NONAMENABLE SUBFORESTS OF MULTI-ENDED QUASI-PMP GRAPHS

RUIYUAN CHEN, GRIGORY TERLOV, AND ANUSH TSERUNYAN

Abstract. We prove the a.e. nonamenability of quasi-pmp locally finite Borel graphs whose every component admits at least three nonvanishing ends with respect to the underlying Radon–Nikodym cocycle. We witness their nonamenability by constructing Borel subforests with at least three nonvanishing ends per component, and then applying Tserunyan and Tucker-Drob's recent characterization of amenability for acyclic quasi-pmp Borel graphs. Our main technique is a weighted cycle-cutting algorithm, which yields a weight-maximal spanning forest. We also introduce a random version of this forest, which generalizes the Free Minimal Spanning Forest, to capture nonunimodularity in the context of percolation theory.

Contents

Date: November 15, 2022.

2020 *Mathematics Subject Classification.* Primary 37A20, 03E15, 60K35; Secondary 37A40, 05C22, 60B99.

Key words and phrases. Borel graphs, amenable, countable Borel equivalence relations, quasi-pmp, nonsingular group actions, Radon–Nikodym cocycle, spanning forest, random forest, percolation.

The first author was partially supported by the NSF grant DMS-2224709, the second and third authors were partially supported by the NSF grant DMS-1855648, and the third author was also partially supported by the NSERC Discovery Grant RGPIN-2020-07120.

1. Introduction

Locally countable Borel graphs[1] on standard probability spaces have been a center of attention in a variety of areas such as Descriptive Graph Combinatorics, Measured Group Theory, Ergodic Theory, and the study of Countable Borel Equivalence Relations[2] (**CBERs**). The interest in them is largely facilitated by the fact that these graphs arise as Schreier graphs of measurable actions of countable groups, enabling graph-theoretic and algorithmic approaches, as well as techniques from probabilistic combinatorics and geometric group theory, for studying such actions and their orbit equivalence relations.

The majority of the developed theory concerns probability measure preserving (**pmp**) actions, where in particular, notions like amenability are well-understood. On the other hand, much less is known about **quasi-pmp** (nonsingular) actions, where points in the same orbit have different "relative weights", whence even free such actions may not reflect the properties of the acting group. For instance, the typical nonamenable group $\mathbb{F}_2$ has amenable free quasi-pmp actions (see Example 1.2). In this article – in the spirit of the Day–von Neumann question – we construct nonamenable subforests of certain locally finite quasi-pmp Borel graphs, in particular witnessing their nonamenability. We do so using the interplay between the geometry of these graphs (the space of ends) and the behavior of the underlying Radon–Nikodym cocycle, which we interpret as a relative weight function on the graph that accounts for the failure of invariance of the underlying measure.

Just like amenability is an important conceptual threshold for groups, **μ-amenability** is an equally fundamental concept for CBERs/group actions/graphs on a standard probability space (X, μ). The concept of μ-amenability was originally introduced in [Zim77], and a very useful equivalent definition was given in [CFW81]; see also [JKL02, 2.4]. By the Connes–Feldman–Weiss theorem [CFW81] (see also [KM04, 10.1] and [Mar17]), μ-amenability of CBERs is equivalent to **μ-hyperfiniteness**, i.e., being a countable increasing union of finite Borel equivalence relations[2] off of a μ-null set. In the present paper, we are only concerned with the measured context, so we use the terms **amenable** and **hyperfinite** interchangeably, dropping μ from notation when it is clear from the context. We also use these terms for Borel actions of countable groups and for locally countable Borel graphs on (X, μ), calling

[1]A **Borel graph** on a standard Borel space X is a graph with vertex set X and edge set a symmetric Borel subset of X^2.

[2]An equivalence relation E on a standard Borel space X is **countable** (resp. **finite**) **Borel** if it is Borel as subsets of X^2 and each equivalence class is countable (resp. finite).

them amenable or hyperfinite, when the induced orbit equivalence relation or connectedness relation, respectively, are so.

1.A. **Overview of nonamenability in the pmp setting.** A fundamental question in this line of research is to determine whether a given Borel action of a countable group on (X, μ) is amenable. By [JKL02, 2.5(i)], every such action is amenable when the group is amenable. For free probability-measure-preserving (**pmp**) actions, we also have the converse: the action is nonamenable when the group is nonamenable [JKL02, 2.5(ii)]. More generally, free pmp actions of countable groups tend to reflect the properties of the group. A big part of this is due to the theory of cost for pmp CBERs[3] – an analogue for CBERs of the free rank for groups – which is only available in the pmp setting. The **cost** of a pmp CBER E is defined as the infimum of the cost (i.e. half of expected degree) of its **graphings**, i.e. Borel graphs G on X whose connectedness relation $\mathbb{E}_G$ is equal to E a.e.

Fundamental results of Gaboriau [Gab98] and Hjorth [Hjo06] establish a strong analogy between free groups and and **treeable** pmp CBERs (those which admit acyclic graphings). In particular, an ergodic[4] treeable CBER is nonamenable exactly when its cost is > 1. By [Gab98], any acyclic graphing achieves the cost of a pmp CBER, so in the pmp context, acyclic Borel graphs of cost > 1 play the same role as free groups of rank > 1 among groups. In the spirit of the Day–von Neumann question as to whether every nonamenable group contains a copy of $\mathbb{F}_2$ (the free group on 2 generators), one tries to detect the nonamenability of a given pmp graph G by exhibiting a nonamenable acyclic Borel subforest of G, or at least of its connectedness relation $\mathbb{E}_G$. The following is a corollary of a theorem of Ghys [Ghy95], due to Gaboriau [Gab00, IV.24]:

Theorem 1.1 (Gaboriau–Ghys; 2000)**.** *The connectness relation $\mathbb{E}_G$ of an ergodic pmp graph G whose a.e. component has ≥ 3 ends is of cost > 1, hence $\mathbb{E}_G$ is nonamenable. In fact, there is an ergodic subforest $F \subseteq \mathbb{E}_G$ of cost > 1 witnessing the nonamenability of $\mathbb{E}_G$.*

This result is a key ingredient in one of the proofs of the celebrated Gaboriau–Lyons theorem [GL09], which gives a positive answer to the Day–von Neumann question for Bernoulli shifts: if Γ is a nonamenable countable group, then the orbit equivalence relation of its shift action on $([0, 1]^\Gamma, \lambda^\Gamma)$, where λ is the Lebesgue measure, admits an ergodic subequivalence relation induced by a free action of $\mathbb{F}_2$. Indeed, in this proof, percolation theory yields a subgraph of the orbit equivalence relation to which Theorem 1.1 applies, yielding an ergodic subforest of cost > 1, which is then upgraded to a free action of $\mathbb{F}_2$ by [Hjo06].

1.B. **Nonamenability in the quasi-pmp setting and our main result.** As for general CBERs on standard probability spaces, much less is known. By an argument of Kechris and Woodin, see [Mil04, Proposition 2.1] or [TZ22, 2.2], these equivalence relations are quasi-pmp[3] after discarding a null set. However, most of the aforementioned theory of pmp actions and CBERs fails in the quasi-pmp setting, starting with the statement that nonamenability of the group implies nonamenability of the orbit equivalence relations of its free actions.

Example 1.2. Consider the action of $\mathbb{F}_2$ on its boundary $\partial\mathbb{F}_2$, which we identify with the set of infinite reduced words in the symmetric set $\{a^{\pm 1}, b^{\pm 1}\}$ of generators of $\mathbb{F}_2$. This action

[3] A CBER E on a standard probability space (X, μ) is **pmp** (resp. **quasi-pmp** if for any Borel automorphism γ on X, that maps every point to an E-equivalent point, preserves μ (resp. μ-null sets)).

[4]An equivalence relation is **ergodic** if every invariant measurable set is either null or conull.

is free except at countably many points (which we discard), and it is (Borel) hyperfinite, by [DJK94, 8.2], because its orbit equivalence relation is equal to that induced by the one-sided shift map on $\partial\mathbb{F}_2$. This implies that there is no invariant probability measure on $\partial\mathbb{F}_2$, but there are certainly many quasi-invariant probability measures, e.g. the one with value $\frac{1}{4}\cdot(\frac{1}{3})^{n-1}$ on each cylindrical set based on a word of length n.

That properties of the group are not reflected by its free quasi-pmp actions is due to the fact that in the latter setting, points in the same orbit have different "relative weights." This is made precise by the **Radon–Nikodym cocycle** $(x, y) \mapsto \mathfrak{w}^y(x) : E \to \mathbb{R}^+$ of the orbit equivalence relation E with respect to the underlying probability measure μ, as defined in [KM04, Section 8]. In the present paper, we think of $\mathfrak{w}^y(x)$ as the weight of x relative to y, so we call cocycles to $\mathbb{R}^+$ **relative weight functions**, hence the notation $\mathfrak{w}$ (fraktur w). The Radon–Nikodym cocycle "corrects" the failure of invariance of the measure μ, enabling the (**tilted**) **mass transport principle**: for each $f : E \to [0, \infty]$,

$$\int \sum_{y\in[x]_E} f(x, y)d\mu(x) = \int \sum_{x\in[y]_E} f(x, y)\mathfrak{w}^y(x)d\mu(y).$$

In the absence of the theory of cost in the quasi-pmp setting, we look at the geometry of quasi-pmp graphs and the behaviour of the Radon–Nikodym cocycle along them. One has to first understand amenability for the simplest class of quasi-pmp graphs, namely, acyclic ones. This is done by Tserunyan and Tucker-Drob in [TTD22]. To present their characterization result in analogy with the pmp setting, we first note that for acyclic ergodic pmp graphs, we can replace cost by geometry, namely, the cost of such a graph is ≤ 1 if and only if it has ≤ 2 ends a.e. Thus, an acyclic ergodic pmp graph is amenable exactly when it has ≤ 2 ends a.e. (proven directly by Adams in [Ada90]). In [TTD22], this is generalized to the quasi-pmp setting as follows:

Theorem 1.3 (Tserunyan–Tucker-Drob; 2022)**.** *An acyclic quasi-pmp graph G is amenable exactly when it has ≤ 2 $\mathfrak{w}$-nonvanishing ends a.e., where $\mathfrak{w}$ is the Radon–Nikodym cocycle of $\mathbb{E}_G$ with respect to the underlying measure.*

Here, an end η of G is said to be **$\mathfrak{w}$-vanishing** if the cocycle $\mathfrak{w}$ converges to 0 along any sequence (x_n) of vertices converging to η, i.e. $\lim_n \mathfrak{w}^{x_0}(x_n) = 0$. Notice that in Example 1.2, each connected component of the canonical Schreier graph has exactly one $\mathfrak{w}$-nonvanishing end, namely, the forward direction of the one-sided shift map on $\partial\mathbb{F}_2$.

Theorem 1.3 indeed generalizes the pmp situation because $\mathfrak{w} \equiv 1$ in that case, so every end is nonvanishing. In light of it, we prove the following generalization and strengthening of Theorem 1.1:

Theorem 1.4. *Let G be a quasi-pmp locally countable Borel graph and let $\mathfrak{w}$ denote its Radon–Nikodym cocycle with respect to the underlying probability measure. If a.e. connected component of G has ≥ 3 $\mathfrak{w}$-nonvanishing ends, then there is a Borel subforest $F \subseteq G$ whose a.e. connected component has perfectly many $\mathfrak{w}$-nonvanishing ends. In particular, G is nowhere amenable. Moreover, F can be made ergodic if G is.*

Remark 1.5. If each connected component of G has *exactly* 2 $\mathfrak{w}$-nonvanishing ends, then G is hyperfinite by [Mil08a, 5.1], or by a simple geometric argument using a maximal disjoint set of $\mathfrak{w}$-bifurcations. As in the pmp case, we are unable to conclude anything if each component of G has exactly 1 $\mathfrak{w}$-nonvanishing end.

Remark 1.6. Theorem 1.4 is a strengthening of Theorem 1.1 even for pmp graphs because our ergodic forest F is a subgraph of G and not just of $\mathbb{E}_G$. However, this is only due to the fact that we now know, by a theorem of Tserunyan [Tse22] (which generalizes the analogous theorem of Tucker-Drob for pmp graphs), that every quasi-pmp ergodic graph admits an ergodic hyperfinite subgraph.

Remark 1.7. A significant part of the proof of Theorem 1.1, namely that by Ghys, involves an analogue for pmp graphs of the Stallings analysis of ends of groups [Sta68]. In contrast, our subforest F in Theorem 1.4 is constructed via a much simpler cycle-cutting algorithm, which runs simultaneously on all G-components and cuts the $\mathfrak{w}$-lightest edge in each simple cycle, using a fixed Borel linear ordering on edges as a tiebreaker.

In Section 5, we give concrete applications of Theorem 1.4 to **coinduced group actions** (Example 5.6) and **cluster graphings** of nonunimodular graphs (Corollary 5.7 and Example 5.9).

1.C. **Application to percolation theory: Free $\mathfrak{w}$-Maximal Spanning Forest.** Besides Theorem 1.4, our cycle-cutting algorithm described in Remark 1.7 has other applications, in particular to **random forests** in probability theory. Indeed, this algorithm works abstractly on any countable graph G equipped with a relative weight function $\mathfrak{w}$ and a linear ordering $<$ (tiebreaker) on the edges of G, yielding what we call the **$\mathfrak{w}$-maximal subforest** of G (with respect to the tiebreaker $<$). This is a generalization to the (relatively) weighted setting of the **minimal subforest** algorithm used in probability, which simply deletes the $<$-least edge in each cycle, regardless of its $\mathfrak{w}$-weight. In particular, just like the minimal subforest splits each cluster (i.e. connected component) of G into infinite trees, the $\mathfrak{w}$-maximal subforest does the same with **$\mathfrak{w}$-infinite** clusters, i.e. those whose total $\mathfrak{w}$-weight is infinite:

Proposition 1.8. *Let G be a graph with a relative weight function $\mathfrak{w}$ and a linear ordering $<$ on the edges. Every $\mathfrak{w}$-infinite component of G splits into $\mathfrak{w}$-infinite trees in the $\mathfrak{w}$-maximal subforest of G (with respect to the tiebreaker $<$).*

A crucial strengthening of Proposition 1.8 is proven in Lemma 3.4 and Observation 3.5.

Taking a uniformly random linear ordering (tiebreaker) on the edges of G, the minimal subforest algorithm yields the **Free Minimal Spanning Forest** (**FMSF**) of G – a well-known random subforest that has been useful in percolation theory of unimodular graphs. Analogously, taking a uniformly random linear ordering on the edges of G, our $\mathfrak{w}$-maximal subforest algorithm yields a random subforest of G, which we call the **Free $\mathfrak{w}$-Maximal Spanning Forest** and denote by $\mathrm{FMaxSF}_{\mathfrak{w}}(G)$.

FMSF is often applied to random subgraphs of a connected locally finite graph G, in particular, to an invariant (under automorphisms of G) bond or site percolation $\mathbf{P}$ on G, yielding an invariant random subforest $\mathrm{FMSF}(\mathbf{P})$ of G. If the relative weight function $\mathfrak{w}$ on $V(G)$ is invariant, then same is true about $\mathrm{FMaxSF}_{\mathfrak{w}}$: applied to an invariant percolation $\mathbf{P}$ on G, it yields an invariant random subforest $\mathrm{FMaxSF}_{\mathfrak{w}}(\mathbf{P})$ of G. Indeed, any closed subgroup Γ of the automorphism group $\mathrm{Aut}(G)$ (equipped with the pointwise convergence topology) induces a Γ-invariant relative weight function $\mathfrak{w}_\Gamma$ on the vertex set $V(G)$, namely, $\mathfrak{w}_\Gamma^y(x)$ is the ratio of the Haar measures of the stabilizers of x and y.

