ANNALES
HENRI LEBESGUE

LÁSZLÓ MÁRTON TÓTH

INVARIANT SCHREIER DECORATIONS OF UNIMODULAR RANDOM NETWORKS

DÉCORATIONS DE SCHREIER INVARIANTES SUR LES GRAPHES ALÉATOIRES UNIMODULAIRES

ABSTRACT. — We prove that every $2d$-regular unimodular random network carries an invariant random Schreier decoration. Equivalently, it is the Schreier coset graph of an invariant random subgroup of the free group F_d. As a corollary we get that every $2d$-regular graphing is the local isomorphic image of a graphing coming from a p.m.p. action of F_d.

The key ingredients of the analogous statement for finite graphs do not generalize verbatim to the measurable setting. We find a more subtle way of adapting these ingredients and prove measurable coloring theorems for graphings along the way.

RÉSUMÉ. — Nous démontrons que tout graphe aléatoire unimodulaire $2d$-régulier admet une décoration de Schreier aléatoire invariante. De manière équivalent, c'est le graphe de Schreier des classes d'un sous-groupe invariant aléatoire du groupe libre F_d. Comme corollaire, nous obtenons que tout graphage $2d$-régulier est localement l'image isomorphe d'un graphage provenant d'une action de F_d préservant une mesure de probabilité.

Les ingrédients-clé de l'énoncé analogue pour les graphes finis ne se généralisent pas tels quels au contexte mesurable. Nous trouvons une manière plus subtile d'adapter ces ingrédients, et nous obtenons en chemin des théorèmes de coloriages mesurables de graphages.

1. Introduction

It is a nice exercise in combinatorics to show that any 2d-regular finite graph is a Schreier graph of the free group F_d on d generators. That is, the edges can be oriented and colored by $\{1, \ldots, d\} = [d]$ such that every vertex has exactly one incoming and outgoing edge of each color. We call such a coloring and orientation a *Schreier decoration* of our graph. We do not want to spoil the argument for the reader, but will inevitably have to talk about parts of the solution. See [Can13] for a proof.

The main result of this paper generalizes the above result to unimodular random rooted graphs, a notion that arises as a natural generalization of a finite graph in the limit theory of sparse graph sequences.

THEOREM 1.1. — *Every 2d-regular unimodular random rooted graph has an invariant random Schreier decoration.*

We now introduce unimodularity and invariant random Schreier decorations. For the precise definitions omitted from the Introduction see Section 2.

Let $\mathfrak{G}_\circ^{2d}$ denote the space of rooted, connected (possibly infinite) graphs with degrees bounded by $2d$. Similarly, let $\mathfrak{G}_\circ^{\mathrm{Sch}}$ denote the space of rooted, connected Schreier graphs of F_d. There is a natural forgetting map $\Phi : \mathfrak{G}_\circ^{\mathrm{Sch}} \to \mathfrak{G}_\circ^{2d}$.

We denote by $\mathcal{M}(X)$ the set of Borel probability measures on a topological space X. A measure $\mu \in \mathcal{M}(\mathfrak{G}_\circ^{2d})$ is called a *unimodular random rooted graph* or *unimodular random rooted network* if it satisfies an involution invariance property mimicking reversibility of random walks. The "rooted" is dropped to shorten terminology, unless we want to explicitly emphasize the presence of a root. A *random Schreier graph* is a measure $\mu' \in \mathcal{M}(\mathfrak{G}_\circ^{\mathrm{Sch}})$. It is *invariant*, if it is preserved by the natural action $F_d \curvearrowright \mathfrak{G}_\circ^{\mathrm{Sch}}$, where generators $s \in [d]$ act by moving the root along the unique outward edge colored by s adjacent to it.

DEFINITION 1.2. — *Let $\mu \in \mathcal{M}(\mathfrak{G}_\circ^{2d})$ be a unimodular random rooted graph that is 2d-regular with probability 1. An invariant random Schreier decoration of μ is a measure $\mu' \in \mathcal{M}(\mathfrak{G}_\circ^{\mathrm{Sch}})$ that is invariant and has $\Phi_*\mu' = \mu$.*

The condition $\Phi_*\mu' = \mu$ expresses that forgetting both the orientation and the coloring from a random Schreier graph generated by μ' gives the same random graph as μ in distribution.

Cannizzo investigated the Schreier decorations of unimodular random rooted graphs in [Can13]. He showed that for an invariant random Schreier graph μ' the pushforward $\Phi_*\mu'$ is always unimodular. Moreover, under some assumptions, if a unimodular random graph has an invariant random Schreier decoration, then it has uncountably many ergodic ones. Theorem 1.1 complements these results nicely.

Invariant random Schreier graphs are closely related to Invariant Random Subgroups (IRS's) introduced by Abért, Glasner and Virág in [AGV14]. On one hand, the invariance property of the graph translates to conjugation invariance of the random subgroup that is the stabilizer of the root. On the other hand, given an IRS $H \leqslant F_d$, the Schreier graph on the coset space F_d/H rooted at H gives an invariant random Schreier graph.

Keeping this correspondence in mind, Theorem 1.1 can be thought of as a statement on IRS's of F_d. Bowen proved in [Bow12] that F_d has a large "zoo" of ergodic IRS's. Cannizzo's work shows that non-unimodularity is a probabilistic obstruction to finding an IRS with a given factor geometry. Our result shows that it is the only obstruction.

The proof of the finite case relies on two parts. First, a finite $2d$-regular graph has a *balanced* orientation. That is, an orientation with equal in- and outdegrees at all vertices. Second, a d-regular bipartite finite graph can be properly edge-colored with d colors. This is Kőnig's line coloring theorem. Our approach to Theorem 1.1 is to study these questions for *graphings* instead of graphs.

Graphings are measurable, bounded degree graphs on some standard Borel probability space (X, ν) as vertex set with an additional measure preserving property. Their close connection to unimodular random graphs is also explained in Section 2. For now it suffices to know that any unimodular random graph μ can be encoded in a single graphing $\mathcal{G}$. We say $\mathcal{G}$ *represents* μ if the connected component in $\mathcal{G}$ of a ν-random vertex is the same as μ in distribution. Such a representing graphing exists for any unimodular random graph, though the choice is not unique.

Unfortunately neither of the two finite statements directly generalizes to graphings. In [Lac88] Laczkovich constructed a 2-regular graphing that has no measurable balanced orientation. Moreover, there are d-regular bipartite graphings that have no measurable proper edge coloring with d colors [CK13, Section 6]. We substitute these steps with the next two theorems.

First we claim that representing graphings can be chosen smartly.

THEOREM 1.3. — *Let μ be a $2d$-regular unimodular random rooted graph. Then there exists a graphing $\mathcal{G} = (X, E, \nu)$ on a standard Borel space (X, ν) such that $\mathcal{G}$ represents μ, and there is a measurable balanced orientation of the edges of $\mathcal{G}$.*

Second, we claim that the coloring can be done with little error. Coloring questions for Borel graphs and graphings have been extensively studied, see [KM15] for a survey. We build on the work of Csóka, Lippner and Pikhurko in [CLP16] to derive a version of Kőnig's theorem that suits our purposes.

THEOREM 1.4. — *Let $\mathcal{G} = (X_1, X_2, E, \nu)$ be a d-regular bipartite graphing. For every $\varepsilon > 0$ there is a measurable proper edge coloring of $\mathcal{G}$ with $d + 1$ colors, such that the measure of edges of the $(d + 1)^{\text{th}}$ color is at most ε.*

Combining Theorems 1.3 and 1.4 with a subsequential weak limit argument gives the proof of Theorem 1.1.

We can formulate versions of both Theorem 1.1 and 1.4 using local isomorphisms. We say a graphing $\mathcal{G}_1 = (X_1, E_1, \nu_1)$ is the *local isomorphic image* of another graphing $\mathcal{G}_2 = (X_2, E_2, \nu_2)$ if their is a measure preserving map $\varphi : X_2 \to X_1$ that is a rooted isomorphism restricted to the connected component of almost all vertices of $\mathcal{G}_2$. It is clear that in this case $\mathcal{G}_1$ and $\mathcal{G}_2$ represent the same unimodular random graph. Theorem 1.1 can be reformulated as follows.

COROLLARY 1.5. — *Every $2d$-regular graphing $\mathcal{G}$ is the local isomorphic image of some graphing $\mathcal{G}'$ that comes from a p.m.p. action of F_d. That is, the edges of $\mathcal{G}'$ have a measurable balanced orientation and coloring with $[d]$.*

This formulation is a priori stronger, it is easy to see that Corollary 1.5 implies Theorems 1.1 (and 1.3) directly. The other implication is proved in Section 6.

Similarly, Theorem 1.4 has the following version.

COROLLARY 1.6. — Let $\mathcal{G} = (X_1, X_2, E)$ be a bipartite graphing with maximum degree d. Then $\mathcal{G}$ is the local isomorphic image of a bipartite graphing that can be measurably edge-colored with d colors.

