

A full characterization of invariant embeddability of unimodular planar graphs

Ádám Timár and László Márton Tóth

Abstract

When can a unimodular random planar graph be drawn in the Euclidean or the hyperbolic plane in a way that the distribution of the random drawing is isometry-invariant? This question was answered for one-ended unimodular graphs in [BT21], using the fact that such graphs automatically have locally finite (simply connected) drawings into the plane. For the case of graphs with multiple ends the question was left open. We revisit Halin’s graph theoretic characterization of graphs that have a locally finite embedding into the plane. Then we prove that such unimodular random graphs do have a locally finite invariant embedding into the Euclidean or the hyperbolic plane, depending on whether the graph is amenable or not.

1 Introduction

Consider a random planar map embedded in the Euclidean or hyperbolic plane M with an isometry-invariant distribution. Simple examples include a lattice shifted by a suitable random isometry or a Voronoi tessellation coming from some invariant point process in M [Mol12, BS01]. If the expected number of vertices in a unit area is finite, then one can condition on having a vertex in a fixed point 0 of M , and define the Palm version, which is hence a random *rooted* graph embedded in the plane. This rooted graph is always unimodular (in the sense defined by Aldous and Lyons in [AL07]), even if the embedding is also taken into account as a decoration (marking) of the rooted graph. (For Euclidean spaces a proof using the terminology of this present paper can be found in [BT21], but this is a basic result in the theory of point processes, where unimodularity is replaced by point-stationarity [HL05, CSKM13].) It is natural to ask whether the converse is true: *does every unimodular random planar graph (G, o) have an isometry-invariant embedding into M , with a suitable choice of M as the Euclidean or hyperbolic plane?* By a unimodular random planar graph (URPG) we mean a unimodular random graph that is planar (has a proper embedding in the plane) almost surely.

To break the above question into two parts, when does (G, o) have a unimodular embedding into M (where a *unimodular embedding*, to be formally defined later, is one where the embedding as a decoration of the graph makes it a unimodular decorated graph), with the root o is embedded at the origin $0 \in M$? And can such an

Keywords and phrases: unimodular random maps, invariant planar embedding, locally finite embedding, excluded minors

embedding be used to define an isometry-invariant embedding? The present paper fully answers both questions. In [BT21], the question of unimodular embeddings of *one-ended* URPG's was answered, always in the affirmative, with amenable G 's embeddable into the Euclidean plane and nonamenable G 's into the hyperbolic plane. Our main contribution here is the solution of the question of unimodular embeddings of URPG's with two or infinitely many ends, whenever it is possible. We also construct isometry invariant embeddings from unimodular ones, which settles the question for invariant embeddings.

The problem in the present setup is somewhat more delicate than in the one-ended case, because the existence of an embedding depends on the graph structure. We mention that 2-ended unimodular random graphs are always invariantly amenable, while infinitely ended ones are always invariantly non-amenable. (See the next section for the definitions of invariant amenability/non-amenability, which will nevertheless not be needed until Section 6.) This will imply that in the former case embeddings into the Euclidean plane are meant, while in the latter case embeddings into the hyperbolic plane.

For the special situation when G is transitive, the expression of “an isometry-invariant embedding of G ” makes sense right away: we want the random embedded graph to be isomorphic to this fixed graph almost surely. (This is how the term is used in [Tim21].) However, when the given rooted graph (G, o) is in fact random, one has to assign a root to the embedded graph to be able to ask for an almost sure rooted isomorphism. This is why it was essential to introduce the Palm version above, and for that definition to work, we need the set of embedded vertices, as a point process, to have finite intensity. If the vertices of a random embedded graph have some accumulation point, then the corresponding invariant point process has infinite intensity. In particular, if every embedding of a given URPG has some accumulation point, then the above question of isometry-invariant embeddings is not defined for that graph. Let us mention that there exist random graphs embedded in the Euclidean plane in an invariant way such that it is not possible to assign to them a root so that the resulting rooted graph is unimodular. See Example 2.4. This shows that defining an invariant embedding of a URPG through the Palm version and hence the requirement of finite intensity, is not the artifact of a possibly wrongly chosen definition, but seems to be the most general and natural choice.

Say that an embedding into a surface is *locally finite* if every compact set is intersected by only finitely many embedded edges. A point of the surface is an *accumulation point* if all its neighborhoods intersect infinitely many embedded edges. As we will see later, it is easy to find examples of URPG's that do not have a locally finite embedding into the plane, and even for those that do, it is a priori unclear whether such an embedding can be chosen to be unimodular. We are led to the following questions:

1. Give some combinatorial characterization of infinite locally finite graphs that have a locally finite embedding in the plane.
2. Characterize those unimodular random planar graphs where there is a *unimodular* embedding into \mathbb{R}^2 or \mathbb{H}^2 with the above property, and construct such an embedding. Similarly, find an isometry-invariant embedding.

We will fully answer both questions. For the first one, a characterization is given not

only for the plane, but other orientable surfaces as well (Theorem 9). The planar case itself (Theorem 2) was already characterized by Halin [Hal66]. These results are analogues of Kuratowski’s theorem with minors replaced by “minors that can use infinity”.

For the second problem, having a locally finite (not necessarily unimodular) embedding will turn out to be equivalent to having a unimodular *combinatorial* embedding in the plane with no accumulation points (a notion yet to be defined in the context of combinatorial embeddings), see Theorem 5. These are further equivalent to having a locally finite unimodular embedding into the Euclidean plane or the hyperbolic plane, which will be equivalent to having an isometry-invariant embedding (Theorem 1). Another variant of our result says that having a one-ended URPG as a unimodular supergraph is equivalent to having a (locally finite) invariant embedding into the Euclidean or hyperbolic plane (combine Theorem 4 with Theorem 1 for one direction, and Proposition 1.1 for the other).

When there is a positive answer to the second problem, our starting point is the construction in [Tim23] for a unimodular *combinatorial* embedding of a URPG. A combinatorial version of “locally finite” has to be introduced, and we modify the construction of [Tim23] to make sure that the unimodular combinatorial embedding is locally finite whenever G has a locally finite embedding at all (Theorem 5). After obtaining the unimodular locally finite combinatorial embedding, one has to realize it as an actual locally finite embedding into the plane. At this point, the metric imposed on the plane needs to be clarified, to make sense of what is meant by the embedding to be unimodular. The two natural choices are the Euclidean and the hyperbolic metric. Then a dichotomy result holds (Theorem 1), similarly to [BT21].

1.1 Our main results

As before, let M denote either \mathbb{R}^2 or \mathbb{H}^2 . A *drawing* of a graph G on M is a locally finite embedding ι of G to M , viewed up to isometries of M . A *rooted drawing* is a drawing together with a distinguished vertex. We denote by $\mathbb{RD}(M)$ the space of rooted drawings of locally finite graphs on M , with a suitable topology given by “local closeness” in a neighborhood of the root. A probability measure μ on $\mathbb{RD}(M)$ is *unimodular*, if it satisfies the appropriate Mass Transport Principle. We give the detailed definitions in Subsection 2.4.

A *unimodular embedding* of a URPG (G, o) into M is a unimodular measure μ on $\mathbb{RD}(M)$, such that forgetting the embedding of a μ -random element (and taking the rooted-isomorphism class of the resulting rooted graph) gives (G, o) in distribution. This was first defined in [AHNR18]; a related notion of unimodular discrete spaces is studied in [BHMK21].

Let \mathfrak{G} denote a random embedded unrooted graph on M with $\text{Isom}(M)$ -invariant distribution. (Both the graph and its embedding can be random.) The *intensity* of \mathfrak{G} is $p(\mathfrak{G}) = \mathbb{E}[|V(\mathfrak{G}) \cap B|]$, where $B \subset M$ is an arbitrary measurable set of area 1. When the intensity is finite, one can define the *Palm version* \mathfrak{G}^* of \mathfrak{G} , by conditioning on $0 \in V(\mathfrak{G})$, where 0 is the origin in M . By standard theory of point processes this makes sense ([LP17, Chapter 9]), and \mathfrak{G}^* defines a random rooted drawing, with the distinguished vertex at 0 .

An *invariant embedding* of a URPG (G, o) into M is an invariant random locally

finite embedded unrooted graph \mathfrak{G} such that the Palm version of \mathfrak{G} is a unimodular embedding of (G, o) . So when invariant embeddings of URPG's are considered, finite intensity is inherent.

It is clear that in order to find either a unimodular or an invariant embedding of a URPG without accumulation points, we need to assume that it is supported on graphs that have a locally finite planar embedding. We show that under this assumption a URPG has an invariant or unimodular embedding of finite, positive intensity into either the Euclidean or the hyperbolic plane, depending on whether it is (invariantly) amenable or not. (So far intensity only has a meaning for invariant embeddings, but the intensity of a unimodular embedding will also be defined, in Subsection 2.6.) We say that a URPG (G, o) has *finite expected degree* if the expected degree of o is finite.

Theorem 1. *Let (G, o) be a URPG with finite expected degree, and assume that G has a locally finite embedding into the plane with probability one. Then (G, o) has an invariant embedding and also a unimodular embedding with positive intensity and no accumulation point into*

- *the Euclidean plane if and only if (G, o) is invariantly amenable;*
- *the hyperbolic plane if and only if (G, o) is invariantly non-amenable.*

Theorem 1 is an extension of similar results in [BT21]. Here we do not require our URPG to be one-ended, replacing that assumption with having a locally finite embedding. Furthermore, the invariant embedding in the non-amenable case was not investigated in [BT21] (and the method of obtaining an invariant embedding from a unimodular one in the hyperbolic plane is not straightforward).

To better understand the assumptions of Theorem 1, one can ask for a graph theoretic characterization of infinite locally finite graphs that have a locally finite embedding in the plane. Such a characterization is given by Halin in [Hal66], see [BR03, Theorem 1.1] for a statement in English. Suppose G is infinite. Loosely speaking, we say that a finite graph H is a $*$ -minor of G if we can find H as a minor of G with one of the vertices of H being “at infinity” in G . See Definition 3.5.

Theorem 2 (Halin). *An infinite locally finite planar graph has a locally finite planar embedding if and only if it does not have $K_{3,3}$ or K_5 as a $*$ -minor.*

In fact we proved a version of this statement for embedding graphs in orientable surfaces, and only found Halin's theorem upon review of the present paper. See Theorem 9 for our more general result.

For completeness, let us state the $*$ -minor characterizations of invariant/unimodular embeddability as a separate theorem. This is straightforward from Theorem 2 and Theorem 1.

Theorem 3. *Let (G, o) be an invariantly amenable (respectively, invariantly non-amenable) URPG with finite expected degree. Then*

- *(G, o) has a unimodular embedding of finite positive intensity into \mathbb{R}^2 (respectively, into \mathbb{H}^2) if and only if it does not have $K_{3,3}$ or K_5 as a $*$ -minor.*

- It has an invariant embedding of positive intensity into \mathbb{R}^2 (respectively, into \mathbb{H}^2) if and only if it does not have $K_{3,3}$ or K_5 as a $*$ -minor.
- It has no unimodular embedding of positive intensity and no invariant embedding into \mathbb{H}^2 (respectively, into \mathbb{R}^2).

As mentioned at the beginning, most of Theorem 1 was already proved in [BT21] for the case when G is almost surely one-ended. Our next theorem shows that one-endedness is, in a sense, the real reason for a graph to be unimodularly/invariantly embeddable in the plane.

Theorem 4. *Let (G, o) be a URPG with finite expected degree. Then the following are equivalent.*

- (1) G has a locally finite planar embedding almost surely, and hence Theorem 1 applies.
- (2) G has a unimodular one-ended planar supergraph.

Proof. It is a straightforward corollary of Proposition 1.1 below. ■

Our constructions for the embeddings in Theorem 1 will follow the strategy adopted in [BT21], once the following is established.

Proposition 1.1. *Let (G, o) be a URPG with finite expected degree that has a locally finite planar embedding almost surely. Then there is a decorated URPG $(G^+, o^+; S)$, with a connected, positive density subgraph S of G^+ , such that (G^+, o^+) is a planar triangulation of finite expected degree, and G^+ has one end. (Positive density means $\mathbb{P}[o^+ \in S] > 0$.) Furthermore, conditioned on $o^+ \in S$, (S, o^+) is distributed as (G, o) , and (G^+, o^+) is invariantly amenable if and only if (G, o) is.*

The definition of unimodularity for decorated graphs is introduced in Subsection 2.1. Proposition 1.1 is a generalization of Theorem 2.2 in [BT21]. It will follow from Theorem 5 presented below, which is the main novelty and most laborious result of this paper.

Already in the one-ended case a key ingredient of the proof is to construct *combinatorial embeddings* for URPG's in a unimodular way, see [Tim23]. Combinatorial embeddings are thoroughly discussed in Section 2. For now it is enough to know that a combinatorial embedding is a decoration of the vertices by finite sets of labels, where each label describes the cyclic order of edges around the vertex in an embedding, and together the labels encode the embedding of the graph up to homeomorphisms. We will also define the number $\text{acc}(G, \pi)$ of accumulation points of a graph G with a combinatorial embedding π , see Definition 4.1. The difficulty in our case is coming from the fact that the unimodular combinatorial embedding constructed in [Tim23] can have $\text{acc}(G, \pi) > 1$ for graphs with more than one end. Here we give a different, more careful construction.

Theorem 5. *Let (G, o) be a URPG with finite expected degree, which has a locally finite planar embedding almost surely. Then (G, o) has a unimodular random combinatorial embedding π with $\text{acc}(G, \pi) \leq 1$ almost surely.*

The paper is structured as follows. Definitions are introduced in Section 2, together with some notation. Section 3 contains a graph theoretic characterization of infinite graphs that have some embedding into a given surface with only one accumulation point, yielding a proof of Theorem 2 as a special case.

The proof of our main Theorem 1, presented in Section 5, can be broken down into parts. In the Euclidean case we construct an invariant embedding first, which can be trivially used to define a unimodular embedding (similarly to [BT21]). In the hyperbolic case we are able to construct a unimodular embedding first. Obtaining an invariant embedding from the unimodular one is the nontrivial direction, we prove the necessary Theorem 7 in Subsection 2.6. Our first step towards a unimodular embedding in the hyperbolic case is to find unimodular *combinatorial* embeddings into the plane with one combinatorial accumulation point. That is, we prove Theorem 5 in Section 4. This is the lengthiest and most laborious section of our paper.

A unimodular combinatorial embedding as above can then be used to define an actual unimodular embedding of positive intensity into \mathbb{H}^2 when the graph is invariantly non-amenable. This was known before, via circle packings, and we briefly summarize the method in Section 5. Here we also prove Proposition 1.1 as a consequence of Theorem 5, and the construction parts of Theorem 1 using Proposition 1.1 and Theorem 7. Finally in Section 6 we prove the other direction, the necessary conditions of our embeddability theorems.