A graph G is called **unimodular** if its automorphism group $\mathrm{Aut}(G)$ is unimodular; equivalently, the relative weight function induced by $\Gamma := \mathrm{Aut}(G)$ is constant 1. While nonunimodularity makes some questions easier to answer [Hut20], many of the techniques developed for unimodular graphs do not extend to the nonunimodular setting leaving large gaps in the understanding of percolation on such graphs. This includes the techniques that make use of $\mathrm{FMSF}(\mathbf{P})$ because their correct analogues in this setting would need to involve the induced relative weight function, which $\mathrm{FMSF}(\mathbf{P})$ does not account for, while $\mathrm{FMaxSF}_{\mathfrak{w}}(\mathbf{P})$ does. Even when G is unimodular, it is possible to have a closed nonunimodular subgroup $\Gamma \leq \mathrm{Aut}(G)$ whose induced relative weight function $\mathfrak{w}_\Gamma$ is nonconstant. For instance, $\Gamma_\xi \leq \mathrm{Aut}(T_d)$ that fixes some specified end ξ of the d-regular tree T_d with $d \geq 3$. For more examples see [HPS99, Tim06, Hut20]. In such cases $\mathrm{FMaxSF}_{\mathfrak{w}_\Gamma}$ is again more applicable to a Γ-invariant percolation on G. The following result is a part of the statement of Theorem 4.15.

Theorem 1.9. *Let*

- *G be a connected locally finite graph;*
- *Γ be a closed subgroup of $\mathrm{Aut}(G)$ that acts transitively on G;*
- *$\mathfrak{w}_\Gamma$ be the Γ-invariant relative weight function on $V(G)$ induced by Γ;*
- *$\mathbf{P}$ be a Γ-invariant deletion tolerant Γ-ergodic percolation process on G, e.g. a Bernoulli percolation.*

If a $\mathbf{P}$-configuration almost surely has a cluster with ≥ 3 $\mathfrak{w}_\Gamma$-nonvanishing ends, then an $\mathrm{FMaxSF}_{\mathfrak{w}_\Gamma}(\mathbf{P})$-configuration almost surely has a cluster with perfectly many $\mathfrak{w}_\Gamma$-nonvanishing ends.

In Section 5.B, we present concrete examples where Theorem 1.9 applies (Corollary 5.7 and Example 5.9). We also derive a further combinatorial property called **infinite visibility** (introduced in [Tse22, 8.1]) of percolation configurations in such graphs (Theorem 5.11).

Organization. The rest of the paper organized as follows. In Section 2 we discuss preliminaries. In Section 3 we present the construction of the maximal forest in Borel setting and prove Theorem 1.4. Section 4 reviews the significance of random spanning forests in percolation theory and presents our construction in the context of this theory. Finally, in Section 5 we present applications of our main results on concrete examples.

Acknowledgments. We thank Russell Lyons and Robin Tucker-Drob for many insightful conversations.

2. Preliminaries

2.A. **Graphs.** Throughout this paper, by a **graph** we mean a simple undirected graph, represented formally as a symmetric *reflexive* binary edge relation $G \subseteq X^2$ on the vertex set X. We therefore write

$$(x, y) \in G \iff x \, G \, y$$

interchangeably to mean that there is an edge from x to y. We will refer to the graph by (X, G) or, when X is clear from context, simply by G.

Definition 2.1. By a **connected** graph (X, G), we will mean one whose *vertex set X is nonempty* and such that any $x, y \in X$ are joined by a path. A subset $A \subseteq X$ is **G-connected** if the induced subgraph $G|A$ is.

We write $\mathbb{E}_G \subseteq X^2$ for the **induced equivalence relation** relating two vertices iff they are joined by a path, and write $X/G := X/\mathbb{E}_G$ for the quotient set, i.e., connected components of G. Note that since we are considering reflexive graphs, $\mathbb{E}_G$ is itself a graph. For a subset $A \subseteq X$, we write $[A]_G := [A]_{\mathbb{E}_G}$ for its G**-saturation**, i.e. the union of all components intersecting A.

Definition 2.2. A **simple path** $x_0 \, G \, x_1 \, G \cdots G \, x_n$ will mean one with no repeated vertices. A **simple cycle** will mean such a path with $n \geq 3$, $x_0 = x_n$ and no other repeated vertices. (Recall that we are working with reflexive graphs.) **Acyclic** means there are no simple cycles.

A **forest** is an acyclic graph; a **tree** is a connected forest.

Definition 2.3. Given a graph (X, G) and subset $A \subseteq X$,

- the **inner (vertex) boundary** of A is the set of vertices in A adjacent to $X \setminus A$;
- the **outer (vertex) boundary** of A is the inner boundary of $X \setminus A$;
- the **edge boundary** of A is the set of all G-edges between two vertices in $A, X \setminus A$.

Note that if G is locally finite, then one of these notions of boundary of A is finite iff all are, in which case we say that A has **finite boundary** or is **boundary-finite**.

Remark 2.4. A connected locally finite graph has only countably many boundary-finite subsets.

2.B. **End spaces.** As our working definition of "ends of a graph", we find it convenient to take the following point-set topological approach. For the equivalence with the more concrete geometrical approach via rays, see Example 2.14 below.

Definition 2.5. For a connected locally finite graph (X, G), its **end compactification**

$$\widehat{X} = \widehat{X}^G \supseteq X$$

is the zero-dimensional Polish compactification of the discrete space X whose clopen sets are precisely the closures of boundary-finite $A \subseteq X$. Thus these closures, which we denote by $\widehat{A} \subseteq \widehat{X}$, form an open basis for $\widehat{X}$. Formally, we may construct $\widehat{X}$ by identifying points $u \in \widehat{X}$ with their neighborhood filters, hence with ultrafilters of boundary-finite $A \subseteq X$ (in other words, $\widehat{X}$ is the Stone space of the Boolean algebra of all such A, into which X embeds as the principal ultrafilters).

The **end space** of (X, G) is

$$\partial X = \partial^G X := \widehat{X}^G \setminus X,$$

or equivalently the closed subspace of nonisolated points in $\widehat{X}^G$, which are the **ends** of G.

Remark 2.6. A clopen set $\widehat{A} \subseteq \widehat{X}$ contains an end iff $A \subseteq X$ is infinite, by compactness.

Remark 2.7. For boundary-finite $A \subseteq X$, it is easily seen that $B \subseteq A$ has finite boundary in G iff it has finite boundary in the induced subgraph $G|A$. In other words, the notation $\widehat{A}$ may also be consistently interpreted as the end compactification of $(A, G|A)$, which embeds into $\widehat{X}$ as a clopen subspace. Because of this, we will refer to an end $u \in \partial^G X$ which is in $\widehat{A}$ as an **end in** A or say that A **is a neighborhood of** u.

In contrast, *for boundary-infinite $Y \subseteq X$, we must carefully distinguish between $\widehat{Y}^{G|Y}$ and the closure $\overline{Y}$ of Y in $\widehat{X}^G$*. There is always a canonical map from the former space to the latter, but it need not be injective; see Example 2.15 below.

Remark 2.8. Every boundary-finite $A \subseteq X$ is the finite (disjoint) union of its connected components, which are also boundary-finite. Thus, the closures $\widehat{C} \subseteq \widehat{X}$ of *connected* boundary-finite $C \subseteq X$ also form an open basis for $\widehat{X}$.

Lemma 2.9. *The following families of subsets of X are the same:*

(i) connected boundary-finite $C \subseteq X$ such that $X \setminus C$ is also connected;
(ii) connected components C of $X \setminus F$ for some finite connected $F \subseteq X$.

The family of $\widehat{C} \subseteq \widehat{X}$ for all (infinite) such C form a neighborhood basis for each end $u \in \partial X$.

Proof. Clearly (ii) $\implies$ (i); to see the converse, let F be the outer boundary of C together with finitely many paths in $X \setminus C$ to make F connected. That such $\widehat{C}$ form a basis for ends is the trivial $n = 1$ case of Lemma 2.11 below. □

Definition 2.10. For a finite connected $F \subseteq X$, we call the components of $X \setminus F$ the **sides** of F.

For $n \in \mathbb{N}^+$, an n**-furcation** is a finite (nonempty) connected set $F \subseteq X$ with at least n infinite sides. An n**-furcation vertex** is a vertex x such that $\{x\}$ is an n-furcation. When $n = 2, 3$, we say **bifurcation**, **trifurcation** respectively. (Note that a trifurcation is also a bifurcation.)

Lemma 2.11. *$|\partial X| \geq n$ iff there is at least one n-furcation. In that case, for any n distinct ends $u_1, \dots, u_n \in \partial X$ and clopen neighborhoods $u_i \in \widehat{A}_i \subseteq \widehat{X}$, there is an n-furcation F and distinct sides $C_i \subseteq X \setminus F$ of it such that $u_i \in \widehat{C}_i \subseteq \widehat{A}_i$. (In other words, the products of distinct sides of n-furcations form a neighborhood basis for each pairwise distinct $(u_1, \dots, u_n) \in (\partial X)^n$.)*

Proof. If there is an n-furcation, then (the closures of) its $\geq n$ infinite sides each contain an end (by Remark 2.6). Conversely, if $u_1, \dots, u_n$ are distinct ends, each contained in a clopen neighborhood $\widehat{A}_i$, then we may find $u_i \in \widehat{B}_i \subseteq \widehat{A}_i$ such that the $\widehat{B}_i$ are pairwise disjoint, and let F be the union of the outer boundaries of the B_i together with finitely many paths to make F connected; then the u_i must belong to (the closures of) distinct sides C_i of F, whence F is an n-furcation. □

Definition 2.12. Let $f : X \to Y$ be a map between the vertex sets of two connected locally finite graphs (X, G) and (Y, H). We extend f by continuity to a (partial) map

$$\begin{aligned} \widehat{f} : \widehat{X}^G &\longrightarrow \widehat{Y}^H \\ u &\longmapsto \lim_{X \ni x \to u} f(x), \end{aligned}$$

where this limit exists; it clearly exists and equals $f(u)$ for vertices $u \in X$. If it also exists for every end $u \in \partial^G X$, we call $\widehat{f}$ the **map induced by** f (it is then automatically continuous). When (X, G) is a subgraph of (Y, H), we denote the inclusion map by $\iota : X \to Y$ and the induced map by $\widehat{\iota} : \widehat{X}^G \to \widehat{Y}^H$.

Lemma 2.13. *If $f : (X, G) \to (Y, H)$ is a finite-to-one graph homomorphism, in particular if f is the inclusion of a subgraph, then the induced map $\widehat{f}$ exists, and restricts to a map $\partial^G X \to \partial^H Y$.*

Proof. As $X \ni x \to u \in \partial X$, $f(x)$ cannot cluster around a vertex $y \in Y$, since $A := f^{-1}(y) \subseteq X$ is finite and so $\widehat{X \setminus A}$ is a neighborhood of u which f maps to $Y \setminus \{y\}$. It remains to rule out the possibility that $f(x)$ clusters around two distinct ends $v, w \in \partial Y$. Indeed, let $\widehat{A} \subseteq \widehat{Y}$ be a clopen set such that $v \in \widehat{A} \not\ni w$, with $A \subseteq Y$ boundary-finite. Since f is a graph homomorphism, f maps the inner boundary of $f^{-1}(A)$ into that of A; since f is also finite-to-one, $f^{-1}(A)$ thus has finite boundary, and so either $\widehat{f^{-1}(A)} \subseteq \widehat{X}$ or its complement is a neighborhood of u, but not both, which means $f(x)$ cannot cluster around both $v \in \widehat{A}$ and $w \in \widehat{Y} \setminus \widehat{A}$ as $x \to u$. □

Example 2.14. If R is the infinite ray graph $0 — 1 — 2 — \cdots$ on vertices $\mathbb{N}$, then an injective graph homomorphism $f : (\mathbb{N}, R) \to (X, G)$ takes the unique end of R to an end of G. It is easily seen that every end of G can be approached along a ray in this way; thus, ends may be equivalently represented as certain equivalence classes of rays.

Example 2.15. Even for the inclusion $\iota : (X, G) \to (Y, H)$ of a subgraph, with either the same vertex set $X = Y$ and a subset of edges $G \subseteq H$, or the induced subgraph $G = H|X$ on a subset of vertices $X \subseteq Y$, there is no reason for the induced map $\widehat{\iota} : \partial^G X \to \partial^H Y$ to be injective. The square lattice graph on $Y = \mathbb{Z}^2$ is one-ended; by removing either vertices or edges, we can turn it into a tree with $2^{\aleph_0}$ ends.

Lemma 2.16. *Under the assumptions of Lemma 2.13, if also f^{-1} preserves (nonempty) connected subsets (it is enough to check 1- and 2-element subsets), then $\widehat{f}$ restricts to a homeomorphism $\partial^G X \cong \partial^H Y$.*

Proof. Recall from Definition 2.1 that "connected" includes "nonempty"; thus f is surjective, hence so is $\widehat{f}$ by the density of $Y \subseteq \widehat{Y}^H$. To check injectivity: let $u, v \in \partial X$ such that $\widehat{f}(u) = \widehat{f}(v)$. Then for any finite connected $F \subseteq X$, since f is a graph homomorphism, $f(F) \subseteq Y$ is still connected (and finite); and $\widehat{f}(u) = \widehat{f}(v)$ lies on one side $D \subseteq Y \setminus f(F)$ of it. Since f^{-1} preserves connectedness, $f^{-1}(D) \subseteq X \setminus f^{-1}(f(F)) \subseteq X \setminus F$ is contained in one side of F, which thus contains both u, v. So u, v lie on the same side of every finite connected F, whence $u = v$ (by Lemma 2.9). □

Definition 2.17. For a possibly disconnected locally finite graph (X, G), we define its **end compactification**, respectively **end space**, to be the disjoint union of those of its components:

$$\widehat{X} = \widehat{X}^G := \bigsqcup_{C \in X/G} \widehat{C}^G,$$
$$\partial X = \partial^G X := \bigsqcup_{C \in X/G} \partial^G C = \widehat{X}^G \setminus X.$$

Note that these are locally compact spaces. The notions of n-**furcation** and **side** are interpreted as in Definition 2.10 within a single G-component.

For a map $f : X \to Y$ between the vertex sets of two such graphs $(X, G), (Y, H)$, we define the induced map $\widehat{f} : \widehat{X}^G \to \widehat{Y}^H$ componentwise (i.e., on $\widehat{C}^G$ for each $C \in X/G$) as in Definition 2.12. This map is guaranteed to exist everywhere if the conditions of Lemma 2.13 are satisfied componentwise, i.e., $f|C : C \to Y$ is a finite-to-one graph homomorphism for each $C \in X/G$.

2.C. **Weighted graphs and ends.** We denote by $\mathbb{R}^+$ the multiplicative group of positive reals.

Definition 2.18. A **weight function** on a graph (X, G) is an arbitrary function $\mathfrak{w} : X \to \mathbb{R}^+$. We often treat $\mathfrak{w}$ as an atomic measure on X, writing $\mathfrak{w}(A) := \sum_{x \in A} \mathfrak{w}(x)$ for a set $A \subseteq X$.

We call a set $A \subseteq X$ **$\mathfrak{w}$-finite** if $\mathfrak{w}(A) < \infty$; otherwise, we call it **$\mathfrak{w}$-infinite**. These notions are respectively called **$\mathfrak{w}$-light** and **$\mathfrak{w}$-heavy** in percolation theory.

Definition 2.19. Let $\mathfrak{w} : X \to \mathbb{R}^+$ be a weight function.