Both corollaries are instances of a more general correspondence between coloring results for unimodular random networks and graphings. We express this phenomenon in Theorem 6.3 in Section 6.

Remark. — For graphings being bipartite is a stronger assumption than containing no odd cycles. This distinction is not relevant for our results, Theorem 1.4 and Corollary 1.6 hold with the weaker assumption as well. The reason for stating our results this way is that the relevant graphings arising in this paper are automatically bipartite.

It is natural to ask if any of these random decorations can be constructed as factor of i.i.d. (FIID) processes. See for example [CL17] for the definition of FIID for deterministic graphs, the definition for random graphs is analogous.

QUESTION. — Which $2d$-regular unimodular random networks have a factor of i.i.d. Schreier decoration? Which even degree unimodular random networks have a factor of i.i.d. balanced orientations?

Note that the bi-infinite line has no FIID balanced orientation.

The structure of the paper is as follows. In Section 2 we give a detailed introduction to unimodularity, graphings, and explain the connection to Benjamini–Schramm convergence. We also study the finite statement to motivate our question. We consider orientations in Section 3 and prove Theorem 1.3. In Section 4 we turn to coloring the edges and prove Theorem 1.4. We put these together and prove Theorem 1.1 in Section 5. In Section 6 we prove Corollaries 1.5 and 1.6.

Acknowledgements

2. Graphings, unimodularity and local convergence

In this section we will introduce the main notions of the paper and provide background and motivation for Theorem 1.1. A thorough and beautifully written exposition of the notions introduced here can be found in [Lov12].

Throughout this paper graphs will be simple; they do not contain loops or multiple edges. We will use standard letters $F, G, \dots$ to denote countable graphs, calligraphic letters $\mathcal{G}, \mathcal{H}, \dots$ to denote graphings and $\mathfrak{G}$ with certain indices to denote spaces of

graphs. The measures $\mu, \mu', \tilde{\mu}, \ldots$ will denote random graphs, while $\nu, \nu', \tilde{\nu}, \ldots$ will denote measures on vertex or edge sets of graphings. To avoid confusion, decorations of edges will be *colors* and decorations of vertices will be *labels*.

The map Φ will always forget all the decoration of a graph. We do not burden the notation with indicating what type of decoration we are forgetting, this will always be clear from context. We consider Schreier graphs to be rooted and connected, unless explicitly stated otherwise. Properties of random graphs are always understood to hold with probability 1.

2.1. Local convergence

Recall that $\mathfrak{G}_\circ^{2d}$ denotes the space of rooted, connected graphs with degrees bounded by $2d$. For two such graphs $(G_1, o_1), (G_2, o_2) \in \mathfrak{G}_\circ^{2d}$ we define their distance as

$$d\big((G_1, o_1), (G_2, o_2)\big) = \frac{1}{2^r},$$

where r is the largest integer with $B_{G_1}(r, o_1) \cong B_{G_2}(r, o_2)$. $B_G(r, o)$ denotes the rooted ball in G around o of radius r, and the isomorphism has to map o_1 to o_2. Two rooted graphs are close, if they look the same in a large neighborhood of the root. If the balls above are isomorphic for all $r \in \mathbb{N}$, then (G_1, o_1) and (G_2, o_2) are isomorphic, and their distance is defined 0. With this distance $\mathfrak{G}_\circ^{2d}$ becomes a totally disconnected compact metric space [Lov12, Section 18.3.1].

Local (or Benjamini–Schramm [BS01]) convergence takes a connected finite graph F (with degree bound $2d$) and by choosing a uniform random root among its vertices turns it into the random rooted graph (F, o). If the graph is not connected, only the component of the random root is kept. This way we obtain a probability measure μ_F on $\mathfrak{G}_\circ^{2d}$. Formally we define the map $f : V(F) \to \mathfrak{G}_\circ^{2d}$, $f(v) = (C_F(v), v)$, where $C_F(v)$ denotes the component of F containing v. Then μ_F is simply the pushforward of the uniform measure on $V(F)$ onto $\mathfrak{G}_\circ^{2d}$ by f.

A sequence of graphs (F_n) is convergent, if the corresponding μ_{F_n} converge weakly in the space $\mathcal{M}(\mathfrak{G}_\circ^{2d})$ of probability measures on $\mathfrak{G}_\circ^{2d}$. In this case the weak limit μ is considered to be the limit of the finite graphs F_n. Since $\mathfrak{G}_\circ^{2d}$ is totally disconnected, weak convergence is equivalent to the convergence of the measures of the following clopen sets. For a finite rooted graph $(F, v) \in \mathfrak{G}_\circ^{2d}$ of radius r we define the clopen set

$$\mathfrak{G}_{(F,v)} = \left\{ (G, o) \in \mathfrak{G}_\circ^{2d} \,\middle|\, B_G(r, o) \cong (F, v) \right\}.$$

A sequence of measures $\mu_n \in \mathcal{M}(\mathfrak{G}_\circ^{2d})$ converges weakly to μ if and only if $\mu_n(\mathfrak{G}_{(F,v)}) \to \mu(\mathfrak{G}_{(F,v)})$ for all choices of (F, v).

While the μ_{F_n} are supported on finite rooted graphs, their limit μ can be supported on infinite graphs as well.

Throughout the paper various spaces of rooted and decorated graphs will appear. They are all metric spaces with rooted metrics that are analogous to the one on $\mathfrak{G}_\circ^{2d}$, with isomorphisms required to respect the decoration as well as the graph structure and the root. We believe that not defining these metrics explicitly will cause no confusion.

2.2. Sofic and unimodular measures

A central question of this limit theory is to determine which measures are limits of finite graphs. This leads to the notions of *sofic* and *unimodular* measures.

Soficity is easy to define: a measure $\mu \in \mathcal{M}(\mathfrak{G}^{2d}_\circ)$ is sofic if it is the limit of measures μ_{F_n} corresponding to finite graphs.

Unimodularity is more subtle, and we will only define it in the regular case. See for example [Lov12] for a general treatment. Let $\mu \in \mathcal{M}(\mathfrak{G}^{2d}_\circ)$ be $2d$-regular with probability 1. It is unimodular if it satisfies the following involution invariance property. We generate a random instance (G, o) of μ. We then choose a uniform random neighbor o' of o in G. We say that μ is unimodular, if the random birooted graph (G, o, o') is invariant with respect to inverting the order of the roots.

Formally we introduce the space $\mathfrak{G}^{2d}_{\circ\circ}$ of connected graphs with degrees bounded by $2d$, that have two distinguished vertices, the first and the second root. The topology is defined similarly to $\mathfrak{G}^{2d}_\circ$, two such birooted graphs are close, if the look the same in large neighborhoods around both roots. There is a natural involution $\iota : \mathfrak{G}^{2d}_{\circ\circ} \to \mathfrak{G}^{2d}_{\circ\circ}$ swapping the roots: $\iota(G, o, o') = (G, o', o)$. The random rooted graph $\mu \in \mathcal{M}(\mathfrak{G}^{2d}_\circ)$ determines a measure $\widetilde{\mu} \in \mathcal{M}(\mathfrak{G}^{2d}_{\circ\circ})$ as above, by choosing the second root uniformly among the neighbors of the original. We say μ is unimodular if $\iota_*\widetilde{\mu} = \widetilde{\mu}$.

We mention two important facts about unimodularity. First, measures μ_F coming from finite graphs are unimodular. (This is true without assuming F to be regular, but here we only defined unimodularity for regular graphs.) Indeed, unimodularity is just a reformulation of the simple random walk being not only stationary, but also reversible with respect to the uniform measure on the vertices of a finite regular graph.

Second, unimodularity is a closed condition in $\mathcal{M}(\mathfrak{G}^{2d}_\circ)$ with respect to weak convergence. The reason is that involution invariance can be checked using only the local neighborhood probabilities $\mu(\mathfrak{G}_{(F,v)})$.

These two facts imply that every sofic measure is unimodular. The converse however is an open question, it is not known whether there is a non-sofic unimodular measure.

2.3. Invariance in the finite case

To motivate Theorem 1.1 we explore how the invariance property appears in the finite case.

Our finite statement takes a connected finite graph F and turns it into an unrooted Schreier graph $\widetilde{F} = (F, \mathrm{or}, c)$, where or and c signify the orientation and coloring respectively. If we then choose a uniform random root, we get the random rooted Schreier graph (F, o, or, c), where $o \in V(F)$ is uniform random. We denote its distribution by $\mu'_{\widetilde{F}} \in \mathcal{M}(\mathfrak{G}^{\mathrm{Sch}}_\circ)$. The action of any generator s is by a bijection on $V(F)$, so the uniform distribution on $V(F)$ is invariant with respect to the action $F_d \curvearrowright V(F)$. This implies that $\mu'_{\widetilde{F}}$ is invariant.

