Unimodular Hausdorff and Minkowski dimensions[*]

François Baccelli[†]   Mir-Omid Haji-Mirsadeghi[‡]   Ali Khezeli[§ ¶]

Abstract

This work introduces two new notions of dimension, namely the *unimodular Minkowski and Hausdorff dimensions*, which are inspired from the classical analogous notions. These dimensions are defined for *unimodular discrete spaces*, introduced in this work, which provide a common generalization to stationary point processes under their Palm version and unimodular random rooted graphs. The use of unimodularity in the definitions of dimension is novel. Also, a toolbox of results is presented for the analysis of these dimensions. In particular, analogues of Billingsley's lemma and Frostman's lemma are presented. These last lemmas are instrumental in deriving upper bounds on dimensions, whereas lower bounds are obtained from specific coverings. The notions of unimodular Hausdorff size, which is a discrete analogue of the Hausdorff measure, and unimodular dimension function are also introduced. This toolbox allows one to connect the unimodular dimensions to other notions such as volume growth rate, discrete dimension and scaling limits. It is also used to analyze the dimensions of a set of examples pertaining to point processes, branching processes, random graphs, random walks, and self-similar discrete random spaces. Further results of independent interest are also presented, like a version of the max-flow min-cut theorem for unimodular one-ended trees and a weak form of pointwise ergodic theorems for all unimodular discrete spaces.

[†]INRIA Paris and the University of Texas at Austin. E-mail: francois.baccelli@ens.fr
[‡]Sharif University of Technology. E-mail: mirsadeghi@sharif.ir
[§]Tarbiat Modares University. E-mail: khezeli@modares.ac.ir
[¶]Current address: INRIA Paris. E-mail: ali.khezeli@inria.fr

1 Introduction

Infinite discrete random structures are ubiquitous: random graphs, branching processes, point processes, graphs or zeros of discrete random walks, discrete or continuum percolation, to name a few. The large scale and macroscopic properties of such spaces have been thoroughly discussed in the literature. In particular, various notions of dimension have been proposed; e.g., the *mass dimension* and the *discrete (Hausdorff) dimension* defined by Barlow and Taylor [8] for subsets of $\mathbb{Z}^d$.

The main novelty of the present paper is the definition of new notions of dimension for a class of discrete structures that, heuristically, enjoy a form of statistical homogeneity. The mathematical framework proposed to handle such structures is that of *unimodular (random) discrete spaces*, where unimodularity is defined here by a version of the mass transport principle. This framework unifies unimodular random graphs and networks, stationary point processes (under their Palm version) and point-stationary point processes. It does not require more than a metric; for instance, no edges or no underlying Euclidean spaces are needed. The statistical homogeneity of such spaces has been used to define localized versions of global notions such as intensity. The novelty of the present paper is the use of this homogeneity to define the notions of *unimodular Minkowski and Hausdorff dimensions*, which are inspired by the analogous classical notions. The definitions are obtained naturally from the classical setting by replacing the infinite sums pertaining to infinite coverings by the expectation of certain random variables at the origin (which is a distinguished point), and also by considering large balls instead of small balls. These definitions are local but capture macroscopic (large scale) properties of the space.

The definitions are complemented by a toolbox for the analysis of unimodular dimensions. Several analogues of the important results known about the classical Hausdorff and Minkowski dimensions are established, like for instance the comparison of the unimodular Minkowski and Hausdorff dimensions as well as unimodular versions of Billingsley's lemma and Frostman's lemma. These lemmas allow one to connect the dimension to the (polynomial) volume growth rate of the space, which is also called *mass dimension* or *fractal dimension* in the literature. While many ideas in this toolbox are imported from the continuum setting, their adaptation is nontrivial and there is no automatic way to import results from the continuum to the discrete setting. For some results, the statements fundamentally differ from their continuum analog; e.g., the statement of Billingsley's lemma.

These notions of dimension are complemented by further definitions which can be used for a finer study of dimension. An analogue of the Hausdorff measure is defined, which is called the unimodular Hausdorff size here. This can be used to compare sets with the same dimension. The notion of unimodular dimension function is also defined for a finer quantification of the dimension. Such notions are new for discrete spaces to the best of the authors' knowledge. Another new notion introduced in the present paper is that of regularity for unimodular spaces, which is the equality of the unimodular Minkowski and Hausdorff dimensions. Similar notions of regularity exist in the continuum setting (see e.g., the definition of fractals in [14]) and for subsets of $\mathbb{Z}^d$ [9].

