ &= \iiint_{\mathbb{S}^n\times\mathbb{S}^n\times\mathbb{S}^n} (1-\alpha)\|x-z\|^2 + \alpha\|z-y\|^2 \,\mathrm{d}\pi_{\mu,\rho,\nu}(x,z,y) \\ &\geq \iiint_{\mathbb{S}^n\times\mathbb{S}^n\times\mathbb{S}^n} c_\alpha(x,y)\,\mathrm{d}\pi_{\mu,\rho,\nu}(x,z,y) = \iint_{\mathbb{S}^n\times\mathbb{S}^n} c_\alpha(x,y)\,\mathrm{d}\pi_{\mu,\nu}(x,y). \end{aligned} \tag{4.6}$$

The inequality (4.6) is saturated if and only if $z = p_\alpha(x,y)$ for $\pi_{\mu,\rho,\nu}$-a.e. $(x,z,y) \in (\mathbb{S}^n)^3$, that is, if $\rho = \rho^{(\alpha)}_{\pi_{\mu,\nu}}$. Moreover, the right-hand side of (4.6) is minimal if and only if $\pi_{\mu,\nu} = \pi_*$. Consequently,

$$Q_\alpha^{\mu,\nu}(\rho) \geq \min_{\pi\in\Pi(\mu,\nu)} \left\{ \iint_{\mathbb{S}^n\times\mathbb{S}^n} c_\alpha(x,y)\,\mathrm{d}\pi(x,y) \right\} = \iint_{\mathbb{S}^n\times\mathbb{S}^n} c_\alpha(x,y)\,\mathrm{d}\pi_*(x,y)$$

and the only ρ realizing this minimum is $\rho^{(\alpha)}_{\pi_*}$. On the other hand, if π_* puts weight on antipodal points, that is, $\pi_*\left(\{(z,-z) : z \in \mathbb{S}^n\}\right) > 0$, and $\alpha = \frac{1}{2}$, then we have an infinite collection of minimizing measures by the theorems of Thales and Pythagoras — or by a simple direct computation. □

In the next proposition, we consider the case when the first argument of p_α is fixed, and we clarify injectivity/surjectivity properties of p_α as α varies form 0 to 1.

Proposition 4.3. *Let $\alpha \in \left(0,\frac{1}{2}\right) \cup \left(\frac{1}{2},1\right]$ and let $N \in \mathbb{S}^n$ be arbitrary but fixed — it may be considered as the "north pole". Let $p_\alpha(N,\cdot)\colon \mathbb{S}^n \to \mathbb{S}^n$ be the map sending u to*

$$p_\alpha(N,u) = \frac{(1-\alpha)N + \alpha u}{||(1-\alpha)N + \alpha u||}.$$

Then for $\alpha \in \left(\frac{1}{2},1\right]$, $p_\alpha(N,\cdot)$ is bijective. For $\alpha \in \left(0,\frac{1}{2}\right)$, $p_\alpha(N,\cdot)$ is neither surjective nor injective: it is 2-to-1 for almost all points of $\mathbb{S}^n$. Finally, $p_{\frac{1}{2}}(N,\cdot)\colon \mathbb{S}^n\setminus\{-N\} \to \mathbb{S}^n\setminus\{-N\}$ is injective, and its range is the open "upper" hemisphere $\{z \in \mathbb{S}^n \,|\, \langle z,N\rangle > 0\}$.

Proof. When considering $p_\alpha(N,u)$, we can assume without loss of generality that $N = (0,0,\dots,0,1)$ and $u = (\cos\theta, 0, \dots, 0, \sin\theta)$ for some $\theta \in (-\pi,\pi]$. Let $c_u^{(\alpha)}$ be the normalising constant $||(1-\alpha)N + \alpha u||$. Note that $c_u^{(\alpha)} > 0$ if and only if $(\alpha,u) \neq \left(\frac{1}{2}, -N\right)$. Whenever $(\alpha,u) \neq \left(\frac{1}{2}, -N\right)$, we have that

$$c_u^{(\alpha)} p_u(N,u) = (\alpha\cos\theta, 0, \dots, 0, (1-\alpha) + \alpha\sin\theta),$$

and so for a fixed $\alpha \neq \frac{1}{2}$, setting

$$x_\theta := \alpha\cos\theta$$
$$y_\theta := (1-\alpha) + \alpha\sin\theta,$$

we see that (x_θ, y_θ) satisfy $x_\theta^2 + (y_\theta - (1-\alpha))^2 = \alpha^2$, i.e., they lie on the circle of radius α with the centre at $(0, 1-\alpha)$. For any point $u = (\cos\theta, 0, \dots, 0, \sin\theta)$, its image $p_\alpha(N,u)$ is the projection of $(x_\theta, 0, \dots, 0, y_\theta)$ onto $\mathbb{S}^n$, i.e., the point obtained as the intersection of $\mathbb{S}^n$ and the half-line from $(0,\dots,0)$ through $(x_\theta, 0, \dots, 0, y_\theta)$. Now the statements of this proposition are easy to see from Figure 1.