The graph structure of $\mu'_{\widetilde{F}}$ without the orientation and coloring is the same as μ_F. Using our notation, we say $\Phi_*\mu'_{\widetilde{F}} = \mu_F$. These two phenomena, $\mu'_{\widetilde{F}}$ being invariant

and $\Phi_*\mu'_{\widetilde{F}} = \mu_F$ are what motivate the definition of an invariant random Schreier decorations.

With this in mind it is relatively easy to prove that sofic random rooted graphs have invariant random Schreier decorations. Finding Schreier decorations of the finite approximations, then taking a subsequential weak limit of the corresponding random rooted Schreier graphs solves the problem. Theorem 1.1 shows that the property of having an invariant random Schreier decoration does not distinguish soficity and unimodularity.

Remark. — For a fixed connected, $2d$-regular infinite rooted or unrooted graph finding a Schreier decoration is simply a compactness argument once one knows the finite version. In fact rooted (possibly infinite), connected Schreier graphs are in one-to-one correspondence with subgroups of F_d.

If one wants to build an invariant Schreier decoration however, one has to pick the individual Schreier decorations of the random instances of the unimodular random graph in a coherent way, so that the whole process becomes invariant. The additional freedom is that the choice can be random.

2.4. Graphings

Graphings have quite a few equivalent definitions. We are going to stick with [Lov12].

DEFINITION 2.1 (Graphing). — *Let X be a standard Borel space and let ν be a probability measure on the Borel sets in X. A* graphing *with degree bound $2d$ is a graph $\mathcal{G}$ on $V(\mathcal{G}) = X$ with Borel edge set $E(\mathcal{G}) \subset X \times X$ in which all degrees are at most $2d$ and*

$$(2.1) \qquad \int_A \deg(x, B)\, d\nu(x) = \int_B \deg(x, A)\, d\nu(x)$$

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

Formally the graphing consists of the quadruple $\mathcal{G} = (X, \mathcal{B}, E, \nu)$, where $\mathcal{B}$ denotes the Borel σ-algebra on X. We will suppress $\mathcal{B}$ from the notation.

Looking at the connected component of a random vertex will represent a unimodular random rooted graph. Recall that for $x \in X$ its connected component in $\mathcal{G}$ is denoted $C_{\mathcal{G}}(x)$. It is a basic property of Borel graphs that the map $f : X \to \mathfrak{G}_\circ^{2d}$, defined by $f(x) = (C_{\mathcal{G}}(x), x)$ is Borel. This way we can define $\mu_{\mathcal{G}} = f_*\nu$ similarly to the finite graph case.

The measure preserving property (2.1) is closely related to the involution invariance that we used to define unimodularity. In fact (2.1) implies that $\mu_{\mathcal{G}}$ is unimodular.

In the other direction the connection is a bit more subtle. Fix a unimodular random rooted graph $\mu \in \mathcal{M}(\mathfrak{G}_\circ^{2d})$. We say that a graphing $\mathcal{G}$ represents μ if $\mu = \mu_{\mathcal{G}}$.

THEOREM 2.2 ([AL07, Ele07]). — *Every unimodular random rooted graph is represented by some graphing.*

As mentioned in the Introduction, the choice of a representing graphing is not unique. Graphings with distinctly different global measurable structure can represent the same unimodular random rooted graph.

In [Lov12, Theorem 18.37] a construction of a representing graphing is carefully explained and motivated. The representing graphings described there always have the same vertex and edge sets, only the measure changes. The space is a vertex-labeled version of $\mathfrak{G}_\circ^{2d}$, with edges connecting graphs that are the same up to a displacement of the root to a neighbor. Starting from a unimodular random rooted graph μ one uses further random vertex labels to break all potential symmetries of any realizations. This results in a measure ν on the space of vertex-labeled rooted graphs. The unimodularity of μ implies that (2.1) is satisfied by ν. We will adapt this construction to our situation in Subsection 3.3.

2.5. Measuring edges of graphings

The property (2.1) of graphings allows one to measure set of edges in a meaningful way. Let $\tilde{\nu}$ be the measure on $E \subseteq X \times X$ obtained the following way. For a Borel subset $C \subseteq E$ of edges let

$$\tilde{\nu}(C) = \frac{1}{2} \int_X \deg_C(x) \; d\nu(x).$$

Here $\deg_C(x)$ denotes the degree of x using the edge set C, that is the number of edges in C adjacent to x. The normalization accounts for counting edges from both endpoints. Intuitively equation (2.1) expresses that edges have the same weight from both endpoints. For the edge set $C = E(A, B)$ of edges between A and B it says that measuring them from A and B gives the same answer. In our case the measure $\tilde{\nu}$ takes values in $[0, 2d]$.

2.6. Sparse vertex labeling of graphings

We will use vertex labeling to break local symmetries of graphings. Proposition 2.3 below is a basic tool when working with graphings.

We say a labeling of a graph is r-sparse, if vertices of distance at most r always have distinct labels. A 1-sparse labeling is traditionally called a *proper* labeling.

Concerning vertex labelings of Borel graphs Kechris, Solecki and Todorcevic [KST99, Theorem 4.6] show that if the degree is bounded by Δ, then there is a proper Borel labeling of the vertices with $\Delta + 1$ labels.

One can also show that introducing additional edges in a Borel graph connecting vertices at distance at most r gives a Borel graph. This together with the above coloring result has the following useful corollary.

PROPOSITION 2.3. — *For every bounded degree Borel graph $\mathcal{G}$ and $r \in \mathbb{N}$ there is an r-sparse Borel labeling $l : V \to [k]$ for some $k \in \mathbb{N}$.*

See for example [CLP16, Corollary 3.1] for a more detailed exposition.

3. Orientation

In this section we prove Theorem 1.3. At first we exhibit a way to choose a random balanced orientation of rooted, even degree infinite graphs. The important property will be that for any even degree graph G and vertices v_1, v_2 the random balanced orientations of (G, v_1) and (G, v_2) will be the same in distribution on the edges of G. We then explain how this allows one to construct a measurably orientable graphing to represent a unimodular random rooted graph.

3.1. Orienting even degree trees

Let (T, o) be an infinite (connected) rooted tree with even degrees. We describe a random balanced orientation of (T, o). For all vertices $v \in V(T)$ let $\deg(v) = 2d_v$. We start at the root, and orient d_o edges inwards and d_o edges outwards randomly. We proceed to the neighbors of o. If v is a neighbor of o, then one of the $2d_v$ edges at v is already oriented, say towards v. So from the $2d_v - 1$ undirected ones we orient $d_v - 1$ inwards and d_v outwards randomly, independently of the previous choice. We continue this procedure radially outwards from o, independently randomizing at each vertex. As T is a tree we will always have exactly 1 already oriented edge at each vertex. The result of this random procedure is a balanced orientation of (T, o).

LEMMA 3.1. — *The random orientation described above does not depend on the choice of root in T. More precisely, if v_1 and v_2 are two vertices of T, then the random orientations started at v_1 and v_2 are the same in distribution.*

Proof. — Let $F \subseteq T$ be a *saturated* finite subtree containing both v_1 and v_2 as internal vertices. By saturated we mean that vertices of F are either leaves of F, or have all their neighbors in F as well. The internal vertices of F are the non-leaves. Fix any orientation or of the edges of F that is balanced at internal vertices. We compute the probability of getting or of F when starting our procedure at v_1. List the internal vertices of F in a radially expanding order: $\{u_1, u_2, \ldots, u_n\}$, specifically $u_1 = v_1$. At u_1 we have to match 1 out of

$$\binom{2d_{u_1}}{d_{u_1}}$$

possible orientations. At each later u_i we have to match 1 out of the possible

$$\binom{2d_{u_i} - 1}{d_{u_i}}$$

orientations of the remaining edges. Note that this does not depend on whether the edge already oriented at u_i is inwards or outwards. In total, the probability of generating our fixed orientation from v_1 is

$$\frac{1}{\binom{2d_{u_1}}{d_{u_1}}} \prod_{i=2}^{n} \frac{1}{\binom{2d_{u_i}-1}{d_{u_i}}}.$$

Using that $\binom{2d_{u_1}}{d_{u_1}} = 2 \cdot \binom{2d_{u_1}-1}{d_{u_1}}$, we can rewrite this as

$$\frac{1}{2}\prod_{i=1}^{n}\frac{1}{\binom{2d_{u_i}-1}{d_{u_i}}}.$$

Notice that this later expression does not change if we reorder the u_i, so starting from v_2 we will get the same probability.

This holds for all finite subtrees F and orientations $\mathtt{or}$ as above. Let $\mathrm{Cyl}(F, \mathtt{or})$ denote the set of all orientations of T (not necessarily balanced) that coincide with $\mathtt{or}$ at internal vertices of F. We proved that the two random orientations assign equal measure to all such $\mathrm{Cyl}(F, \mathtt{or})$. (This holds also if $\mathtt{or}$ is not balanced at all internal vertices of F; then both random orientations assign measure 0 to $\mathrm{Cyl}(F, \mathtt{or})$.)