The paper also contains new mathematical results of independent interest. A weak version of Birkhoff's pointwise ergodic theorem is stated for all unimodular discrete spaces. A unimodular version of the max-flow min-cut theorem is also proved for unimodular one-ended trees, which is used in the proof of the unimodular Frostman lemma. Also, for unimodular one-ended trees, a relation between the volume growth rate and the height of the root is established as explained below.

The framework is used to derive concrete results on the dimension of several instances of unimodular random discrete metric spaces. This is done for the zeros and the graph of discrete random walks, sets defined by digit restriction, trees obtained from branching processes and drainage network models, etc. Some general results are obtained for all unimodular trees. For instance, a general relation is established between the unimodular dimensions of a unimodular one-ended tree and the tail of the distribution of the height of the root. The dimensions of some unimodular discrete random self-similar sets are also discussed. The latter are defined in this paper as unimodular discrete analogues of self similar sets such as the Koch snowflake, the Sierpinski triangle, etc.

This framework opens several further research directions. Firstly, it might be useful for the study of some discrete examples which are of interest in mathematical physics. Many examples in this domain enjoy some kind of homogeneity and give rise to unimodular spaces, directly or indirectly; e.g., percolation clusters and self-avoiding random walks. A few such examples are studied in detail in this work and in the preprint [6]. One might expect that in these examples, the values of unimodular dimensions match the conjectures or results pertaining to other notions of dimension that are applicable. Secondly, the definitions and many of the results are valid for exponential (or other) gauge functions. The proposed framework might hence have applications in group theory (or other areas), where most interesting examples have super-polynomial growth. A third important line of thoughts is the connections of unimodular dimensions to other notions of dimension. Some first connections are discussed in Subsection 8.1. The preprint [6] discusses ongoing research on these questions as well as further developments of these notions of dimensions.

1.1 Summary of the Main Definitions and Results

Recall that the ordinary Minkowski dimension of a compact metric space X is defined using the minimum number of balls of radius ϵ needed to cover X. Now, consider a (unimodular) discrete space D (it is useful to have in mind the example $D = \mathbb{Z}^k$ to see how the definitions work). It is convenient to consider coverings of D by balls of equal but large radius. Of course, if D is unbounded, then an infinite number of balls is needed to cover D. So one needs another measure to assess how many balls are used in a covering. Let $S \subseteq D$ be the set of centers of the balls in the covering. The idea pursued in this paper is that if D is unimodular, then the *intensity* of S is a measure of *the average number of points of S per points of D* (S should be *equivariant* for the intensity to be defined, as discussed later). This gives rise to the definition of the unimodular Minkowski dimension naturally.

The idea behind the definition of the unimodular Hausdorff dimension is similar. Recall that the α-dimensional *Hausdorff content* of a compact metric space X is defined by considering the infimum of $\sum_i R_i^\alpha$, where the R_i's are the radii of a sequence of balls that cover X. Also, it is convenient to enforce an upper bound on the radii. Now, consider a unimodular discrete space D and a covering of D by balls which may have different radii. Let $R(v)$ be the radius of the ball centered at v. It is convenient to consider a lower bound on the radii, say $R(\cdot) \geq 1$. Again, if D is unbounded, then $\sum_v R(v)^\alpha$ is always infinite. The idea is to leverage the unimodularity of D and to consider *the average of the values $R(\cdot)^\alpha$ per point* as a replacement of the sum. Under the unimodularity assumption, this can be defined by $\mathbb{E}[R(o)^\alpha]$, where o stands for the distinguished point of D (called the origin) and where, by convention, $R(o)$ is zero if there is no ball centered at o. This is used to define the *unimodular Hausdorff dimension* of D in a natural way.