2 Preliminaries

Our focus in this paper is planar graphs and unimodularity. Nevertheless, throughout Sections 2 and 3 we will keep our setup more general, and consider embeddings of arbitrary graphs into orientable surfaces. The proof of Theorem 5 and as a consequence Theorem 1 will not make use of any of this general setup.

All graphs studied are assumed to be connected with at least 2 vertices. Disconnected graphs only appear when we delete subgraphs to study how the components of the remainder behave. $\mathcal{C}_G(x)$ denotes the connected component of the vertex x in the graph G .

Let $H' \subseteq H$ be a subgraph. The graph $H \setminus H'$ is obtained by removing the vertices of H' and all adjacent edges from H . For a vertex $v \in V(H)$ let $N_H(v)$ denote the neighbors of v in H .

To remove a set of edges $E \subseteq E(H)$ we also write $H \setminus E$. In this case all vertices of H are kept, only the edges are deleted.

2.1 Unimodular random graphs

A *random rooted graph* is a probability measure ν on \mathcal{G}_\bullet , the space of rooted, connected, locally finite graphs, considered up to rooted isomorphisms. The set \mathcal{G}_\bullet is a locally compact Polish space, the topology is defined by the rooted distance d_r , where

$$d_r((G_1, o_1), (G_2, o_2)) = \frac{1}{2^{k+1}}, \text{ where } k = \sup \{i \in \mathbb{N} \mid B_{G_1}(o_1, i) \cong B_{G_2}(o_2, i)\}.$$

We say ν is *unimodular*, if it satisfies the Mass Transport Principle [AL07]. That

is, for any measurable function $f : \mathcal{G}_{\bullet\bullet} \rightarrow [0, \infty)$ on the space $\mathcal{G}_{\bullet\bullet}$ of connected, locally finite birooted graphs (up to birooted isomorphisms) we have

$$\int_{(G,o)} \sum_{u \in V(G)} f(G, o, u) d\nu(G, o) = \int_{(G,o)} \sum_{u \in V(G)} f(G, u, o) d\nu(G, o).$$

The definition can be repeated when vertices or edges of the graph are decorated by elements of a complete separable metric space. See [AL07] or [Lov12, Chapter 18.3] for details. Usually we write (G, o) for a unimodular random rooted graph, meaning implicitly that its distribution is a unimodular measure. When a decoration, like a function f on edges or vertices, or a subset U of edges or vertices is present, we write $(G, o; f)$ or $(G, o; U)$. An equivalent formulation of unimodularity is by involution invariance. Assume (G, o) has finite expected degree, and let (G', o') result from (G, o) after biasing by the degree of o . Let (G', o', o'') denote the random birooted graph obtained by taking a uniform random neighbor o'' of o' in G' . The random rooted graph (G, o) is unimodular if and only if the distribution of (G', o', o'') is the same as the distribution of (G', o'', o') . The same equivalence holds for decorated graphs.

We say a unimodular random graph (G, o) is *invariantly amenable*, if for every $\varepsilon > 0$ there is a random subset $U \subseteq V(G)$ such that $(G, o; U)$ is unimodular, every component of $G \setminus U$ is finite, and $\mathbb{P}[o \in U] < \varepsilon$. (See [AL07] for equivalent definitions.) Otherwise it is called *invariantly non-amenable*. For simplicity from now on we will refer to these as amenable and non-amenable. As we only consider unimodular random graphs in this paper, this should cause no confusion.

2.2 Embeddings of arbitrary graphs

For the rest of this Section 2 let G be a connected, locally finitely graph and Σ a closed, connected, and orientable surface. An *embedding* of G into Σ is a map ι that maps vertices of G to distinct points of Σ and edges e with endpoints x and y to arcs between $\iota(x)$ and $\iota(y)$ such that no inner point of these arcs is contained in another arc. We denote by $\iota(G)$ the union of all these vertices and arcs in Σ . We call the connected components of $\Sigma \setminus \iota(G)$ the *topological faces* of $\iota(G)$. We say the embedding is *amicable*, if for any finite subgraph F of G , the topological faces of $\iota(F)$ are the interiors of compact orientable surfaces with boundary. All embeddings considered in this paper are assumed to be amicable. For planar embeddings in this paper we also assume amicability, by which we mean amicability of the embedding as viewed to be mapping into the one-point compactification S^2 .

A few remarks are in order:

- We do not require the the embedding to be *cellular*, i.e. the topological faces of F do not have to be disks. In the planar case, which is our main interest, amicable embeddings are automatically cellular. In general however, topological faces of F are assumed to be of the form $\Sigma' \setminus (D_1 \cup \dots \cup D_k)$, where Σ' is a closed orientable surface and the $D_j \subseteq \Sigma'$ are disjoint closed disks. This makes our setup somewhat different from the ones discussed in [LZ13] and [MT01]. The reason for our choice is that we want to make sure that graphs that can be embedded into a fixed surface form a minor-closed family. This is not the

case if we require all topological faces to be disks. The only place where this remark applies is Theorem 9, where we go beyond the planar setup.

- In the definition of amicability, we do not impose our topological requirement on topological faces of infinite subgraphs, because we want to allow examples with “infinite faces”, like the natural embedding of \mathbb{Z} into \mathbb{R}^2 . The complement of \mathbb{Z} on S^2 is connected, but not homeomorphic to an open disk. The requirement on the finite subgraphs will give us the right generality and enough control.
- When considering planar embeddings of infinite graphs on S^2 , the embedding is of course not locally finite anymore. Nevertheless, such locally finite planar embeddings are in one-to-one correspondence with embeddings into S^2 with exactly one accumulation point (which point is not in the image of any edge or vertex by the embedding). We will reserve the term *simply connected*, for such embeddings into S^2 . So “simply connected” and “locally finite planar” embeddings are the same, but we will use the two terminologies to indicate the space we are embedding into.

The *graph boundary* of a topological face of G (with respect to ι) is the subgraph that is mapped onto the topological boundary of the face by ι . So while a topological face of G is a domain in Σ , the graph boundary of a topological face in G is a subgraph of G . (The term “simply connected”, frequently used in the literature, indicates that the union of the closures of the topological faces with finite graph boundary is a simply connected topological space.)

2.3 Combinatorial embeddings

In this subsection we introduce some notions from the theory of graphs on surfaces. See [LZ13, Chapter 1.3.3], [MT01, Chapter 3.1] or [BWGT09] for more details on the topic, but keep in mind that our setup is slightly different, because we do not assume embeddings to be cellular.

Fix an embedding ι of G into Σ . For each edge e of G introduce two oriented edges \vec{e} and \overleftarrow{e} , or *darts* in opposite directions. We denote by $s(\vec{e})$ and $t(\vec{e})$ the source and terminal vertex of a dart. When the edge is a loop it still gives two opposing darts, but in this case the source and terminus are the same.

The orientation of the surface defines a cyclic permutation π_v at every vertex $v \in V(G)$ of the darts leaving that vertex. The collection of all these permutations $(\pi_v)_{v \in V(G)}$ is the *combinatorial embedding* of G corresponding to ι . In general such a collection of cyclic permutations is called a combinatorial embedding, if there is some ι that realizes it.

We go on to explore how this combinatorial data can be used to reconstruct the embedding when G is finite.

Given a graph G with a combinatorial embedding $(\pi_v)_{v \in V(G)}$ into Σ we can define the *combinatorial faces* of G . A combinatorial face is a cyclically ordered tuple of darts $(\vec{e}_1, \dots, \vec{e}_n)$ with $t(\vec{e}_i) = s(\vec{e}_{i+1}) = v_i$ and such that the dart \vec{e}_{i+1} is the successor of the dart \overleftarrow{e}_i according to the permutation π_{v_i} . We understand $i+1$ cyclically, so $\vec{e}_{n+1} = \vec{e}_1$. We can think of a combinatorial face as a returning

walk along oriented edges, always continuing on the unique succeeding dart at every vertex. Note that a combinatorial face might contain both darts belonging to the same edge. When G is finite, for every dart \vec{e} there is a unique combinatorial face that contains it. We will revisit the infinite case in subsection 4.10.

When $\Sigma = S^2$ the embedding is automatically cellular, so each combinatorial face is the graph boundary of a topological face, and these topological faces are all homeomorphic to the open disk. So gluing disks to the combinatorial faces of the combinatorial embedding π reconstructs an embedding that realizes π . In particular if two embeddings ι_1 and ι_2 define the same combinatorial embedding, then they are the same up to a homeomorphism of S^2 . That is, there exists a homeomorphism $\varphi : S^2 \rightarrow S^2$ such that $\iota_1 = \varphi \circ \iota_2$.

When Σ is some other orientable surface the picture is more complicated. The cyclic graph C_n can be embedded in the torus T^2 two different ways that are not the same up to homeomorphisms, yet define the same (and only) combinatorial embedding of C_n . (Neither of these embeddings is cellular.) To be able to reconstruct an embedding from the combinatorial data we further need to record the genus of the topological faces of G , as well as the number of closed disks removed (see Example 2.1). Each topological face F_i is of the form $\Sigma_i \setminus (D_1^i \cup \dots \cup D_{k_i}^i)$, where Σ_i is an orientable surface and the $D_j^i \subseteq \Sigma_i$ are disjoint closed disks ($k_i \geq 1$).

The graph boundary ∂F_i is a union of combinatorial faces of G . Each combinatorial face is part of the graph boundary of exactly one topological face. So we get a partition of the combinatorial faces $F_c(G, \pi)$ of (G, π) into sets $P_i \subseteq F_c(G, \pi)$. We have $P_i \cap P_j = \emptyset$ for $i \neq j$ and $\cup_i P_i = F_c(G, \pi)$. To each subset P_i we record the genus g_i of Σ_i . Note that $g_i \leq g(\Sigma)$.

To summarize, we collected the data $\pi = (\pi_v)_{v \in V(G)}$, the partition (P_i) of $F_c(G, \pi)$ and for each P_i the number g_i . This information is sufficient to reconstruct the embedding ι by gluing $\Sigma_i \setminus (D_1^i \cup \dots \cup D_{k_i}^i)$ to G along the combinatorial faces collected in P_i , where Σ_i is the surface of genus g_i . Here $k_i = |P_i|$. If two embeddings give the same combinatorial data, they are the same up to homeomorphisms. Note that the correspondence between the combinatorial faces in P_i and the D_j^i need not be recorded for the reconstruction of the embedding, as any bijective mapping gives the same embedding up to homeomorphisms.

Corollary 6. *Let the finite graph G and the orientable surface Σ be fixed. There are finitely many embeddings of G into Σ up to homeomorphisms.*

Proof. There are finitely many choices for the combinatorial data π , (P_i) and (g_i) , with $g_i \leq g(\Sigma)$. ■

Example 2.1. For the cycle C_n let π denote its only combinatorial embedding (which is an involution of the two outgoing darts at every vertex). Let E and E' denote the two combinatorial faces. When embedding C_n into T^2 the two distinct embeddings produce the combinatorial data $(\pi, P_1 = \{E, F\}, g_1 = 0)$ and $(\pi, P_1 = \{E\}, P_2 = \{F\}, g_1 = 0, g_2 = 1)$.

Remark 2.2. Not all possible choices of π , (P_i) and (g_i) with $g_i \leq g(\Sigma)$ will give an embedding into Σ . Even in the cellular case (meaning each P_i is a singleton and all g_i are 0), the genus of the surface created by the gluing procedure might vary depending on π . See [BWGT09, Chapter 3] where the enumeration of cellular maps into surfaces of different genus is discussed.

2.4 Unimodular embeddings

We study invariant and unimodular embeddings in the planar case, that is graphs embeddable into $\Sigma = S^2$. In the previous subsections we considered embeddings up to homeomorphisms. As we mentioned before, the difference between amenable and non-amenable URPG's becomes apparent when we specify the metric on the punctured sphere, and consider embeddings up to isometries.

As in the Introduction, let M denote either \mathbb{R}^2 or \mathbb{H}^2 . We denote by $\mathbb{G}(M)$ the space of graphs amicably embedded into M . Elements of $\mathbb{G}(M)$ are denoted by typewriter font letters like \mathbf{G} . $\mathbb{G}(M)$ carries a standard Borel structure by requiring all edge- and vertex-counting functions over Borel subsets of M to be measurable. The group $\Gamma = \text{Isom}(M)$ acts on $\mathbb{G}(M)$ measurably by shifting the vertices and edges. The shift of $\mathbf{G} \in \mathbb{G}(M)$ by $\varphi \in \Gamma$ is denoted $\varphi.\mathbf{G}$. We denote by $\mathbb{G}_0(M)$ the subspace of embedded graphs that have a vertex at the origin $0 \in M$; this is again a standard Borel space. We denote by $\mathbb{G}_\bullet(M) = \{(\mathbf{G}, \mathfrak{o}) \mid \mathbf{G} \in \mathbb{G}(M), \mathfrak{o} \in V(\mathbf{G})\}$ the set of rooted embedded graphs.

A *drawing* of a graph on M is an *equivalence class* of locally finite, amicably embedded graphs, where two such embeddings are equivalent if they are the same up to an element of Γ . The set of such drawings is denoted $\mathbb{D}(M)$. The equivalence class of an embedded graph $\mathbf{G} \in \mathbb{G}(M)$ is denoted by $[\mathbf{G}]$.

A *rooted drawing* is a drawing together with a distinguished vertex. More precisely, the space of rooted drawings $\mathbb{RD}(M)$ is $\mathbb{G}_\bullet(M)/\Gamma$. The equivalence class of $(\mathbf{G}, \mathfrak{o}) \in \mathbb{G}_\bullet(M)$ is denoted by $[\mathbf{G}, \mathfrak{o}]$. Equivalently, a rooted drawing can be thought of as a rooted graph embedded into M with the root at the origin $0 \in M$, up to isometries of M fixing 0. Therefore $\mathbb{RD}(M)$ is in fact in bijection with $\mathbb{G}_0(M)/\text{Stab}_\Gamma(0)$. We use this correspondence to define the Borel structure: as $\mathbb{RD}(M)$ is a factor of a standard Borel space by a compact subgroup it is itself standard Borel. The embedded graph $\mathbf{G} \in \mathbb{G}_0(M)$ represents the class $[\mathbf{G}, 0]$.