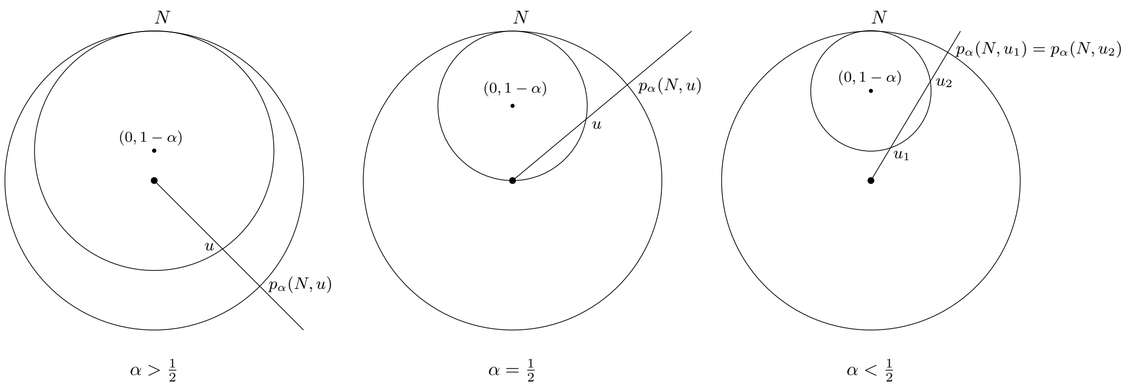

FIGURE 1. $p_\alpha(N,u)$ lies on the $\mathbb{S}^1$ spanned by N and u, at the spherical projection of $c_u^{(\alpha)} p_\alpha(N,u)$. The bigger circle displayed in each of the cases is $\mathbb{S}^1$, the smaller one is $c_\cdot^{(\alpha)} p_\alpha(N,\cdot)$.

□

Now we turn to the proof of Theorem 1.1 which, for the sake of clarity, we divide into six steps.

Step 1. Similarly as in the proof of Theorem 3.1, we first understand the action of Ψ on the set of Dirac measures. The maximal distance in $\mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ is 2, and is attained only on pairs of Dirac measures that are concentrated on antipodal points. Since

$$2 = d_{\mathcal{W}_2}(\mu,\nu) = d_{\mathcal{W}_2}(\Psi(\mu), \Psi(\nu)),$$

we get that $\Psi(\delta_x)$ is a Dirac measure for all $x \in \mathbb{S}^n$. Since Ψ and Ψ^{-1} are both isometries, the map $\psi\colon \mathbb{S}^n \to \mathbb{S}^n$ defined by $\Psi(\delta_x) = \delta_{\psi(x)}$ is a bijection, and furthermore, since

$$\|\psi(x) - \psi(y)\| = d_{\mathcal{W}_2}\left(\delta_{\psi(x)}, \delta_{\psi(y)}\right) = d_{\mathcal{W}_2}\left(\Psi(\delta_x), \Psi(\delta_y)\right) = d_{\mathcal{W}_2}\left(\delta_x, \delta_y\right) = \|x-y\|,$$

it is in fact an isometry. Just as before, we will be done once we prove that an isometry of $\mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ which fixes all Dirac measures is necessarily the identity, because then in particular, $\left(\psi^{-1}\right)_\# \circ \Psi = \mathrm{id}_{\mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)}$, and so $\Psi = \psi_\#$ as claimed. From now on, we assume that Φ is an isometry of $\mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ such that $\Phi(\delta_x) = \delta_x$, and our aim is to show that $\Phi(\mu) = \mu$ for all $\mu \in \mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$.

Step 2. Next we claim that Φ preserves the barycenter of measures.