For an arbitrary subset $A \subseteq X$, we put

$$\limsup_A \mathfrak{w} = \limsup_{x \in A} \mathfrak{w}(x) := \inf_{\text{finite } F \subseteq A} \sup_{x \in A \setminus F} \mathfrak{w}(x) \in [0, \infty].$$

If this quantity is 0, we say A is **($\mathfrak{w}$-)vanishing**; otherwise A is **($\mathfrak{w}$-)nonvanishing**.

For an end $u \in \partial X$, we put

$$\widehat{\mathfrak{w}}(u) := \limsup_{x \to u} \mathfrak{w}(x) = \inf_{\widehat{A} \ni u} \sup_{x \in A} \mathfrak{w}(x) \in [0, \infty]$$

(where $\widehat{A}$ ranges over clopen neighborhoods of u). In other words, $\widehat{\mathfrak{w}} : \widehat{X} \to [0, \infty]$ is the minimal upper semicontinuous extension of $\mathfrak{w}$. If $\widehat{\mathfrak{w}}(u) = 0$, we say that the end u is **($\mathfrak{w}$-)vanishing**; otherwise u is **($\mathfrak{w}$-)nonvanishing**. Let

$$\partial_{\mathfrak{w}} X = \partial^G_{\mathfrak{w}} X := \{u \in \partial X \mid u \text{ is } \mathfrak{w}\text{-nonvanishing}\}.$$

Remark 2.20. By upper semicontinuity, $\partial_{\mathfrak{w}} X = \bigcup_n \widehat{\mathfrak{w}}^{-1}([1/n, \infty]) \subseteq \partial X$ is an F_σ subset. It may not be G_δ, as in Figure 2.21 where it is a countable dense set. In other words, $\partial_{\mathfrak{w}} X \subseteq \partial X$ *with the subspace topology may not be Polish!*

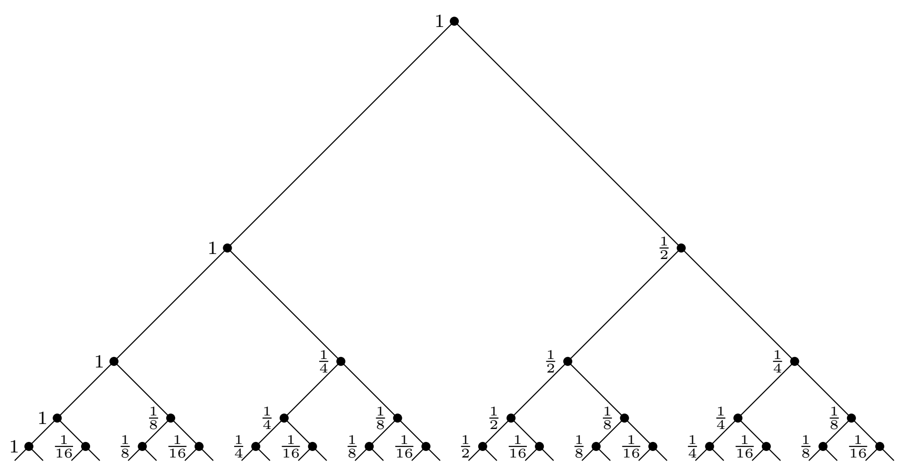

FIGURE 2.21. A weighted graph whose set of nonvanishing ends is F_σ but not G_δ

The notions of vanishing for sets and ends are related as follows (this is analogous to Remark 2.6):

Lemma 2.22.

(a) For an end $u \in \partial X$,

$$\widehat{\mathfrak{w}}(u) = \inf_{\widehat{A} \ni u} \limsup_A \mathfrak{w}.$$

Thus if u is nonvanishing, then every boundary-finite $A \subseteq X$ containing u is nonvanishing.

(b) For an infinite boundary-finite $A \subseteq X$ contained in a single G-component,

$$\limsup_A \mathfrak{w} = \max_{u \in \widehat{A}} \widehat{\mathfrak{w}}(u).$$

Thus A is nonvanishing iff it contains a nonvanishing end.

Proof. (a) Clearly $\widehat{\mathfrak{w}}(u) = \inf_{\widehat{A} \ni u} \sup_A \mathfrak{w} \geq \inf_{\widehat{A} \ni u} \limsup_A \mathfrak{w}$; conversely, for each $\widehat{A} \ni u$, we have $\widehat{\mathfrak{w}}(u) = \inf_{\widehat{B} \ni u} \sup_B \mathfrak{w} \leq \limsup_A \mathfrak{w}$, since A minus any finite set contains a neighborhood B of u.

(b) First, note that by the compactness of $\widehat{A}$ (because A is contained in a single G-component) and upper semicontinuity of $\widehat{\mathfrak{w}}$, the maximum is achieved. Now $\geq$ follows from (a). To show $\leq$: if $\limsup_A \mathfrak{w} > 0$, then there is a sequence of distinct vertices $x_0, x_1, \ldots \in A$ such that $\lim_{n \to \infty} \mathfrak{w}(x_n) = \limsup_A \mathfrak{w}$; a subsequence of these converges to some end $u \in \widehat{A}$ with $\widehat{\mathfrak{w}}(u) \geq \limsup_A \mathfrak{w}$. $\square$

Remark 2.23. The converse of the last statement of Lemma 2.22(a) is false, as shown by the example in Figure 2.21. In fact, Lemma 2.22(b) shows that every neighborhood of an end u is nonvanishing iff u belongs to the *closure* $\overline{\partial_{\mathfrak{w}} X} \subseteq \partial X$ of the set of nonvanishing ends.

Definition 2.24 (cf. Definition 2.10)**.** A **$\mathfrak{w}$-n-furcation** is a finite connected $F \subseteq X$ with at least n nonvanishing sides $C_1, \ldots, C_n \subseteq [F]_G \setminus F$. A **$\mathfrak{w}$-$n$-furcation vertex** is a singleton $\mathfrak{w}$-n-furcation. When $n = 2, 3$, we say **$\mathfrak{w}$-bifurcation**, **$\mathfrak{w}$-trifurcation** respectively.

Lemma 2.25 (cf. Lemma 2.11)**.** *For connected G, $|\partial_{\mathfrak{w}} X| \geq n$ iff there is at least one $\mathfrak{w}$-n-furcation. In that case, for any n distinct $\mathfrak{w}$-nonvanishing ends $u_1, \ldots, u_n \in \partial_{\mathfrak{w}} X$ and clopen neighborhoods $u_i \in \widehat{A}_i \subseteq \widehat{X}$, there is a $\mathfrak{w}$-n-furcation F and distinct sides $C_i \subseteq X \setminus F$ of it such that $u_i \in \widehat{C}_i \subseteq \widehat{A}_i$. (In other words, the products of distinct sides of $\mathfrak{w}$-n-furcations form a neighborhood basis for each pairwise distinct $(u_1, \ldots, u_n) \in (\partial_{\mathfrak{w}} X)^n$.)*

Proof. Analogous to Lemma 2.11, using Lemma 2.22 in place of Remark 2.6. $\square$

Lemma 2.26. *Let a finite-to-one homomorphism between connected locally finite graphs $f : (X, G) \to (Y, H)$ induce $\widehat{f} : \widehat{X}^G \to \widehat{Y}^H$ as in Lemma 2.13. For a weight function $\mathfrak{w} : X \to \mathbb{R}^+$, put*

$$\begin{aligned} \sup\nolimits_f \mathfrak{w} : Y &\longrightarrow \mathbb{R}^+ \\ y &\longmapsto \sup_{x \in f^{-1}(y)} \mathfrak{w}(x). \end{aligned}$$

Then for an end $v \in \partial^H Y$,

$$(*) \qquad \widehat{\sup_f \mathfrak{w}}(v) = \sup_{u \in \widehat{f}^{-1}(v)} \widehat{\mathfrak{w}}(u).$$

Thus for $u \in \partial^G X$,

$$\widehat{\mathfrak{w}}(u) \leq \widehat{\sup_f \mathfrak{w}}(\widehat{f}(u)).$$

In particular, $\widehat{f}$ restricts to a map $\partial^G_{\mathfrak{w}} X \to \partial^H_{\sup_f \mathfrak{w}} Y$ between spaces of nonvanishing ends.

Proof. Both sides of $(*)$ define the least upper semicontinuous map $\widehat{Y}^H \to [0,\infty]$ whose composite with $f : X \to Y \subseteq \widehat{Y}^H$ is $\geq \mathfrak{w}$. $\square$

Corollary 2.27. *Let a componentwise finite-to-one homomorphism between (possibly disconnected) locally finite graphs $f : (X,G) \to (Y,H)$ induce $\widehat{f} : \widehat{X}^G \to \widehat{Y}^H$ as in Definition 2.17. For weight functions $\mathfrak{w}_G : X \to \mathbb{R}^+$ and $\mathfrak{w}_H : Y \to \mathbb{R}^+$ such that $\mathfrak{w}_G \leq \mathfrak{w}_H \circ f$, we have for each $u \in \partial^G X$,*

$$\widehat{\mathfrak{w}_G}(u) \leq \widehat{\mathfrak{w}_H}(\widehat{f}(u)).$$

In particular, $\widehat{f}$ restricts to a map $\partial^G_{\mathfrak{w}_G} X \to \partial^H_{\mathfrak{w}_H} Y$ between spaces of nonvanishing ends.

Proof. For each component $C \in X/G$ mapping into $D := [f(C)]_H \in Y/H$, since $\mathfrak{w}_G \leq \mathfrak{w}_H \circ f$, we have $\sup_{f|C}(\mathfrak{w}_G|C) \leq \mathfrak{w}_H|D$, so we may apply Lemma 2.26 to $f|C : C \to D$. $\square$

In light of Lemma 2.13 and Corollary 2.27, we use the following notion:

Definition 2.28. Let (Y,H) be a graph equipped with a weight function $\mathfrak{w}$. For a subgraph (X,G) of (Y,H), the **canonical maps** $\partial^G X \to \partial^H Y$ and $\partial^G_{\mathfrak{w}} X \to \partial^H_{\mathfrak{w}} Y$ are the restrictions of the map $\widehat{\iota} : \widehat{X}^G \to \widehat{Y}^H$ induced by the inclusion $\iota : X \to Y$. We also refer to the ι-images of ends of (X,G) as **canonical images**.

Remark 2.29. It will be important in what follows that all of the notions considered in this subsection are **homogeneous in $\mathfrak{w}$**, meaning preserved under scaling $\mathfrak{w}$ by any constant in $\mathbb{R}^+$.

2.D. **Relative weight functions (cocycles) on graphs.** In the sequel, we use the notion of $\mathfrak{w}$-vanishing sets and ends for graphs equipped with a *relative* weight function $\mathfrak{w}$, which we now define. Let (X,G) be a locally finite possibly disconnected graph.

Definition 2.30. An $\mathbb{R}^+$**-valued cocycle** or a **relative weight function** on G is a map $\mathfrak{w} : G \to \mathbb{R}^+$ satisfying the **cocycle identity**

$$\mathfrak{w}(x_0,x_1)\mathfrak{w}(x_1,x_2)\cdots\mathfrak{w}(x_{n-1},x_n) = 1 \quad \text{for any cycle } x_0 \mathrel{G} x_1 \mathrel{G} \cdots \mathrel{G} x_n = x_0.$$

Such $\mathfrak{w}$ then extends uniquely to a cocycle on the induced equivalence relation $\mathbb{E}_G$, which we also denote $\mathfrak{w}$, namely $\mathfrak{w}(x,y) := \mathfrak{w}(x_0,x_1)\cdots\mathfrak{w}(x_{n-1},x_n)$ for any path $x = x_0 \mathrel{G} x_1 \mathrel{G} \cdots \mathrel{G} x_n = y$.

For vertices x, y in the same component of G, we think of $\mathfrak{w}^y(x) := \mathfrak{w}(x,y)$ as the weight of x relative to y. Indeed, the map

$$\mathfrak{w}^y := \mathfrak{w}(-,y) : [y]_G \longrightarrow \mathbb{R}^+$$

is simply a weight function on the G-component of y, and these weight functions $\mathfrak{w}^y$ and $\mathfrak{w}^z$ for different basepoints y, z in the same G-component are constant multiples of each other by the cocycle identity $\mathfrak{w}^z = \mathfrak{w}^z(y)\mathfrak{w}^y$. Because of this, for a fixed G-component C, **homogeneous** statements about $\mathfrak{w}^b$ do not depend on the choice of the basepoint $b \in C$; for example:

- the definitions of $\mathfrak{w}^b$-finite, $\mathfrak{w}^b$-vanishing, $\mathfrak{w}^b$-n-furcation for sets and ends in C,
- $\mathfrak{w}^b(x) < \mathfrak{w}^b(y)$ for $x, y \in C$,
- $\min\{\mathfrak{w}^b(x), \mathfrak{w}^b(y)\} \leq \min\{\mathfrak{w}^b(u), \mathfrak{w}^b(v)\}$ for $x, y, u, v \in C$.

We drop c from the superscript in such ($\mathfrak{w}$-homogeneous) statements and simply write $\mathfrak{w}$, e.g. $\mathfrak{w}$-nonvanishing. In particular, per Remark 2.31, we may use the notions and statements of Section 2.C for a *relative* weight function $\mathfrak{w}$ on G.

2.E. **Borel and quasi-pmp graphs and equivalence relations.** Let (X, G) be a locally finite Borel graph, i.e., the vertex set X is a standard Borel space, and $G \subseteq X^2$ is Borel as a set of pairs.

Remark 2.31. In general, notions of **end space**, etc., for (X, G) are to be understood in the general sense of disconnected locally finite graphs, as in Definition 2.17. Thus for example, $\widehat{X}^G$ is typically a nonseparable locally compact Hausdorff space. Note that the topology on $\widehat{X}^G$ has nothing to do with any compatible Polish topology on X.

Definition 2.32. Let μ be a probability measure on X.

We say that μ is **(G-)quasi-invariant** (or that G is a **quasi-pmp** graph) if for every Borel μ-null $A \subseteq X$, $[A]_G$ is still μ-null.

For a Borel cocycle $\mathfrak{w} : G \to \mathbb{R}^+$, we say that μ is **$\mathfrak{w}$-invariant** if for any Borel sets $A, B \subseteq X$ and Borel bijection $\gamma : A \cong B$ with graph contained in G (i.e., perfect G-matching between A, B),

$$\mu(B) = \int_A \mathfrak{w}^x(\gamma(x))\, d\mu(x).$$

It follows that the same holds for γ with graph contained merely in $\mathbb{E}_G$.

In fact, it is enough to require this equation only for countably many Borel bijections γ whose graphs cover G. For instance, if G is the Schreier graph of a Borel action of a countable group $\Gamma \curvearrowright X$, then it is enough to require this for γ among the generators of Γ. See [KM04, 8.1, 2.1].

If μ is $\mathfrak{w}$-invariant, then it is clearly quasi-invariant. Conversely, every quasi-invariant μ is $\mathfrak{w}$-invariant for an essentially unique (mod μ-null) Borel cocycle $\mathfrak{w} : G \to \mathbb{R}^+$, called the **Radon–Nikodym cocycle** of $\mathbb{E}_G$ with respect to μ; see [KM04, 8.3].

A **countable Borel equivalence relation** (CBER) $E \subseteq X^2$ is a Borel equivalence relation with countable classes; see [Kec22] for general background. These are exactly the connectedness relations $\mathbb{E}_G$ of locally finite Borel graphs G [JKL02, remark after proof of 3.12].

A CBER E on X is called

- **smooth** if it has a **Borel transversal** $A \subseteq X$, meaning a Borel set containing exactly one element from each E-class.
- **hyperfinite** if it is an increasing union of finite Borel equivalence relations.
- **amenable** if there is a sequence of Borel functions $\lambda_n : E \to [0, 1]$ that are summable to 1 on each equivalence class and for all $(x, y) \in E$ we have that $\|\lambda_n(x, \cdot) - \lambda_n(y, \cdot)\|_1 \to 0$ as n tends to infinity.
- **treeable** if it admits an acyclic graphing, where a **graphing** of E is a Borel graph G on X whose connectedness relation $\mathbb{E}_G$ is E.