We say a probability measure μ on $\mathbb{RD}(M)$ is *unimodular*, if it satisfies the appropriate Mass Transport Principle. Namely, one defines the space $\mathbb{BRD}(M)$ of *birooted drawings* of graphs on M similarly to $\mathbb{RD}(M)$, the equivalence class of a $(\mathbf{G}, \mathfrak{u}, \mathfrak{v})$ is denoted by $[\mathbf{G}, \mathfrak{u}, \mathfrak{v}]$. A *payment function* is a measurable function $f : \mathbb{BRD}(M) \rightarrow [0, \infty)$. We say μ is unimodular if the expected income of the root of a μ -random rooted drawing is equal to the expected outpay, for any payment function f :

$$\int_{\mathbb{RD}(M)} \sum_{\mathfrak{u} \in V(\mathbf{G})} f([\mathbf{G}, \mathfrak{o}, \mathfrak{u}]) \, d\mu([\mathbf{G}, \mathfrak{o}]) = \int_{\mathbb{RD}(M)} \sum_{\mathfrak{u} \in V(\mathbf{G})} f([\mathbf{G}, \mathfrak{u}, \mathfrak{o}]) \, d\mu([\mathbf{G}, \mathfrak{o}]).$$

In the formula we have to choose a representative $(\mathbf{G}, \mathfrak{o})$ of $[\mathbf{G}, \mathfrak{o}]$ to compute the sum over $V(\mathbf{G})$, even though we integrate over $[\mathbf{G}, \mathfrak{o}]$. This is still well-defined because the value of the sum does not depend on the choice of representative. This problem of having to choose representatives for formulas to make sense is a recurring phenomenon in this section; our intention is to keep the notation reasonably clean and intuitive. We will point this out whenever it might cause confusion.

When defining unimodular and invariant embeddings of URPG's, we are interested in random variables with values in the spaces $\mathbb{G}(M)$ and $\mathbb{RD}(M)$. We will use

gothic letters like \mathfrak{G} to denote them. A *unimodular embedding* of a URPG (G, o) into M is a random rooted drawing $[\mathfrak{G}, \mathfrak{o}]$ such that its distribution μ is unimodular, and forgetting the embedding of any representative $(\mathfrak{G}, \mathfrak{o})$ gives (G, o) in distribution.

Remark 2.3. It is standard to take random rooted graphs themselves as equivalence classes of rooted graphs up to rooted isomorphism. (As we implicitly did in subsection 2.1.) One could ask why we never highlighted this in our notation, yet keep emphasizing the difference between classes and representatives when it comes to rooted drawings. The reason is that, in the case of drawings, equivalence is up to an isometry of the ambient space. This is much less intuitive from the name “rooted drawing” than the usual identification of isomorphic rooted graphs.

2.5 Invariant embeddings

In a moment we will define invariant embeddings through unimodularity of their Palm versions. Before that we would like to note that properties of Palm versions of invariant point processes are thoroughly studied in the literature, but conditions on point processes that inherently characterize them to be the Palm version of some invariant point process are much less frequent. Nevertheless, in the Euclidean setting such techniques have been elaborated (see [HL05], and also [Nev77, Proposition II.12]). The theory was also extended to Abelian locally compact groups (see [LT09]), and then to homogeneous spaces of general locally compact groups, in particular covering the hyperbolic case. The work of Last [Las10] covers the theory introduced in the present and consequent subsections in this generality, while much of what we need here is already present in [RZ90].

Recall that $\Gamma = \text{Isom}(M)$, and $\mathbb{G}(M)$ is the space of amicably embedded graphs. A *random embedded unrooted graph* \mathfrak{G} is a random variable with values in $\mathbb{G}(M)$. When \mathfrak{G} is Γ -invariant in distribution, the *intensity* is $p(\mathfrak{G}) = \mathbb{E}[|V(\mathfrak{G}) \cap B|]$, where $B \subset M$ is an arbitrary measurable set of area 1. The value $p(\mathfrak{G})$ does not depend on the choice of B because of the Γ -invariance. Notice that the intensity cannot be 0, but it can be ∞ . For \mathfrak{G} with finite intensity, let \mathfrak{G}^* denote its Palm version. Then the distribution of $[\mathfrak{G}^*, 0]$ is a unimodular measure on $\mathbb{RD}(M)$, see e.g. Example 9.5 in [AL07]. We say \mathfrak{G} is an *invariant embedding* of the URPG (G, o) into M , if $[\mathfrak{G}^*, 0]$ is a unimodular embedding of (G, o) . To sum up, an invariant embedding of a given URPG is a Γ -invariant random embedded unrooted graph, whose Palm version rooted at 0 gives back the URPG after forgetting the embedding.

Notice that we assume finite intensity of \mathfrak{G} to be able to choose a root in the graph, giving rise to a unimodular random rooted graph: the choice is provided by the Palm version. The next example indicates that the requirement of finite intensity is necessary.

Example 2.4. Consider an invariant point process of finite intensity in the Euclidean plane; let ω be the corresponding random point set. For every $x \in \omega$, let $p_{x,i}$ be a uniformly chosen point from the ball of radius 2^{-i} around x . Define the Voronoi tessellation for the set $\{p_{x,i}, x \in \omega, i = 1, 2, \dots\}$. Connect two points if their Voronoi cells are adjacent. The resulting graph has countably infinitely many ends, hence it is not unimodular with any choice of root. (Like transitive graphs, unimodular random rooted graphs have 0, 1, 2 or uncountably many ends, see [AL07,

Proposition 6.10].) On the other hand it is a random graph drawn in the plane with an isometry-invariant distribution.

2.6 Invariant embeddings from unimodular ones

Assume as before that \mathfrak{G} is a Γ -invariant embedded unrooted graph with $p(\mathfrak{G}) < \infty$, and therefore $[\mathfrak{G}^*, 0]$ is unimodular. If λ denotes the area in M and $\text{Vor}(\mathfrak{G}^*, 0)$ the Voronoi cell of 0, then

$$E_{\mathfrak{G}^*}[\lambda(\text{Vor}(\mathfrak{G}^*, 0))] = \frac{1}{p(\mathfrak{G})}. \quad (1)$$

Equation (1) is very intuitive, and a basic corollary of the *Voronoi inversion formula*, a basic tool in the theory of point processes, at least in the Euclidean case ([LP17, Section 9.4]). For homogeneous spaces [Las10] covers the theory.

Inspired by (1) we introduce the *intensity* of a unimodular embedding $[\mathfrak{G}_\circ, \circ]$ as

$$p([\mathfrak{G}_\circ, \circ]) = \mathbb{E}_{[\mathfrak{G}_\circ, \circ]}[\lambda(\text{Vor}(\mathfrak{G}_\circ, \circ))]^{-1}.$$

The reason for this less direct definition is that a priori fixing a measurable subset B of the space is not possible (as it was in the definition of intensity of an invariant embedding), since in this context everything is understood up to isometries. The value $\lambda(\text{Vor}(\mathfrak{G}_\circ, \circ))$ is well defined, it does not depend on the choice of representative of $[\mathfrak{G}_\circ, \circ]$. Notice that we allow the intensity of a unimodular embedding to be 0 (when $\mathbb{E}_{[\mathfrak{G}_\circ, \circ]}[\lambda(\text{Vor}(\mathfrak{G}_\circ, \circ))]$ is infinite), but not ∞ .

Moreover, the intensity of \mathfrak{G} can be measured using any factor allocation (to be defined next), not just the Voronoi. An *allocation scheme* is a method of partitioning the space M (up to measure 0) into measurable pieces, given a discrete subset $\omega \subset M$, such that the measurable pieces are bijectively associated to the points in ω . One also assumes that the partition depends on ω in a measurable way. For example in the Voronoi allocation scheme each point of ω is associated to its Voronoi cell. We denote the piece associated to a point $v \in \omega$ by $\Psi_\omega(v)$.

The allocation is a *factor allocation* if it is equivariant, i.e., for any isometry g of M we have $\Psi_{g\omega}(g.v) = g.\Psi_\omega(v)$. If the random point set ω is defined as the embedded vertices of some URPG by a unimodular embedding, which is by definition only defined up to rooted isometries of M , the property of equivariance guarantees that such an allocation scheme provides a well-defined allocation to the points of these ω almost surely. Hence factor allocations are automatically well-defined for unimodular embedded graphs (even the 0 intensity ones). We denote the scheme itself by Ψ .

If \mathfrak{G} is Γ -invariant random, then for any factor allocation Ψ we have the analogue of (1),

$$\mathbb{E}_{\mathfrak{G}^*}[\lambda(\Psi_{V(\mathfrak{G}^*)}(0))] = \frac{1}{p(\mathfrak{G})}.$$

The following theorem allows us to reverse this correspondence.

Theorem 7 (Theorem 7.1 in [Las10] and Theorem 3 in [RZ90]). *Let $[\mathfrak{G}_\circ, \circ]$ be a unimodular random rooted drawing, and Ψ a factor allocation scheme such that the expected area of the cell of \circ is finite, that is $\mathbb{E}_{[\mathfrak{G}_\circ, \circ]}[\lambda(\Psi_{V(\mathfrak{G}_\circ)}(\circ))] < \infty$. Then there*

is a Γ -invariant random embedded unrooted graph \mathfrak{G} of finite intensity such that $[\mathfrak{G}^*, 0]$ is the same as $[\mathfrak{G}_\circ, \circ]$ in distribution.

In particular if a URPG (G, \circ) has a unimodular embedding of positive intensity, it also has an invariant embedding of finite intensity.

For the sake of completeness and full agreement with our terminology and notation, we include the proof of Theorem 7 in the Appendix.

The Palm version of an invariant process determines the process, therefore as a corollary of Theorem 7 one gets the following.

Corollary 8. *Let $[\mathfrak{G}_\circ, \circ]$ be a unimodular random rooted drawing into M . Then for any factor allocation schemes Ψ and Ψ' we have*

$$\mathbb{E}_{[\mathfrak{G}_\circ, \circ]} [\lambda(\Psi_{V(\mathfrak{G}_\circ)}(\circ))] = \mathbb{E}_{[\mathfrak{G}_\circ, \circ]} [\lambda(\Psi'_{V(\mathfrak{G}_\circ)}(\circ))].$$

Remark 2.5. Again, there is a slight abuse of notation in the corollary above. A unimodular random rooted drawing only gives an equivalence class, so the cell of the root is not defined as a subset of M . Nevertheless, as both allocation schemes are factors, the area of the cell of the root is well-defined.

3 Surface graphs with 1 accumulation point

Fix a compact orientable surface Σ . Let \mathcal{G}_Σ denote the family of locally finite, connected graphs that can be embedded in Σ . Let \mathcal{G}_Σ^f denote the family of finite graphs in \mathcal{G}_Σ . As \mathcal{G}_Σ^f is minor closed, by the Robertson-Seymour graph minor theorem (see [RS04] and references therein, or Chapter 12.5 in [Die17]) it can be characterized by a finite set of forbidden minors $M_\Sigma = \{H_1, \dots, H_n\}$.

$$\mathcal{G}_\Sigma^f = \{G \text{ finite, connected} \mid H_i \text{ is not a minor of } G, \forall 1 \leq i \leq n\}.$$

This characterization extends to \mathcal{G}_Σ almost automatically:

Lemma 3.1. *Let G be a locally finite graph, and $G_1 \subseteq G_2 \subseteq \dots$ an exhaustion of G by finite graphs. If all G_i can be embedded in Σ , then G can also be embedded in Σ .*

Proof. By Corollary 6 each G_i has finitely many embeddings into Σ up to homeomorphisms. Hence, by Konig's lemma, we can find a coherent sequence of embeddings of the G_i into Σ . That is, embeddings ι_i of G_i and homeomorphisms $\varphi_i : \Sigma \rightarrow \Sigma$ such that $\iota_i = \varphi_i \circ \iota_{i+1}|_{G_i}$.

Then we define an embedding ι of G by setting $\iota|_{G_i} = \varphi_1 \circ \dots \circ \varphi_{i-1} \circ \iota_i$. This is well defined since it stabilizes for every edge, and indeed an embedding, because any unwanted intersection of arcs would be witnessed in some G_i . ■

We can now characterize \mathcal{G}_Σ using minors. The possibly infinite, locally finite graph H is said to be a minor of the locally finite graph G if we can contract (possibly infinitely many) edges and delete edges and isolated vertices of G to obtain H .

Remark 3.2. Note that contracting infinitely many edges of G can result in graphs that are not locally finite. We however restrict our definition to the case where the minor H is also locally finite. When H is actually finite the situation is more straightforward, H is a minor of G if and only if it is a minor of a finite subgraph of G in the traditional sense.

Lemma 3.3. *The family \mathcal{G}_Σ consists of all locally finite graphs not having a minor in M_Σ .*

$$\mathcal{G}_\Sigma = \{G \text{ locally finite, connected} \mid H \text{ is not a minor of } G, \forall H \in M_\Sigma\}.$$

Proof. Let G be a locally finite graph, $G_1 \subseteq G_2 \subseteq \dots$ be an exhaustion of G by finite graphs. A finite graph H is a minor of G if and only if H is a minor of some G_i . By Lemma 3.1 $G \in \mathcal{G}_\Sigma$ if and only if $G_i \in \mathcal{G}_\Sigma^f$ for all i . ■

We now turn to embeddings with 1 accumulation point.

Definition. Let \mathcal{G}_Σ^* denote the family of graphs in \mathcal{G}_Σ that can be embedded in Σ with at most 1 accumulation point.

Remark 3.4. It is clear that \mathcal{G}_Σ^* is minor closed, yet this does not imply it can be characterized by a finite set of forbidden minors, because the Robertson-Seymour theorem applies to finite graphs, and it is an open question whether it extends to countable graphs [Die17, Chapter 12.7].

Recall that for finite graphs, H is a minor of G if and only if we can find $|H|$ -many pairwise disjoint subtrees of G such that contracting the subtrees gives a supergraph of H . Our definition of $*$ -minor follows this formulation.

Definition 3.5. Let $|H| = n$. H is a $*$ -minor of G if we can find $n - 1$ finite subtrees X_1, \dots, X_{n-1} and a subforest X_n with all connected components infinite such that

1. The vertex sets $V(X_i)$, $1 \leq i \leq n$, partition $V(G)$.
2. H is isomorphic to a subgraph of the graph we get from G by contracting all edges of the X_i and identifying all the points obtained from distinct connected components of X_n to a single one.

The graph obtained in part 2 might have infinitely many loops at the vertex coming from X_n . This makes no difference for us.

Remark 3.6. Note that a minor of G is not necessarily a $*$ -minor of G . In fact, finite graphs have no $*$ -minors at all. See Figure 1 for an infinite graph with $K_{3,3}$ as a $*$ -minor. The blue $*$ indicates the point of $K_{3,3}$ “at infinity”.

Lemma 3.7. *H is a $*$ -minor of G if and only if G has an infinite, locally finite minor G' such that*

1. G' has a subgraph isomorphic to $H \setminus v$ for some $v \in V(H)$. By slight abuse of notation $(H \setminus v) \subset G'$.
2. For all vertices $u \in N_H(v) \subseteq (H \setminus v)$ there is an infinite path starting at u in the graph $G' \setminus (H \setminus v)$.