For any $\mu \in \mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ and $x \in \mathbb{S}^n$, we have

$$d^2_{\mathcal{W}_2}(\mu, \delta_x) = \int_{\mathbb{S}^n} \|y - x\|^2 \mathrm{d}\mu(y) = 2 - 2\left\langle x, \int_{\mathbb{S}^n} y \,\mathrm{d}\mu(y)\right\rangle = 2(1 - \langle x, m(\mu)\rangle). \tag{4.7}$$

Therefore,

$$\operatorname{argmin}\left\{d^2_{\mathcal{W}_2}(\mu, \delta_x) \colon x \in \mathbb{S}^n\right\} = \begin{cases} \frac{m(\mu)}{||m(\mu)||} & \text{if } m(\mu) \neq 0 \\ \mathbb{S}^n & \text{if } m(\mu) = 0 \end{cases}$$

and

$$\min\left\{d^2_{\mathcal{W}_2}(\mu, \delta_x) \colon x \in \mathbb{S}^n\right\} = 2(1 - ||m(\mu)||).$$

Since $\Phi(\delta_x) = \delta_x$ and hence $d^2_{\mathcal{W}_2}(\Phi(\mu), \delta_x) = d^2_{\mathcal{W}_2}(\mu, \delta_x)$ for all $x \in \mathbb{S}^n$, we get that

$$\begin{aligned} 2(1 - ||m(\Phi(\mu))||) &= \min\left\{d^2_{\mathcal{W}_2}(\Phi(\mu), \delta_x) \colon x \in \mathbb{S}^n\right\} \\ &= \min\left\{d^2_{\mathcal{W}_2}(\mu, \delta_x) \colon x \in \mathbb{S}^n\right\} = 2(1 - ||m(\mu)||). \end{aligned}$$

This implies that $||m(\Phi(\mu))|| = ||m(\mu)||$, and in particular, $m(\Phi(\mu)) = 0 \in \mathbb{R}^{n+1}$ whenever $m(\mu) = 0$. Moreover, if $m(\mu) \neq 0$ then

$$\begin{aligned} \frac{m(\mu)}{||m(\mu)||} &= \operatorname{argmin}\left\{d^2_{\mathcal{W}_2}(\mu, \delta_x) \colon x \in \mathbb{S}^n\right\} \\ &= \operatorname{argmin}\left\{d^2_{\mathcal{W}_2}(\Phi(\mu), \delta_x) \colon x \in \mathbb{S}^n\right\} = \frac{m(\Phi(\mu))}{||m(\Phi(\mu))||} \end{aligned}$$

which implies the desired equality $m\left(\Phi\left(\mu\right)\right) = m\left(\mu\right)$ for all $\mu \in \mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$.

Step 3. Now we prove that measures supported on two points are mapped to measures supported on two points. We first show that for all $\mu, \nu \in \mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$,

$$\operatorname{affspan}\left(\operatorname{supp}\big(\Phi(\mu)\big)\right) \perp \operatorname{affspan}\left(\operatorname{supp}\big(\Phi(\nu)\big)\right)$$

holds if and only if

$$\operatorname{affspan}\left(\operatorname{supp}(\mu)\right) \perp \operatorname{affspan}\left(\operatorname{supp}(\nu)\right).$$

Kloeckner proved in [18, Lemma 6.2] that orthogonality of supports can be characterized by the metric in the ambient space $\mathcal{W}_2\left(\mathbb{R}^{n+1}, \varrho_{\|\cdot\|}\right)$. Namely,

$$d^2_{\mathcal{W}_2(\mathbb{R}^{n+1})}(\mu, \nu) = ||m(\mu) - m(\nu)||^2 + d^2_{\mathcal{W}_2(\mathbb{R}^{n+1})}\left(\mu, \delta_{m(\mu)}\right) + d^2_{\mathcal{W}_2(\mathbb{R}^{n+1})}\left(\nu, \delta_{m(\nu)}\right)$$

holds if and only if there exist two orthogonal affine subspaces $L, M \subset \mathbb{R}^{n+1}$ such that $\operatorname{supp}(\mu) \subseteq L$ and $\operatorname{supp}(\nu) \subseteq M$. We proceed by showing that the isometries of $\mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$

leave the $\mathcal{W}_2(\mathbb{R}^{n+1})$-distance of a measure from the Dirac mass concentrated on its barycenter invariant, that is,

$$d_{\mathcal{W}_2(\mathbb{R}^{n+1})}\left(\mu, \delta_{m(\mu)}\right) = d_{\mathcal{W}_2(\mathbb{R}^{n+1})}\left(\Phi(\mu), \delta_{m(\Phi(\mu))}\right)$$

for any $\mu \in \mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$. Indeed, a direct computation very similar to (4.7) shows that

$$d^2_{\mathcal{W}_2(\mathbb{R}^{n+1})}\left(\mu, \delta_{m(\mu)}\right) = 1 - ||m(\mu)||^2 \text{ and } d^2_{\mathcal{W}_2(\mathbb{R}^{n+1})}\left(\Phi(\mu), \delta_{m(\Phi(\mu))}\right) = 1 - ||m(\Phi(\mu))||^2,$$

which implies our statement as we have shown $m(\Phi(\mu)) = m(\mu)$ in Step 2. Hence for any $\mu, \nu \in \mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$,