In the presence of a Borel probability measure μ on X, the notions of hyperfinite, amenable, and treeable are relaxed to μ-hyperfinite, μ-amenable, and μ-treeable by demanding that the the corresponding property holds off of a μ-null set. We also the notions of μ-amenability and μ-hyperfiniteness interchangeably because they are equivalent by the Connes–Feldman–Weiss theorem [CFW81, KM04, Mar17]. Finally, we often omit μ before these terms when it is clear from the context.

Lemma 2.33. *If (X, E) is a smooth countable Borel equivalence relation, $\mathfrak{w} : E \to \mathbb{R}^+$ is a Borel cocycle, and each E-class is $\mathfrak{w}$-infinite, then there are no $\mathfrak{w}$-invariant probability measures.*

This fact is well-known; see [Mil08b, 2.1], [Tse22, 5.6]. For the reader's convenience, we include the easy proof.

Proof. Let $A \subseteq X$ be a Borel transversal. By Lusin–Novikov uniformization [Kec95, 18.10], there are Borel maps $\gamma_0, \gamma_1, \ldots : A \to X$ such that for each $x \in A$, $(\gamma_i(x))_i$ is an injective enumeration of $[x]_E$. Then for any $\mathfrak{w}$-invariant μ,

$$\mu(X) = \sum_i \mu(\gamma_i(A)) = \sum_i \int_A \mathfrak{w}^x(\gamma_i(x))\, d\mu(x) = \int_A \sum_{y \in [x]_E} \mathfrak{w}^x(y)\, d\mu(x) = \int_A \infty\, d\mu(x). \quad \square$$

3. $\mathfrak{w}$-maximal subforests

In this section, we present our main cycle-cutting algorithm mentioned in Remark 1.7. We do so in several stages, starting with a connected locally finite graph equipped with a relative weight function and building our way up to locally finite quasi-pmp graphs.

3.A. **For a connected graph with enough trifurcation vertices.** Throughout this subsection, we let (X, G) be a connected locally finite graph with a relative weight function $\mathfrak{w} : G \to \mathbb{R}^+$. Fixing a basepoint $b \in X$, we get a genuine weight function $\mathfrak{w}^b$ on X, which we use below, omitting the superscript b from $\mathfrak{w}$-homogeneous statements as they do not depend on the choice of the basepoint b. Especially in this subsection, the reader can think of $\mathfrak{w}$ as a weight function on X without any harm.

Definition 3.1. We extend the weight function $\mathfrak{w}^b$ from X to (the edge-set of) G by setting

$$\mathfrak{w}^b(e) := \min\{\mathfrak{w}^b(x), \mathfrak{w}^b(y)\}$$

for an edge $e = \{x, y\} \in G$. Fix also an arbitrary linear ordering $<$ on the *undirected* G-edges. Define a new linear ordering on the undirected G-edges as follows: for $e_1, e_2 \in G$,

$$e_1 <_{\mathfrak{w}} e_2 \;:\Longleftrightarrow\; \mathfrak{w}(e_1) < \mathfrak{w}(e_2) \text{ or } \big[\mathfrak{w}(e_1) = \mathfrak{w}(e_2) \;\&\; e_1 < e_2\big].$$

We emphasize that the definition of $<_{\mathfrak{w}}$ does not depend on the basepoint b.

Definition 3.2. Let $<$ be a linear ordering on the set of edges of G, and let $H \subseteq G$ be an acyclic subgraph. The **$\mathfrak{w}$-maximal subforest of G (fixing H, with tiebreaker $<$)** is the subforest $H \subseteq M \subseteq G$ obtained from G by deleting the $<_{\mathfrak{w}}$-least edge not in H from each simple cycle.

All of our analysis of this subforest will be based on an abstract property it obeys, Lemma 3.4 below, which relates the subforest to the following notion:

Definition 3.3. A subset $Y \subseteq X$ is **(G-)cycle-invariant** if whenever it contains an edge in a simple G-cycle, it also contains the entire cycle.

For example, for any bifurcation vertex $x \in X$ and side $C \subseteq X \setminus \{x\}$ of x (cf. Definition 2.10), the subset $C \cup \{x\}$ is cycle-invariant. (This also trivially holds for non-bifurcation vertices.)

Lemma 3.4. *For any G-connected cycle-invariant $Y \subseteq X$, if the $\mathfrak{w}$-maximal subforest M is such that $M|Y$ is disconnected, then every $M|Y$-component is $\mathfrak{w}$-nonvanishing. In particular, if Y is $\mathfrak{w}$-nonvanishing, then so is every $M|Y$-component.*

Proof. Note the following key property of the $\mathfrak{w}$-maximal subforest construction: if we restrict both G and H to a G-cycle-invariant set Y (keeping the same relative weights $\mathfrak{w}$ and tiebreaker $<$), the maximal subforest we obtain is $M|Y$. Thus we may assume that $Y = X$.

We will show that if M is disconnected, then the *G-edge boundary of every M-component C is $\mathfrak{w}$-nonvanishing*. This will imply that C is itself $\mathfrak{w}$-nonvanishing, by local finiteness and our definition of the weight of an edge as the *minimum* of the weights of the incident vertices.

Since M is disconnected, there is an edge e in $G \setminus M$ between C and another M-component. For any such edge e, since e was deleted in M, it is the $<_{\mathfrak{w}}$-least edge not in H in a simple G-cycle, which must thus contain another edge between C and another M-component, which is $>_{\mathfrak{w}} e$ and also not in H. Hence, there is a strictly $<_{\mathfrak{w}}$-increasing sequence $e_0 <_{\mathfrak{w}} e_1 <_{\mathfrak{w}} \cdots$ on the G-edge boundary of C. Passing to a subsequence (using local finiteness), we may assume these edges are pairwise disjoint (nonadjacent). Then the endpoints of these edges in C are infinitely many vertices $x_0, x_1, \ldots$ on the inner boundary of C with $\mathfrak{w}(x_i) \geq \mathfrak{w}(e_i) \geq \mathfrak{w}(e_0)$, where the first inequality is again due to the weight of e_i is defined to be the minimum of that of its endpoints. Whence C is $\mathfrak{w}$-nonvanishing. ∎

Observation 3.5. Note that the above proof of Lemma 3.4 exhibits something stronger, namely if our cycle-cutting algorithm (Definition 3.2) disconnects a cycle-invariant G-connected set Y then there is a $<_{\mathfrak{w}}$-increasing sequence of pairwise disjoint edges on the boundary of each $M|Y$-component.

In the rest of this subsection, we will prove various combinatorial properties of the $\mathfrak{w}$-maximal subforest M; these proofs will only make use of Lemma 3.4, and not any other specific features of our construction. We therefore make the following

Hypothesis 3.6. Let $M \subseteq G$ be *any* subforest for which Lemma 3.4 holds.

One benefit of isolating this abstract property is:

Observation 3.7. If Lemma 3.4 holds for M, then it continues to hold if we replace $\mathfrak{w}$ by a different relative weight function $\widetilde{\mathfrak{w}}$ such that every $\mathfrak{w}$-nonvanishing subset is also $\widetilde{\mathfrak{w}}$-nonvanishing. In particular, we may take $\widetilde{\mathfrak{w}} \equiv 1$, yielding that the following results also hold for unweighted ends.

Lemma 3.8. *If (X, G) has a $\mathfrak{w}$-trifurcation vertex x, then the M-component of x has at least 3 $\mathfrak{w}$-nonvanishing M-ends.*

Proof. For each of the at least 3 $\mathfrak{w}$-nonvanishing sides $C \subseteq X \setminus \{x\}$ of x, we have a G-connected cycle-invariant set $C \cup \{x\} \subseteq X$, whence the $M|(C \cup \{x\})$-component of x is $\mathfrak{w}$-nonvanishing by Lemma 3.4, whence x is a $\mathfrak{w}$-trifurcation in its M-component, which therefore has at least 3 $\mathfrak{w}$-nonvanishing ends (by Lemma 2.25). ∎

Lemma 3.9. *Suppose every boundary-finite $A \subseteq X$ containing a $\mathfrak{w}$-bifurcation (of (X, G)) also contains a $\mathfrak{w}$-bifurcation vertex (of (X, G)). Then the canonical map $\partial_{\mathfrak{w}}^M X \to \partial_{\mathfrak{w}}^G X$ (Definition 2.28) has dense image.*

Proof. If $\partial^G_{\mathfrak{w}} X \neq \varnothing$, then $\partial^M_{\mathfrak{w}} X \neq \varnothing$ by Lemma 3.4 (and Lemma 2.22); this proves the case $|\partial^G_{\mathfrak{w}} X| = 1$. Now suppose $|\partial^G_{\mathfrak{w}} X| \geq 2$. Then a basic open set in $\partial^G_{\mathfrak{w}} X$ is given by $\partial^G_{\mathfrak{w}} X \cap \widehat{A}$ for a side A of a $\mathfrak{w}$-bifurcation (Lemma 2.25). Let $F \subseteq A$ be finite connected and containing the inner boundary of A; then each side of F is contained in either A or $X \setminus A$, and so F is a $\mathfrak{w}$-bifurcation. So A contains a $\mathfrak{w}$-bifurcation, hence also contains a $\mathfrak{w}$-bifurcation vertex x. At most one nonvanishing side $D \subseteq X \setminus \{x\}$ of x can contain the G-connected set $X \setminus A$; thus at least one nonvanishing side C of x is disjoint from $X \setminus A$, hence contained in A. So A contains the nonvanishing G-connected cycle-invariant set $C \cup \{x\}$, which has a nonvanishing M-end by Lemma 3.4 whose canonical image is in $\widehat{A}$. ☐

Definition 3.10. The **Cantor–Bendixson derivative** of a topological space X is the closed subspace $X' \subseteq X$ of nonisolated points.

Lemma 3.11. *Suppose every boundary-finite $A \subseteq X$ containing a $\mathfrak{w}$-trifurcation (of (X, G)) also contains a $\mathfrak{w}$-trifurcation vertex (of (X, G)). Then every neighborhood of a nonisolated $\mathfrak{w}$-nonvanishing G-end contains the canonical images of two distinct $\mathfrak{w}$-nonvanishing M-ends from a single M-component with at least 3 $\mathfrak{w}$-nonvanishing M-ends. In particular, if Y denotes the union of M-components with at least 3 $\mathfrak{w}$-nonvanishing M-ends, then the canonical image of $\partial^M_{\mathfrak{w}} Y$ is a dense subset of $(\partial^G_{\mathfrak{w}} X)'$.*

Proof. We may again assume that A is a side of a $\mathfrak{w}$-bifurcation. Since v is nonisolated in $\partial^G_{\mathfrak{w}} X$, $G|A$ has infinitely many nonvanishing ends. By applying Lemma 2.25 to three distinct nonvanishing ends of $G|A$ and clopen neighborhoods of them disjoint from the inner G-boundary of A, we get a finite connected $F \subseteq A$ containing the inner G-boundary of A and with at least 3 nonvanishing sides in A, hence also in X since F contains the inner G-boundary of A. Thus A contains a $\mathfrak{w}$-trifurcation F, hence also contains a $\mathfrak{w}$-trifurcation vertex x. Now as in the preceding lemma, at most one nonvanishing side of x can contain $X \setminus A$, hence at least two nonvanishing sides are contained in A, each of which has a nonvanishing M-end whose canonical image is in $\widehat{A}$. ☐

3.B. **For general connected graphs.** Given a connected locally finite graph (X, G) with a relative weight function $\mathfrak{w}$ and with many $\mathfrak{w}$-nonvanishing ends, there may not be any $\mathfrak{w}$-(bi/tri)furcation vertices. Our goal now is to "collapse" enough $\mathfrak{w}$-(bi/tri)furcation *sets* into $\mathfrak{w}$-(bi/tri)furcation *vertices*, and then apply the analysis of the preceding subsection to the resulting quotient graph.

The construction below is $\mathfrak{w}$-homogeneous, so we present it for a genuine weight function $\mathfrak{w} : X \to \mathbb{R}^+$ instead of a relative weight function, to avoid notational complications. Formally, the construction is done for $\mathfrak{w}^b$, where $b \in X$ is a fixed basepoint, observing that it does not depend on the choice of b.

Definition 3.12. Let (X, G) be a connected locally finite graph, $\mathcal{F}$ be a pairwise disjoint family of finite connected subsets $F \subseteq X$. Let $X/\mathcal{F}$ denote the quotient of X identifying all vertices in a single $F \in \mathcal{F}$; formally,

$$X/\mathcal{F} := \mathcal{F} \cup \{\{x\} \mid x \in X \setminus \bigcup \mathcal{F}\}.$$

Let $G/\mathcal{F}$ denote the G-adjacency graph on $X/\mathcal{F}$: for $F, F' \in X/\mathcal{F}$,

$$F \; G/\mathcal{F} \; F' \;\; :\iff \;\; \exists x \in F,\, y \in F' \, (x \; G \; y).$$

We call $(X/\mathcal{F}, G/\mathcal{F})$ the **quotient graph** of (X, G) by $\mathcal{F}$.

Given a weight function $\mathfrak{w} : X \to \mathbb{R}^+$, let $\mathfrak{w}/\mathcal{F} : X/\mathcal{F} \to \mathbb{R}^+$ be $\sup_\pi \mathfrak{w}$ as defined in Lemma 2.26, where $\pi : X \twoheadrightarrow X/\mathcal{F}$ is the quotient map; that is,

$$(\mathfrak{w}/\mathcal{F})(F) := \max_{x \in F} \mathfrak{w}(x).$$

By Lemma 2.16, π induces a homeomorphism

$$\widehat{\pi} : \partial^G X \cong \partial^{G/\mathcal{F}}(X/\mathcal{F}),$$

which by Lemma 2.26 takes $\widehat{\mathfrak{w}} : \partial^G X \to [0, \infty]$ to $\widehat{\mathfrak{w}/\mathcal{F}} : \partial^{G/\mathcal{F}}(X/\mathcal{F}) \to [0, \infty]$, thus restricts to

$$\widehat{\pi} : \partial^G_{\mathfrak{w}} X \cong \partial^{G/\mathcal{F}}_{\mathfrak{w}/\mathcal{F}}(X/\mathcal{F}). \tag{3.13}$$

Definition 3.14. Let (X, G) be a connected locally finite graph with a weight function $\mathfrak{w} : X \to \mathbb{R}^+$. Consider the following method for choosing a family $\mathcal{F}$ as above:

(1) First, take a maximal disjoint family of $\mathfrak{w}$-trifurcations.
(2) Next, enlarge to a maximal disjoint family of $\mathfrak{w}$-*bi*furcations.
(3) Finally, enlarge to a maximal disjoint family of (ordinary) bifurcations.

Let $M_{\mathcal{F}} \subseteq G/\mathcal{F}$ be a $\mathfrak{w}/\mathcal{F}$-maximal subforest constructed according to Definition 3.2, with respect to some (unspecified) fixed subforest $H \subseteq G/\mathcal{F}$ and tiebreaker linear ordering $<$ on the undirected $G/\mathcal{F}$-edges.

Finally, let $M \subseteq G$ be a subgraph defined by arbitrarily choosing a spanning tree on each (bi/tri)furcation $F \in \mathcal{F}$, and for each $M_{\mathcal{F}}$-edge between two different $F, F' \in X/\mathcal{F}$, arbitrarily choosing a single G-edge between them (which exists by the definition of $G/\mathcal{F} \supseteq M_{\mathcal{F}}$).