The volume growth rate of the space is the polynomial growth rate of $\#N_r(o)$, where $N_r(o)$ represents the closed ball of radius r centered at the origin and $\#N_r(o)$

is the number of points in this ball. It is shown that the upper and lower volume growth rates of $\#N_r(o)$ (i.e., limsup and liminf of $\log(\#N_r(o))/\log r$ as $r \to \infty$) provide upper and lower bound for the unimodular Hausdorff dimension, respectively. This is a discrete analogue of Billingsley's lemma (see e.g., [14]). A discrete analogue of the *mass distribution principle* is also provided, which is useful to derive upper bounds on the unimodular Hausdorff dimension. In the Euclidean case (i.e., for point-stationary point processes equipped with the Euclidean metric), it is shown that the unimodular Minkowski dimension is bounded from above by the polynomial decay rate of $\mathbb{E}[1/\#N_n(o)]$. Weighted versions of these inequalities, where a weight is assigned to each point, are also presented. As a corollary, a weak form of Birkhoff's pointwise ergodic theorem is established for all unimodular discrete spaces. These results are very useful for calculating the unimodular dimensions in many examples. An important result is an analogue of Frostman's lemma. Roughly speaking, this lemma states that the mass distribution principle is sharp if the weights are chosen appropriately. This lemma is a powerful tool to study the unimodular Hausdorff dimension. In the Euclidean case, another proof of Frostman's lemma is provided using a version of the max-flow min-cut theorem for unimodular one-ended trees, which is of independent interest.

Depending on whether one defines the unimodular Minkowski dimension as the decay rate or the growth rate of the optimal intensity of the coverings by balls of radius r, one gets positive or negative dimensions. The present paper adopts the convention of positive dimensions for the definitions of both the unimodular Minkowski and Hausdorff dimensions, despite some mathematical arguments in favor of negative dimensions. Further discussion on the matter is provided in Subsection 8.3.

1.2 Organization of the Material

Section 2 defines unimodular discrete spaces and *equivariant processes*, which are needed throughout. Section 3 presents the definitions of the unimodular Minkowski and Hausdorff dimensions and the unimodular Hausdorff size. It also provides some basic properties of these unimodular dimensions as part of the toolbox for the analysis of unimodular dimensions. Various examples are discussed in Section 4. These examples are used throughout the paper. Section 5 is focused on the connections with volume growth rates and contains the statements and proofs of the unimodular Billingsey lemma and of the mass distribution principle. The unimodulat Frostman lemma is discussed in Section 7. Section 6 completes the analysis of the examples discussed in Section 4 and also discusses new examples for further illustration of the results. Section 8 discusses further topics on the matter. This includes a discussion of the connections to earlier notions of dimensions for discrete sets, in particular those proposed by Barlow and Taylor in [8, 9], as well as a discussion on negative dimensions. A collection of conjectures and open problems are also listed in this section.

Throughout the paper, some easier proofs as well as some extra details (e.g., measure-theoretic requirements) are skipped for the sake of brevity. These proofs and details are nevertheless available in the arXiv version [7] of the present paper. Precise indications on where to find them in these preprints are given in the text.

2 Unimodular Discrete Spaces

The main objective of this section is the definition of unimodular discrete spaces as a common generalization of unimodular graphs, Palm probabilities and point-stationary point processes. If the reader is familiar with unimodular random graphs, he or she can restrict attention to the case of unimodular graphs and jump to Subsection 2.5 at first reading.

2.1 Notation and Definitions

The following notation will be used throughout. The set of nonnegative real (resp. integer) numbers is denoted by $\mathbb{R}^{\geq 0}$ (resp. $\mathbb{Z}^{\geq 0}$). The minimum and maximum binary operators are denoted by $\wedge$ and $\vee$ respectively. The number of elements in a set A is denoted by $\#A$, which is a number in $[0, \infty]$. If $P(x)$ is a property about x, the indicator $1_{\{P(x)\}}$ is equal to 1 if $P(x)$ is true and 0 otherwise.

Discrete metric spaces (discussed in details in Subsection 2.2) are denoted by D, D', etc. Graphs are an important class of discrete metric spaces. So the symbols and notations are mostly borrowed from graph theory.