$$\begin{aligned}||m(\mu) - m(\nu)||^2 &+ d^2_{\mathcal{W}_2}\left(\mu, \delta_{m(\mu)}\right) + d^2_{\mathcal{W}_2}\left(\nu, \delta_{m(\nu)}\right)\\ &= ||m(\Phi(\mu)) - m(\Phi(\nu))||^2 + d^2_{\mathcal{W}_2}\left(\Phi(\mu), \delta_{m(\Phi(\mu))}\right) + d^2_{\mathcal{W}_2}\left(\Phi(\nu), \delta_{m(\Phi(\nu))}\right),\end{aligned}$$

meaning that orthogonally supported measures must be mapped to orthogonally supported measures by Φ.

A maximal set of measures whose supports are one-dimensional and pairwise orthogonal must therefore be mapped to a set of measures whose supports are zero- or one-dimensional. But zero-dimensionally supported measures are exactly the Dirac masses, to which only Dirac masses can be mapped by Φ, and so one-dimensionally supported measures must be mapped to one-dimensionally supported measures. Continuing similarly, we would see more generally that the affine dimension of the support is preserved by Φ, but since on the sphere, one-dimensionally supported measures are exactly the two-point supported measures, the one-dimensional case is enough to prove our statement.

Step 4. We proceed with showing that measures supported on two points are fixed by Φ. Let us introduce the notation $\Delta'_2(\mathbb{S}^n)$ for the set of all elements in $\mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$ with a two-point support, set $\widetilde{\mu} := \left(t_{-m(\mu)}\right)_{\#}\mu$ for all $\mu \in \Delta'_2(\mathbb{S}^n)$, and $\Delta'_{2,0}(\mathbb{S}^n) := \left\{\widetilde{\mu} \in \mathcal{P}(\mathbb{R}^{n+1}) : \mu \in \Delta'_2(\mathbb{S}^n)\right\}$. By Step 3, $\Phi|_{\Delta'_2(\mathbb{S}^n)} : \Delta'_2(\mathbb{S}^n) \to \Delta'_2(\mathbb{S}^n)$ is an isometric embedding. By Proposition 4.1 we know that for all $\mu, \nu \in \Delta'_2(\mathbb{S}^n)$,

$$\begin{aligned}d^2_{\mathcal{W}_2}\left(\widetilde{\mu}, \widetilde{\nu}\right) &= d^2_{\mathcal{W}_2}\left(\mu, \nu\right) - ||m(\mu) - m(\nu)||^2\\ &= d^2_{\mathcal{W}_2}\left(\Phi(\mu), \Phi(\nu)\right) - ||m(\Phi(\mu)) - m(\Phi(\nu))||^2 = d^2_{\mathcal{W}_2}\left(\widetilde{\Phi(\mu)}, \widetilde{\Phi(\nu)}\right).\end{aligned}$$

Consequently, $\widetilde{\mu} = \widetilde{\nu}$ holds if and only if $\widetilde{\Phi(\mu)} = \widetilde{\Phi(\nu)}$, in other words, $\Phi(\nu)$ is a translate of $\Phi(\mu)$ if and only if ν is a translate of μ.

Let a measure $\mu \in \Delta'_2(\mathbb{S}^n)$ be fixed. We can assume without loss of generality that

$$\operatorname{supp}(\mu) = \{(\cos\theta, 0, \ldots, 0, \sin\theta), (\cos\theta, 0, \ldots, 0, -\sin\theta)\}$$

for some $\theta \in (0, \pi/2]$. In this case,

$$\frac{1}{2}\sum_{x \in \operatorname{supp}(\mu)} x = (\cos\theta, 0, \ldots, 0)$$

and

$$\left(\operatorname{affspan}\left(\operatorname{supp}(\mu)\right) - \frac{1}{2}\sum_{x \in \operatorname{supp}(\mu)} x\right)^{\perp} = \{(v_1, \ldots, v_n, 0) \colon v_1, \ldots, v_n \in \mathbb{R}\}.$$

Define $\vec{\mu} := \left\{v \in \mathbb{R}^{n+1} : (t_v)_{\#}\mu \in \mathcal{P}\left(\mathbb{S}^n\right)\right\}$, and observe that $\vec{\mu}$ is the set of those vectors $(v_1, \ldots, v_{n+1})$ such that

$$\|(v_1 + \cos\theta, v_2, \ldots, v_n, v_{n+1} + \sin\theta)\| = \|(v_1 + \cos\theta, v_2, \ldots, v_n, v_{n+1} - \sin\theta)\| = 1.$$

Then we get that $v \in \vec{\mu}$ if and only if

$$(v_1 + \cos\theta)^2 + v_2^2 + \cdots + v_n^2 + v_{n+1}^2 \pm 2v_{n+1}\sin\theta = 1 - \sin^2\theta.$$