It is easily seen that M is then a forest, and that $M_{\mathcal{F}} = M/\mathcal{F}$ (with each $F \in \mathcal{F}$ an M-tree). The respective spaces of $\mathfrak{w}$-nonvanishing ends are related as follows:

$$\begin{CD}
\partial^M_{\mathfrak{w}} X @>{\widehat{\iota}}>> \partial^G_{\mathfrak{w}} X \\
@V{\widehat{\pi}^M}V{\cong}V @V{\cong}V{\widehat{\pi}^G}V \\
\partial^{M/\mathcal{F}}_{\mathfrak{w}}(X/\mathcal{F}) @>{\widehat{\iota/\mathcal{F}}}>> \partial^{G/\mathcal{F}}_{\mathfrak{w}}(X/\mathcal{F})
\end{CD} \tag{3.15}$$

Here, the horizontal maps are the canonical maps induced by the subgraph inclusions $\iota : (X, M) \to (X, G)$ and $\iota/\mathcal{F} : (X/\mathcal{F}, M/\mathcal{F}) \to (X/\mathcal{F}, G/\mathcal{F})$ (which preserve $\mathfrak{w}$-nonvanishing ends by Corollary 2.27), while the vertical homeomorphisms are induced by the quotient map $\pi : X \twoheadrightarrow X/\mathcal{F}$ as in (3.13). Since clearly $(\iota/\mathcal{F}) \circ \pi = \pi \circ \iota$, this square commutes.

We now have the following main result, summarizing the end-preservation properties of the maximal subforest construction for a single connected graph:

Theorem 3.16. *Let (X, G) be a connected locally finite graph with positive weight function $\mathfrak{w} : X \to \mathbb{R}^+$. The "collapsed maximal subforest" $M \subseteq G$ constructed as in Definition 3.14 has the following properties, where $\iota : (X, M) \to (X, G)$ is the inclusion:*

(a) $\widehat{\iota} : \partial^M X \to \partial^G X$ has dense image, as does its restriction $\widehat{\iota} : \partial^M_{\mathfrak{w}} X \to \partial^G_{\mathfrak{w}} X$.
(b) If G has ≥ 3 $\mathfrak{w}$-nonvanishing ends, then so does at least one component of M.
(c) Every neighborhood of a nonisolated $\mathfrak{w}$-nonvanishing G-end contains the canonical images of two distinct $\mathfrak{w}$-nonvanishing M-ends from a single M-component with at least 3 $\mathfrak{w}$-nonvanishing M-ends.

Proof. By Lemmas 3.8, 3.9 and 3.11, $\widehat{\iota/\mathcal{F}}$ has the claimed properties, given our choice of $\mathcal{F}$ in Definition 3.14; hence so does $\widehat{\iota}$ since the above square commutes. (To see the first part of (a), apply Lemma 3.9 with $\mathfrak{w}$ replaced by the constant function 1 (as noted in Observation 3.7); the hypothesis of Lemma 3.9 is still satisfied, by Definition 3.14(3).) □

One might expect that the properties stated in Theorem 3.16 can be strengthened in various ways; for instance, perhaps one could demand more of the M-components than merely "at least 3 nonvanishing ends". Indeed, we will show below that more can be said for almost every component of a quasi-pmp graph (see Corollary 3.22). However, the following shows that there are limitations to such strengthenings.

Example 3.17 (Windmill graph)**.** Let (X, G) be the graph depicted in Figure 3.18. Each "blade" of the windmill is a quadrant of the square lattice graph on $\mathbb{Z}^2$. The weight function $\mathfrak{w}$ is constant 1; thus all ends are nonvanishing. The big dot vertices are trifurcations, and already form a maximal disjoint family of bifurcations $\mathcal{F}$ as in Definition 3.14; thus there is no need to collapse. The tiebreaker linear ordering $<$ is chosen so that each "row" of dotted edges is strictly increasing, and each dotted edge is $<$ each solid edge. Then the solid edges are precisely those that belong to the maximal subforest M.

Now the original end space $\partial^G X$ of this graph is **perfect** (has no isolated points). But each M-component is just 3 rays joined at their basepoint, hence has exactly 3 ends. This shows that Theorem 3.16(c) is best possible in some sense. Moreover, by removing some of the "blades" from G, we can cause M to have infinitely many 2-ended components, thereby showing that Theorem 3.16(b) cannot be strengthened to "every component of M".

3.C. **For Borel and quasi-pmp graphs.** Let (X, G) be a locally finite Borel graph equipped with a Borel cocycle $\mathfrak{w} : G \to \mathbb{R}^+$. We recall from Remark 2.31 that $\partial^G X$, $\partial^G_{\mathfrak{w}} X$, etc. are interpreted as the (uncountable) disjoint unions of the end spaces of all components.

Theorem 3.19. *Let (X, G) be a locally finite Borel graph, $\mathfrak{w} : G \to \mathbb{R}^+$ be a Borel cocycle. There is a Borel subforest $M \subseteq G$ with the following properties, where $\iota : (X, M) \to (X, G)$ is the inclusion:*

- *(a) The induced $\widehat{\iota} : \partial^M X \to \partial^G X$ has dense image, as does its restriction $\widehat{\iota} : \partial^M_{\mathfrak{w}} X \to \partial^G_{\mathfrak{w}} X$.*
- *(b) Each G-component $C \in X/G$ with ≥ 3 nonvanishing G-ends contains at least one M-component with ≥ 3 nonvanishing M-ends.*
- *(c) For every nonisolated nonvanishing G-end u, every clopen neighborhood $\widehat{A}$ of u contains the canonical image of at least two distinct nonvanishing M-ends from a single M-component with at least 3 nonvanishing M-ends.*

Proof. This follows from implementing the algorithm of Definition 3.14 in a Borel manner on each G-component. In detail, the maximal family $\mathcal{F}$ of furcations in that algorithm may be chosen in a Borel manner (see [KM04, 7.3]), since the notions of "$\mathfrak{w}$-trifurcation", etc., are clearly Borel. This yields a finite, hence smooth, Borel subequivalence relation $\sim_{\mathcal{F}} \subseteq \mathbb{E}_G$, whose standard Borel quotient $X/{\sim_{\mathcal{F}}}$ yields on each G-component the quotient $X/\mathcal{F}$ from Definition 3.12.

Let $Y \subseteq X$ be a Borel transversal for $\sim_{\mathcal{F}}$, choosing from each $F \in X/\mathcal{F}$ a single element with maximum $\mathfrak{w}$-weight (i.e., maximum $\mathfrak{w}^x$-weight for any $x \in F$). Define now the cocycle $\mathfrak{w}/\mathcal{F}$ on $G/\mathcal{F} \subseteq (X/\mathcal{F})^2$, by identifying $X/\mathcal{F}$ with Y and then taking the restriction of

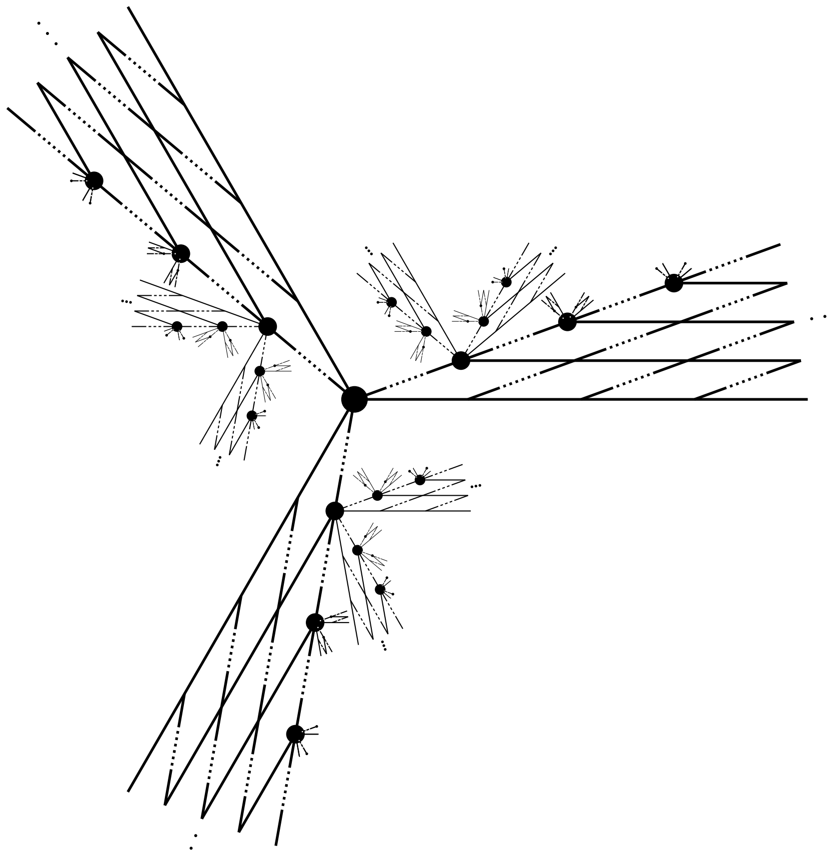

FIGURE 3.18. Windmill graph described in Example 3.17

$\mathfrak{w} : \mathbb{E}_G \to \mathbb{R}^+$ to Y. In other words, for $F, F' \in X/\mathcal{F}$, we define $(\mathfrak{w}/\mathcal{F})(F, F')$ to be $\mathfrak{w}(x, y)$ for $\mathfrak{w}$-heaviest elements $x \in F$ and $y \in F'$. Then for $\mathfrak{w}$-heaviest x in F, the weight function $(\mathfrak{w}/\mathcal{F})^F : [F]_{G/\mathcal{F}} \to \mathbb{R}^+$ will be exactly the quotient weight function $\mathfrak{w}^x/\mathcal{F}$ from Definition 3.12.

So we have defined a quotient Borel graph $(X/\mathcal{F}, G/\mathcal{F})$ with cocycle $\mathfrak{w}/\mathcal{F} : G/\mathcal{F} \to \mathbb{R}^+$, which on each G-component is exactly the quotient graph from Definition 3.12. We may now construct the $\mathfrak{w}/\mathcal{F}$-maximal subforest $M/\mathcal{F} \subseteq G/\mathcal{F}$ in a Borel manner as in Definition 3.2 (with any Borel subforest H of $G/\mathcal{F}$, e.g., $H := \varnothing$, and any Borel tiebreaker linear ordering $<$ on the undirected $G/\mathcal{F}$-edges). Finally, lift $M/\mathcal{F}$ to $M \subseteq G$ as in Definition 3.14, choosing the finite spanning trees and liftings of $G/\mathcal{F}$-edges in a Borel manner using Lusin–Novikov uniformization [Kec95, 18.10]. The desired properties of this M are then given by Theorem 3.16. □

For a quasi-pmp Borel graph, the above properties of the subforest M may be significantly strengthened on a conull set, due to the following fact (whose analogue in percolation on a unimodular graph is [LS99, Proposition 3.9]):

Lemma 3.20. *Let (X, G) be a locally finite Borel graph, $\mathfrak{w} : G \to \mathbb{R}^+$ be a Borel cocycle, and μ be a $\mathfrak{w}$-invariant probability measure on X. For a.e. G-component, the space of $\mathfrak{w}$-nonvanishing ends either has ≤ 2 elements or is perfect*[5] *(has no isolated points).*

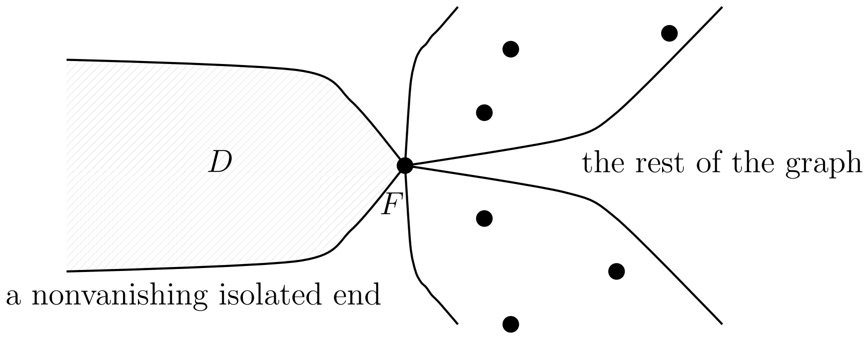

FIGURE 3.21. This is an illustration of the proof of Lemma 3.20, where the large dots represent the maximal disjoint Borel family $\mathcal{F}$ of $\mathfrak{w}$-trifurcations.

Proof. Suppose some G-component C has at least 3 $\mathfrak{w}$-nonvanishing ends, at least one of which is isolated (among $\mathfrak{w}$-nonvanishing ends). Then any such isolated end $u \in \partial_{\mathfrak{w}} C$ belongs to a side $D \subseteq C \setminus F$ of a $\mathfrak{w}$-trifurcation F with no other $\mathfrak{w}$-nonvanishing ends: to see this, apply Lemma 2.25 to u, any neighborhood isolating it, and two other nonvanishing ends. Furthermore, such D then cannot also contain a $\mathfrak{w}$-trifurcation F' (of C), or else at least two nonvanishing sides of F' would be disjoint from F, yielding at least two nonvanishing ends in D.

Now take a maximal disjoint Borel family $\mathcal{F}$ of $\mathfrak{w}$-trifurcations $F \subseteq X$ which have at least one side D with exactly one nonvanishing end (which is hence isolated); see Figure 3.21. Let $Y \subseteq X$ be the union of all such sides D of all $F \in \mathcal{F}$. Then each $y \in Y$ belongs to a unique such D for a unique $F \in \mathcal{F}$, since if it also belonged to a one-ended side D' of another $F' \in \mathcal{F}$, then either $F \subseteq D'$ or $F' \subseteq D$ which is impossible as noted above. Let $E \subseteq \mathbb{E}_G$ be the equivalence relation on Y whose classes are exactly all such sides D of $F \in \mathcal{F}$, hence are nonvanishing. Then E is smooth, since we may choose in a Borel way a nonempty finite subset of each class $D \in Y/E$, namely the inner boundary of D (i.e., the vertices adjacent to F). So by Lemma 2.33, $\mu(Y) = 0$, and hence $\mu([Y]_G) = 0$ by quasi-invariance of μ. But by maximality of $\mathcal{F}$, $[Y]_G$ is precisely the union of the G-components with at least 3 $\mathfrak{w}$-nonvanishing ends, at least one of which is isolated (among $\mathfrak{w}$-nonvanishing ends). ☐

Corollary 3.22. *The subforest $M \subseteq G$ from Theorem 3.19 additionally obeys the following for every $\mathfrak{w}$-invariant probability measure μ:*

- *($*$) For a.e. G-component $C \in X/G$ with ≥ 3 nonvanishing G-ends, the space $\partial^M_{\mathfrak{w}} D$ of nonvanishing ends in each M-component $D \subseteq C$ is perfect*[5].

Moreover, given a $\mathfrak{w}$-invariant probability measure μ on X which is G-ergodic, then we may choose the subforest M to also make μ ergodic.

[5]This does not in general imply that it has continuum-many elements because the space may not be Polish; see Remark 2.20.

Proof. To conclude (∗) notice that by Lemma 3.20, the union Y of all M-components with at least 3, and at least one isolated, nonvanishing ends is μ-null, whence so is $[Y]_G$ by quasi-invariance of μ; by Theorem 3.19(b), $[Y]_G$ contains all G-components C not satisfying (∗).