For $r > 0$, $N_r(v) := N_r(D, v)$ denotes the closed r-neighborhood of $v \in D$; i.e., the set of points of D with distance less than or equal to r from v. An exception is made for $r = 0$ (Subsection 3.3), where $N_0(v) := \emptyset$. The diameter of a subset A is denoted by $\mathrm{diam}(A)$. For a function $f : [1, \infty) \to \mathbb{R}^{\geq 0}$, the *polynomial growth rates* and *polynomial decay rates* are defined by the following formulas:

$$
\begin{aligned}
\underline{\mathrm{growth}}\,(f) &:= -\overline{\mathrm{decay}}\,(f) := \liminf_{r \to \infty} \log f(r)/\log r, \\
\overline{\mathrm{growth}}\,(f) &:= -\underline{\mathrm{decay}}\,(f) := \limsup_{r \to \infty} \log f(r)/\log r, \\
\mathrm{growth}\,(f) &:= -\mathrm{decay}\,(f) := \lim_{r \to \infty} \log f(r)/\log r.
\end{aligned}
$$

Definition 2.1. *Let μ be a probability measure on a measurable space X and $w : X \to \mathbb{R}^{\geq 0}$ be a measurable function. Assume $0 < c := \int_X w(x)d\mu(x) < \infty$. By **biasing** μ **by** w we mean the probability measure ν on X defined by*

$$
\nu(A) := \frac{1}{c} \int_A w(x)d\mu(x).
$$

2.2 The Space of Pointed Discrete Spaces

Throughout the paper, the metric on any metric space is denoted by d, except when explicitly mentioned. In this paper, it is always assumed that the discrete metric spaces under study are **boundedly finite**; i.e., every set included in a ball of finite radius in D is finite (note that this is stronger than being locally-finite). This implies that the metric space is indeed discrete; i.e., every point is isolated. The term **discrete space** will always refer to boundedly finite discrete metric space. A **pointed set** (or a *rooted set*) is a pair (D, o), where D is a set and o a distinguished point of D called the **origin** (or the *root*) of D. Similarly, a **doubly-pointed** set is a triple (D, o_1, o_2), where o_1 and o_2 are two distinguished points of D.

Let Ξ be a complete separable metric space called the **mark space**. A **marked pointed discrete space** is a tuple $(D, o; m)$, where (D, o) is a pointed discrete space and m is a function $m : D \times D \to \Xi$. The mark of a single point x may also be defined by $m(x) := m(x, x)$, where the same symbol m is used for simplicity. An **isomorphism** (or *rooted isomorphism*) between two such spaces $(D, o; m)$ and $(D', o'; m')$ is an isometry $\rho : D \to D'$ such that $\rho(o) = o'$ and $m'(\rho(u), \rho(v)) = m(u, v)$ for all $u, v \in D$. An isomorphism between doubly-pointed marked discrete spaces is defined similarly. An isomorphism from a space to itself is called an **automorphism**.

Most of the examples of discrete spaces in this work are graphs or discrete subsets of the Euclidean space. More precisely, connected and locally-finite simple graphs equipped with the graph-distance metric [2] are instances of discrete spaces. Similarly, *networks*; i.e., graphs equipped with marks on the edges [2], are instances of marked discrete spaces.

Let $\mathcal{D}_*$ (resp. $\mathcal{D}_{**}$) be the set of equivalence classes of pointed (resp. doubly-pointed) discrete spaces under isomorphism. Let $\mathcal{D}'_*$ and $\mathcal{D}'_{**}$ be defined similarly for marked

discrete spaces with mark space Ξ (which is usually given). The equivalence class containing (D, o), $(D, o; m)$ etc., is denoted by brackets $[D, o]$, $[D, o; m]$, etc.

Every element of $\mathcal{D}_*$ can be regarded as a *boundedly-compact* measured metric space (where the measure is the counting measure). Therefore, the generalization of the *Gromov-Hausdorff-Prokhorov* metric in [37] defines a metric on $\mathcal{D}_*$. By using the results of [37], one can show that $\mathcal{D}_*$ is a Borel subset of some complete separable metric space. The proof of this result is skipped for brevity. It can be found in [7]. Similarly, one can show that $\mathcal{D}_{**}, \mathcal{D}'_*$ and $\mathcal{D}'_{**}$ are Borel subsets of some complete separable metric spaces (see also [36]). This enables one to define *random pointed discrete spaces*, etc., which are discussed in the next subsection.

2.3 Random Pointed Discrete Spaces

Definition 2.2. *A **random pointed discrete space** is a random element in $\mathcal{D}_*$ and is denoted by bold symbols $[\boldsymbol{D}, \boldsymbol{o}]$. Here, $\boldsymbol{D}$ and $\boldsymbol{o}$ represent the discrete space and the origin respectively.*

In this paper, the probability space is not referred to explicitly[1]. The main reason is that the notions of dimension, to be defined, depend only on the distribution of the random object. Also, extra randomness will be considered frequently and it is easier to forget about the probability space. By an abuse of notation, the symbols $\mathbb{P}$ and $\mathbb{E}$ are used for all random objects, possibly living in different spaces. They refer to probability and expectation with respect to the random object under consideration.