Since $\sin\theta \neq 0$, this holds exactly when $v_{n+1} = 0$ and $(v_1 + \cos\theta)^2 + v_2^2 + \cdots + v_n^2 = \cos^2\theta$, i.e., the first n coordinates span an $n-1$-dimensional sphere with radius $|\cos\theta|$ centered at $(-\cos\theta, 0, \ldots, 0)$, or they are just the singleton containing $0 \in \mathbb{R}^n$ in the case $\cos\theta = 0$. In other words,

$$\begin{aligned}\vec{\mu} &= -(\cos\theta, 0, \ldots, 0) + |\cos\theta| \cdot \left(\mathbb{S}^n \cap \{(v_1, \ldots, v_n, 0) : v_1, \ldots, v_n \in \mathbb{R}\}\right) \\ &= -\left(\frac{1}{2}\sum_{x \in \operatorname{supp}(\mu)} x\right) + \left\|\frac{1}{2}\sum_{x \in \operatorname{supp}(\mu)} x\right\| \cdot \left(\mathbb{S}^n \cap \left(\operatorname{affspan}\left(\operatorname{supp}(\mu)\right) - \frac{1}{2}\sum_{x \in \operatorname{supp}(\mu)} x\right)^{\perp}\right).\end{aligned}$$

As Φ maps the translates of μ to the translates of $\Phi(\mu)$, there is an $\eta \in \Delta'_{2,0}(\mathbb{S}^n) \subset \mathcal{P}\left(\mathbb{R}^{n+1}\right)$ such that

$$\Phi\left((t_v)_{\#}\mu\right) = (t_{v+m(\mu)})_{\#}\eta \qquad (v \in \vec{\mu}). \tag{4.8}$$

We emphasize that η does not depend on v. It follows that $\vec{\mu} + m(\mu) + \operatorname{supp}(\eta) \subset \mathbb{S}^n$. But by plugging $v = 0 \in \mathbb{R}^{n+1}$ to (4.8), we get that $\operatorname{supp}\left(\Phi(\mu)\right) = m(\mu) + \operatorname{supp}(\eta)$, and so the previous line becomes $\vec{\mu} + \operatorname{supp}\left(\Phi(\mu)\right) \subset \mathbb{S}^n$. By the definition of $\vec{\Phi(\mu)}$, this means that $\vec{\mu} \subseteq \vec{\Phi(\mu)}$.

For any μ with $\operatorname{diam}(\operatorname{supp}(\mu)) < 2$, we have that $\cos\theta \neq 0$, and so $\vec{\mu}$ is an $n-1$-dimensional sphere, implying that $\vec{\mu} = \vec{\Phi(\mu)}$ and $\operatorname{supp}(\mu) = \operatorname{supp}\left(\Phi(\mu)\right)$. Now μ and $\Phi(\mu)$ are probability measures with the same 2-point support and the same barycenter, and so $\mu = \Phi(\mu)$. Finally, $\mu = \Phi(\mu)$ for all $\mu \in \Delta'_2\left(\mathbb{S}^n\right)$ by continuity.

Step 5. Now assume that $\mu = \sum_{i=1}^m \lambda_i \delta_{x_i}$ where $x_i \neq -x_j$ for all $1 \leq i < j \leq m$. Such measures form a dense subset of $\mathcal{W}_2\left(\mathbb{S}^n, \varrho_{\|\cdot\|}\right)$. We claim that

$$\operatorname{supp}\big(\Phi(\mu)\big) \subseteq \{x_1, \ldots, x_m\} \cup \{-x_1, \ldots, -x_m\}$$

and $\big(\Phi(\mu)\big)(\{x_i, -x_i\}) = \mu(\{x_i\})$ for all $1 \leq i \leq m$.

The proof of this claim relies on preserving the mass of *bisectors* which are defined as follows: for $u, v \in \mathbb{S}^n$, the corresponding bisector is

$$B(u, v) := \{y \in \mathbb{S}^n : ||u - y|| = ||v - y||\} \cong \mathbb{S}^{n-1}.$$