Proof. Contracting the finite trees in the definition gives G' and vice versa. If $G' \setminus (H \setminus v)$ has some finite connected components, then we can contract them into $(H \setminus v)$ to ensure that we get a partition of the vertices. ■

We are now ready to state the main result of this section. Let Σ_g denote the orientable closed surface with genus g (which is unique up to homeomorphisms). Clearly $\mathcal{G}_{\Sigma_g} \subset \mathcal{G}_{\Sigma_{g+1}}$.




Figure 1: An infinite graph that has $K_{3,3}$ as a $*$ -minor

Theorem 9. *A planar graph is in $\mathcal{G}_{\Sigma_2}^*$ if and only if it does not have $K_{3,3}$ or K_5 as a $*$ -minor. For a positive genus $g \geq 1$ a graph $G \in (\mathcal{G}_{\Sigma_g} \setminus \mathcal{G}_{\Sigma_{g-1}})$ is in $\mathcal{G}_{\Sigma_g}^*$ if and only if it has no $*$ -minors in M_{Σ_g} .*

Remark 3.8. Note that because of the $G \in (\mathcal{G}_{\Sigma_g} \setminus \mathcal{G}_{\Sigma_{g-1}})$ condition the theorem only deals with the case when an infinite graph is embedded into a surface of minimal genus possible. This assumption is indeed necessary. The planar graph $K_3 \times \mathbb{Z}$ has no embedding into the torus T^2 with a single accumulation point. Still it does not have $*$ -minors in M_{T^2} , because all its $*$ -minors are embeddable into T^2 . To see this, consider the embedding β to the torus $[0, 1)^2$ (with $[0, 1)$ understood with 0 and 1 identified) that maps $(x, n) \in V(K_3 \times \mathbb{Z})$ to $(b_1(x), b_2(n))$, where b_1 is some bijection between $V(K_3)$ and $\{0, 1/3, 2/3\}$, and $b_2 : \mathbb{Z} \rightarrow (0, 1)$ is a strictly increasing function with $\lim_{n \rightarrow -\infty} b_2(n) = 0$ and $\lim_{n \rightarrow \infty} b_2(n) = 1$. Edges are embedded into geodesics between their endpoints. If H is a finite graph that arises as a $*$ -minor of G , then, using Definition 3.5, we can make $\beta(X_n)$ connected by adding some broken line segments to it that are disjoint from $\cup_{e \in E(K_3 \times \mathbb{Z})} \beta(e) \setminus \cup_{v \in V(K_3 \times \mathbb{Z})} \beta(v)$. This set together with the $\beta(X_i)$, $i = 1, \dots, n - 1$ can be used to represent the vertices of H (see the equivalent formulation of minors before Definition 3.5) for an embedding, with edges between them coming from $\beta(K_3 \times \mathbb{Z})$.

Remark 3.9. Halin's original theorem [Hal66] is not stated using $*$ -minors, but instead lists the 4 infinite graphs that, if minors of G , can be witnesses of $K_{3,3}$ or K_5 being a $*$ -minor of G . Such a translation is theoretically possible in the higher genus case as well. If one knew the graphs in M_{Σ_g} , one could construct the finitely many infinite graphs that are minor-obstructions to G being in $\mathcal{G}_{\Sigma_g}^*$.

Proof. Suppose that $G \in \mathcal{G}_{\Sigma_g}^*$, and assume towards contradiction that some $H \in M_{\Sigma_g}$ is a $*$ -minor of G . Carrying out the minor operations on the embedded G we can find an embedded G' with $(H \setminus v) \subset G'$ for some $v \in V(H)$. This gives an embedding of $H \setminus v$ into Σ_g .

We claim that with this embedding, all $u \in N_H(v)$ are on the topological boundary of the same topological face of $(H \setminus v)$, namely the one containing the accumulation point of the embedded G' . Indeed, all these u are part of an infinite path in $G' \setminus (H \setminus v)$, and in the embedding of G' all these infinite paths have to accumulate at the unique accumulation point.

Since all the vertices of $N_H(v)$ are on the topological boundary of the same topological face, we can place a vertex v anywhere on that topological face, and connect it with $N_H(v)$ without intersections. We found an embedding of $H \in M_{\Sigma_g}$ into Σ_g which is a contradiction.

For the other implication, suppose that G has no $*$ -minors in M_{Σ_g} . Let $G_1 \subset G_2 \subset \dots \subset G$ be an exhaustion of G by finite graphs. By increasing each G_i if

necessary, we can assume that all components of $G \setminus G_i$ are infinite. We also assume that $N_G(V(G_i)) \subseteq G_{i+1}$.

For each G_i let $V^*(G_i)$ denote the set of vertices $v \in V(G_i)$ that are adjacent to (an infinite component of) $G \setminus G_i$. By our second assumption $V^*(G_{i+1}) \cap V(G_i) = \emptyset$. We construct the graph G_i^+ by wiring together all vertices in $V^*(G_i)$ using an additional vertex. That is $V(G_i^+) = V(G_i) \cup \{p_i\}$ where p_i is an auxiliary point. For the edges we have $E(G_i^+) = E(G_i) \cup \{(p_i, u) \mid u \in V^*(G_i)\}$.

The graph G_i^+ can be embedded in Σ_g . Indeed, if $H \in M_{\Sigma_g}$ were a minor of G_i^+ , then H would be a $*$ -minor of G . Equivalently, G_i can be embedded in a way that all vertices in $V^*(G_i)$ are on the topological boundary of the same topological face F_i of G_i . This holds for all i .

For each G_i there are finitely many such embeddings up to homeomorphisms. By compactness we can choose consistent embeddings for the G_i . That is, restricting the embedding of G_i to G_j ($j \leq i$) agrees with the embedding of G_j . The embeddings of the G_i together give an embedding ι of G , as shown in Lemma 3.1. Since all vertices of G_i connected to infinity are along the same topological face F_i , we will be able to make sure that there is only one accumulation point, by choosing F_i such that $\lim F_i = \{x\}$ for some point $x \in \Sigma_g$.

In the case of planar graphs the intuition is very clear: each finite part gets embedded such that all vertices connected to infinity are along the outer face. We then claim that we can keep extending this embedding in a way that no accumulation point (besides infinity) is created.

To be precise, we argue as follows. For some i large enough, G_i cannot be embedded in a surface of smaller genus, which implies that all topological faces must be disks. In particular the topological face F_i is a closed disk D_i . Fix a homeomorphism between D_i and the unit disk in the Euclidean plane. This allows us to metrize D_i such that it has radius 1.

By changing the embedding of G_{i+1} by a homeomorphism only inside D_i , we can assume that all vertices in $V(G_{i+1}) \setminus V(G_i)$ lie in the region $D_i \setminus D_{i+1}$, where D_{i+1} is the disk concentric to D_i with radius $1/2$. Moreover we make sure that $V^*(G_{i+1})$ lies on the topological boundary of D_{i+1} .

Repeating this for all $k \geq 0$, we achieve that $V(G_{i+k+1}) \setminus V(G_{i+k})$ is mapped in the region $D_{i+k} \setminus D_{i+k+1}$, where D_{i+k} is concentric to D_i with radius $1/2^k$, while making sure that $V^*(G_{i+k})$ lies on the topological boundary of D_{i+k} . Note that at each step we only change the embedding inside D_{i+k} so for every vertex and edge the embedding eventually stabilizes. The resulting embedding has 1 accumulation point, the center of D_i . ■

Proof of Theorem 2. This is a special case of the first claim in Theorem 9. ■

4 Unimodular combinatorial embeddings of URPG's

From now on we only consider embeddings into S^2 .

In this section we prove Theorem 5. The titles of the subsections are meant to provide a sketch of the proof. The first three subsections give the necessary definitions and earlier results that we will be relying on.

4.1 Accumulation points of planar combinatorial embeddings

Let G be an infinite, locally finite graph. Given a combinatorial embedding of G , can we identify accumulation points that will be present for any actual embedding representing this combinatorial embedding?

In Section 2.3 we argued that for general surfaces a combinatorial embedding is not enough to reconstruct an embedding even up to homeomorphisms. In the case of S^2 , however, the reconstruction is possible when the graph is finite, because all combinatorial faces bound a topological face homeomorphic to a disk.

Unfortunately this does not hold for infinite graphs. Indeed, there are embeddings of \mathbb{Z} in S^2 with 1 or with 2 accumulation points, while both embeddings define the same combinatorial embedding. This means that for a combinatorial embedding $(\pi_v)_{v \in V(G)}$ the number of accumulation points of an embedding ι realizing π is not well defined. We are, however, interested in embeddings with as few accumulation points as possible.

We proceed to define $\text{acc}(G, \pi)$, the number of accumulation points of a combinatorial embedding, using only the combinatorial data. It will turn out that $\text{acc}(G, \pi)$ is exactly the minimal number of accumulation points needed to realize π .

Given a cycle on darts we define the inside of that cycle. Let C be a cyclically ordered tuple of distinct darts $(\vec{e}_1, \dots, \vec{e}_n)$ such that $t(\vec{e}_i) = s(\vec{e}_{i+1}) = v_{i+1}$ for $1 \leq i \leq n$. Let $\vec{e} \notin C$ be a dart with $s(\vec{e}) = v_i$. We say that \vec{e} is *inside* C if \vec{e} is somewhere between \overleftarrow{e}_{i-1} and \vec{e}_i in the cyclic order π_{v_i} . If $t(\vec{e}) \notin \{v_1, \dots, v_n\}$ (i.e. \vec{e} is not a chord of C), then the whole connected component of $t(\vec{e})$ in the graph $G \setminus C$ is said to be inside C . This definition does not depend on the choice of \vec{e} connecting the component to C .

The outside of a cycle is defined analogously. Note that inverting all darts along the cycle also inverts the notion of inside and outside. When the graph is finite, a combinatorial face is a cycle that has nothing inside it. We will say a part of a graph is *surrounded* by a cycle if it is inside that cycle.

Consider a finite connected subgraph $H \subset G$. It is endowed with the combinatorial embedding inherited from G . The cyclical order π_v^H at a vertex $v \in V(H)$ on the adjacent H -darts is the restriction of the cyclical order π_v on the adjacent G -darts. Hence one can also define the faces of H , inherited from the combinatorial embedding of G . These may not be faces of G , of course.

For a connected component of $G \setminus H$ there is a unique face of H that surrounds it. Let a_H^π denote the number of faces of H that surround at least one infinite component of $G \setminus H$. It is straightforward to check that if $H \subseteq H'$, where H' is also finite, then $a_H^\pi \leq a_{H'}^\pi$. This leads us to the following.

Definition 4.1. Let $H_1 \subset H_2 \subset \dots \subset G$ be an exhaustion of G by finite graphs. We define $\text{acc}(G, \pi) = \sup a_{H_i}^\pi$. The value of $\text{acc}(G, \pi)$ does not depend on the particular exhaustion because of the monotonicity. When the combinatorial embedding is clear from context, we write $\text{acc}(G)$.

Any embedding of G into S^2 that realizes π has at least $\text{acc}(G, \pi)$ accumulation points. On the other hand when $\text{acc}(G, \pi)$ is finite, there exists an embedding realizing π with exactly $\text{acc}(G, \pi)$ accumulation points. Indeed, suppose that $a_{H_i}^\pi$

stabilizes at i_0 . We can make sure that accumulation points only occur at the centers of the $\text{acc}(G, \pi)$ -many topological faces of H_{i_0} that surround infinitely many points of $G \setminus H_{i_0}$. The reasoning is the same as the last part of the proof of Theorem 9.

4.2 The general strategy

In [Tim23] the first author shows that URPG's have unimodular combinatorial embeddings into the plane. Notice that when the URPG has one end, *any* combinatorial embedding into S^2 gives $\text{acc}(G) = 1$.

Here, in Theorem 5, we show that as long as the URPG (G, o) is supported on $\mathcal{G}_{S^2}^*$, it has a unimodular random combinatorial embedding into the plane with $\text{acc}(G) \leq 1$ almost surely.

Recall that an embedding with one topological accumulation point is called *simply connected*. For example the d -regular tree T_d ($d \geq 2$) has continuum many ends, but admits simply connected embeddings into S^2 . On the other hand $K_2 \times T_3$ does not.

Throughout this Section we consider graphs of increasing generality and build up to a proof of Theorem 5. As in [Tim23], we construct the random embedding by decomposing our graph into 2-connected, and then further into 3-connected components. We choose random combinatorial embeddings for those parts and glue them together to form a combinatorial embedding of the whole graph.

The main reason why our construction works is that multi-ended infinite 2-connected components essentially have a unique combinatorial embedding with 1 accumulation point whenever they have at least one. (See Lemma 4.6.) Technical difficulties arise because we have to keep track of where ends get embedded.

4.3 Building blocks of graphs: Tutte and block-cut decompositions

The following are the main tools for constructing a random combinatorial embedding in [Tim23, Section 2]. References and a more detailed introduction can be found there, here we try to tighten the exposition.

4.3.1 3-connected components

The starting point of the construction is the fact (due to Whitney and Imrich) that a 3-connected graph has a unique combinatorial embedding, up to inverting all permutations, even in the infinite case.

Example 4.2. To find examples of 3-connected graphs with infinitely many ends that have a locally finite embedding do the following. Start with any (possibly unimodular random) graph with infinitely many ends and a locally finite planar embedding, and assume for simplicity that it has no leaves (e.g. regular-, or unimodular Galton-Watson trees). Perform the local operation shown in Figure 2 around every vertex to turn it into a 3-connected graph while preserving the other properties.




Figure 2: Local operation ensuring 3-connectedness

4.3.2 Tutte-decomposition

Any 2-connected graph G has a unique Tutte-decomposition into 3-connected components, 3-links and cycles. A 3-link is a graph consisting of 3 parallel edges.

By the *amalgam* of two graphs along an edge we mean taking their disjoint union, picking an edge in both (together with a bijection between their endpoints), deleting those edges, and identifying the endpoints. Amalgamation of graphs A and B is denoted by $A + B$, the edge will be clear from context.

Informally speaking, the Tutte decomposition describes a treelike way of obtaining G by amalgamating 3-connected components and cycles, possibly keeping (1 copy of) the edge we are amalgamating along. We can achieve this by inserting a 3-link between two graphs during the amalgamating procedure whenever we want to keep the common edge. See Figure 3 for the Tutte decomposition of the infinite ladder.