The sets $\mathrm{Cyl}(F, \mathtt{or})$ form a base of the product topology on the set of (not necessarily balanced) orientations of T. Therefore the two random processes are the same in distribution. □

3.2. Orienting even degree graphs

We now define a random balanced orientation of rooted, bounded even degree infinite graphs. Again, the random orientation will not depend on the choice of root in distribution.

Given such a graph (G, o), we will first orient cycles as long as there are any. Our procedure will consist of countably many stages, with each stage consisting of countably many steps. During the procedure, we will have both oriented and undirected edges. We say that at a certain point a cycle is undirected, if all its edges are undirected. Two cycles of the same length are neighbors, if they share an edge.

After stage n, with probability 1 there will be no undirected cycles of length at most n. At every step of stage n we consider currently undirected cycles of length n, and randomize a uniform $[0, 1]$ label for each. If an undirected cycle has a higher label than all of its neighbors, then it gets oriented, randomly in one of the two possible directions.

Since our graph has bounded degree, the number of neighbors of an undirected cycle considered during stage n is bounded. So as long as it remains undirected, it has a positive probability (bounded away from 0) to get a higher label than all its neighbors at every step. Therefore staying undirected after countably many steps happens with probability 0.

At every step the partial orientation is balanced. That is, the indegrees and outdegrees of vertices are the same. Clearly the undirected degree of each vertex is always even. By the end of our procedure with probability 1 the undirected edges contain no cycle, so they span a disjoint union of even degree trees. We can randomly orient these trees by choosing a root in each component: for example the closest one to o in G, or if there are multiple ones, then picking one uniformly at random. Then we use the random orientation described in Subsection 3.1 for the rooted trees.

We call the random balanced orientation of (G, o) described above the *canonical* random balanced orientation of (G, o). The distribution on the edges does not depend

on the choice of o. This is obvious for the cycle removal, and the random orientation of the trees was shown to be invariant under changing the root. In the next subsection we show how the canonical random orientation allows us to build a measurably orientable graphing for any unimodular random network.

3.3. Constructing orientable graphings

Theorem 2.2 states that any unimodular random network can be represented by a graphing. The construction thoroughly explained in [Lov12, Theorem 18.37] can be adapted to our situation. We outline our construction and the key ideas, but do not burden the reader with lengthy exposition.

The set of vertices of our graphing will be the space of $2d$-regular, oriented, connected rooted graphs with labels from $[0, 1]$ on each vertex. That is, every node is a $2d$-regular rooted graph (H, v) with an orientation or $: E(H) \to V(H)$ and a vertex labeling $l : V(H) \to [0, 1]$. Let $\mathfrak{G}_\circ^{\mathrm{or}, l}$ denote the set of such quadruples.

The Borel structure is generated by the following cylinder sets: for any $r \geqslant 0$, we fix the oriented isomorphism type of the ball with radius r about the root, and also for every vertex in the ball we specify a Borel set in $[0, 1]$ from which the weight is to be chosen.

For a node (H, v, or, l) we introduce an edge to the nodes (H, v', or, l), where v' runs through the neighbors of v in H. In other words, (H, v, or, l) and $(H', v', \mathrm{or}', l')$ are connected if there is an isomorphism (of unrooted graphs) $\eta : H \to H'$, that preserves the orientation and the vertex labels, and $\eta(v) \sim_{H'} v'$. (The $\sim_G$ symbol indicates being neighbors in the graph G.)

Starting from a unimodular measure μ on $\mathfrak{G}_\circ^{2d}$ we construct a measure $\bar{\mu}$ on $\mathfrak{G}_\circ^{\mathrm{or}, l}$. We first choose (G, o) according to μ, then put i.i.d. $[0, 1]$ labels on the vertices, and randomize an orientation according to the canonical balanced orientation of (G, o).

PROPOSITION 3.2. — *The Borel graph $\mathfrak{G}_\circ^{\mathrm{or}, l}$ with the measure $\bar{\mu}$ is a graphing, and it represents the unimodular random rooted graph μ. Moreover, it has a measurable balanced orientation.*

Proof. — The measure $\bar{\mu}$ will be involution invariant on $\mathfrak{G}_\circ^{\mathrm{or}, l}$, because μ is involution invariant on $\mathfrak{G}_\circ^{2d}$, and both the orientation an the labels are chosen in a way that is in distribution the same regardless of the position of the root. This makes sure that $(\mathfrak{G}_\circ^{\mathrm{or}, l}, \bar{\mu})$ is indeed a graphing.

The labeling breaks all symmetries of (G, o, or) almost surely, which makes sure that $(\mathfrak{G}_\circ^{\mathrm{or}, l}, \bar{\mu})$ represents μ.

Finally by construction the edges of $\mathfrak{G}_\circ^{\mathrm{or}, l}$ can be measurably oriented. The edge between (H, v, or, l) and (H, v', or, l) is oriented towards (H, v, or, l) if and only if (v, v') is oriented towards v in (H, v, or). $\qquad\square$

This finishes the proof of Theorem 1.3.

4. Measurable colorings of bipartite graphings

In this section we prove Theorem 1.4. We will repeatedly make use of two results of Csóka, Lippner and Pikhurko [CLP16, Theorems 1.9 and 4.2]. Their main result is the following

THEOREM 4.1. — *For every $d \geqslant 1$ there is $r_0 = r_0(d)$ such that if $\mathcal{G} = (X_1, X_2, E, \nu)$ is a bipartite graphing with maximum degree at most $d + 1$ such that the set J of vertices of degree $d + 1$ is r_0-sparse, then there is a measurable proper edge coloring of $\mathcal{G}$ with $d + 1$ colors.*

During the proof of Theorem 4.1 they use the following theorem as an inductive step.

THEOREM 4.2. — *For every $d > 2$ and $r > 1$, there is $r_1 = r_1(d, r)$ such that the following holds. Let $\mathcal{G} = (X_1, X_2, E, \nu)$ be a bipartite graphing with degree bound $d + 1$ that has no finite components. If the set $J \subseteq X_1 \cup X_2$ of vertices of degree $d + 1$ is r_1-sparse, then there is a Borel matching M such that, up to removing a null-set, $\mathcal{G} \setminus M$ has maximum degree at most d and its set of degree-d vertices is r-sparse.*

We will follow the induction and make sure the color $(d + 1)$ is used rarely. Note that vertices of high degree being sparse does not immediately imply that they have small density, there might be small connected components. However, if a component contains at least 2 such vertices, then the density can be bound in terms of the sparsity. Ruling out components with exactly one high degree vertex is key to our proof of Theorem 1.4.

We start by proving two lemmas to take care of finite components in graphings.

LEMMA 4.3. — *Let $G = (S, T, E)$ be a finite bipartite graph with degrees bounded by $(d + 1)$, and such that*
  (1) *the vertices of degree $(d + 1)$ are 3-sparse, and*
  (2) *there are at most $\rho \cdot |V(G)|$ vertices of degree $(d + 1)$ in total, $(0 \leqslant \rho \leqslant 1)$.*
Then there is a proper edge coloring of G with $(d + 1)$ colors such that there are at most $\rho \cdot |V(G)|$ edges of color $(d + 1)$.

Proof. — For each vertex of degree $(d + 1)$ choose an arbitrary adjacent edge and color it with $(d + 1)$. As these points are 3-sparse, this coloring is so far proper. We colored at most $\rho \cdot |V(G)|$ edges with color $(d + 1)$. The remaining graph has degrees at most d, so by Kőnig's line coloring theorem it can be properly colored with d colors. Notice that if $\rho = 0$ the statement is simply Kőnig's theorem. □

LEMMA 4.4. — *Let $\mathcal{G} = (X_1, X_2, E, \nu)$ be a bipartite graphing with all connected components finite and satisfying the conditions of Lemma 4.3. Then it has a Borel proper edge coloring with $(d + 1)$ colors such that the $\tilde{\nu}$-measure of edges of color $(d + 1)$ is at most ρ.*

Proof. — For $i = 1, 2, \ldots$ we edge color the components with exactly $i + 1$ vertices. For each i we fix an i-sparse Borel labeling $l : X_1 \cup X_2 \to [k_i]$. In components of size $i + 1$ all vertices have different label. For each isomorphism type of such labeled

graphs choose a way to color its edges as in Lemma 4.3, and apply this coloring consistently everywhere. At each step the color classes are Borel, and there are countably many steps, so the color classes in the end are also Borel.

At each step the measure of edges colored by $(d+1)$ is at most ρ times the measure of vertices with components of size $i+1$. As $\mu(X_1 \cup X_2) = 1$, the measure of edges colored $(d+1)$ in the end is at most ρ. $\qquad\square$

Remark. — Lemma 4.4 can also be proved using uniformization, specificly the Lusin-Novikov Uniformization Theorem.