To start, we apply Lemma 3.17 from [14] with $E = \mathbb{R}^{n+1}$, $p = 2$, $x \in \mathbb{S}^n$, $a = 1$ and $b = -1$ to obtain that

$$\begin{aligned}\mu\left(B(x, -x)\right) = \max\left\{\alpha : d_{\mathcal{W}_2}\left(\mu, \alpha\delta_x + (1-\alpha)\delta_{-x}\right) = m_\mu\right\} \\ - \min\left\{\alpha : d_{\mathcal{W}_2}\left(\mu, \alpha\delta_x + (1-\alpha)\delta_{-x}\right) = m_\mu\right\},\end{aligned}$$

where $m_\mu := \min\left\{d_{\mathcal{W}_2}\left(\mu, \alpha\delta_x + (1-\alpha)\delta_{-x}\right) : 0 \leq \alpha \leq 1\right\}$ and $B(x, -x) \cong \mathbb{S}^{n-1}$ is the bisector between x and $-x$, i.e., the set of all points equidistant from x and $-x$. But since $\Phi\left(\alpha\delta_x + (1-\alpha)\delta_{-x}\right) = \alpha\delta_x + (1-\alpha)\delta_{-x}$ for all $\alpha \in [0, 1]$ by Steps 1 and 4, we get that $m_\mu = m_{\Phi(\mu)}$, and subsequently $\mu(B(x, -x)) = \left(\Phi(\mu)\right)\left(B(x, -x)\right)$ for all $x \in \mathbb{S}^n$. Since for every $x \in \mathbb{S}^n$, $B(x,-x)$ is an $n-1$-dimensional subsphere of $\mathbb{S}^n$, and every $n-1$-dimensional subsphere of $\mathbb{S}^n$ is of the form $B(x,-x)$ for some $x \in \mathbb{S}^n$, the previous sentence says that $\mu(S) = (\Phi(\mu))(S)$ for every subsphere S of codimension 1.

For every $\widetilde{x} \in \{x_1, \dots, x_m\} = \mathrm{supp}(\mu)$, there exists a sequence $(S_j)_{j\in\mathbb{N}}$ of $n-1$-dimensional subspheres of $\mathbb{S}^n$ such that $S_j \cap \mathrm{supp}(\mu) = \{\widetilde{x}\}$ for every j, and the intersection of any n subspheres is trivial, that is, $\bigcap_{k=1}^n S_{j_k} = \{\tilde{x}, -\tilde{x}\}$ for any choice of $j_1 < j_2 < \cdots < j_n$. Therefore,

$$\mu(\{\widetilde{x}\}) = \mu(S_j) = \big(\Phi(\mu)\big)(S_j) \qquad (j \in \mathbb{N}),$$

and we are in the right position to prove that

$$\Phi(\mu)(\{\widetilde{x}, -\widetilde{x}\}) = \mu(\{\widetilde{x}\}).$$

The inequality $\Phi(\mu)(\{\widetilde{x}, -\widetilde{x}\}) \le \big(\Phi(\mu)\big)(S_j) = \mu(\{\widetilde{x}\})$ holds because $\{\widetilde{x}, -\widetilde{x}\} \subseteq S_j$. If $n = 1$ then in fact $\{\tilde{x}, -\tilde{x}\} = S_j$, and we are done. If $n \ge 2$, assume indirectly that $\Phi(\mu)(\{\widetilde{x}, -\widetilde{x}\}) < \mu(\{\widetilde{x}\})$, and let $\varepsilon > 0$ denote the gap between the two sides of this strict inequality. Now we have

$$\big(\Phi(\mu)\big)(S_j \setminus \{\widetilde{x}, -\widetilde{x}\}) = \big(\Phi(\mu)\big)(S_j) - \Phi(\mu)(\{\widetilde{x}, -\widetilde{x}\}) = \varepsilon$$

for every $j \in \mathbb{N}$. The fact that $\bigcap_{k=1}^n S_{j_k} = \{\tilde{x}, -\tilde{x}\}$ for every $j_1 < j_2 < \cdots < j_n$ implies that the family of sets $(S_j \setminus \{\widetilde{x}, -\widetilde{x}\})_{j=1}^\infty$ covers any point of $\mathbb{S}^n$ at most n times. This means that $\sum_{j=1}^\infty \big(\Phi(\mu)\big)(S_j \setminus \{\widetilde{x}, -\widetilde{x}\})$ is bounded from above by $n \cdot \big(\Phi(\mu)\big)(\mathbb{S}^n) = n$, which is a contradiction as $\big(\Phi(\mu)\big)(S_j \setminus \{\widetilde{x}, -\widetilde{x}\}) = \varepsilon$ for every $j \in \mathbb{N}$ and $\sum_{j=1}^\infty \varepsilon = \infty$.

Step 6. A crucial consequence of the claim made in Step 5 is that the isometry Φ fixes all measures that are supported within an open hemisphere of $\mathbb{S}^n$. Indeed, we learned from Step 2 that Φ preserves the barycenter of measures, and from Step 5 that the only possible action Φ can do is to send some mass from a point to its antipodal point. But if a measure is supported on an open hemisphere then the transport of any mass to its antipodal point would change the barycenter.