Figure 3: The Tutte decomposition of the infinite ladder

More precisely, the Tutte decomposition gives a tree T with vertices $\alpha \in V_T$ labelled by graphs G_α , and edges (α, β) labelled by functions $f(\alpha, \beta)$. The G_α are 3-connected, 3-links, or cycles, and the function $f(\alpha, \beta)$ picks an edge from G_α and G_β , and a bijection between their endpoints. Moreover, each edge of each G_α is picked by at most one $f(\alpha, \beta)$. The order of amalgamations is arbitrary, and we denote the graph obtained in the end by $\Gamma(T)$. We say an edge in some G_α is *virtual*, if it is selected by some $f(\alpha, \beta)$. An edge of some G_α is present in $\Gamma(T)$ if and only if it is not virtual. As amalgams of cycles are cycles, we also assume that cycles cannot be neighbors in the Tutte tree. This guarantees uniqueness of the decomposition.

4.3.3 Block-cut decompositions

Finally, if G is a locally finite, connected planar graph, it has a block-cut tree decomposition into 2-connected components. The 2-connected components (blocks) G_a are indexed by some set A . The cut vertices form the set C . The vertex set of the block-cut tree \mathcal{T} is $A \cup C$, and $a \in A$ is connected to $v \in C$ by an edge if $v \in V(G_a)$. In \mathcal{T} the vertices $a, b \in A$ are at distance 2 if and only if G_a and G_b have a (unique) common vertex.

4.4 2-connected graphs with full core

In this subsection we define the *core* of 2-connected infinite graphs, and investigate combinatorial embeddings of such graphs with *full core*.

Definition. Let G be an infinite, 2-connected graph and $T = (V_T, E_T)$ the Tutte tree with $\Gamma(T) = G$. Let $V_\infty \subseteq V_T$ denote the set of vertices that represent infinite 3-connected Tutte components. Let $\text{core}(T)$ denote the convex hull of $V_\infty \cup \text{Ends}(T)$. That is, a vertex $\alpha \in V_T$ is in $\text{core}(T)$ if and only if α is either contained in a path between two vertices in V_∞ (allowing also the case $\alpha \in V_\infty$), or contained in an infinite path starting at a vertex in V_∞ , or contained in a bi-infinite path. Let $\text{core}(G) = \Gamma(\text{core}(T))$.

Remark 4.3. Note that $\text{core}(T)$ can be obtained by deleting all finite parts of T that contain no vertices of V_∞ and can be separated from T by deleting an edge.

Example 4.4. We have seen the Tutte decomposition of the infinite ladder in Figure 3. In this case the Tutte tree T is a bi-infinite path, so $T = \text{core}(T)$, and consequently $G = \text{core}(G)$.

Example 4.5. Another 2-connected example with infinitely many ends can be seen in Figure 4. The core is the 4-regular-tree-like part drawn with black, while the rest of the graph, consisting of finite components amalgamated at some places, is drawn with red.

Lemma 4.6. *Let G be an infinite, 2-connected planar graph such that $\text{core}(G) = G$. Assume also that G has a simply connected embedding into S^2 . Then (up to inverting all permutations) it has a unique combinatorial embedding with $\text{acc}(G) = 1$.*

Proof. Let T denote the Tutte tree of G , so $G = \Gamma(T)$. Since G is planar, all the Tutte components are planar. Also any amalgam of these obtained by performing some amalgamations as prescribed by T is planar.

Most importantly, as it is proved in [Tim23, Lemma 12], a combinatorial embedding π^G of G is uniquely determined by its restrictions to the Tutte components. We denote these restrictions $(\pi^\alpha)_{\alpha \in V(T)}$. For 3-connected components and 3-links we have 2 possible combinatorial embeddings, while for cycles we have 1.

We know that G has a simply connected embedding. Pick one and let π_+^G denote the corresponding combinatorial embedding. Accordingly, the restrictions are denoted π_+^α . When G_α is a 3-connected component or a 3-link, let π_-^α denote the other combinatorial embedding of G_α . If G_α is a cycle, we put $\pi_-^\alpha = \pi_+^\alpha$. Let π_-^G

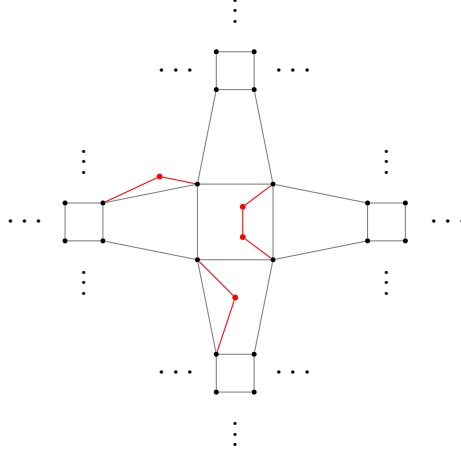


Figure 4: A 2-connected graph with infinitely many ends and a simply connected embedding

denote the combinatorial embedding of G with restrictions π_-^α for all $\alpha \in V(T)$. Since π_-^G is obtained from π_+^G by inverting all permutations, it also belongs to a simply connected embedding of G .

As we have said, to specify a combinatorial embedding π^G of G one has to make a choice between π_+^α and π_-^α for all $\alpha \in V(T)$. (For cycles this means no choice.) We claim that from these possible collections of choices only two give $\text{acc}(G, \pi^G) = 1$, namely π_+^G where we choose π_+^α for all α , and π_-^G , where we choose π_-^α for all α .

Suppose π^G is neither π_+^G nor π_-^G . We aim to find a cycle in (G, π^G) which has infinitely many points both on the inside and outside. Such a cycle shows $\text{acc}(G, \pi^G) \geq 2$. Now, we can find G_α and G_β , 3-connected components or 3-links in the Tutte decomposition such that $\pi^\alpha = \pi_+^\alpha$ and $\pi^\beta = \pi_-^\beta$. Moreover, we can pick α and β to be neighbors in T , or at worst at distance two in T with some γ in between, with G_γ a cycle. Assume first that α and β are neighbors. The proof will not be essentially different when there is γ in between.

Consider the case when both G_α and G_β are finite. Let e denote the virtual edge of G_α that we use to amalgamate it to G_β . As $\alpha \in \text{core}(T)$, we have another virtual edge of G_α , say f . Let T_e and T_f denote the components of $T \setminus \alpha$ corresponding to e and f . We know that $\Gamma(T_e)$ and $\Gamma(T_f)$ are both infinite by our assumption that $G = \text{core}(G)$ and consequently $T = \text{core}(T)$. Let $E = (\vec{e}, \vec{e}_2, \dots, \vec{e}_n)$ and $F = (\overleftarrow{e}, \vec{f}_2, \dots, \vec{f}_k)$ denote the two faces of G_α adjacent to e . Set $v_1 = s(\vec{e})$ and $v_2 = t(\vec{e})$.

The only edge that E and F share is e , hence

$$C := E + F = (\vec{e}_2, \dots, \vec{e}_n, \vec{f}_2, \dots, \vec{f}_k)$$

is a cycle in G_α . The cycle C surrounds e , therefore when G_β is amalgamated to G_α along e , the whole of $G_\beta \setminus e$ will be inside C .

Similar to G_α , the Tutte component G_β also has at least one other virtual edge f' apart from e . Clearly $f \neq f'$, as the only edge shared by G_α and G_β is e . What happens to the faces E and F when we amalgamate along e ? Because G_β is either 3-connected or a 3-link we can find three edge-disjoint paths in G_β between the




Figure 5: The faces E and F in G_α, π_+^α

endpoints of e . At least one of these avoids both e and f' , let us denote it by $P = (\vec{p}_1, \dots, \vec{p}_l)$. With slight abuse of notation we use $v_1 = s(\vec{p}_1)$ and $v_2 = t(\vec{p}_l)$ to denote the endpoints of e that are both in G_β and G_α . Note that P might contain other virtual edges of β , different from e and f' , but that will not cause problems for us.

In $G_\alpha + G_\beta$ the edge e is deleted, so the G_α -faces E and F merge into the cycle $C = E + F$, but the path P splits it into two cycles E' and F' in $G_\alpha + G_\beta$. (Recall that $P \subseteq G_\beta$ is inside C .) To be precise let $E' = (\vec{p}_1, \dots, \vec{p}_l, \vec{e}_2, \dots, \vec{e}_n)$ and $F' = (\overleftarrow{p}_l, \dots, \overleftarrow{p}_1, \vec{f}_2, \dots, \vec{f}_k)$.

We know that f' is inside C and not a part of P , so it is inside either E' or F' , depending on the choice of combinatorial embedding π^β . Assume that when we amalgamate according to π_+^β the β -edge f' falls inside the E' cycle, and when we amalgamate according to π_-^β it falls inside F' .




Figure 6: The amalgam $(G_\alpha + G_\beta, \pi_+^\alpha, \pi_+^\beta)$

When we amalgamate according to π_+^G , in $G_\alpha + G_\beta$ the virtual edges f and f' must not be separated by a cycle, because that would indicate $\text{acc}(G, \pi_+^G) \geq 2$. This means that f also has to be along the face E in G_α , otherwise E' would separate them in $G_\alpha + G_\beta$. See Figure 6.

We claim that f cannot be part of the face F in G_α . Indeed, in a 3-connected planar graph one cannot find two edges that have two common faces. This also holds for the 3-link. If f was incident to F , then e and f would be such edges.

We use the fact that f is not incident to F in G_α to establish $\text{acc}(G, \pi^G) \geq 2$ when we amalgamate according to π_+^α and π_-^β . In this case we have f' surrounded by the cyclic walk F' , see Figure 7. Therefore it is separated from f in $G_\alpha + G_\beta$, since f is along E' , outside F' . Irrespective of the other choices of π^G the two separated virtual edges already indicate $\text{acc}(G, \pi^G) \geq 2$. This finishes the proof in the finite case.




Figure 7: The amalgam $(G_\alpha + G_\beta, \pi_+^\alpha, \pi_-^\beta)$

If at least one of G_α and G_β is infinite, we can assume without loss of generality that G_β is infinite. In G_α , using the 3-connectedness (without assuming finiteness) we find two disjoint directed paths \vec{P}_1 and \vec{P}_2 evading e from v_2 to v_1 , and form the two cycles $E = (\vec{e}, \vec{P}_1)$ and $F = (\overleftarrow{e}, \overleftarrow{P}_2)$. We proceed the same way as in the finite case. We either find a virtual edge f of G_α , or G_α is infinite as well. In this case it has infinitely many points, almost all of which have to fall inside $C = (\vec{P}_1, \overleftarrow{P}_2)$. The argument leads to F' separating infinitely many points of G_β from infinitely many points of G_α (or f).

Finally if α and β are at distance two with γ in between, we proceed almost identically. We consider G_α and $G_\beta + G_\gamma$, and since G_γ is a cycle, $G_\beta + G_\gamma$ is a subdivision of G_β . Let e be the common virtual edge of G_α and G_γ , and e' the common virtual edge of G_β and G_γ . When choosing the path P in $G_\beta + G_\gamma$ to replace e in the G_α -cycles E and F we first find the path P_0 in G_β evading e' and f' , and then complement it with $G_\gamma \setminus \{e, e'\}$ to get a path connecting the endpoints of e . ■

4.5 General 2-connected planar graphs

Let G be a 2-connected planar graph with at least 2 ends that has a simply connected embedding. Let T denote its Tutte tree. Then $\text{core}(T)$ is nonempty, and $\text{core}(G)$ also has a simply connected embedding. By Lemma 4.6 there are exactly two combinatorial embeddings with $\text{acc}(\text{core}(G)) = 1$, using the fact that $\text{core}(G)$ is

always connected. As before, we denote the induced combinatorial embeddings on the Tutte components by π_+^α and π_-^α where $\alpha \in \text{core}(T)$.

In general a combinatorial embedding π^G gives $\text{acc}(G, \pi^G) = 1$ if and only if $\pi^\alpha = \pi_+^\alpha$ for all $\alpha \in \text{core}(T)$, or $\pi^\alpha = \pi_-^\alpha$ for all $\alpha \in \text{core}(T)$. For the vertices $\beta \in V(T) \setminus \text{core}(T)$, in which case G_β is always finite, the choice of π^β does not influence $\text{acc}(G, \pi^G) = 1$.

Example 4.7. Consider the 2-connected graph shown in Figure 4. Lemma 4.6 implies that the combinatorial embedding (with $\text{acc}(G) = 1$) is essentially unique for the core. On the other hand at every red component we have a choice between two possible ways of embedding, see Figure 8.




Figure 8: A 2-connected graph with infinitely many ends and a simply connected embedding

Remark 4.8. Notice that when G has 1 end we have $\text{core}(T) = \emptyset$ or it is a single point. And indeed, any choice of combinatorial embeddings for the Tutte components results in a combinatorial embedding of G with $\text{acc}(G) = 1$.

4.6 2-connected graphs with special vertices

When we later consider arbitrary planar graphs decomposed into 2-connected blocks, there will be a need to treat certain special vertices of G , informally speaking, as if they were ends of G . That is, as if they alone counted as infinitely many vertices, and by themselves gave rise to an accumulation point in one of the incident faces. The reason we did not include this technicality so far is that Lemma 4.6 - which is the key ingredient of our construction - is simpler to present without it, and the proof of Lemma 4.6 is heavy in notation as it is. Moreover, the need for these special vertices arises further down the line, and their presence does not change the proof, only complicates the exposition.

Nevertheless, to be precise, we now explain how the definition of $\text{core}(G)$, $\text{acc}(G)$, and the statement and proof of Lemma 4.6 need to be modified in order to incorporate the special vertices. See the upcoming subsection 4.7 for an example where the special vertices are crucial.

Assume that a set of vertices $W \subseteq V(G)$ is specified to be *special*. Earlier V_∞ stood for the set of Tutte components α where $|G_\alpha| = \infty$, now we redefine V_∞ so that every α with $V(G_\alpha) \cap W \neq \emptyset$ is also included. Recall that $\text{core}(G)$ is defined as $\Gamma(\text{core}(T))$, where $\text{core}(T)$ is the convex hull of $V_\infty \cup \text{Ends}(T)$ in the Tutte tree T . This definition remains the same, but with the updated V_∞ . Note that we automatically have $W \subseteq \text{core}(G)$.

The definition of $\text{acc}(G, \pi)$ changes as well. When considering a finite subgraph $H \subset G$, a_H^π denoted the number of combinatorial faces of H which surrounded infinitely many vertices of $G \setminus H$. Now we update that definition by counting combinatorial faces surrounding vertices from W as well. (Note that only vertices from $W \setminus V(H)$ can be surrounded by combinatorial faces of H .) We set $\text{acc}(G, \pi) = \sup a_{H_i}^\pi$ as before, with $H_1 \subset H_2 \subset \dots \subset G$ an exhaustion by finite graphs as before.