It remains to show that if μ is G-ergodic then we can choose M to be μ ergodic as well. Recall the identification $X/\mathcal{F} \cong Y \subseteq X$ from the proof of Theorem 3.19. Since $Y \subseteq X$ is a complete $\mathbb{E}_G$-section, by quasi-invariance of μ, $\mu(Y) > 0$. Since μ was G-ergodic and $\mathfrak{w}$-invariant, $\mu|Y$ is $\mathbb{E}_G|Y$-ergodic (again by quasi-invariance of μ) and $\mathfrak{w}|Y$-invariant, hence corresponds to a $G/\mathcal{F}$-ergodic and $\mathfrak{w}/\mathcal{F}$-invariant measure on $X/\mathcal{F}$, which we denote $\mu/\mathcal{F}$. Now in the proof of Theorem 3.19, when constructing the maximal subforest $M/\mathcal{F} \subseteq G/\mathcal{F}$ as in Definition 3.2, take the fixed subforest H to be a $\mu/\mathcal{F}$-ergodic hyperfinite subforest, which exists by [Tse22, Theorem 1.3] (which gives an ergodic hyperfinite subgraph, but every such subgraph contains a Borel treeing by [Mil08a, Lemma 2.4]). This ensures that $M/\mathcal{F}$ is also $\mu/\mathcal{F}$-ergodic, which easily implies that its lift M (as in the proof of Theorem 3.19) is μ-ergodic. □

Corollary 3.22 together with Theorem 1.3 yields:

Corollary 3.23 (Theorem 1.4)**.** *Let G be a quasi-pmp locally finite Borel graph on a standard probability space (X, μ) and let $\mathfrak{w} : \mathbb{E}_G \to \mathbb{R}^+$ be the Radon–Nikodym cocycle of $\mathbb{E}_G$ with respect to μ. If every connected component of G admits ≥ 3 $\mathfrak{w}$-nonvanishing ends, then G is μ-nowhere amenable. In fact, G contains a μ-nowhere amenable Borel subforest.*

4. Maximal forest as a random subgraph

We connect the construction of a maximal subforest of a connected locally finite graph from Section 3 to the study of random spanning forests, in particular the Free Minimal Spanning Forest (FMSF) [LPS06]. We show that our construction of a maximal subforest yields the natural extension of FMSF for nonunimodular graphs. We first give a brief introduction to random spanning forests and their importance in percolation theory, after which we quickly review nonunimodular graphs and present our construction of the Free Maximal Spanning Forest for such graphs, as well as its properties.

Throughout this section, let $G := (V, E)$ denote a locally finite graph.

4.A. **Random spanning forests and percolation theory.** Classically, the **Free Minimal Spanning Forest** $\mathrm{FMSF}(G)$ on the graph $G := (V, E)$ is a random subforest of G constructed as follows:

- Let $\{U_e\}_{e \in E}$ be a collection of independent random variables with Uniform$[0, 1]$ distribution. Notice that almost surely for any two distinct edges e and e' we have $U_e \neq U_{e'}$.
- For each cycle in G, delete an edge e with largest value of the label U_e.

In other words, for each $e \in E$,

$$e \in \mathrm{FMSF}(G) \Leftrightarrow \text{for each cycle } C \text{ in } G \text{ that contains } e \text{ there is } e' \in C \text{ such that } U_e < U_{e'}.$$

The **Wired Minimal Spanning Forest** $\mathrm{WMSF}(G)$ is constructed similarly, but the bi-infinite paths are also considered to be cycles.

A **bond percolation** process on G is a probability measure $\mathbf{P}$ on 2^E. We refer to elements $\omega \in 2^E$ as **configurations** and we say that an edge $e \in E$ is **present** (or **open**) in ω if $e \in \omega$

(we think of ω as a subset of E). The connected components of ω are referred to as **clusters**. Finally, for a subgroup $\Gamma \leq \mathrm{Aut}(G)$ we say that percolation $\mathbf{P}$ is Γ-**invariant** if the measure $\mathbf{P}$ is invariant under the diagonal action of Γ on G.

For $p \in [0,1]$, a bond percolation process on G is called Bernoulli(p) if every edge is present in a configuration independently with probability p. We denote the measure associated with Bernoulli(p) bond percolation by $\mathbb{P}_p$. Henceforth, we will drop the word "bond" as we never use other kinds of percolations.

The study of FMSF and WMSF is closely connected to percolation theory. For example, [LPS06, Proposition 3.6] stated below as Theorem 4.3 relates these random forests to a famous conjecture of Benjamini and Schramm [BS96, Conjecture 6]. This conjecture says that a quasi-transitive[6] graph G is amenable if and only if $p_c(G) = p_u(G)$, where

$$p_c(G) := \inf\{p \in [0,1] : \mathbb{P}_p(\text{there is an infinite cluster}) = 1\}, \tag{4.1}$$

$$p_u(G) := \inf\{p \in [0,1] : \mathbb{P}_p(\text{there is exactly one infinite cluster}) = 1\}. \tag{4.2}$$

The fact that amenability of G implies $p_c = p_u$ was proven in [GKN92]; however the other direction of this conjecture is only known in particular cases and is one of the central questions of the subject. In light of this, the following characterisation of $p_c(G) = p_u(G)$ is particularly interesting.

Theorem 4.3 ([LPS06, Proposition 3.6])**.** *For any graph G $p_c(G) = p_u(G)$ if and only if* $\mathrm{FMSF}(G) = \mathrm{WMSF}(G)$.

4.B. **Unimodular and nonunimodular graphs.** We recall that locally compact group is called **unimodular** if its left Haar measure is also right-invariant. The automorphism group $\mathrm{Aut}(G)$ of a *connected* locally finite graph $G := (V, E)$ is a locally compact group when equipped with the topology of pointwise convergence. Thus, the graph G is called **unimodular** if $\mathrm{Aut}(G)$ is unimodular. We refer to [BLPS99, LP16] for a survey of unimodular automorphism groups and their significance for random subgraphs.

Let Γ be a closed subgroup of $\mathrm{Aut}(G)$ and m be a Haar measure on Γ. For $x, y \in V$ that are in the same Γ-orbit, we define the weight of x relative to y by

$$\mathfrak{w}_\Gamma(x, y) := m(\Gamma_x)/m(\Gamma_y), \tag{4.4}$$

where $\Gamma_v := \{\gamma \in \Gamma \mid \gamma v = v\}$ is the stabilizer of $v \in V$. The map $\mathfrak{w}_\Gamma : (x, y) \mapsto \mathfrak{w}_\Gamma(x, y)$ is an $\mathbb{R}^+$-valued cocycle on the orbit equivalence relation of the action of Γ on V.

It is a well-known fact, proven in [Tro85], that a closed subgroup $\Gamma \leq \mathrm{Aut}(G)$ is unimodular if and only if for all $x, y \in V$ in the same Γ-orbit

$$|\Gamma_x y| = |\Gamma_y x| \tag{4.5}$$

where $\Gamma_x y$ denotes the (finite) set $\{\gamma y \mid \gamma \in \Gamma,\, \gamma x = x\}$. Furthermore, the cocycle $\mathfrak{w}$ is invariant under the action of Γ, i.e., for any $\gamma \in \Gamma$ and $x, y \in V$ in the same Γ-orbit, we have

$$\mathfrak{w}_\Gamma^y(x) = \frac{m(\Gamma_x)}{m(\Gamma_y)} = \frac{|\Gamma_x y|}{|\Gamma_y x|} = \frac{|\Gamma_{\gamma x}\gamma y|}{|\Gamma_{\gamma y}\gamma x|} = \frac{m(\Gamma_{\gamma x})}{m(\Gamma_{\gamma y})} = \mathfrak{w}_\Gamma^{\gamma y}(\gamma x). \tag{4.6}$$

In particular, equations (4.5) and (4.6) imply that if Γ is unimodular then the function $x \mapsto m(\Gamma_x)$ is constant on each Γ-orbit. Moreover, if Γ acts quasi-transitively, then, by definition, there are only finitely many Γ-orbits and hence the set $\{m(\Gamma_x) \mid x \in V\}$ is finite.

[6] A graph G is called **quasi-transitive** if the natural action of $\mathrm{Aut}\,(G)$ group on it has finitely many orbits.

The converse is also true (see [Tan19, Lemma 2.3] for a proof): if $\{m(\Gamma_x) \mid x \in V\}$ is finite, then Γ is unimodular.

4.C. **The Random Maximal Spanning Forest.** Let $G := (V, E)$ be a countable locally finite graph and let $\mathfrak{w}$ be a weight function on V.

Definition 4.7. Let $<$ be a uniformly random linear ordering of E defined by setting, for any $e, e' \in E$,

$$e < e' :\Leftrightarrow U_e > U_{e'},$$

where $(U_e)_{e \in E}$ is a sequence of independent random variables with Uniform$[0,1]$ distribution. Then the $\mathfrak{w}$-maximal subforest of G with the random tiebreaker $<$ is a random subforest of G, which we call the **Free $\mathfrak{w}$-Maximal Spanning Forest** of G and denote it by $\mathrm{FMaxSF}_{\mathfrak{w}}(G)$.

In other words, for every cycle in G select the set of edges that are adjacent to the vertices with the smallest $\mathfrak{w}$-weight in that cycle, and delete the one that has the largest U_e associated with it.

Remark 4.8. It is immediate that $\mathrm{FMaxSF}_{\mathfrak{w}}(G) = \mathrm{FMSF}(G)$ when $\mathfrak{w}$ is constant 1. Thus, $\mathrm{FMaxSF}_{\mathfrak{w}}(G)$ is a natural generalization of $\mathrm{FMSF}(G)$ suitable to the nonunimodular setting, where we take $\mathfrak{w}$ to be the weight function $\mathfrak{w}_\Gamma$ induced by a nonunimodular closed subgroup $\Gamma \leq \mathrm{Aut}(G)$.

Below, we assume that G is *connected*, fix a closed subgroup $\Gamma \leq \mathrm{Aut}(G)$, and let $\mathfrak{w} := \mathfrak{w}_\Gamma$ be defined as in (4.4). We then consider $\mathrm{FMaxSF}_{\mathfrak{w}}(\omega)$ for a subgraph ω of G. Often, ω will itself be a random subgraph of G and thus, $\mathrm{FMaxSF}_{\mathfrak{w}}(\omega)$ will have two sources of randomness: one from ω and the other from the random linear ordering $<$.

Similarly to FMSF, the random subforest $\mathrm{FMaxSF}_{\mathfrak{w}}(\omega)$ of a Γ-invariant percolation configuration ω on G has the following properties due to the fact that the weight function $\mathfrak{w}_\Gamma$ is Γ-invariant (see (4.6)).

Proposition 4.9. *Let Γ be a closed subgroup of $\mathrm{Aut}(G)$ that acts transitively on G and let $\mathbf{P}$ be a Γ-invariant measure on $\{0,1\}^E$. If $\omega \sim \mathbf{P}$ then the random forest $\mathrm{FMaxSF}_{\mathfrak{w}}(\omega)$ is a Γ-equivariant factor of $(\mathfrak{w}, (U_e)_{e\in E})$. In particular:*

(a) The distribution of $\mathrm{FMaxSF}_{\mathfrak{w}}(\omega)$ is invariant under the action of Γ.
(b) If $\mathbf{P}$ is ergodic then $\mathrm{FMaxSF}_{\mathfrak{w}}(\omega)$ is weakly mixing (hence, ergodic) under the action of Γ.

In particular, all this applies to $\mathrm{FMaxSF}_{\mathfrak{w}}(G)$ because taking $\mathbf{P} := 1_E$, we have $\omega = G$ a.s.

Proof. Since $\mathfrak{w}$ is Γ-invariant and any cycle in ω and the tiebreaker $<$ are the same as their γ-images for any $\gamma \in \Gamma$ the map $(\omega, (U_e)_{e\in E}) \mapsto \mathrm{FMaxSF}_{\mathfrak{w}}(\omega)$ is Γ-equivariant. Since $\{(U_e)_{e\in E}\}$ are i.i.d. and percolation is ergodic the product of corresponding measures is weakly mixing, and thus any Γ-equivariant factor also weakly mixes. In particular, this implies that $\mathrm{FMaxSF}_{\mathfrak{w}}(\omega)$ is ergodic. □

Observation 4.10. In the setting of Proposition 4.9, if $\mathbf{P}$ is the Bernoulli(p) percolation, for any $p \in [0,1]$ then $\mathrm{FMaxSF}_{\mathfrak{w}}(\omega)$ is a factor of i.i.d..

The following lemma is helpful in proving a variety of statements in percolation theory, for instance the continuity of the percolation phase transition in unimodular nonamenable graphs [LP16, Theorem 8.21].

Lemma 4.11 ([LP16, Lemma 7.7])**.** *Let* $\mathbf{P}$ *be a* Γ*-invariant percolation process on a graph* G*. If with positive probability (w.p.p.) there is a component of* $\mathbf{P}$*-configuration* ω *with at least three ends, then the joint distribution of the pair* $(\mathrm{FMSF}(\omega), \omega)$ *is* Γ*-invariant and w.p.p. there is a component of* $\mathrm{FMSF}(\omega)$ *that has at least three ends.*

Recall, from Section 2.C, that we say that a component C is ($\mathfrak{w}_\Gamma$-)**heavy** if $\sum_{x \in C} \mathfrak{w}_\Gamma^y(x) = \infty$ for some/any $y \in V$; otherwise, we call it ($\mathfrak{w}_\Gamma$-)**light**. Thus in addition to the critical parameter p_c and the uniqueness threshold p_u defined in (4.1) and (4.2) we consider the heaviness transition parameter p_h, defined as follows:

$$p_h := p_h(G, \Gamma) := \inf\{p \in [0,1] : \mathbb{P}_p(\text{there is a } (\mathfrak{w}_\Gamma\text{-})\text{heavy cluster}) = 1\}. \tag{4.12}$$

It is easy to see that a unique infinite cluster has to be heavy, yielding that

$$p_c(G) \le p_h(G) \le p_u(G) \le 1.$$

Moreover, each of these inequalities could be strict, for details and examples see [HPS99, Tim06, Hut20].

Our Theorem 4.15 is an analog of Lemma 4.11 for the maximal forest and $\mathfrak{w}$-nonvanishing ends under additional assumptions of the deletion tolerance.

Definition 4.13 (Insertion and deletion tolerance)**.** Given a set of configurations $A \subseteq 2^E$ and an edge $e \in E$, let $\Pi_e A = \{\omega \cup \{e\} \mid \omega \in A\}$ and $\Pi_{\neg e} A = \{\omega \setminus \{e\} \mid \omega \in A\}$. A bond percolation process $\mathbf{P}$ is called an **insertion (deletion) tolerant** if $\mathbf{P}(\Pi_e A) > 0$ ($\mathbf{P}(\Pi_{\neg e} A) > 0$) for every $e \in E$ and every measurable $A \subseteq 2^E$ such that $\mathbf{P}(A) > 0$.

Example 4.14. Let $\mathbb{P}_p$ be Bernoulli percolation on G. Then for any edge $e \in E$ and any $A \subseteq 2^E$ we have that

$$\mathbb{P}_p(\Pi_e A) \ge p\mathbb{P}_p(A) \qquad \text{and} \qquad \mathbb{P}_p(\Pi_{\neg e} A) \ge (1-p)\mathbb{P}_p(A).$$

In particular, this implies that Bernoulli bond percolation is both insertion and deletion tolerant.

Theorem 4.15 (Maximal forest in percolation)**.** *Let* G *be a countable locally finite graph,* Γ *be a closed subgroup of* $\mathrm{Aut}(G)$ *that acts transitively on* G*, and* $\mathfrak{w} := \mathfrak{w}_\Gamma$ *be as in* (4.4)*. Let* $\mathbf{P}$ *be a* Γ*-invariant deletion tolerant percolation process on* G*. Then the following holds:*

- *(a) The (*$\mathfrak{w}$*-nonvanishing) ends of* $\mathrm{FMaxSF}_\mathfrak{w}(\omega)$ *are canonically mapped (Definition 2.28) onto a dense subset of the space of (*$\mathfrak{w}$*-nonvanishing) ends of* ω*.*
- *(b) If w.p.p. there is a component of* ω *with at least three* $\mathfrak{w}$*-nonvanishing ends, then w.p.p. there is a tree in* $\mathrm{FMaxSF}_\mathfrak{w}(\omega)$ *that has perfectly many* $\mathfrak{w}$*-nonvanishing ends.*

Proof. Proof of (a) follows directly from Theorem 3.19(a).