Suppose that $\mu(\{-N\}) = 0$. Then the spherical projection $\rho^{(\frac{1}{2})}_{\delta_N \otimes \mu}$ of the displacement convex combination of δ_N and μ is well-defined, and since it is supported on the upper hemisphere, $\Phi\left(\rho^{(\frac{1}{2})}_{\delta_N \otimes \mu}\right) = \rho^{(\frac{1}{2})}_{\delta_N \otimes \mu}$. Let us now consider the sets

$$\begin{aligned} A := &\left\{ Q_{\frac{1}{2}}^{\delta_N, \mu}(\rho) : \rho \in \mathcal{P}(\mathbb{S}^n) \right\} \\ = &\left\{ Q_{\frac{1}{2}}^{\Phi(\delta_N), \Phi(\mu)}(\Phi(\rho)) : \rho \in \mathcal{P}(\mathbb{S}^n) \right\} \\ = &\left\{ Q_{\frac{1}{2}}^{\delta_N, \Phi(\mu)}(\kappa) : \kappa \in \mathcal{P}(\mathbb{S}^n) \right\} =: B \end{aligned}$$

where the last equality follows from the surjectivity of the isometry Φ. Since $A = B$, necessarily

$$\min B = \min A = Q_{\frac{1}{2}}^{\delta_N, \mu}\left(\rho^{(\frac{1}{2})}_{\delta_N \otimes \mu}\right) = Q_{\frac{1}{2}}^{\Phi(\delta_N), \Phi(\mu)}\left(\Phi\left(\rho^{(\frac{1}{2})}_{\delta_N \otimes \mu}\right)\right) = Q_{\frac{1}{2}}^{\delta_N, \Phi(\mu)}\left(\Phi\left(\rho^{(\frac{1}{2})}_{\delta_N \otimes \mu}\right)\right).$$

The fact that B has a unique minimizer implies by the second statement of Proposition 4.2 that $\Phi(\mu)(\{-N\}) = 0$. Consequently — let us use now the first statement of Proposition

4.2 — the unique minimizer of $Q_\alpha^{\delta_N,\Phi(\mu)}$ is $\rho^{(\frac{1}{2})}_{\delta_N\otimes\Phi(\mu)}$ and hence

$$\rho^{(\frac{1}{2})}_{\delta_N\otimes\Phi(\mu)} = \Phi\left(\rho^{(\frac{1}{2})}_{\delta_N\otimes\mu}\right) = \rho^{(\frac{1}{2})}_{\delta_N\otimes\mu}. \tag{4.9}$$

By the injectivity of $p_{\frac{1}{2}}(N,\cdot)$ on $\mathbb{S}^n\backslash\{-N\}$, see Proposition 4.3, for every measure $\nu\in\mathcal{P}\left(\mathbb{S}^n\right)$ supported within the upper hemisphere, there is a unique measure $\kappa\in\mathcal{P}\left(\mathbb{S}^n\right)$ such that ν is the spherical projection $\rho^{(\frac{1}{2})}_{\delta_N\otimes\kappa}$ of the displacement convex combination of δ_N and κ. Therefore, (4.9) implies that $\Phi(\mu)=\mu$, which completes the proof.

References

[1] L. Ambrosio, N. Gigli, *A user's guide to optimal transport. Modelling and optimisation of flows on networks*, Lecture Notes in Math., 2062, Fond. CIME/CIME Found. Subser., Springer, Heidelberg, 2013.

[2] Z. M. Balogh, T. Titkos, D. Virosztek, *Isometries and isometric embeddings of Wasserstein spaces over the Heisenberg group*, manuscript, arXiv:2303.15095 (2023).

[3] Z. M. Balogh, T. Titkos, D. Virosztek, *Isometric rigidity of the Wasserstein space* $\mathcal{W}_1(\mathbf{G})$ *over Carnot groups*, manuscript, arXiv:2305.05492 (2023).

[4] J. Bertrand, B. Kloeckner, *A geometric study of Wasserstein spaces: isometric rigidity in negative curvature*, Int. Math. Research Notices **5** (2016), 1368–1386.

[5] D. Burago, Y. Burago, S. Ivanov, *A Course in Metric Geometry* Graduate Studies in Mathematics vol. 33, American Mathematical Society, Providence, RI, 2001.

[6] G. Dolinar, B. Kuzma, D. Mitrovic, *Isometries of probability measures with respect to the total variation distance* J. Math. Anal. Appl. **507** (2021), Paper No. 125829.

[7] G. Dolinar, L. Molnár, *Isometries of the space of distribution functions with respect to the Kolmogorov–Smirnov metric*, J. Math. Anal. Appl. **348** (2008), 494–498.