Earlier in Lemma 4.6 we assumed that G has a simply connected embedding, which means a unique accumulation point $x \in S^2$. Let D denote the connected component of x in $S^2 \setminus \iota(G)$. In the plane $D \setminus \{x\}$ is the infinite topological face of the graph (which might be empty). When taking the special vertices W into consideration, we need a stronger assumption in Lemma 4.6, namely that there exists a simply connected embedding such that vertices in W are embedded on the topological boundary of D . The implication remains that G has a unique (up to inverting all permutations) combinatorial embedding with $\text{acc}(G) = 1$.

Lemma 4.9. *Let G be an infinite, 2-connected planar graph with special vertices $W \subseteq V(G)$. Assume that $\text{core}(G) = G$, and that G has a simply connected embedding into S^2 which maps W onto ∂D . Then (up to inverting all permutations) G has a unique combinatorial embedding with $\text{acc}(G) = 1$.*

The proof of Lemma 4.9 is essentially the same as the proof of Lemma 4.6. Earlier, contradictions arose because we could establish $\text{acc}(G, \pi) \geq 2$ by finding cycles with either infinitely many vertices, or virtual edges representing infinitely many vertices both on the inside and the outside of the cycle. When the special vertices W are present, finding such a vertex inside or outside the cycle also suffices to show $\text{acc}(G, \pi) \geq 2$.

4.7 An example with special vertices

Consider the graph G in Figure 9. The vertex v is a cut vertex, and let G_a denote the infinite 2-connected component of G . We aim to illustrate that v needs to be treated as a special vertex in G_a in order to construct a combinatorial embedding of G with $\text{acc}(G) = 1$.

Let us first explore what happens if we do not treat v as special. Let C denote the 4-cycle in G_a containing v , and P the infinite path that v separates from G_a . Notice, that C is not in $\text{core}(G_a)$, in fact $\text{core}(G_a)$ is the infinite ladder from Figure 3. Lemma 4.6 says that $\text{core}(G_a)$ has an essentially unique combinatorial embedding with $\text{acc}(\text{core}(G_a)) = 1$ (which corresponds to the way it is drawn in these pictures, the only accumulation point being at infinity). In the Tutte decomposition of G_a , the cycle C is amalgamated to the cycle C' in $\text{core}(G_a)$ through a 3-link L . When we will later choose a random combinatorial embedding of G_a as described in subsection

4.11, the combinatorial embeddings of C and L are going to be chosen uniformly randomly, because they are not in the core. The cycle C has only 1, but L has 2 choices. And this choice of π^L determines whether C will be embedded inside C' , or outside it. In case C is mapped inside C' we will get $\text{acc}(G) = 2$ no matter how we reattach the infinite path P . So we will not manage to build the right combinatorial embedding this way.

Now let us explore what happens if we treat v as special in G_a . In this case C is also included in $\text{core}(G_a)$, so $G_a = \text{core}(G_a)$. We see that the embedding places v on the boundary of the infinite face, so Lemma 4.9 applies. It provides us with the essentially unique way of combinatorially embedding G_a such that we can then reattach P while keeping $\text{acc}(G) = 1$.




Figure 9: An example with special vertices

4.8 Towards combinatorial embeddings of arbitrary connected planar graphs

Having found (essentially unique) combinatorial embeddings with 1 combinatorial accumulation point for 2-connected graphs, we now turn our attention to general connected graphs.

Let G be any locally finite, connected planar graph, with block-cut tree \mathcal{T} . We write G_a , $a \in A$ for the 2-connected components and C for the cut vertices of G .

A combinatorial embedding π^G induces a combinatorial embedding π^{G_a} on all blocks. In order to be able to reconstruct π^G we also have to record how the π^{G_a} are to be glued at the cut vertices.

For a block G_a we denote by $\text{Cut}(G_a)$ the set of vertices $C \cap V(G_a)$. Similarly, we write $\text{Cut}_\infty(G_a)$ for those vertices in $\text{Cut}(G_a)$ that separate G_a from infinitely many points in G .

$$\text{Cut}_\infty(G_a) = \{v \in V(G_a) \mid |\mathcal{C}_{G \setminus E(G_a)}(v)| = \infty\}.$$

Our aim is to construct π^G randomly such that $\text{acc}(G, \pi^G) = 1$. We will do this by first choosing the π^{G_a} randomly for all blocks G_a , using Subsection 4.5 in the infinite case. Secondly, when making choices on how to glue the π^{G_a} at the cut vertices, we will make sure that in the end $\text{acc}(G, \pi^G) = 1$ holds.

4.9 Finite blocks with special vertices

Let G_a be a block, and assume first that G_a is finite. Let $v \in \text{Cut}_\infty(G_a)$, and let G_{b_1}, \dots, G_{b_i} denote the other blocks containing v . For any combinatorial embedding of G each G_{b_i} is surrounded by one of the combinatorial faces of G_a that are incident to v . If $\mathcal{C}_{G \setminus G_a}(G_{b_i})$ is infinite, this produces an accumulation point inside that combinatorial face.

So in order to achieve $\text{acc}(G, \pi^G) = 1$ we need all vertices in $\text{Cut}_\infty(G_a)$ to be adjacent to the same combinatorial face of G_a . Moreover, all neighboring blocks G_b with $\mathcal{C}_{G \setminus G_a}(G_{b_i})$ infinite have to fall in this distinguished combinatorial face.

Remark 4.10. Note that fixing π^{G_a} for each $a \in A$, and also fixing the choice of the combinatorial face of each (G_a, π^{G_a}) where each neighboring G_b should be embedded still does not determine the global combinatorial embedding π^G . When there are several blocks joined by a single cut-vertex we still need to choose a cyclic order of the blocks. This choice however does not impact $\text{acc}(G, \pi^G) = 1$.

We choose a random combinatorial embedding π^{G_a} such that π^{G_a} places all of $\text{Cut}_\infty(G_a)$ on the same distinguished face as follows. The simply connected embedding of G (that we assumed it has) induces such a combinatorial embedding on G_a . We wire $\text{Cut}_\infty(G_a)$ together. That is, add an auxiliary vertex v_a to $V(G_a)$ and connect v_a with all $v \in \text{Cut}_\infty(G_a)$. This graph G_a^w is still planar, and as it is finite it has finitely many combinatorial embeddings. We choose a (uniform) random combinatorial embedding of G_a^w , and choose the G_a -face that surrounds v_a to be distinguished.

4.10 Infinite blocks with special vertices

For an infinite block G_a the role of the distinguished combinatorial face will be played by the “infinite face” of G_a . We will not define the infinite face, but rather say that a point “sees infinity”. We will then make sure that $\text{acc}(G_a, \pi^{G_a}) = 1$ and all of $\text{Cut}_\infty(G_a)$ can see infinity.

Since G_a is infinite it is no longer true that any dart \vec{e} is part of a unique finite combinatorial face. Let $W(\vec{e})$ denote the unique bi-infinite walk containing \vec{e} with the property that at every vertex the next step uses the dart succeeding the opposite of the previous. When G_a was finite, $W(\vec{e})$ kept looping the finite combinatorial face containing \vec{e} infinitely many times in both directions. When G_a is infinite however, $W(\vec{e})$ might be a bi-infinite path of darts. We say a vertex $x \in V(G_a)$ *sees infinity* in the combinatorial embedding π^{G_a} if there is a dart \vec{e} with $t(\vec{e}) = x$ such that $W(\vec{e})$ is bi-infinite. Let \vec{f} denote the dart after \vec{e} in $W(\vec{e})$, i.e. $\vec{f} = \pi_x^{G_a}(\overleftarrow{e})$. (Recall that $\pi_x^{G_a}$ denotes the cyclic permutation on the outgoing darts at the vertex x in the combinatorial embedding π^{G_a} .)

We call the pair $(\overleftarrow{e}, \vec{f})$ the *infinite region* at x . We claim that this pair, if it exists, is unique for x . Indeed, as G_a is 2-connected, we can find a directed path \vec{P} from $t(\vec{f})$ to $s(\overleftarrow{e})$, and form the finite cycle $\vec{C} = (\vec{P}, \overleftarrow{e}, \vec{f})$. The walk $W(\vec{e})$ has infinitely many points inside \vec{C} . If there was another pair $\overleftarrow{e}_1, \vec{f}_1$ giving an infinite region at x , then $W(\overleftarrow{e}_1)$ would have infinitely many points outside \vec{C} contradicting $\text{acc}(G_a, \pi^{G_a}) = 1$.

Lemma 4.11. *Let G_a be an infinite block with a fixed combinatorial embedding π^{G_a} such that $\text{acc}(G_a, \pi^{G_a}) = 1$. A vertex $x \in V(G_a)$ sees infinity if and only if there is no cycle separating it from infinitely many vertices of G_a .*

Proof. First assume towards contradiction that x sees infinity and there is a cycle \vec{C} surrounding x with infinitely many points on the outside. The bi-infinite path $W(\vec{e})$ cannot cross \vec{C} , so there are infinitely many vertices inside \vec{C} as well, contradicting $\text{acc}(G_a, \pi^{G_a}) = 1$.

On the other hand assume that x does not see infinity, that is all darts leaving x are part of finite combinatorial faces. Form the finite subgraph H consisting of the vertices and edges of these combinatorial faces, together with the combinatorial embedding π^H restricted from π^{G_a} . This is a finite graph, and none of its combinatorial faces adjacent to x surround any vertices of G_a (when considered as oriented cycles in G_a). However, all vertices in $V(G_a) \setminus V(H)$ are surrounded by some combinatorial face of H . As G_a is infinite, there is some other combinatorial face of H , not adjacent to x that surrounds infinitely many points of G_a . This, as a cycle in G_a , separates x from infinitely many vertices of G_a . ■

4.11 Random combinatorial embeddings of infinite blocks

We now construct a random combinatorial embedding of an infinite block G_a using Lemma 4.9, by setting $W = \text{Cut}_\infty(G_a)$.

If G has at least 2 ends, then any infinite G_a has at least 2 ends (where vertices from $\text{Cut}_\infty(G_a)$ are also regarded as “ends”), so $\text{core}(G_a)$ is not empty. Also, if G has a simply connected embedding ι , then the restriction $\iota|_{G_a}$ is a simply connected embedding of G_a placing $\text{Cut}_\infty(G_a)$ on the topological boundary of the infinite topological face of G_a . Thus the block G_a satisfies the assumptions of Lemma 4.9.

In subsection 4.6 we emphasized that special vertices always belong to the core, so $\text{Cut}_\infty(G_a) \subseteq V(\text{core}(G_a))$. We claim that with respect to either of the two combinatorial embeddings $\pi_+^{\text{core}(G_a)}$ and $\pi_-^{\text{core}(G_a)}$, all vertices $x \in \text{Cut}_\infty(G_a)$ have to see infinity in $\text{core}(G_a)$ (where $\pi_+^{\text{core}(G_a)}$ and $\pi_-^{\text{core}(G_a)}$ stand for the unique combinatorial embedding and its inverse, as in Lemma 4.6). Indeed, by Lemma 4.11, if x did not see infinity, it would be separated from infinitely many points by a cycle, witnessing $\text{acc}(\text{core}(G)) \geq 2$, using that x is special.

Remark 4.12. If some G_α is 2-connected and has no vertices that see infinity then we see that it cannot contain Cut_∞ -vertices at all, since that would give rise to at least 2 accumulation points of the whole graph G .

Note that amalgamating the remaining Tutte components $\alpha \notin \text{core}(T_a)$ to $\text{core}(G_a)$ cannot create new cycles separating vertices of $\text{Cut}_\infty(G_a)$ from infinitely many points. So the choices of the π^α for $\alpha \notin \text{core}(T_a)$ are arbitrary, as before in subsection 4.5. These choices do not influence $\text{Cut}_\infty(G_a)$ seeing infinity.

We make independent uniform random choices between $\pi_+^{\text{core}(G_a)}$ and $\pi_-^{\text{core}(G_a)}$ for the core, and also to choose each π^α for $\alpha \notin \text{core}(T_a)$. The chosen $\pi^{\text{core}(G_a)}$ and π^α together form the random combinatorial embedding π^{G_a} .

4.12 Random embedding of G from its embedded blocks

Recall that at this point G is a deterministic, infinite, locally finite graph that has a locally finite embedding. We now construct a random combinatorial embedding π^G of G with $\text{acc}(G, \pi^G) = 1$.

For finite blocks G_a , as described in subsection 4.9, we choose a uniform random combinatorial embedding π^{G_a} with the additional property that the vertices in $\text{Cut}_\infty(G_a)$ are incident to a combinatorial face that we consider distinguished.

For infinite blocks G_a , we chose π^{G_a} as in subsection 4.11, with $\text{Cut}_\infty(G_a)$ considered as special points representing infinitely many vertices (as explained in subsection 4.6, and therefore being placed so that they see infinity).

We then construct the combinatorial embedding π^G by putting together the π^{G_a} . At cut vertices $v \in \text{Cut}_\infty(G_a)$ we define π_v^G so that all edges leading to all other blocks G_b intersecting G_a at v are placed inside the distinguished face of (G_a, π^{G_a}) . In the infinite case this means that in the cyclic order we place these edges between the two darts $(\overleftarrow{e}, \overrightarrow{f})$ that form the unique infinite region of G_a at v .

There is some freedom still in the choice of π_v when more than two blocks are glued at v , see the Remark in subsection 4.9. In that case we choose uniformly randomly among the finitely many choices satisfying the above.

We claim that $\text{acc}(G, \pi^G) = 1$. Indeed, any cycle C has to belong to a block G_a . The choice of the distinguished face (or the infinite region in case G_a is infinite) makes sure that C cannot have infinitely many points on both sides.

Remark 4.13. Note that so far in the present Section 4 the graph G was deterministic and unrooted (but assumed to have a locally finite embedding).

4.13 Proof of Theorem 5

The following proposition completes the proof of Theorem 5.

Proposition 4.14. *Let (G, o) be a URPG. Given a random instance of (G, o) , choose the combinatorial embedding π of G randomly as described in subsection 4.12. Then the random triple $(G, o; \pi)$ is a unimodular combinatorial embedding of (G, o) .*

Proof. Let $(G, o_1, o_2; \pi)$ denote the birooted (decorated) random graph induced from $(G, o; \pi)$ by taking a random step in the degree-biased version of $(G, o; \pi)$. We have to check that the distribution of $(G, o_1, o_2; \pi)$ is invariant with respect to swapping the o_1 and o_2 .

First the unimodularity of (G, o) says that (G, o_1, o_2) is invariant under swapping o_1 and o_2 . We have two operations on the random graph (G, o) : one is taking a random step on a degree-biased sample of (G, o) , the second is decorating (G, o) randomly. These two operations commute, their order is irrelevant. The bias and the random step never considers the combinatorial embedding and the random combinatorial embedding never considers any roots. ■

5 Embeddings of URPG's into \mathbb{H}^2 and \mathbb{R}^2

In this section we present the proof of Theorem 1. Much of it is adopted from [BT21], so for brevity we will refer to them as much as possible, and emphasize the contribution of the present work instead.