Proof of (b). Since ω contains a cluster with at least three $\mathfrak{w}$-nonvanishing ends, it must contain a $\mathfrak{w}$-k-furcation F for some $k \ge 3$. Denote the components that F disconnects that have $\mathfrak{w}$-nonvanishing ends by $\{C_i\}_{i \le k}$. By Fubini's theorem the set of configurations such that F forms a furcation has a positive measure. Hence, by deletion tolerance w.p.p. such furcation is a tree and is connected to each C_i by a single edge. On this event edges of F are not contained in any cycles of ω and hence $F \subset \mathrm{FMaxSF}_\mathfrak{w}(\omega)$. It remains to show that the tree in $\mathrm{FMaxSF}_\mathfrak{w}(\omega)$ that contains F has a $\mathfrak{w}$-nonvanishing end in each of C_i. This follows from the argument of Lemma 3.4. Hence w.p.p. there is a component of $\mathrm{FMaxSF}_\mathfrak{w}(\omega)$ that has at least three $\mathfrak{w}$-nonvanishing ends.

Each tree in $\mathrm{FMaxSF}_{\mathfrak{w}}(\omega)$ that has at least three $\mathfrak{w}$-nonvanishing ends a.s. has perfectly many of such ends. This follows from an argument analogous to the one in the proof of Lemma 3.20 that uses smoothness. It can also be proved by following the argument of [LS99, Proposition 3.9] and applying Tilted Mass Transport Principle (see [LP16, Theorem 8.7] or [Hut20, Section 2.1] for details). □

Remark 4.16 (Assumption of deletion tolerance)**.** While Theorem 4.15 is a natural extension of Lemma 4.11 to the setting where the relative weight function is not constant, one can see that we have an additional assumption of the deletion tolerance. The proofs of both of these statements rely on the fact that one can force (w.p.p.) a particular finite tree to be present in the forest. When the forest is defined purely in terms of the linear order induced by $\{U_e\}_{e\in E}$, one can do so by restricting to the event where particular finite set of labels are less or greater than $1/2$. On the other hand, when the weight function is nonconstant this technique would not apply and hence we require the the deletion tolerance property to "cut" all the cycles around the tree of interest and hence force it to be present in the forest.

Corollary 4.17. *Under assumptions of Theorem 4.15(b), if* $\mathbf{P}$ *is ergodic under the action of* Γ *then almost surely at least one tree in* $\mathrm{FMaxSF}_{\mathfrak{w}}(\omega)$ *has perfectly many* $\mathfrak{w}$*-nonvanishing ends.*

If $\mathbf{P}$ is also insertion tolerant then verifying the assumptions of Theorem 4.15 becomes significantly easier as we show in the following lemma.

Lemma 4.18. *Let G be a countable locally finite graph,* Γ *is a closed subgroup of* $\mathrm{Aut}(G)$ *that acts transitively on G, and m be the Haar measure on it. Let* $\mathbf{P}$ *be a* Γ*-invariant insertion and deletion tolerant percolation process on G.*

- *(a) If w.p.p. there are at least three heavy clusters then w.p.p. there is a cluster with at least three* $\mathfrak{w}$*-nonvanishing ends (i.e., assumptions of Theorem 4.15 are satisfied).*
- *(b) Then each heavy tree of* $\mathrm{FMaxSF}_{\mathfrak{w}}(\omega)$ *contains at least one nonvanishing end.*

Proof of Lemma 4.18. Each heavy tree of $\mathrm{FMaxSF}_{\mathfrak{w}}(\omega)$ is contained in some heavy cluster of ω. Deletion and insertion tolerance of $\mathbf{P}$, following argument of [Tim06, Lemma 5.2], yields that any heavy cluster has a sequence of vertices with $\mathfrak{w}$-weights going to infinity. Indeed, if w.p.p. there was a heavy cluster such that all weights of vertices in it is bounded from above, then by deletion and insertion tolerance w.p.p. there would be a cluster with the unique vertex that has maximum $\mathfrak{w}$-weight. This would contradict Tilted Mass Transport Principle. Since the set of ends of a cluster is compact it has to contain at least one $\mathfrak{w}$-nonvanishing end. Hence each heavy cluster has to have at least one $\mathfrak{w}$-nonvanishing end. Lemma 3.4 implies that the same is true for each tree of $\mathrm{FMaxSF}_{\mathfrak{w}}(\omega)$.

If with w.p.p. ω contains at least three heavy components then by insertion tolerance they can be connected yielding a cluster with at least three $\mathfrak{w}$-nonvanishing ends. This happens w.p.p., hence the assumptions of Theorem 4.15 are satisfied and by Theorem 4.15(b) w.p.p. $\mathrm{FMaxSF}_{\mathfrak{w}}(\omega)$ contains a tree with perfectly many $\mathfrak{w}$-nonvanishing ends. □

Corollary 4.19 ($\mathfrak{w}$-nonvanishing ends in Bernoulli(p) percolation)**.** *For any locally finite graph G, suppose* $\Gamma \leq \mathrm{Aut}(G)$ *is a nonunimodular closed subgroup that acts transitively on G and is such that such that* $p_h(G,\Gamma) < p_u(G)$*. Moreover, suppose the orbit of each vertex under action of* Γ *is infinite, then* Bernoulli(p) *bond percolation is insertion and deletion tolerant and is ergodic* [LP16, Proposition 7.3] *and by* [Tan19] *the heavy clusters of* Bernoulli(p)

percolation are indistinguishable. Therefore, it satisfies the assumptions of Corollary 4.17 and for almost all configurations ω that contains infinitely many heavy components, any heavy tree in $\mathrm{FMaxSF}_{\mathfrak{w}}(\omega)$ *has perfectly many* $\mathfrak{w}$*-nonvanishing ends.*

Indistinguishably of FMSF on unimodular graphs was shown in [Tim18, Theorem 1.2]. One consequence of this result is a simplification of the construction of the treeable ergodic subrelation in Gaboriau–Lyons theorem [GL09, Proposition 13]. For analogous reasons it would be of interest to investigate if FMaxSF is also indistinguishable.

Question 4.20. Are $\mathfrak{w}_\Gamma$-heavy trees of $\mathrm{FMaxSF}_{\mathfrak{w}_\Gamma}(G)$ indistinguishable? In the sense that for any Γ-invariant Borel property $\mathcal{A}$ either almost surely all $\mathfrak{w}_\Gamma$-heavy trees of $\mathrm{FMaxSF}_{\mathfrak{w}_\Gamma}(G)$ satisfy $\mathcal{A}$ or almost surely all such trees do not satisfy $\mathcal{A}$.

5. Applications

In this section we present several concrete applications of our results.

5.A. **Coinduced actions.**

Let $\Gamma \leq \Delta$ be countably infinite groups and let $\Gamma \curvearrowright X$ be a Borel action on a standard Borel space X. Let X_Γ^Δ be the set of all Γ-equivariant maps from Δ to X, where Γ acts on Δ by left translation. Let Δ act on X_Γ^Δ (on the left) by right shift, namely, for any $\pi \in X_\Gamma^\Delta$ and $\delta, \delta' \in \Delta$, we set

$$(\delta \cdot \pi)(\delta') := \pi(\delta' \cdot \delta).$$

It is straightforward to verify that $\delta \cdot \pi$ is a Γ-equivariant map, so the action is well-defined.

An isomorphic description of this action may be given as follows. Note that for any $\delta \in \Delta$, the values of a Γ-equivariant $\pi : \Delta \to X$ on the right coset $\Gamma\delta$ are uniquely determined by $\pi(\delta)$. Thus for any family of coset representatives $(\delta_C \in C)_{C \in \Gamma\backslash\Delta}$, we have a bijection

$$\begin{aligned} X_\Gamma^\Delta &\cong X^{\Gamma\backslash\Delta} \\ \pi &\mapsto (\widetilde{\pi} : C \mapsto \pi(\delta_C)). \end{aligned} \tag{5.1}$$

Transferring the action of Δ from X_Γ^Δ to $X^{\Gamma\backslash\Delta}$, we get

$$(\delta \cdot \widetilde{\pi})(C) := (\delta_C \cdot \delta) \cdot \delta_{C\cdot\delta}^{-1} \cdot \widetilde{\pi}(C \cdot \delta), \tag{5.2}$$

for each $\delta \in \Delta$, $\widetilde{\pi} \in X^{\Gamma\backslash\Delta}$ and $C = \Gamma\delta_C \in \Gamma\backslash\Delta$.

Using such a bijection (5.1), we may put a measure on X_Γ^Δ by transferring the product measure $\mu^{\Gamma\backslash\Delta}$ on $X^{\Gamma\backslash\Delta}$ for any base measure μ on X.

Remark 5.3. If μ is a Γ-invariant measure on X, then the measure on X_Γ^Δ induced in this way does not depend on the choice of coset representatives $(\delta_C)_C$, and is Δ-invariant; see [KQ19]. However, the same does not hold for quasi-invariant measures, which are our main interest.

In spite of this remark, we will denote the measure on X_Γ^Δ induced as above by a measure μ on X by μ_Γ^Δ, leaving the choice of coset representatives to be implied by the context.

Lemma 5.4. *If Γ acts freely on X, and μ is an atomless probability measure on X, then there is a μ_Γ^Δ-conull Δ-invariant subset $Y \subseteq X_\Gamma^\Delta$ on which Δ acts freely.*

Proof. Suppose $\pi \in X_\Gamma^\Delta$ and $\delta \in \Delta$ such that $\delta \cdot \pi = \pi$. We have a Γ-equivariant map $X_\Gamma^\Delta \to X$, namely the projection $\pi \mapsto \pi(1)$; thus if $\delta \in \Gamma$, then since $\Gamma \curvearrowright X$ is free, $\delta = 1$. If $\delta \notin \Gamma$, then we have $\pi(\gamma\delta) = (\delta \cdot \pi)(\gamma) = \pi(\gamma)$ for all $\gamma \in \Gamma$, whence in particular, there are two distinct cosets $C \neq D \in \Gamma\backslash\Delta$ (namely $C := \Gamma\delta$ and $D := \Gamma$) such that $\pi(C) = \pi(D) \subseteq X$. The set $Y \subseteq X_\Gamma^\Delta$ of all π for which there exist such $C \neq D$ with $\pi(C) = \pi(D)$ is clearly Δ-invariant, and it is contained in the set of π such that there exist $C \neq D$ with $\pi(\delta_C) \in \pi(D) = \Gamma \cdot \pi(\delta_D)$, which is null since its image under the bijection (5.1) is a countable union of diagonals, which is null since μ (hence also all Γ-translates of μ) are atomless. Thus Y works. □

We are particularly interested in the case $\Delta = \Gamma * \Lambda$ for another countable group Λ. In this case, there is a canonical choice of coset representatives $(\delta_C)_{C\in\Gamma\backslash(\Gamma*\Lambda)}$, namely those elements of $\Gamma * \Lambda$ whose normal form does not start with a nonidentity element of Γ. Note that the right translation actions of Γ and Λ on $\Gamma\backslash(\Gamma * \Lambda)$ affect these coset representatives as follows: for $\lambda \in \Lambda$, $\gamma \in \Gamma$, and $C \in \Gamma\backslash(\Gamma * \Lambda)$,

$$\begin{aligned}
\delta_{C\cdot\lambda} &= \delta_C \cdot \lambda, \\
\delta_{C\cdot\gamma} &= \begin{cases} \delta_C \cdot \gamma & \text{if } \delta_C \neq 1, \text{ i.e., } C \neq \Gamma, \\ 1 & \text{if } \delta_C = 1, \text{ i.e., } C = \Gamma. \end{cases}
\end{aligned}$$

Thus, the formula (5.2) for the action of Δ on $X^{\Gamma\backslash\Delta}$ becomes

$$\begin{aligned}
(\lambda \cdot \widetilde{\pi})(C) &:= \widetilde{\pi}(C \cdot \lambda), \\
(\gamma \cdot \widetilde{\pi})(C) &:= \begin{cases} \widetilde{\pi}(C \cdot \gamma) & \text{if } \delta_C \neq 1, \text{ i.e., } C \neq \Gamma, \\ \gamma \cdot \widetilde{\pi}(\Gamma) & \text{if } \delta_C = 1, \text{ i.e., } C = \Gamma. \end{cases}
\end{aligned}$$

Using this, we have

Lemma 5.5. *If μ is a Γ-quasi-invariant measure on X, with Radon–Nikodym cocycle $\mathfrak{w} : \mathbb{E}_\Gamma \to \mathbb{R}^+$, then $\mu^{\Gamma\backslash(\Gamma*\Lambda)}$ is a $(\Gamma * \Lambda)$-quasi-invariant measure on $X^{\Gamma\backslash(\Gamma*\Lambda)}$ (with the above action), with Radon–Nikodym cocycle $\widetilde{\mathfrak{w}}$ defined on generators $\lambda \in \Lambda$ and $\gamma \in \Gamma$ by*

$$\begin{aligned}
\widetilde{\mathfrak{w}}(\widetilde{\pi}, \lambda \cdot \widetilde{\pi}) &:= 1, \\
\widetilde{\mathfrak{w}}(\widetilde{\pi}, \gamma \cdot \widetilde{\pi}) &:= \mathfrak{w}(\widetilde{\pi}(\Gamma), \gamma \cdot \widetilde{\pi}(\Gamma)).
\end{aligned}$$

In particular, the action of Λ on $X^{\Gamma\backslash(\Gamma*\Lambda)}$ is $\mu^{\Gamma\backslash(\Gamma*\Lambda)}$-preserving.

Proof. λ acts via right shift $X^{\Gamma\backslash(\Gamma*\Lambda)} \to X^{\Gamma\backslash(\Gamma*\Lambda)}$, which preserves the product measure; while γ acts via the composite of right shift followed by acting on the Γth coordinate via $\gamma : X \to X$, the latter of which clearly has Radon–Nikodym cocycle $\mathfrak{w}(\widetilde{\pi}(\Gamma), \gamma \cdot \widetilde{\pi}(\Gamma))$. □

Example 5.6. Let Γ, Λ be infinite finitely generated groups, such that Λ is Kazhdan, and let $\Gamma \curvearrowright (X, \mu)$ be a free quasi-pmp action. For example, we may take $\Lambda := \mathrm{SL}_3(\mathbb{Z})$, $\Gamma := \mathbb{F}_2$, and X to be the boundary action; see Example 1.2. By the above, we get an a.e. free quasi-pmp coinduced action $\Gamma * \Lambda \curvearrowright (X^{\Gamma\backslash(\Gamma*\Lambda)}, \mu^{\Gamma\backslash(\Gamma*\Lambda)})$. Since Λ is nonamenable and its action on $X^{\Gamma\backslash(\Gamma*\Lambda)}$ is pmp by the above lemma, $\mathbb{E}_\Lambda$ and hence also $\mathbb{E}_{\Gamma*\Lambda}$ is nowhere amenable. However, it is also nowhere treeable, by [AS90]. Using our construction, we may produce a subforest of $\mathbb{E}_{\Gamma*\Lambda}$ witnessing its nonamenability: namely, since the action of Λ is pmp, it is easily seen that each Λ-orbit yields a distinct $\widetilde{\mathfrak{w}}$-nonvanishing end in the Schreier graph with respect to a union of finite generating sets for Γ, Λ, which thus contains a subforest with a.e. perfectly many $\widetilde{\mathfrak{w}}$-nonvanishing ends by Corollary 3.22.