The following proposition is the first step of our inductive procedure to prove Theorem 1.4.

PROPOSITION 4.5. — *Let $\mathcal{G} = (X_1, X_2, E, \nu)$ be a bipartite graphing with degrees 2 or 3. Assume that the vertices of degree 3 are r-sparse for $r \geqslant r_0(3)$, where r_0 is the constant from Theorem 4.1. Then it has a measurable proper edge coloring with 3 colors such that the $\tilde{\nu}$-measure of edges of the 3rd color is at most $4/r$.*

Proof. — We will use 3 colors: red, blue and purple, with as few purple edges as possible. It suffices to prove the theorem when $\mathcal{G}$ only has finite components, and also when it only has infinite components. In general $\mathcal{G}$ can be split into two measurable parts according to the components being finite or infinite, and measurably coloring each side with few purple edges.

Suppose all components are finite. Finite components that contain no vertices of degree 3 are even cycles, and can be properly edge-colored with red and blue, no purple edges needed.

We also claim that finite components that contain a degree 3 vertex contain at least two. Indeed, the sum of degrees has to be even. This implies that the size of such a finite component is at least $r+1$, as degree 3 vertices are r-sparse. We claim that the density of degree 3 vertices is at most $2/r$. Indeed, because of the sparsity the balls of radius $r/2$ around the degree 3 points are disjoint, and these balls contain at least $r/2$ points each. By Lemma 4.4 we have a Borel proper edge coloring with only $2/r$ proportion of edges colored purple.

Now assume that all components are infinite. We simply apply Theorem 4.1. We get a measurable proper edge coloring $c : E \rightarrow \{\text{red}, \text{blue}, \text{purple}\}$. We will modify this coloring in finitely many steps so that purple edges become sparse.

Call a purple edge *standard*, if both of its endpoints have degree 2. Our procedure will consist of $r+1$ steps. In step 0 we will make sure that the neighbors of standard purple edges have different color. Standard purple edges that are between two blue edges can be recolored red. We will now argue that this recoloring maintains measurability.

Let R, B, and P denote the edge sets $c^{-1}(\{\text{red}\})$, $c^{-1}(\{\text{blue}\})$, and $c^{-1}(\{\text{purple}\})$ respectively. For a measurable set of edges F let V_F denote the set of endpoints of F, that is

$$V_F = \left\{ x \in X_1 \cup X_2 \mid \exists\, y \in X_1 \cup X_2 \text{ s.t. } (x,y) \in F \right\}.$$

If F is Borel, then so is V_F since degrees in $\mathcal{G}$ are finite. Moreover, $\tilde{\nu}(F) = 0$ implies $\nu(V_F) = 0$ by the definition of $\tilde{\nu}$. Consequently if F is $\tilde{\nu}$-measurable, then V_F is ν-measurable.

Let $A \subseteq X_1 \cup X_2$ denote the set of vertices that are the endpoints of a blue and a standard purple edge. Measurability of A is shown by $A = V_P \cap V_B \setminus V_R$. The set of standard purple edges between two blue edges is $A^2 \cap P$, which is measurable. Therefore the recoloring preserves measurability. Similarly, standard purple edges between two red edges can be recolored blue. This finishes step 0.

Generally for $n \geqslant 1$, during step n we make sure that there are no standard purple edges at distance at most n. Since steps 1 through $n-1$ are already done, a path between two standard purple edges at distance n is colored by red and blue in an alternating way. Step 0 has made sure that the neighbors of these purple edges are colored differently.
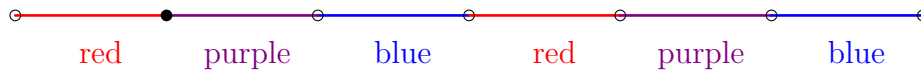

Figure 4.1. *Path between standard purple edges at distance 2.*

So swapping the colors on the path between the purple edges results in them having neighbors of the same color, and then we can recolor them red or blue according to the color of their neighbors. The case of $n = 2$ is illustrated in Figures 4.1 and 4.2.

Figure 4.2. *Path recolored.*

As before, we do not loose measurability. We execute steps $0, 1, \ldots, r$, and estimate the measure of purple edges at the end. Purple edges are either standard, or have an endpoint of degree 3. We divide the set $P = \{e \in E \mid c(e) = \text{"purple"}\}$ accordingly: $P_1 = \{e \in P \mid e \text{ is standard}\}$, $P_2 = P \setminus P_1 = \{e = (x, y) \in P \mid \deg(x) = 3 \text{ or } \deg(y) = 3\}$.

As points of degree 3 are r-sparse in infinite components they have measure at most $2/r$. Each such point has exactly one edge from P_2, so $\tilde{\nu}(P_2) \leqslant 2/r$. Edges in P_1 are also r-sparse, so $\tilde{\nu}(P_1) \leqslant 2/r$. We indeed proved that $\tilde{\nu}(P) \leqslant 4/r$. $\qquad\square$

Proof of Theorem 1.4. — The idea is to use Theorem 4.2 $(d-2)$ times to remove matchings from the edge set, coloring each one with a unique color. Whenever removing those matchings creates finite components we color those using Lemma 4.4. In the end we will apply Proposition 4.5 to the remaining graphing.

Instead of re-normalizing every time we split our graphings, we will talk about *density* of $(d+1)$-color edges. This means their $\tilde{\nu}$-measure divided by the ν-measure

of the vertex set of the graphing in question. At each step the density of $(d+1)$-color edges will be at most ε, so in the end their density will still be at most ε, and this coincides with their $\tilde{\nu}$-measure.

To be precise we start by picking $r_2 \geqslant r_0(3)$ such that $4/r_2 \leqslant \varepsilon$. Looking to apply Theorem 4.2 repeatedly we define $r_i \geqslant r_1(i, r_{i-1})$ for $i = 3, \dots d$, also assuming $r_{i-1} \leqslant r_i$.

The graphing $\mathcal{G}$ has no vertices of degree $(d+1)$, so in particular vertices of degree $(d+1)$ are r_d-sparse. The finite components can be colored with $[d]$ using Lemma 4.4 (choosing $\rho = 0$). So without loss of generality we can assume that all components are infinite. Then $\mathcal{G}$ satisfies the conditions of Theorem 4.2, so we find a measurable matching $M_1 \subseteq E(\mathcal{G})$, such that the set of vertices of degree d in $\mathcal{G}_1 = \mathcal{G} \setminus M_1$ is r_{d-1}-sparse. We color M_1 with color d, and turn our attention to $\mathcal{G}_1$. Notice that all degrees in $\mathcal{G}_1$ are either d or $d-1$.

We claim that all finite components of $\mathcal{G}_1$ that contain a vertex of degree d contain at least two. Assume towards contradiction that there was a finite component $G = (S, T, E')$, with a unique vertex x of degree d. Without loss of generality we can assume $x \in S$. Double counting the edges we get $(d-1)|S| + 1 = (d-1)|T|$, which is not possible if $d \geqslant 3$, as the right hand side is divisible by $(d-1)$, while the left hand side is not. This proves our claim.

This rules out small components with vertices of degree d. We use this to argue that in the finite components vertices of degree d are not only sparse, but have small density as well. Components having degree d vertices have size at least $r_{d-1} + 1$, since they have at least two vertices at distance at least r_{d-1}. The $r_{d-1}/2$ balls around degree d vertices are disjoint, and have size at least $r_{d-1}/2$ as the component has at least $r_{d-1} + 1$ points. So the density of degree d points in finite components of $\mathcal{G}_1$ is at most $2/r_{d-1}$.

We split $\mathcal{G}_1$ into the disjoint union $\mathcal{G}_1 = \mathcal{F}_1 \cup \mathcal{G}_1'$, where $\mathcal{F}_1$ is the part with finite components, and $\mathcal{G}_1'$ is the one with infinite components. To $\mathcal{F}_1$ we apply Lemma 4.4 with $\rho = 2/r_{d-1} \leqslant 4/r_2 \leqslant \varepsilon$. We use the colors $\{1, \dots, d-1\}$ for most edges, and $(d+1)$ with density at most ε.

To $\mathcal{G}_1'$ we apply Theorem 4.2, and find a measurable matching M_2 such that $\mathcal{G}_2 = \mathcal{G}_1' \setminus M_2$ has maximum degree $(d-1)$, and degree $(d-1)$ vertices in $\mathcal{G}_2$ are r_{d-2}-sparse. We color M_2 with color $d-1$. Notice that all vertices in $\mathcal{G}_2$ are of degree either $d-2$ or $d-1$.