We tackle the case of more than one ends through Theorem 5 and Proposition 1.1. Besides these, the case not covered by [BT21], namely an invariant embedding of a non-amenable URPG into \mathbb{H}^2 will be dealt with using Theorem 7.

Proof of Proposition 1.1. By Theorem 5 we can start with a unimodular random combinatorial embedding π of G with $\text{acc}(G, \pi) \leq 1$. Based on this we can construct the triangulation (G^+, o^+) as in [BT21, Theorem 2.2], the resulting G^+ has finite expected degree, and invariant amenability or non-amenableity is preserved, and positive density of G in G^+ is guaranteed. Moreover the construction provides a unimodular combinatorial embedding π^+ with $\text{acc}(G^+, \pi^+) = 1$, where the restriction of π^+ to $G \leq G^+$ is π .

We claim G^+ has one end. If G^+ had more than one end, there would be a finite induced subgraph F with $G^+ \setminus F$ having at least two infinite components. These infinite components would also have to be surrounded by distinct combinatorial faces of F , because G^+ is triangulated (no edge can be added without violating planarity). In fact all vertices of $G^+ \setminus F$ that are surrounded by the same combinatorial face of F belong to the same connected component of $G^+ \setminus F$. Consequently F would have two distinct faces surrounding infinitely many points, implying $\text{acc}(G^+, \pi^+) \geq 2$. ■

Proof of Theorem 1. For amenable URPG's one in fact first finds an invariant embedding of finite intensity, and takes the Palm version to produce an unimodular embedding. By Proposition 1.1 one can find (G, o) as a positive density subgraph of an amenable planar triangulation. By [BT21, Theorem 4.2] such a triangulation can be embedded into \mathbb{R}^2 invariantly with vertices mapped onto some invariant point process of finite intensity.

For non-amenable URPG's the construction works in the other direction. We construct a unimodular embedding of positive intensity, and use Theorem 7 to deduce the existence of an invariant embedding. The triangulation (G^+, o^+) has a circle packing representation on \mathbb{H}^2 which is unique up to isometries [HS93, AHN16]. Therefore the embedding defined by the circle packing (by connecting centers of tangent circles by straight line segments) is a well defined rooted drawing of (G^+, o^+) (as defined in Subsection 2.4). As (G^+, o^+, S) is unimodular, restricting the rooted drawing to (S, o^+) on the event that $o^+ \in S$, we get a unimodular embedding of (G, o) . Moreover, by a Mass Transport argument the intensity is simply multiplied by $\mathbb{P}[o^+ \in S] > 0$ when restricting the embedding. The payment function is $f([\mathfrak{G}^+, \mathfrak{o}^+, \mathfrak{u}, \mathfrak{S}]) = \mathbb{1}_{\mathfrak{u} \in \mathfrak{S}} \cdot \lambda(\text{Vor}(\mathfrak{G}^+, \mathfrak{o}^+) \cap \text{Vor}(\mathfrak{S}, \mathfrak{u}))$. (The reader who uses gothic letters as rarely as the authors may want to note that \mathfrak{S} is the gothic version of the letter S .)

We now prove that this unimodular embedding has positive intensity. All edges of (G^+, o^+) are mapped to geodesics, and the embedding gives a triangulation of \mathbb{H}^2 . We consider the following factor allocation: take the barycentric subdivision of each triangle into 6 smaller triangles, and assign each to the one original vertex it is incident to. The number of triangles assigned to o^+ is $2 \deg_{G^+}(o^+)$, so the expected area of these pieces is at most $(\pi/2) \cdot 2\mathbb{E}[\deg_{G^+}(o^+)] < \infty$. Consequently

the intensity of the embedding of (G^+, o^+) is positive. As S has positive density in G^+ , the embedding restricted to S also has positive intensity. Therefore we get an invariant embedding by Theorem 7.

Finally, URPG's have no invariant or unimodular embeddings into the wrong space as proven in Propositions 6.1, 6.2 and 6.3 below. ■

6 Cases with no embedding because of isoperimetry

The final pieces of the proof of Theorem 1 are the following propositions.

Proposition 6.1. *A non-amenable URPG G has no invariant embedding into \mathbb{R}^2 .*

Proof as in Theorem 1.1 in [BT21]. Suppose that G had an isometry-invariant embedding into \mathbb{R}^2 without accumulation points. Then one could use the invariant random partitions of \mathbb{R}^2 to $2^n \times 2^n$ squares to define a unimodular finite exhaustion of G . Thus G is amenable, a contradiction. ■

Proposition 6.2. *An amenable URPG G has no invariant embedding into \mathbb{H}^2 .*

Proof. Suppose that G has an invariant embedding into \mathbb{H}^2 . Take the Voronoi partition of \mathbb{H}^2 corresponding to the embedded vertices. For the random configuration ω and $x, y \in \mathbb{H}^2$, define $f(x, y; \omega)$ to be 1 if x is in the Voronoi cell of an embedded vertex that is at distance ≤ 1 from y , and for $A, B \subset \mathbb{H}^2$ Borel define

$$\mu(A, B) = \mathbb{E} \left[\int_{x \in A} \int_{y \in B} f(x, y; \omega) dx dy \right].$$

The continuous version of the Mass Transport Principle (Theorem 5.2 in [BS01]) shows that the expected area of the Voronoi cell of the origin has to be finite.

Consider some unimodular finite exhaustion of G , i.e. some unimodular random $(G, o; (P_n)_{n \in \mathbb{N}})$ with the following properties almost surely. Each P_n is a partition of $V(G)$ into finite subsets, P_{n+1} is a coarsening of P_n , and for every $u, v \in V(G)$ there exists some n for which u and v are in the same part in P_n . The existence of such a unimodular finite exhaustion is equivalent to the amenability of (G, o) .

The P_n can be used to define an invariant random sequence of coarser and coarser partitions K_n of \mathbb{H}^2 , by taking the unions of the Voronoi cells of vertices that belong to the same part in P_n . The parts in K_n have finite area. Any two points of \mathbb{H}^2 end up in the same part of K_n for high enough n , which contradicts the non-amenability of \mathbb{H}^2 . (To see this using our setup, take an invariant random non-amenable tiling in \mathbb{H}^2 with a transitive underlying graph Γ . The above sequence of partitions would generate a unimodular finite exhaustion of Γ , a contradiction.) ■

Propositions 6.1 and 6.2 covered invariant embeddings; the next observations are their counterparts addressing unimodular embeddings.

Proposition 6.3. *Let (G, o) be a URPG.*

- (1) *If (G, o) is amenable, then it has no unimodular embedding of positive intensity into \mathbb{H}^2 .*

- (2) *There exists an amenable (G, o) with a unimodular embedding (of 0 intensity) into \mathbb{H}^2 .*
- (3) *If (G, o) is non-amenable then it has no unimodular embedding of positive intensity into \mathbb{R}^2 .*

Proof. If an embedding as in (1) or (3) existed, one could apply Theorem 7 to obtain an invariant embedding, contradicting Propositions 6.2 and 6.1, respectively.

Example 1.5 in [BT21] proves (2). That is, consider a bi-infinite geodesic in \mathbb{H}^2 through the origin 0, and take the points at integer distance from 0 on this geodesic to be the embedded vertices.

Part (3) is part of Theorem 1.2 in [BT21]. ■

Appendix

Proof of Theorem 7. We denote by ν_H the Haar measure on Γ , normalized such that $\nu_H(\text{Stab}_\Gamma(0)) = 1$. Any factor allocation scheme, in particular Ψ can be used to sample the Palm version of an invariant random embedded graph \mathfrak{G} with finite intensity as follows. First pick \mathfrak{G} randomly, then let $\mathfrak{o} \in V(\mathfrak{G})$ be the vertex whose cell contains the origin, i.e. $0 \in \Psi_{V(\mathfrak{G})}(\mathfrak{o})$. Now pick a Haar-random element $\varphi \in \Gamma$ such that $\varphi \cdot \mathfrak{o} = 0$ (the set of such φ is compact, so it has finite Haar measure), and let $\mathfrak{G}' = \varphi \cdot \mathfrak{G}$. Let \mathfrak{G}^* denote the Palm version of \mathfrak{G} . Then in distribution \mathfrak{G}' is exactly \mathfrak{G}^* biased by the area of the cell of the origin, i.e. $\lambda(\Psi_{V(\mathfrak{G}^*)}(0))$. To be precise, if μ' and μ^* denote the distributions of \mathfrak{G}' and \mathfrak{G}^* respectively, for any $A \subseteq \mathbb{G}_0(M)$ measurable we have

$$\mu'(A) = \int_A \lambda(\Psi_{V(\mathfrak{G})}(0)) d\mu^*(\mathfrak{G}) \Big/ \int_{\mathbb{G}_0(M)} \lambda(\Psi_{V(\mathfrak{G})}(0)) d\mu^*(\mathfrak{G}).$$

For a reference see [HP05], where a balanced allocation scheme (assigning equal sized cells for all centers) is constructed. In that case there is no bias happening, there is no difference between \mathfrak{G}' and \mathfrak{G}^* .

We will reverse these steps by starting from a unimodular embedding $[\mathfrak{G}_o, \mathfrak{o}]$, sample it biased by $\lambda(\Psi_{V(\mathfrak{G}_o)}(\mathfrak{o}))$, taking a representative of the class $[\mathfrak{G}_o, \mathfrak{o}]$ from $\mathbb{G}_0(M)$ with the root at 0, picking a uniform random point v from $\Psi_{V(\mathfrak{G}_o)}(0)$ (the cell has finite area almost surely), picking a random isometry $\varphi \in \Gamma$ with $\varphi \cdot 0 = v$, and shifting the embedding with φ^{-1} . We will show that the resulting embedding is Γ -invariant in distribution. This will complete the proof of Theorem 7, as the original unimodular embedding is the Palm version of the invariant one, because our construction is the reverse of the one in the previous paragraph.

Throughout this paper this is the only point where fixing an arbitrary embedding of each rooted drawing $[\mathfrak{G}, \mathfrak{o}]$ is necessary. As we have said before, $\mathbb{RD}(M)$ is in bijection with $\mathbb{G}_0(M)/\text{Stab}_\Gamma(0)$, so let $\theta : \mathbb{RD}(M) \rightarrow \mathbb{G}_0(M)$ be a measurable selecting function (as both $\mathbb{G}_0(M)$ and $\mathbb{RD}(M)$ are standard Borel spaces, such a function exists, see for example [Kec12, Exercise 18.3]). To shorten notation we write $\theta([\mathfrak{G}_o, \mathfrak{o}]) = \overline{[\mathfrak{G}, \mathfrak{o}]}$. Note that $\overline{[\mathfrak{G}, \mathfrak{o}]}$ is no longer a rooted object, it is an embedded graph that happens to have a vertex at $0 \in M$.

Let $a : \mathbb{G}(M) \times M \rightarrow [0, \infty)$ be defined as follows:

$$a(\mathbb{G}, y) = \begin{cases} 1 & \text{if } y \in V(\mathbb{G}) \text{ and } 0 \in \Psi_{V(\mathbb{G})}(y); \\ 0 & \text{otherwise.} \end{cases}$$

Note that if $0 \in M$ is not in the measure zero subset of M where $\Psi_{V(\mathbb{G})}$ fails to be a partition, then it is contained in exactly one cell, so we have $\sum_{y \in V(\mathbb{G})} a(\mathbb{G}, y) = 1$.

Let $\mu_0 \in \text{Prob}(\mathbb{RD}(M))$ denote the distribution of $(\mathfrak{G}_o, \mathfrak{o})$. We define the measure μ on $\mathbb{G}(M)$ as follows.

$$\mu(A) = \int_{[\mathbb{G}, \mathfrak{o}] \in \mathbb{RD}(M)} \int_{\varphi \in \Gamma} \mathbb{1}_A(\varphi \cdot \overline{[\mathbb{G}, \mathfrak{o}]}) \cdot a(\varphi \cdot \overline{[\mathbb{G}, \mathfrak{o}]}, \varphi \cdot 0) \, d\nu_H(\varphi) \, d\mu_0([\mathbb{G}, \mathfrak{o}]) \quad (2)$$

Remark 6.4. In the definition of μ it is not apparent that any biasing with the size of the cell of 0 took place. But in fact it did, otherwise there would be a $1/\lambda(\Psi_{V(\overline{[\mathbb{G}, \mathfrak{o}]})}(0))$ term as well in the integral. The role of the biasing is exactly to cancel this term out. Also note that we are not normalizing μ , as that would lengthen the expressions in the coming calculation. The measure of the whole space $\mu(\mathbb{G}(M))$ is $\mathbb{E}_{[\mathfrak{G}_o, \mathfrak{o}]}[\lambda(\Psi_{V(\mathfrak{G}_o)}(\mathfrak{o}))]$, this is ensured by having picked the right normalization of ν_H . By assumption, this is finite, so we get a probability measure after normalization.