Note that the whole symbol $[\boldsymbol{D}, \boldsymbol{o}]$ represents one random object, which is a random equivalence class of pointed discrete spaces. Therefore, any formula using $\boldsymbol{D}$ and $\boldsymbol{o}$ should be well defined for equivalence classes; i.e., it should be invariant under pointed isomorphisms.

The following convention is helpful throughout.

Convention 2.3. In this paper, bold symbols are usually used in the random case or when extra randomness is used. For example, $[D, o]$ refers to a deterministic element of $\mathcal{D}_*$ and $[\boldsymbol{D}, \boldsymbol{o}]$ refers to a random pointed discrete space.

Note that the distribution of a random pointed network $[\boldsymbol{D}, \boldsymbol{o}]$ is a probability measure on $\mathcal{D}_*$ defined by $\mu(A) := \mathbb{P}\left[[\boldsymbol{D}, \boldsymbol{o}] \in A\right]$ for events $A \subseteq \mathcal{D}_*$.

Definition 2.4. *A **random pointed marked discrete space** is a random element in $\mathcal{D}'_*$ and is denoted by bold symbols $[\boldsymbol{D}, \boldsymbol{o}; \boldsymbol{m}]$. Here, $\boldsymbol{D}$, $\boldsymbol{o}$ and $\boldsymbol{m}$ represent the discrete space, the origin and the mark function respectively.*

Most examples in this work are either random rooted graphs (or networks) [2] or point processes (i.e., random discrete subset of $\mathbb{R}^k$) and marked point processes that contain 0, where 0 is considered as the origin. Other examples are also studied by considering different metrics on such objects.

2.4 Unimodular Discrete Spaces

Once the notion of random pointed discrete space is defined, the definition of unimodularity is a straight variant of [2]. In what follows, the notation is similarly to [5]. Here, the symbol $g[D, o, v]$ is used as a short form of $g([D, o, v])$. Similarly, brackets $[\cdot]$ are used as a short form of $([\cdot])$.

---

[1]Indeed, one may regard $\mathcal{D}_*$, equipped with a probability measure, as the canonical probability space. The last paragraph of Subsection 2.2 ensures that this is a standard probability space, and hence, the classical tools of probability theory are available. One may also define a random pointed discrete space as a measurable function from some standard probability space to $\mathcal{D}_*$.

Definition 2.5. *A **unimodular discrete space** is a random pointed discrete space, namely* $[\boldsymbol{D}, \boldsymbol{o}]$*, such that for all measurable functions* $g : \mathcal{D}_{**} \to \mathbb{R}^{\geq 0}$*,*

$$\mathbb{E}\left[\sum_{v \in \boldsymbol{D}} g[\boldsymbol{D}, \boldsymbol{o}, v]\right] = \mathbb{E}\left[\sum_{v \in \boldsymbol{D}} g[\boldsymbol{D}, v, \boldsymbol{o}]\right]. \tag{2.1}$$

*Similarly, a **unimodular marked discrete space** is a random pointed marked discrete space* $[\boldsymbol{D}, \boldsymbol{o}; \boldsymbol{m}]$ *such that for all measurable functions* $g : \mathcal{D}'_{**} \to \mathbb{R}^{\geq 0}$*,*

$$\mathbb{E}\left[\sum_{v \in \boldsymbol{D}} g[\boldsymbol{D}, \boldsymbol{o}, v; \boldsymbol{m}]\right] = \mathbb{E}\left[\sum_{v \in \boldsymbol{D}} g[\boldsymbol{D}, v, \boldsymbol{o}; \boldsymbol{m}]\right]. \tag{2.2}$$

Note that the expectations may be finite or infinite.

When there is no ambiguity, the term $g[D, o, v]$ is also denoted by $g_D(o, v)$ or simply $g(o, v)$. The sum in the left (respectively right) side of (2.1) is called the **outgoing mass from o** (respectively **incoming mass into o**) and is denoted by $g^+(o)$ (respectively $g^-(o)$). The same notation can be used for the terms in (2.2). So (2.1) and (2.2) can be summarized by

$$\mathbb{E}\left[g^+(\boldsymbol{o})\right] = \mathbb{E}\left[g^-(\boldsymbol{o})\right].$$

These equations are called the **mass transport principle** in the literature. The reader will find further discussion on the mass transport principle and unimodularity in [2] and the examples therein.