As long as $i < d-2$ we will split the graphing $\mathcal{G}_i$ into the finite and infinite component parts $\mathcal{F}_i$ and $\mathcal{G}_i'$. We argue that degree $d-i+1$ vertices in $\mathcal{F}_i$ are not only sparse, but also of small density. The divisibility argument works, as we have $(d-i)|S| + 1 = (d-i)|T|$ being impossible. We color $\mathcal{F}_i$ using Lemma 4.4 with colors $\{1, \dots, d-i\}$ and $(d+1)$ used with density at most $2/r_{d-i} \leqslant 4/r_2 \leqslant \varepsilon$. Then we apply Theorem 4.2 to $\mathcal{G}_i'$ to find a measurable matching M_{i+1} such that $\mathcal{G}_{i+1} = \mathcal{G}_i \setminus M_{i+1}$ has maximum degree $d-1$ and its degree $(d-i)$ vertices are r_{d-i}-sparse. We color M_{i+1} with $(d-i)$.

After coloring $d-2$ matchings we end up with $\mathcal{G}_{d-2}$. This satisfies the conditions of Proposition 4.5, which finishes the coloring using colors $1, 2$ and $(d+1)$, the later

with density at most $4/r_2 \leqslant \varepsilon$. Color conflicts cannot arise as the M_i are matchings, and Proposition 4.5 also provides a proper edge coloring. $\qquad\square$

5. Proof of the main Theorem

In this section we complete the the proof of Theorem 1.1. Take any unimodular random rooted graph (G, o) with distribution $\mu \in \mathcal{M}(\mathfrak{G}_\circ^{2d})$ that is $2d$-regular with probability one. By Theorem 1.3 we choose a graphing $\mathcal{G} = (X, E, \nu)$ representing μ that can be measurably oriented in a balanced way. This balanced orientation is a measurable way of choosing a terminus of each edge, that is a symmetric measurable map or $: E \to X$ such that for all edges $e = (u, v) \in E$ we have $\mathrm{or}(e) \in \{u, v\}$.

Using the orientation we construct a d-regular bipartite graphing by doubling the vertex set and placing edges according to their orientation. To be precise let $Y = X \times \{1, 2\}$, $X_i = X \times \{i\}$ and consider the product measure ν' on Y, that is $\nu' = \nu \times u_{\{1,2\}}$, where $u_{\{1,2\}}$ is the uniform distribution on $\{1, 2\}$. We build a bipartite, undirected graphing $\mathcal{H}$ on the space (X_1, X_2, ν'). We set $((u, 1), (v, 2)) \in E(\mathcal{H})$ if and only if $(u, v) \in E(\mathcal{G})$ and $\mathrm{or}(u, v) = v$.

We set $1/n$ as the amount of error we allow with our coloring. By Theorem 1.4 we have a measurable proper edge coloring $c_0^n : E(\mathcal{H}) \to [d+1]$ using the last color infrequently, i.e. $\tilde{\nu}'(c_0^{-1}(d+1)) \leqslant 1/n$. The measure $\tilde{\nu}'$ is the edge measure on the graphing $\mathcal{H}$. The coloring c_0^n gives a coloring $c^n : E(\mathcal{G}) \to [d+1]$, as every oriented edge of $\mathcal{G}$ is represented as exactly one edge of H. For an edge $e = (u, v) \in E(\mathcal{G})$ let $\mathrm{or}^-(e)$ denote its source, that is $\mathrm{or}^-(e) = \{u, v\} \setminus \mathrm{or}(e)$. We set $c^n(e) = c_0^n((\mathrm{or}^-(e), 1), (\mathrm{or}(e), 2))$. Note that c^n is not a proper edge coloring. Instead it has the property that vertices have at most one incoming and one outgoing edge of each color.

We managed to measurably orient and color the edges of $\mathcal{G}$, although the number of colors is $(d+1)$ instead of the desired d. We now look at the connected component of a ν-random point x in $\mathcal{G}$, together with the orientation and coloring of edges, we get a random rooted, oriented and colored graph. Our next two lemmas claim that these are almost Schreier graphs and by passing to a weakly convergent subsequence we get an invariant random Schreier graph.

LEMMA 5.1. — *Let $(G_n, o_n, \mathrm{or}_n, c_n)$ denote the rooted, oriented, edge-colored component of a ν-random vertex o_n in $\mathcal{G}_n = (X, \nu, \mathrm{or}, c^n)$. Then*

$$\mathbb{P}\Big[\text{all edges around the root } o_n \text{ are colored by } [d]\Big] \geqslant 1 - \frac{4}{n}.$$

Proof. — Let $Y_{d+1} \subseteq Y$ denote vertices of $\mathcal{H}$ that encounter an edge colored $d+1$.

$$Y_{d+1} = \Big\{ y \in X_1 \cup X_2 \ \Big| \ \exists \, y' \sim_\mathcal{H} y \text{ with } c_0^n\big((y, y')\big) = d+1 \Big\}.$$

Our assumption of $(d+1)$-colored edges being infrequent implies $\nu(Y_{d+1}) \leqslant 2/n$.
Similarly, set

$$X_{d+1} = \Big\{ x \in X \ \Big| \ \exists \, x' \sim_\mathcal{G} x \text{ with } c^n\big((x, x')\big) = d+1 \Big\}.$$

By construction we have $x \in X_{d+1}$ if and only if $(x, 1) \in Y_{d+1}$ or $(x, 2) \in Y_{d+1}$. This gives $\nu(X_{d+1}) \leqslant 4/n$.

The set $X \setminus X_{d+1}$ is exactly the vertices that have all edges connected to them colored by $[d]$, so we have

$$\mathbb{P}\Big[\text{all edges around the root } o_n \text{ are colored by } [d]\Big] = \nu(X \setminus X_{d+1}) \geqslant 1 - \frac{4}{n}. \quad \square$$

By passing to a subsequence we can assume that $(G_n, o_n, \mathrm{or}_n, c_n)$ weakly converges to some random rooted, oriented, colored graph. Note that for all n the undirected, uncolored rooted graph (G_n, o_n) is the same as (G, o) in distribution, so the same holds for the limit. We will denote the limit by (G, o, or, c), which highlights the fact that it is a random orientation and coloring of the random rooted graph (G, o) we began with.

Remark. — In our notation logically $\mu'_{\mathcal{G}_n}$ should denote the distribution of $(G_n, o_n, \mathrm{or}_n, c_n)$, and μ' the subsequential weak limit. Notice however, that $\mu'_{\mathcal{G}_n} \notin \mathcal{M}(\mathfrak{G}_\circ^{\mathrm{Sch}})$. The weak convergence $\mu'_{\mathcal{G}_n} \to \mu'$ takes place in the larger space $\mathcal{M}(\mathfrak{G}_\circ^{2d, \mathrm{or}, c})$. Here $\mathfrak{G}_\circ^{2d, \mathrm{or}, c}$ denotes the space of rooted, connected graphs with degree bound $2d$, with a balanced orientation and a coloring of the edges by $d + 1$ colors, each vertex having at most one in- and outedge of each color. The forgetting function $\Phi : \mathfrak{G}_\circ^{2d, \mathrm{or}, c} \to \mathfrak{G}_\circ^{2d}$ still makes sense, and our construction ensured $\Phi_*(\mu'_{(\mathcal{G}_n)}) = \mu$. This implies $\Phi_* \mu' = \mu$, which we shorthanded into the notation (G, o, or, c).

A priori we do not know that $\mu' \in \mathcal{M}(\mathfrak{G}_\circ^{\mathrm{Sch}})$, and F_d only partially acts on $\mathfrak{G}_\circ^{2d, \mathrm{or}, c}$. This makes the statement of our next lemma somewhat cumbersome.

LEMMA 5.2. — *With probability 1 the root in (G, o, or, c) has one incoming and one outgoing edge of each color in $[d]$. For any $s \in [d]$ we define $s.o$ as the unique neighbor v of o with the edge (o, v) colored s and oriented towards v. Such a vertex exists almost surely. We claim that (G, o, or, c) and $(G, s.o, \mathrm{or}, c)$ are the same in distribution.*

In the lemma we only claim that edges at the root have the right colors. Together with the invariance however this implies that all edges have the right colors, so Lemma 5.2 finishes the proof of Theorem 1.1.

Proof. — The first part of the statement follows easily from Lemma 5.1. The probability of seeing one incoming and one outgoing edge of each color in $[d]$ adjacent to o in (G, o, or, c) is the limit of the probabilities of the same event for $(G_n, o_n, \mathrm{or}_n, c_n)$, which is clearly 1.

To show invariance with respect to moving the root, let α denote a finite (connected) rooted, oriented, $[d + 1]$ colored graph of radius r. Assume also that the $2d$ edges at the root are properly oriented an colored.