[8] A. Figalli and F. Glaudo, *An Invitation to Optimal Transport, Wasserstein Distances, and Gradient Flows*, EMS Textbooks in Mathematics, EMS Press, Berlin, 2021.

[9] Gy.P. Gehér, *Surjective Kuiper isometries*, Houston J. Math. **44** (2018), 263–281.

[10] Gy.P. Gehér, T. Titkos, *A characterisation of isometries with respect to the Lévy-Prokhorov metric*, Annali della Scuola Normale Superiore di Pisa - Classe di Scienze, Vol. **XIX** (2019), 655–677.

[11] Gy.P. Gehér, T. Titkos, D. Virosztek, *On isometric embeddings of Wasserstein spaces – the discrete case*, J. Math. Anal. Appl. **480** (2019), 123435.

[12] Gy.P. Gehér, T. Titkos, D. Virosztek, *Isometric sudy of Wasserstein spaces – the real line*, Trans. Amer. Math. Soc. **373** (2020), 5855-–5883.

[13] Gy.P. Gehér, T. Titkos, D. Virosztek, *Dirac masses and isometric rigidity*, RIMS Kokyuroku **2125** (2020), 34–41.

[14] Gy. P. Gehér, T. Titkos, D. Virosztek, *The isometry group of Wasserstein spaces: the Hilbertian case*, J. London Math. Soc. **106** (2022), 3865–3894.

[15] Gy. P. Gehér, T. Titkos, D. Virosztek, *Isometric rigidity of Wasserstein tori and spheres*, Mathematika **69** (2023), 20–32.

[16] Gy. P. Gehér, T. Titkos, D. Virosztek, *On the exotic isometry flow of the quadratic Wasserstein space over the real line*, Linear Algebra Appl., In Press, https://doi.org/10.1016/j.laa.2023.02.016.

[17] G. Kiss, T. Titkos, *Isometric rigidity of Wasserstein spaces: the graph metric case*, Proc. Amer. Math. Soc. **150** (2022), 4083–4097.

[18] B. Kloeckner, *A geometric study of Wasserstein spaces: Euclidean spaces*, Annali della Scuola Normale Superiore di Pisa - Classe di Scienze **IX** (2010), 297–323.

[19] L.Molnár *Kolmogorov-Smirnov isometries and affine automorphisms of spaces of distribution functions*, Central European J. of Math. **9** (2011), 789–796.

[20] L. Molnár, P. Szokol, *Kolmogorov-Smirnov isometries of the space of generalized distribution functions*, Math. Slovaca **64** (2014), 433–444.

[21] L. Molnár, *Lévy isometries of the space of probability distribution functions*, J. Math. Anal. Appl. **380** (2011), 847–852.

[22] W. Rudin, *Fourier analysis on groups*, Interscience Tracts in Pure and Applied Mathematics, No. 12 Interscience Publishers, New York-London, (1962).
[23] F. Santambrogio, *Optimal Transport for Applied Mathematicians*, Progress in Nonlinear Differential Equations and Their Applications 87, Birkhauser Basel (2015).
[24] J. Santos-Rodríguez, *On isometries of compact L^p–Wasserstein spaces*, Advances in Math. **409** (2022), Article No. 108632.
[25] C. Villani, *Topics in Optimal Transportation*, Graduate Studies in Mathematics vol. 58, American Mathematical Society, Providence, RI, 2003.
[26] C. Villani, *Optimal Transport: Old and New*, Grundlehren der mathematischen Wissenschaften, Springer, 2009.
[27] D. Virosztek, *Maps on probability measures preserving certain distances — a survey and some new results*, Acta Sci. Math. (Szeged) **84** (2018), 65–80.

György Pál Gehér, Department of Mathematics and Statistics, University of Reading, Whiteknights, P.O. Box 220, Reading RG6 6AX, United Kingdom
Email address: gehergyuri@gmail.com

Aranka Hrušková, Alfréd Rényi Institute of Mathematics, Reáltanoda u. 13-15., Budapest 1053, Hungary, and Central European University, Nádor u. 9, Budapest 1051, Hungary
Email address: aranka@renyi.hu

Tamás Titkos, Alfréd Rényi Institute of Mathematics, Reáltanoda u. 13-15., Budapest 1053, Hungary, and BBS University of Applied Sciences, Alkotmány u. 9., Budapest 1054, Hungary
Email address: titkos.tamas@renyi.hu

Dániel Virosztek, Alfréd Rényi Institute of Mathematics, Reáltanoda u. 13-15., Budapest 1053, Hungary
Email address: virosztek.daniel@renyi.hu

Appendix F

Asymptotically commuting measures have the same harmonic functions

by Aranka Hrušková and Omer Segev