As a basic example, every finite metric space D, equipped with a random root $\boldsymbol{o} \in D$ chosen uniformly, is unimodular. Also, the lattices of the Euclidean space rooted at 0; e.g., $[\mathbb{Z}^k, 0]$ and $[\delta\mathbb{Z}^k, 0]$, are unimodular. In addition, unimodularity is preserved under weak convergence, as observed in [13] for unimodular graphs.

The following two examples show that unimodular discrete spaces unify unimodular graphs and point-stationary point processes. Most of the examples in this work are of these types.

Example 2.6 (Unimodular Random Graphs). In the case of random rooted graphs and networks, the concept of unimodularity in Definition 2.5 coincides with that of [2] (see also Remark A.5 of [7] regarding the topologies). Therefore, unimodular random graphs and networks are special cases of unimodular (marked) discrete spaces.

Example 2.7 (Point-Stationary Point Processes). **Point-stationarity** is defined for point processes Φ in $\mathbb{R}^k$ such that $0 \in \Phi$ a.s. (see e.g., [39]). This definition is equivalent to (2.1), except that g is required to be invariant under translations only (and not under all isometries). This implies that $[\Phi, 0]$ is unimodular. In addition, by considering the mark $\boldsymbol{m}(x, y) := y - x$ on pairs of points of Φ, point-stationarity of Φ will be equivalent to the unimodularity of $[\Phi, 0; \boldsymbol{m}]$ (see also Remark A.5 of [7] regarding the topologies). Note also that Φ can be recovered from $[\Phi, 0; \boldsymbol{m}]$.
For example, if Φ is a stationary point process in $\mathbb{R}^k$ (i.e., its distribution is invariant under all translations), with finite intensity (i.e., a finite expected number of points in the unit cube), then the *Palm version* of Φ is a point-stationary point process, where the latter is heuristically obtained by conditioning Φ to contain the origin (see e.g., Section 13 of [18] for the precise definition). Also, if $(X_n)_{n \in \mathbb{Z}}$ is a stochastic process in $\mathbb{R}^k$ with stationary increments such that $X_0 = 0$, then the graph of this random walk is a point-stationary point process in $\mathbb{R}^{k+1}$. The image of this random walk is also point-stationary provided that it is discrete (i.e., the random walk is transient) and $X_i \neq X_j$ a.s. for every $i \neq j$. See Subsection 4.3.

2.5 Equivariant Process on a Unimodular Discrete Space

In many cases in this paper, an unmarked unimodular discrete space $[\boldsymbol{D}, \boldsymbol{o}]$ is given and various ways of assigning marks to $\boldsymbol{D}$ are considered. Intuitively, an *equivariant process* on $\boldsymbol{D}$ is an assignment of (random) marks to $\boldsymbol{D}$ such that the new marked space is unimodular. Formally, it is

a unimodular marked discrete space $[\boldsymbol{D}', \boldsymbol{o}'; \boldsymbol{m}]$ such that the space $[\boldsymbol{D}', \boldsymbol{o}']$, obtained by forgetting the marks, has the same distribution as $[\boldsymbol{D}, \boldsymbol{o}]$.

In this paper, it is more convenient to work with a disintegrated form of this heuristic, defined below. It can be proved that the two notions are equivalent, but the proof is skipped for brevity (this claim is similar to *invariant disintegration* for group actions). The easy part of the claim is Lemma 2.12 below. For the other direction, see Proposition B.1 of [7].

In the following, the mark space Ξ is fixed as in Subsection 2.2.

Definition 2.8. *Let D be a deterministic discrete space which is boundedly-finite. A* **marking** *of D is a function from $D \times D$ to Ξ; i.e., an element of $\Xi^{D \times D}$. A* **random marking** *of D is a random element of $\Xi^{D \times D}$.*

Definition 2.9. *An* **equivariant process** *$\boldsymbol{Z}$ with values in Ξ is a map that assigns to every deterministic discrete space D a random marking $\boldsymbol{Z}_D$ of D satisfying the following properties:*

(i) $\boldsymbol{Z}$ is compatible with isometries in the sense that for every isometry $\rho : D_1 \to D_2$, the random marking $\boldsymbol{Z}_{D_1} \circ \rho^{-1}$ of D_2 has the same distribution as $\boldsymbol{Z}_{D_2}$.