We claim μ is Γ -invariant, which is equivalent to showing that for any $f : \mathbb{G}(M) \rightarrow [0, \infty)$ and fixed $\eta \in \Gamma$ we have

$$\int_{\mathbb{G}(M)} \eta \cdot f \, d\mu = \int_{\mathbb{G}(M)} f \, d\mu. \quad (3)$$

By the definition of μ we have

$$\int_{\mathbb{G}(M)} \eta \cdot f \, d\mu = \int_{\mathbb{RD}(M)} \int_{\Gamma} f(\eta^{-1} \varphi \cdot \overline{[\mathbb{G}, \mathfrak{o}]}) \cdot a(\varphi \cdot \overline{[\mathbb{G}, \mathfrak{o}]}, \varphi \cdot 0) \, d\nu_H(\varphi) \, d\mu_0([\mathbb{G}, \mathfrak{o}]) = \dots$$

For μ_0 -almost all $[\mathbb{G}, \mathfrak{o}] \in \mathbb{RD}(M)$ and $\varphi \in \Gamma$ we have $\sum_{x \in \eta^{-1} \varphi \cdot V(\overline{[\mathbb{G}, \mathfrak{o}]})} a(\eta^{-1} \varphi \cdot \overline{[\mathbb{G}, \mathfrak{o}]}, x) = 1$, so

$$\begin{aligned} \dots &= \int_{\mathbb{RD}(M)} \int_{\Gamma} \sum_{x \in \eta^{-1} \varphi \cdot V(\overline{[\mathbb{G}, \mathfrak{o}]})} a(\eta^{-1} \varphi \cdot \overline{[\mathbb{G}, \mathfrak{o}]}, x) \cdot f(\eta^{-1} \varphi \cdot \overline{[\mathbb{G}, \mathfrak{o}]}) \\ &\quad \cdot a(\varphi \cdot \overline{[\mathbb{G}, \mathfrak{o}]}, \varphi \cdot 0) \, d\nu_H(\varphi) \, d\mu_0([\mathbb{G}, \mathfrak{o}]) = \dots \end{aligned}$$

By substituting $x = \eta^{-1} \varphi \cdot y$ we have

$$\begin{aligned} \dots &= \int_{\mathbb{RD}(M)} \int_{\Gamma} \sum_{y \in V(\overline{[\mathbb{G}, \mathfrak{o}]})} a(\eta^{-1} \varphi \cdot \overline{[\mathbb{G}, \mathfrak{o}]}, \eta^{-1} \varphi \cdot y) \cdot f(\eta^{-1} \varphi \cdot \overline{[\mathbb{G}, \mathfrak{o}]}) \\ &\quad \cdot a(\varphi \cdot \overline{[\mathbb{G}, \mathfrak{o}]}, \varphi \cdot 0) \, d\nu_H(\varphi) \, d\mu_0([\mathbb{G}, \mathfrak{o}]) = \dots \end{aligned}$$

We will use the Mass Transport Principle with the function g that assigns to $[\mathbf{G}, \mathbf{u}, \mathbf{v}] \in \mathbb{BRD}(M)$ the value

$$g([\mathbf{G}, \mathbf{u}, \mathbf{v}]) = a\left(\eta^{-1}\varphi.\overline{[\mathbf{G}, \mathbf{u}]}, \eta^{-1}\varphi.\overline{\mathbf{v}}\right) \cdot f\left(\eta^{-1}\varphi.\overline{[\mathbf{G}, \mathbf{u}]}\right) \cdot a\left(\varphi.\overline{[\mathbf{G}, \mathbf{u}]}, \varphi.0\right),$$

where $\overline{\mathbf{v}} \in M$ is some point corresponding to \mathbf{v} in the embedded $\overline{[\mathbf{G}, \mathbf{u}]}$, that is to say $[\overline{[\mathbf{G}, \mathbf{u}]}, 0, \overline{\mathbf{v}}] = [\mathbf{G}, \mathbf{u}, \mathbf{v}]$. If there are multiple such points, any choice is sufficient, the value of g does not depend on the choice. We get

$$\begin{aligned} \dots &= \int_{\Gamma} \int_{\mathbb{RD}(M)} \sum_{\mathbf{y} \in V(\mathbf{G})} g([\mathbf{G}, \mathbf{o}, \mathbf{y}]) \, d\mu_0([\mathbf{G}, \mathbf{o}]) \, d\nu_H(\varphi) = \\ &\stackrel{\text{MTP}}{=} \int_{\Gamma} \int_{\mathbb{RD}(M)} \sum_{\mathbf{y} \in V(\mathbf{G})} g([\mathbf{G}, \mathbf{y}, \mathbf{o}]) \, d\mu_0([\mathbf{G}, \mathbf{o}]) \, d\nu_H(\varphi) = \\ &= \int_{\mathbb{RD}(M)} \int_{\Gamma} \sum_{\mathbf{y} \in V(\mathbf{G})} a\left(\eta^{-1}\varphi.\overline{[\mathbf{G}, \mathbf{y}]}, \eta^{-1}\varphi.\overline{\mathbf{o}}\right) \cdot f\left(\eta^{-1}\varphi.\overline{[\mathbf{G}, \mathbf{y}]}\right) \\ &\quad \cdot a\left(\varphi.\overline{[\mathbf{G}, \mathbf{y}]}, \varphi.0\right) \, d\nu_H(\varphi) \, d\mu_0([\mathbf{G}, \mathbf{o}]) = \dots \end{aligned} \quad (4)$$

As we have emphasized before, the sum only makes sense once a representative of $[\mathbf{G}, \mathbf{o}]$ is chosen. For us this will be $\overline{[\mathbf{G}, \mathbf{o}]}$. This way $\overline{[\mathbf{G}, \mathbf{y}]}$ (as \mathbf{y} runs through $V(\mathbf{G})$) is exactly $\overline{[\overline{[\mathbf{G}, \mathbf{o}]}, \mathbf{y}]}$ as \mathbf{y} runs through $V(\overline{[\mathbf{G}, \mathbf{o}]})$. For an $\mathbf{y} \in V(\overline{[\mathbf{G}, \mathbf{o}]})$ let $\varphi_{\mathbf{y}} \in \Gamma$ be such that $\varphi_{\mathbf{y}}.\overline{[\overline{[\mathbf{G}, \mathbf{o}]}, \mathbf{y}]} = \overline{[\mathbf{G}, \mathbf{o}]}$ and $\varphi_{\mathbf{y}}.0 = \mathbf{y}$. We note that as we sum over \mathbf{y} , the term $\overline{\mathbf{o}}$ in (4) also changes. As \mathbf{y} runs through $V(\overline{[\mathbf{G}, \mathbf{o}]})$ in our chosen representative $\varphi_{\mathbf{y}}^{-1}.0$ is a valid choice of $\overline{\mathbf{o}}$. For \mathbf{y} fixed let $\psi = \eta^{-1}\varphi\varphi_{\mathbf{y}}^{-1}$, and we calculate

$$\begin{aligned} \dots &= \int_{\mathbb{RD}(M)} \int_{\Gamma} \sum_{\mathbf{y} \in V(\overline{[\mathbf{G}, \mathbf{o}]})} a\left(\eta^{-1}\varphi.\overline{[\overline{[\mathbf{G}, \mathbf{o}]}, \mathbf{y}]}, \eta^{-1}\varphi\varphi_{\mathbf{y}}^{-1}.0\right) \cdot f\left(\eta^{-1}\varphi.\overline{[\overline{[\mathbf{G}, \mathbf{o}]}, \mathbf{y}]}\right) \\ &\quad \cdot a\left(\varphi.\overline{[\overline{[\mathbf{G}, \mathbf{o}]}, \mathbf{y}]}, \varphi.0\right) \, d\nu_H(\varphi) \, d\mu_0([\mathbf{G}, \mathbf{o}]) = \\ &= \int_{\mathbb{RD}(M)} \sum_{\mathbf{y} \in V(\overline{[\mathbf{G}, \mathbf{o}]})} \int_{\Gamma} a\left(\psi\varphi_{\mathbf{y}}.\overline{[\overline{[\mathbf{G}, \mathbf{o}]}, \mathbf{y}]}, \psi.0\right) \cdot f\left(\psi\varphi_{\mathbf{y}}.\overline{[\overline{[\mathbf{G}, \mathbf{o}]}, \mathbf{y}]}\right) \\ &\quad \cdot a\left(\eta\psi\varphi_{\mathbf{y}}.\overline{[\overline{[\mathbf{G}, \mathbf{o}]}, \mathbf{y}]}, \eta\psi\varphi_{\mathbf{y}}.0\right) \, d\nu_H(\varphi) \, d\mu_0([\mathbf{G}, \mathbf{o}]) = \dots \end{aligned}$$

The Haar measure ν_H is bi-invariant, therefore (for fixed \mathbf{y}) we can integrate

with respect to ψ instead of φ , so we get

$$\begin{aligned}
\dots &= \int_{\mathbb{RD}(M)} \sum_{y \in V(\overline{[\mathbf{G}, \mathbf{o}]})} \int_{\Gamma} a\left(\psi \cdot \overline{[\mathbf{G}, \mathbf{o}]}, \psi \cdot 0\right) \cdot f\left(\psi \cdot \overline{[\mathbf{G}, \mathbf{o}]}\right) \\
&\quad \cdot a\left(\eta\psi \cdot \overline{[\mathbf{G}, \mathbf{o}]}, \eta\psi \cdot y\right) d\nu_H(\psi) d\mu_0([\mathbf{G}, \mathbf{o}]) = \\
&= \int_{\mathbb{RD}(M)} \int_{\Gamma} a\left(\psi \cdot \overline{[\mathbf{G}, \mathbf{o}]}, \psi \cdot 0\right) \cdot f\left(\psi \cdot \overline{[\mathbf{G}, \mathbf{o}]}\right) \\
&\quad \cdot \sum_{y \in V(\overline{[\mathbf{G}, \mathbf{o}]})} a\left(\eta\psi \cdot \overline{[\mathbf{G}, \mathbf{o}]}, \eta\psi \cdot y\right) d\nu_H(\psi) d\mu_0([\mathbf{G}, \mathbf{o}]) = \dots
\end{aligned}$$

We have $\sum_{y \in V(\overline{[\mathbf{G}, \mathbf{o}]})} a\left(\eta\psi \cdot \overline{[\mathbf{G}, \mathbf{o}]}, \eta\psi \cdot y\right) = 1$ for almost all ψ and $[\mathbf{G}, \mathbf{o}]$, therefore

$$\dots = \int_{\mathbb{RD}(M)} \int_{\Gamma} a\left(\psi \cdot \overline{[\mathbf{G}, \mathbf{o}]}, \psi \cdot 0\right) \cdot f\left(\psi \cdot \overline{[\mathbf{G}, \mathbf{o}]}\right) d\nu_H(\psi) d\mu_0([\mathbf{G}, \mathbf{o}]) = \int_{\mathbb{G}(M)} f d\mu.$$

The last equality holds by the definition of μ in (2). Reading the calculation from start to finish yields equation (3), proving that μ is indeed Γ -invariant, which completes the proof. ■

Remark 6.5. We never used that the function a is $\{0, 1\}$ -valued, only that it sums to 1 almost always. For this reason one can extend Theorem 7 and Corollary 8 to *fractional* allocations as well, where instead of partitioning M we cover it by L^1 functions (indexed by the centers) that sum to $\mathbb{1}_M$. The size of a cell is the integral of the function. Considering such fractional allocations has a utility in homogeneous spaces of general locally compact groups. In this generality the Voronoi allocation and balanced allocation of [HP05] can only be defined as fractional allocations, because the measure of points equidistant from two given points can be positive.

Acknowledgement. The first author was partially supported by ERC Consolidator Grant 648017 and by Icelandic Research Fund Grant 185233-051. The second author was partially supported by NKFIH (National Research, Development and Innovation Office, Hungary) grant KKP-139502, “Groups and graph limits”.

References

- [AHNR16] Omer Angel, Tom Hutchcroft, Asaf Nachmias, and Gourab Ray. Unimodular hyperbolic triangulations: circle packing and random walk. *Inventiones mathematicae*, 206(1):229–268, 2016.
- [AHNR18] Omer Angel, Tom Hutchcroft, Asaf Nachmias, and Gourab Ray. Hyperbolic and parabolic unimodular random maps. *Geometric and Functional Analysis*, 28(4):879–942, 2018.
- [AL07] David Aldous and Russell Lyons. Processes on unimodular random networks. *Electronic Journal of Probability*, 12:1454–1508, 2007.

- [BHMK21] François Baccelli, Mir-Omid Haji-Mirsadeghi, and Ali Khezeli. Unimodular Hausdorff and Minkowski dimensions. *Electronic Journal of Probability*, 26:1–64, 2021.
- [BR03] C Paul Bonnington and R Bruce Richter. Graphs embedded in the plane with a bounded number of accumulation points. *Journal of Graph Theory*, 44(2):132–147, 2003.
- [BS01] Itai Benjamini and Oded Schramm. Percolation in the hyperbolic plane. *Journal of the American Mathematical Society*, 14(2):487–507, 2001.
- [BT21] Itai Benjamini and \acute{A} dám Timár. Invariant embeddings of unimodular random planar graphs. *Electronic Journal of Probability*, 26:1–18, 2021.
- [BWGT09] Lowell W Beineke, Robin J Wilson, Jonathan L Gross, and Thomas W Tucker. *Topics in topological graph theory*. Cambridge University Press Cambridge, 2009.
- [CSKM13] Sung Nok Chiu, Dietrich Stoyan, Wilfrid S Kendall, and Joseph Mecke. *Stochastic geometry and its applications*. John Wiley & Sons, 2013.
- [Die17] Reinhard Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer-Verlag Berlin Heidelberg, 5th edition, 2017.
- [Hal66] R Halin. Zur häufungspunktfreien darstellung abzählbarer graphen in der ebene. *Archiv der Mathematik*, 17(3):239–243, 1966.
- [HL05] Matthias Heveling and Günter Last. Characterization of palm measures via bijective point-shifts. *Annals of probability*, pages 1698–1715, 2005.
- [HP05] Alexander E Holroyd and Yuval Peres. Extra heads and invariant allocations. *Annals of probability*, 33(1):31–52, 2005.
- [HS93] Zheng-Xu He and Oded Schramm. Fixed points, koebe uniformization and circle packings. *Annals of Mathematics*, pages 369–406, 1993.
- [Kec12] Alexander Kechris. *Classical descriptive set theory*, volume 156. Springer Science & Business Media, 2012.
- [Las10] Günter Last. Stationary random measures on homogeneous spaces. *Journal of Theoretical Probability*, 23(2):478–497, 2010.
- [Lov12] László Lovász. *Large networks and graph limits*, volume 60. American Mathematical Soc., 2012.
- [LP17] Günter Last and Mathew Penrose. *Lectures on the Poisson process*, volume 7. Cambridge University Press, 2017.
- [LT09] Günter Last and Hermann Thorisson. Invariant transports of stationary random measures and mass-stationarity. *The Annals of Probability*, 37(2):790–813, 2009.

- [LZ13] Sergei K. Lando and Alexander K. Zvonkin. *Graphs on surfaces and their applications*, volume 141. Springer Science & Business Media, 2013.
- [Mol12] Jesper Moller. *Lectures on random Voronoi tessellations*, volume 87. Springer Science & Business Media, 2012.
- [MT01] Bojan Mohar and Carsten Thomassen. *Graphs on surfaces*. Johns Hopkins University Press, 2001.
- [Nev77] Jacques Neveu. Processus ponctuels. In *Ecole d'Été de Probabilités de Saint-Flour VI-1976*, pages 249–445. Springer, 1977.
- [RS04] Neil Robertson and Paul D Seymour. Graph minors. xx. wagner’s conjecture. *Journal of Combinatorial Theory, Series B*, 92(2):325–357, 2004.
- [RZ90] W Rother and M Zähle. Palm distributions in homogeneous spaces. *Mathematische Nachrichten*, 149(1):255–263, 1990.
- [Tim21] Ádám Timár. A nonamenable “factor” of a Euclidean space. *The Annals of Probability*, 49(3):1427 – 1449, 2021.
- [Tim23] Ádám Timár. Unimodular random one-ended planar graphs are sofic. *Combinatorics, Probability and Computing*, page 1–8, 2023.

Ádám Timár

Division of Mathematics, The Science Institute, University of Iceland
and

Alfréd Rényi Institute of Mathematics, Budapest
madaramit[at]gmail.com

László Márton Tóth

École Polytechnique Fédérale de Lausanne
laszломarton.toth[at]epfl.ch