We have to show

$$\mathbb{P}_{(G, o, \mathrm{or}, c)}\Big[B_G(r, o) \cong \alpha\Big] = \mathbb{P}_{(G, s.o, \mathrm{or}, c)}\Big[B_G(r, s.o) \cong \alpha\Big].$$

The event $B_G(r, s.o) \cong \alpha$ can be seen by the $r + 1$ neighborhood of the root in (G, o). So we collect all connected, rooted, oriented and $[d + 1]$-colored finite graphs $\beta = (F_\beta, o_\beta, \mathrm{or}_\beta, c_\beta)$ of radius at most $r + 1$, where $B_{F_\beta}(r, s.o_\beta) \cong \alpha$. We implicitly

assumed that $s.o_\beta$ makes sense. Let $\mathbf{B}$ denote the set of all such β. Now we can express the right hand side of the above as

$$\mathbb{P}_{(G,\,s.o,\,\mathrm{or},\,c)}\Big[B_G(r,s.o) \cong \alpha\Big] = \mathbb{P}_{(G,o,\,\mathrm{or},\,c)}\Big[B_G(r+1,o) \in \mathbf{B}\Big].$$

Since these events correspond to clopen sets in the space of rooted, oriented, colored graphs we have

$$\mathbb{P}_{(G,o,\,\mathrm{or},\,c)}\Big[B_G(r,o) \cong \alpha\Big] = \lim_{n\to\infty} \mathbb{P}_{(G_n,o_n,\,\mathrm{or}_n,\,c_n)}\Big[B_{G_n}(r,o_n) \cong \alpha\Big],$$

$$\mathbb{P}_{(G,o,\,\mathrm{or},\,c)}\Big[B_G(r+1,o) \in \mathbf{B}\Big] = \lim_{n\to\infty} \mathbb{P}_{(G_n,o_n,\,\mathrm{or}_n,\,c_n)}\Big[B_{G_n}(r+1,o_n) \in \mathbf{B}\Big].$$

To compare the events $B_{G_n}(r,o_n) \cong \alpha$ and $B_{G_n}(r+1,o_n) \in \mathbf{B}$ we introduce the sets

$$X_n^\alpha = \Big\{x \in X \mid B_{(\mathcal{G},\,c^n)}(r,x) \cong \alpha\Big\},$$

$$X_n^s = \{x \in X \mid \text{ there is an outward } s\text{-edge of } (\mathcal{G},\mathrm{or},c^n) \text{ at } x\},$$

$$X_n^{\mathbf{B}} = \Big\{x \in X \mid B_{(\mathcal{G},\,c^n)}(r+1,x) \in \mathbf{B}\Big\}.$$

We also introduce the partial bijection φ_n^s defined on X_n^s, setting $\varphi_n^s(x) = y$ if (x,y) is the s-colored edge of $(\mathcal{G},\mathrm{or},c^n)$ oriented towards y. As $\mathcal{G}$ is a graphing, this partial bijection is measure preserving from X_n^s to $\varphi_n^s(X_n^s)$. Also $X_n^{\mathbf{B}} = \{x \in X_n^s \mid \varphi_n^s(x) \in X_n^\alpha\}$, so

$$\mu\left(X_n^{\mathbf{B}}\right) = \mu\Big((\varphi_n^s)^{-1}\left(\mathrm{ran}\,\varphi_n^s \cap X_n^\alpha\right)\Big) = \mu\left(\mathrm{ran}\,\varphi_n^s \cap X_n^\alpha\right).$$

We can rewrite the probabilities as follows:

$$\mathbb{P}_{(G_n,o_n,\,\mathrm{or}_n,\,c_n)}\Big[B_{G_n}(r,o_n) \cong \alpha\Big] = \mu\left(X_n^\alpha\right),$$

$$\mathbb{P}_{(G_n,o_n,\,\mathrm{or}_n,\,c_n)}\Big[B_{G_n}(r+1,o_n) \in \mathbf{B}\Big] = \mu\left(X_n^{\mathbf{B}}\right) = \mu\left(\mathrm{ran}\,\varphi_n^s \cap X_n^\alpha\right).$$

We have $\mu(X_n^\alpha) - \mu(\mathrm{ran}\,\varphi_n^s \cap X_n^\alpha) \leqslant 1 - \mu(\mathrm{ran}\,\varphi_n^s) = 1 - \mu(X_n^s) \to 0$, so

$$\lim_{n\to\infty} \mathbb{P}_{(G_n,o_n,\,\mathrm{or}_n,\,c_n)}\Big[B_{G_n}(r,o_n) \cong \alpha\Big] = \lim_{n\to\infty} \mathbb{P}_{(G_n,o_n,\,\mathrm{or}_n,\,c_n)}\Big[B_{G_n}(r+1,o_n) \in \mathbf{B}\Big].$$

By the equalities above this implies

$$\mathbb{P}_{(G,o,\,\mathrm{or},\,c)}\Big[B_G(r,o) \cong \alpha\Big] = \mathbb{P}_{(G,o,\,\mathrm{or},\,c)}\Big[B_G(r+1,o) \in \mathbf{B}\Big]$$

$$= \mathbb{P}_{(G,\,s.o,\,\mathrm{or},\,c)}\Big[B_G(r,s.o) \cong \alpha\Big],$$

which finishes the proof of Lemma 5.2. $\qquad\square$

6. Graphings from F_d actions

In this section we consider probability measure preserving (p.m.p.) actions of F_d, and prove Corollary 1.5.

Given a p.m.p. action of F_d on a standard Borel probability space (X,ν) the associated Schreier graph $\mathrm{Sch}(F_d \curvearrowright X, S)$ with the standard generating set S is a $2d$-regular graphing. The action of F_d provides a measurable orientation and coloring of the edges.

Not all $2d$-regular graphings come from p.m.p. actions of F_d. In fact finding a measurable orientation and coloring for the edges of a $2d$-regular graphing $\mathcal{G} = (X, E, \nu)$ is equivalent to finding such a p.m.p. action $F_d \curvearrowright (X, \nu)$ with $\mathrm{Sch}(F_d \curvearrowright X, S) = \mathcal{G}$ as unlabeled graphings.

When two graphings $\mathcal{G}_1$ and $\mathcal{G}_2$ represent the same unimodular random rooted graph, that is $\mu_{\mathcal{G}_1} = \mu_{\mathcal{G}_2}$ we say that they are *locally equivalent*. Recall that a graphing $\mathcal{G}_1 = (X_1, E_1, \nu_1)$ is the *local isomorphic image* of another graphing $\mathcal{G}_2 = (X_2, E_2, \nu_2)$ if there is a measure preserving map $\varphi : X_2 \to X_1$ such that $\varphi_* \nu_2 = \nu_1$, and $\left(C_{\mathcal{G}_2}(x), x \right) \cong \left(C_{\mathcal{G}_1}(\varphi(x)), \varphi(x) \right)$ for ν_2-almost all $x \in X_2$. It is clear that if $\mathcal{G}_1$ is a local isomorphic image of $\mathcal{G}_2$, then they are locally equivalent. However, local isomorphisms are not invertible in general. To this end we introduce a symmetric relation, *bi-local isomorphism*. We say $\mathcal{G}_1$ and $\mathcal{G}_2$ are bi-locally isomorphic if they are both local isomorphic images of some graphing $\mathcal{G}_3$.

Clearly, bi-local isomorphism implies local equivalence. We will exploit the fact that the converse also holds.

THEOREM 6.1. — *Two graphings are locally equivalent if and only if they are bi-locally isomorphic.*

See [Lov12, Theorem 18.59] for a detailed proof. Most of the argument is also present in [HLS14].

Proof of Corollary 1.5. — Let $\mathcal{G}$ be an arbitrary $2d$-regular graphing. By Theorem 1.1 the corresponding unimodular random rooted graph $\mu_{\mathcal{G}}$ has an invariant random Schreier decoration μ'. We use μ' to build a graphing that is locally equivalent to $\mathcal{G}$ with measurable orientation and coloring of the edges.

Let $\mathfrak{G}_\circ^{\mathrm{Sch},l}$ denote the space of Schreier graphs of F_d with vertices labeled by $[0, 1]$. The Borel structure and the edges are defined as in Subsection 3.3. By choosing uniform i.i.d. vertex labels for a μ'-random Schreier graph, we get a measure $\bar{\mu}$ on $\mathfrak{G}_\circ^{\mathrm{Sch},l}$ that turns it into a graphing locally equivalent to $\mathcal{G}$. The graphing property holds because μ' is invariant, and $(\mathfrak{G}_\circ^{\mathrm{Sch},l}, \bar{\mu})$ is locally equivalent to $\mathcal{G}$ because $\Phi_* \bar{\mu} = \Phi_* \mu' = \mu_{\mathcal{G}}$. It is clear that $(\mathfrak{G}_\circ^{\mathrm{Sch},l}, \bar{\mu})$ has a measurable orientation and coloring, as by construction it carries a p.m.p. action of F_d.