5.B. **Cluster graphings for nonunimodular graphs.** We recall the general cluster graphing construction from [Gab05, Section 2.2 and 2.3]. Let $G \subseteq V^2$ be a connected locally finite rooted graph on vertex set V and let $\rho \in V$ be the root. Let $\Gamma \leq \mathrm{Aut}(G)$ be a closed subgroup of the automorphism group of G. Given a free pmp action $\Gamma \curvearrowright (X, \mu)$ and a factor map $\pi : X \to 2^G$, the cluster graphing construction produces a quasi-pmp graph G^{cl} on a standard probability space (Y, ν), which is a certain quotient of $X \times V$ identified with a Borel subset of X. Each component of G^{cl} is a copy of the corresponding percolation cluster; more precisely, the G^{cl}-component of a point $x \in Y$ is isomorphic to the cluster of the root ρ in the percolation configuration $\pi(x)$. Furthermore, while [Gab05, Theorem 2.5] states that the cluster connectedness relation[7] $E^{\mathrm{cl}} := \mathbb{E}_{G^{\mathrm{cl}}}$ is pmp if and only if the graph G is unimodular, its proof actually shows that the Radon–Nikodym cocycle of E^{cl} with respect to ν corresponds (via the aforementioned isomorphism[8]) to the cocycle on G induced by Γ as in (4.4). Hence the cocycle on G is takes infinitely many values E^{cl} is quasi-pmp.

Corollary 5.7. *Let G be a countable locally finite graph, Γ is a closed subgroup of $\mathrm{Aut}(G)$ that acts transitively on G. Let $\mathfrak{w}_\Gamma$ be the relative weight function induced by Γ as in (4.4). For any $p \in [0, 1]$ such that there are infinitely many $\mathfrak{w}_\Gamma$-heavy clusters under Bernoulli(p) percolation $\mathbb{P}_p(G)$. Then a cluster graphing G^{cl} with the measure induced by $\mathbb{P}_p(G)$ is nonamenable.*

Proof. Corollary 4.19 yields that G^{cl} is an ergodic graphing that has at least three $\mathfrak{w}_\Gamma$-nonvanishing ends. Notice that it is ergodic as a consequence of the indistinguishability of heavy clusters shown in [Tan19]. By Corollary 3.22 it contains an ergodic treeable subrelation with at least three $\mathfrak{w}_\Gamma$-nonvanishing ends and thus, by Theorem 1.3, G^{cl} is nonamenable. □

Remark 5.8. If indistinguishability of the Free Maximal Forest $\mathrm{FMaxSF}_{\mathfrak{w}_\Gamma}(G)$ holds then its cluster connectedness relation would serve as an ergodic treeable subrelation with at least three $\mathfrak{w}$-nonvanishing ends. (See Question 4.20.)

We now give a concrete example of such a graph.

Example 5.9 (Free product of $\mathrm{GP}(2)$ and $\mathbb{Z}^2$)**.** Let $\mathrm{GP}(k)$ be the grandparent graph, originally introduced in [Tro85]. Such a graph is constructed as follows: start with a $(k+1)$-regular tree with a distinguished end, this defines a unique parent and k children, for each vertex. Connect every vertex to its grandparent. It is easy to check that $\mathrm{GP}(k)$ is nonunimodular. Moreover, equation (4.4) implies that if x is a parent of y then $\mathfrak{w}^x(y) = \frac{1}{k}$. Therefore, any vertices x and y have the same weight if and only if they are in the same generation.

Let the graph G be the free product[9] of the grandparent graph $\mathrm{GP}(2)$ and $\mathbb{Z}^2$. G is still nonunimodular. Notice that the induced cocycle is equal to 1 on each edge that comes

[7] In [Gab05], this relation is called the *reduced equivalence relation* and denoted by $\mathcal{R}$, while the term "cluster equivalence relation" is used for a lift of $\mathcal{R}$ to $X \times V$.

[8] The isomorphisms from the perspective of different points $x, y \in Y$ in the same G^{cl}-component cohere, see [Gab05, Section 2.3].

[9] The free product of graphs is similar to the Cayley graph of the free product of groups. There are several definitions of such products for graphs present in literature; we use the one in, for example, [PT02, Section 4]. This definition and various others are compared in [CTW20], where it is shown that they are all equivalent for vertex-transitive graphs.

from $\mathbb{Z}^2$, while it is $\frac{1}{2}$ or 2 on the edges that come from GP(2). This together with the fact that $p_c(\mathbb{Z}^2) = \frac{1}{2}$ [Kes80] implies that $p_h(G) \leq \frac{1}{2}$. It is also straight forward to check that $p_u(G) = 1$. For $p \in \left(\frac{1}{2}, 1\right)$ each subgraph of G isomorphic to $\mathbb{Z}^2$ almost surely contains an infinite cluster. Since cocycle is constant on the edges in such a subgraph, it is also contained in a heavy cluster and by compactness such a cluster must contain a $\mathfrak{w}$-nonvanishing end. Insertion tolerance and ergodicity of Bernoulli bond percolation imply that almost surely there is a cluster that contains at least three such infinite clusters from disjoint subgraphs of G isomorphic to $\mathbb{Z}^2$. Therefore, for $p \in \left(\frac{1}{2}, 1\right)$ under Bernoulli(p) percolation on G a.s. there is an infinite cluster with at least three $\mathfrak{w}$-nonvanishing ends.

Witnessing nonamenability for the cluster graphing can be used to show purely graph-theoretic or percolation-theoretic properties. For example, $\mathfrak{w}$-visibility of a graph equipped with a relative weight $\mathfrak{w}$.

Definition 5.10. Let $G = (V, E)$ be a graph. Given a relative weight function (cocycle) $\mathfrak{w} : E \to \mathbb{R}^+$ on edges, we say that $u \in V$ is **$\mathfrak{w}$-visible** from $v \in V$ is there is a path $(v = v_0, v_1, v_2, \ldots, v_k = u)$ such that $(v_i, v_{i+1}) \in E$ and $\mathfrak{w}(v_i, v) \leq 1$ for all $i \in \{0, 1, \ldots, k-1\}$. Let $N_{\mathfrak{w}}(v)$ be the set of all points $\mathfrak{w}$-visible from u. We say that G has **$\mathfrak{w}$-finite visibility** if for all $v \in V$ the set $N_{\mathfrak{w}}(v)$ is $\mathfrak{w}$-light.

Finite visibility is introduced in [Tse22] as a sufficient condition of amenability for quasi-pmp graphs. This notion is particularly useful because nonamenable Borel graphs have infinite visibility, which in tern gives room for combinatorial techniques such as (tilted) mass transport principle.

Theorem 5.11. *Let G be a countable locally finite graph, Γ is a closed subgroup of* $\mathrm{Aut}(G)$ *that acts transitively on G. Let $\mathfrak{w}_\Gamma$ be the relative weight function induced by Γ as in* (4.4)*. For any $p \in [0, 1]$ such that there are infinitely many $\mathfrak{w}_\Gamma$-heavy clusters under Bernoulli(p) percolation $\mathbb{P}_p$. Any configuration $\omega \in 2^G$ under $\mathbb{P}_p$ has $\mathfrak{w}_\Gamma$-infinite visibility almost surely.*

Proof. By Corollary 5.7 G^{cl} is nonamenable. Hence by [Tse22, Theorem 1.8] it has to have $\mathfrak{w}_\Gamma$-infinity visibility. This yields that ω has $\mathfrak{w}_\Gamma$-infinity visibility $\mathbb{P}_p$-almost surely. ☐

References

[Ada90] Scot Adams. Trees and amenable equivalence relations. *Ergodic Theory and Dynamical Systems*, 10(1):1–14, 1990. doi:10.1017/S0143385700005368.

[AS90] S. R. Adams and R. J. Spatzier. Kazhdan groups, cocycles and trees. *Amer. J. Math.*, 112(2):271–287, 1990. doi:10.2307/2374716.

[BLPS99] I. Benjamini, R. Lyons, Y. Peres, and O. Schramm. Group-invariant percolation on graphs. *Geom. Funct. Anal.*, 9(1):29–66, 1999. doi:10.1007/s000390050080.

[BS96] I. Benjamini and O. Schramm. Percolation beyond $\mathbf{Z}^d$, many questions and a few answers. volume 1, pages no. 8, 71–82. 1996. doi:10.1214/ECP.v1-978.

[CFW81] A. Connes, J. Feldman, and B. Weiss. An amenable equivalence relation is generated by a single transformation. *Ergodic Theory Dynam. Systems*, 1(4):431–450 (1982), 1981. doi:10.1017/s014338570000136x.

[CTW20] Max Carter, Stephan Tornier, and George Willis. On free products of graphs. *Australas. J. Combin.*, 78:154–176, 2020.

[DJK94] R. Dougherty, S. Jackson, and A. S. Kechris. The structure of hyperfinite Borel equivalence relations. *Trans. Amer. Math. Soc.*, 341(1):193–225, 1994. doi:10.2307/2154620.

[Gab98] Damien Gaboriau. Mercuriale de groupes et de relations. *C. R. Acad. Sci. Paris Sér. I Math.*, 326(2):219–222, 1998. doi:10.1016/S0764-4442(97)89474-8.

[Gab00] Damien Gaboriau. Coût des relations d'équivalence et des groupes. *Invent. Math.*, 139(1):41–98, 2000. doi:10.1007/s002229900019.

[Gab05] D. Gaboriau. Invariant percolation and harmonic Dirichlet functions. *Geom. Funct. Anal.*, 15(5):1004–1051, 2005. doi:10.1007/s00039-005-0539-2.

[Ghy95] Étienne Ghys. Topologie des feuilles génériques. *Ann. of Math. (2)*, 141(2):387–422, 1995. doi:10.2307/2118526.

[GKN92] A. Gandolfi, M. S. Keane, and C. M. Newman. Uniqueness of the infinite component in a random graph with applications to percolation and spin glasses. *Probab. Theory Related Fields*, 92(4):511–527, 1992. doi:10.1007/BF01274266.

[GL09] D. Gaboriau and R. Lyons. A measurable-group-theoretic solution to von Neumann's problem. *Invent. Math.*, 177(3):533–540, 2009. doi:10.1007/s00222-009-0187-5.

[Hjo06] Greg Hjorth. A lemma for cost attained. *Ann. Pure Appl. Logic*, 143(1-3):87–102, 2006. doi:10.1016/j.apal.2005.05.034.

[HPS99] O. Häggström, Y. Peres, and R. H. Schonmann. Percolation on transitive graphs as a coalescent process: relentless merging followed by simultaneous uniqueness. In *Perplexing problems in probability*, volume 44 of *Progr. Probab.*, pages 69–90. Birkhäuser Boston, Boston, MA, 1999.

[Hut20] T. Hutchcroft. Nonuniqueness and mean-field criticality for percolation on nonunimodular transitive graphs. *J. Amer. Math. Soc.*, 33(4):1101–1165, 2020. doi:10.1090/jams/953.

[JKL02] S. Jackson, A. S. Kechris, and A. Louveau. Countable Borel equivalence relations. *J. Math. Log.*, 2(1):1–80, 2002. doi:10.1142/S0219061302000138.

[Kec95] A. S. Kechris. *Classical descriptive set theory*, volume 156 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995. doi:10.1007/978-1-4612-4190-4.

[Kec22] A. S. Kechris. *The theory of countable Borel equivalence relations.* 2022.

[Kes80] Harry Kesten. The critical probability of bond percolation on the square lattice equals $\frac{1}{2}$. *Comm. Math. Phys.*, 74(1):41–59, 1980. URL: http://projecteuclid.org/euclid.cmp/1103907931.

[KM04] A. S. Kechris and B. D. Miller. *Topics in orbit equivalence*, volume 1852 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 2004. doi:10.1007/b99421.

[KQ19] Alexander S. Kechris and Vibeke Quorning. Co-induction and invariant random subgroups. *Groups Geom. Dyn.*, 13(4):1151–1193, 2019. doi:10.4171/ggd/517.

[LP16] R. Lyons and Y. Peres. *Probability on trees and networks*, volume 42 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, New York, 2016. doi:10.1017/9781316672815.

[LPS06] R. Lyons, Y. Peres, and O. Schramm. Minimal spanning forests. *Ann. Probab.*, 34(5):1665–1692, 2006. doi:10.1214/009117906000000269.

[LS99] R. Lyons and O. Schramm. Indistinguishability of percolation clusters. *Ann. Probab.*, 27(4):1809–1836, 1999. doi:10.1214/aop/1022677549.

[Mar17] A. Marks. A short proof of the connes-feldman-weiss theorem. 2017. preprint. URL: https://www.math.ucla.edu/~marks/notes/cfw.pdf.

[Mil04] B. D. Miller. *Full groups, classification, and equivalence relations.* ProQuest LLC, Ann Arbor, MI, 2004. Thesis (Ph.D.)–University of California, Berkeley. URL: http://gateway.proquest.com/openurl?url_ver=Z39.88-2004&rft_val_fmt=info:ofi/fmt:kev:mtx:dissertation&res_dat=xri:pqdiss&rft_dat=xri:pqdiss:3155578.

[Mil08a] B. D. Miller. Ends of graphed equivalence relations, i. *Israel J. Math.*, 169(1):375, Nov 2008. doi:10.1007/s11856-009-0015-z.

[Mil08b] B. D. Miller. The existence of measures of a given cocycle. II. Probability measures. *Ergodic Theory Dynam. Systems*, 28(5):1615–1633, 2008. doi:10.1017/S0143385707001125.

[PT02] Tomaž Pisanski and Thomas W. Tucker. Growth in products of graphs. *Australas. J. Combin.*, 26:155–169, 2002.

[Sta68] John R. Stallings. On torsion-free groups with infinitely many ends. *Ann. of Math. (2)*, 88:312–334, 1968. doi:10.2307/1970577.

[Tan19] P. Tang. Heavy Bernoulli-percolation clusters are indistinguishable. *Ann. Probab.*, 47(6):4077–4115, 2019. doi:10.1214/19-aop1354.

[Tim06] Á. Timár. Percolation on nonunimodular transitive graphs. *Ann. Probab.*, 34(6):2344–2364, 2006. doi:10.1214/009117906000000494.

[Tim18] Á. Timár. Indistinguishability of the components of random spanning forests. *Ann. Probab.*, 46(4):2221–2242, 2018. doi:10.1214/17-AOP1225.

[Tro85] V. I. Trofimov. Groups of automorphisms of graphs as topological groups. *Mat. Zametki*, 38(3):378–385, 476, 1985.

[Tse22] A. Tserunyan. Pointwise ergodic theorem for locally countable quasi-pmp graphs. *J. Mod. Dyn.*, 18:575–621, 2022. doi:10.3934/jmd.2022018.

[TTD22] Anush Tserunyan and Robin Tucker-Drob. Geography of hyperfinite equivalence relations in quasi-pmp acyclic graphs. 2022+.

[TZ22] Anush Tserunyan and Jenna Zomback. A backward ergodic theorem along trees and its consequences for free group actions. 2022+. URL: https://arxiv.org/abs/2012.10522.

[Zim77] Robert J. Zimmer. Hyperfinite factors and amenable ergodic actions. *Invent. Math.*, 41(1):23–31, 1977. doi:10.1007/BF01390162.

Ruiyuan Chen, Mathematics Department, University of Michigan, Ann Arbor, MI, USA
Email address: ruiyuan@umich.edu
URL: https://rynchn.github.io

Grigory Terlov, Mathematics Department, University of Illinois Urbana-Champaign, Urbana, IL, USA
Email address: gterlov2@illinois.edu
URL: https://sites.google.com/a/illinois.edu/gterlov

Anush Tserunyan, Mathematics and Statistics Department, McGill University, Montréal, QC, Canada
Email address: anush.tserunyan@mcgill.ca
URL: https://www.math.mcgill.ca/atserunyan