By Theorem 6.1 the graphings $\mathcal{G}$ and $(\mathfrak{G}_\circ^{\mathrm{Sch},l}, \bar{\mu})$ are bi-locally isomorphic, so they are both local isomorphic images of some graphing $\mathcal{G}' = (X', E', \nu')$. The measurable orientation and coloring can be pulled back by the local isomorphism $\varphi : X' \to \mathfrak{G}_\circ^{\mathrm{Sch},l}$, so $\mathcal{G}'$ satisfies the requirements of the Corollary 1.5. $\qquad\square$

6.1. Unimodular decorations

The argument in the previous subsection can be generalized to arbitrary decorations of graphs. To define an abstract decoration, let $D_1, \ldots D_n$ be compact standard Borel spaces serving as labels; this allows finite sets as well. Let $k_1, \ldots, k_n$ be natural numbers. We use D_i to label k_i-tuples of vertices.

A *decoration* of the graph $G = (V, E)$ is a finite set of functions $\mathcal{C} = \{c_1, \ldots, c_n\}$ with $c_i : V(G)^{k_i} \to D_i$. The decorated graph consists of the pair $(G, \mathcal{C})$.

Fix the sets of labels $\mathcal{D} = (D_1, \ldots, D_n)$, the size of tuples $\mathbf{k} = (k_1, \ldots, k_n)$ and a degree bound Δ. The space $\mathfrak{G}_\circ^{\Delta, \mathcal{D}, \mathbf{k}}$ of rooted, connected, decorated graphs $(G, o, \mathcal{C})$ with degree bound by Δ is a compact standard Borel space as before.

The Borel structure is generated by the following cylinder sets. For any $r \geqslant 0$ we fix the isomorphism type of the ball with radius r about the root. Also for every i and every k_i-tuple of vertices in the fixed ball we specify a Borel set in D_i from which its label is to be chosen.

Unimodular measures $\mu \in \mathcal{M}(\mathfrak{G}_\circ^{\Delta, \mathcal{D}, \mathbf{k}})$ can be defined exactly as before, by the involution invariance of the bi-rooted measure $\tilde{\mu} \in \mathcal{M}(\mathfrak{G}_{\circ\circ}^{\Delta, \mathcal{D}, \mathbf{k}})$ obtained by taking a step with the simple random walk on a μ-random sample of $\mathfrak{G}_\circ^{\Delta, \mathcal{D}, \mathbf{k}}$.

DEFINITION 6.2. — *Let $\mu \in \mathfrak{G}_\circ^\Delta$ be a unimodular random rooted graph. A unimodular random decoration of μ is a measure $\mu' \in \mathcal{M}(\mathfrak{G}_\circ^{\Delta, \mathcal{D}, \mathbf{k}})$ that is unimodular, and $\Phi_* \mu' = \mu$.*

This notion has been studied in the context of Schreier decorations by Biringer and Tamuz in [BT17]. They show that for a random rooted Schreier graph unimodularity is equivalent to invariance. Cannizzo also proves this in [Can13]. The equivalence shows how Definition 6.2 is a generalization of Definition 1.2 in the Introduction.

When decorating a graphing $\mathcal{G}$ the only formal difference is requiring measurability of the maps $c_i : V(\mathcal{G})^{k_i} \to D_i$. Each connected component becomes a decorated graph when restricting the decoration.

Since we are only interested in the connected components, the values of the c_i only matter on k_i-tuples where the vertices are in the same component. As the graph is measurable, changing the value on all tuples that are not contained in connected components does not change measurability of the c_i. So we can assume the c_i to be constant on such irrelevant tuples.

The following theorem translates between the two languages.

THEOREM 6.3. — *Let P be a property of rooted graphs, and Q be a property of rooted, decorated graphs. The following are equivalent.*
  (i) *Every unimodular random rooted graph μ that has property P almost surely has a unimodular random decoration μ' that has property Q almost surely.*
  (ii) *Every graphing $\mathcal{G}$ with almost every component having property P is the local isomorphic image of a graphing $\mathcal{G}'$ that has a measurable decoration with almost every component having property Q.*

Proof. — The (ii) implies (i) part is the easy one. We pick $\mathcal{G}$ to represent μ, then $\mathcal{G}'$ also represents μ. Together with the measurable decoration it defines $\mu' = \mu'_{(\mathcal{G}, \mathcal{C})}$.

The (i) implies (ii) part is proved exactly like Corollary 1.5. Starting from $\mathcal{G}$ we use (i) for $\mu_\mathcal{G}$ to find a locally equivalent graphing $\mathcal{G}_0$ that can be measurable decorated, and use Theorem 6.1 to find $\mathcal{G}'$ that has local isomorphisms into both. The decoration can be pulled back from $\mathcal{G}_0$ to $\mathcal{G}'$. $\square$

We now proceed to prove Corollary 1.6.

Proof of Corollary 1.6. — Our idea is to show that the unimodular random graph $\mu_\mathcal{G}$ has a unimodular random proper edge coloring with d colors, and then use the correspondence in Theorem 6.3.

First we add vertices and edges to our bipartite graphing $\mathcal{G}$, so that it remains bipartite and becomes d-regular. This can always be done. Then we color the edges using Theorem 1.4 with smaller and smaller weight of $(d + 1)$-color edges. The colorings c_n restricted to $\mathcal{G}$ give unimodular random colored graphs $\mu_{(\mathcal{G}, c_n)}$ when looking at the component of a random point. A subsequential weak limit will be a unimodular random edge coloring of $\mu_{\mathcal{G}}$.

We also want to make sure that the graphing we get in the end is bipartite. This is equivalent to the existence of a measurable proper vertex labeling with 2 labels. So we start with such a labeling $l : V(\mathcal{G}) \to \{1, 2\}$, and retain the information in the random graph, that is we consider $\mu_{(\mathcal{G}, c_n, l)}$. Now a subsequential weak limit will still be unimodular random, carrying both edge and vertex decoration. This makes sure that the graphing $\mathcal{G}'$ that has $\mathcal{G}$ as a local isomorphic image is both bipartite and the edges can be measurably colored. $\qquad\square$

BIBLIOGRAPHY

[AGV14]   Miklós Abért, Yair Glasner, and Bálint Virág, *Kesten's theorem for invariant random subgroups*, Duke Math. J. **163** (2014), no. 3, 465–488. ↑1706

[AL07]   David J. Aldous and Russell Lyons, *Processes on unimodular random networks*, Electron. J. Probab. **12** (2007), 1454–1508. ↑1711

[Bow12]   Lewis Bowen, *Invariant random subgroups of the free group*, `https://arxiv.org/abs/1204.5939v1`, 2012. ↑1707

[BS01]   Itai Benjamini and Oded Schramm, *Recurrence of Distributional Limits of Finite Planar Graphs*, Electron. J. Probab. **6** (2001), article no. 23. ↑1709

[BT17]   Ian Biringer and Omer Tamuz, *Unimodularity of invariant random subgroups*, Trans. Am. Math. Soc. **369** (2017), no. 6, 4043–4061. ↑1724

[Can13]   Jan Cannizzo, *On invariant Schreier structures*, `https://arxiv.org/abs/1309.5163v1`, 2013. ↑1706, 1724

[CK13]   Clinton T. Conley and Alexander S. Kechris, *Measurable chromatic and independence numbers for ergodic graphs and group actions*, Groups Geom. Dyn. **7** (2013), no. 1, 127–180. ↑1707

[CL17]   Endre Csóka and Gabor Lippner, *Invariant random perfect matchings in Cayley graphs*, Groups Geom. Dyn. **11** (2017), no. 1, 211–244. ↑1708

[CLP16]   Endre Csóka, Gabor Lippner, and Oleg Pikhurko, *Kőnig's line coloring and Vizing's theorems for graphings*, Forum Math. Sigma **4** (2016), article no. e27. ↑1707, 1712, 1716

[Ele07]   Gábor Elek, *On limits of finite graphs*, Combinatorica **27** (2007), no. 4, 503–507. ↑1711

[HLS14]   Hamed Hatami, László Lovász, and Balázs Szegedy, *Limits of locally–globally convergent graph sequences*, Geom. Funct. Anal. **24** (2014), no. 1, 269–296. ↑1723

[KM15]   Alexander S. Kechris and Andrew S. Marks, *Descriptive graph combinatorics*, 2015, preprint available at `http://www.math.caltech.edu/~kechris/papers/combinatorics20book.pdf`. ↑1707

[KST99]   Alexander S. Kechris, Slawomir Solecki, and Stevo Todorcevic, *Borel chromatic numbers*, Adv. Math. **141** (1999), no. 1, 1–44. ↑1712

[Lac88]   Miklós Laczkovich, *Closed sets without measurable matching*, Proc. Am. Math. Soc. **103** (1988), no. 3, 894–896. ↑1707

[Lov12]   László Lovász, *Large networks and graph limits*, Colloquium Publications, vol. 60, American Mathematical Society, 2012. ↑1708, 1709, 1710, 1711, 1712, 1715, 1723

László Márton TÓTH
Chair of Ergodic and
Geometric Group Theory,
EPFL SB MATH EGG, Station 8,
1015 Lausanne, (Switzerland)
laszlomarton.toth@epfl.ch