(ii) For every measurable subset $A \subseteq D'_*$, the following function on $\mathcal{D}_*$ is measurable:

$$[D, o] \mapsto \mathbb{P}\left[[D, o; \boldsymbol{Z}_D] \in A\right].$$

In addition, given a unimodular discrete space $[\boldsymbol{D}, \boldsymbol{o}]$, such a map is also called an **equivariant process on $\boldsymbol{D}$**. *In this case, one can also let $\boldsymbol{Z}_{(\cdot)}$ be undefined for a class of discrete spaces, as long as it is defined for almost all realizations of $\boldsymbol{D}$. It is important that extra randomness be allowed here.*

Convention 2.10. *If D is clear from the context, $\boldsymbol{Z}_D(\cdot)$ is also denoted by $\boldsymbol{Z}(\cdot)$ for simplicity.*

Note that in the above definition, D is deterministic and is not an equivalence class of discrete spaces. However, for an equivariant process on $[\boldsymbol{D}, \boldsymbol{o}]$, one can define $[\boldsymbol{D}, \boldsymbol{o}; \boldsymbol{Z}_{\boldsymbol{D}}]$ as a random pointed marked discrete space with distribution $\mathcal{Q}$ (on $\mathcal{D}'_*$), defined by

$$\mathcal{Q}(A) := \int \int 1_A[D, o; m] \mathrm{d}\mathcal{P}_D(m) d\mu([D, o]), \quad \forall A \subseteq \mathcal{D}'_*, \tag{2.3}$$

where $\mathcal{P}_D$ is the distribution of $\boldsymbol{Z}_D$ (for every D) and μ is the distribution of $[\boldsymbol{D}, \boldsymbol{o}]$ (note that only the distribution of $\boldsymbol{Z}_D$ is important here and it doesn't matter which probability space is used for $\boldsymbol{Z}_D$). It can be seen that $\mathcal{Q}(A)$ is indeed well defined and is a probability measure on $\mathcal{D}'_*$. As mentioned before, the probabilities and expectations to be used for $\boldsymbol{Z}_D$ and $[\boldsymbol{D}, \boldsymbol{o}; \boldsymbol{Z}_{\boldsymbol{D}}]$ will be denoted by the same symbols $\mathbb{P}$ and $\mathbb{E}$; e.g., $\mathbb{P}\left[\boldsymbol{Z}_D \in B\right]$, $\mathbb{P}\left[[D, o; \boldsymbol{Z}_D] \in A\right]$ and $\mathbb{E}\left[f[\boldsymbol{D}, \boldsymbol{o}; \boldsymbol{Z}_{\boldsymbol{D}}]\right]$. The symbols $\mathcal{P}_D$ and $\mathcal{Q}$ will not be used in what follows.

The following basic examples help to illustrate the definition.

Example 2.11. By choosing the marks of points (or pairs of points) in an i.i.d. manner, one obtains an equivariant process. Also, the following periodic marking of $\mathbb{Z}$ is an

equivariant process on $\mathbb{Z}$: Choose $\boldsymbol{U} \in \{0, 1, \ldots, n-1\}$ uniformly at random and let $\boldsymbol{Z}_{\mathbb{Z}}(x) := 1$ if $x \in n\mathbb{Z} + \boldsymbol{U}$ and $\boldsymbol{Z}_{\mathbb{Z}}(x) := 0$ otherwise. Moreover, given a measurable function $z : \mathcal{D}_{**} \to \Xi$, one can define $\boldsymbol{Z}_D(u, v) := z[D, u, v]$, which will be called a *deterministic process* here.

Lemma 2.12. *Let $[\boldsymbol{D}, \boldsymbol{o}]$ be a unimodular discrete space. If $\boldsymbol{Z}$ is an equivariant process on $\boldsymbol{D}$, then $[\boldsymbol{D}, \boldsymbol{o}; \boldsymbol{Z}_D]$ is also unimodular.*

The proof is straightforward and skipped for brevity. The converse of this claim also holds (see [7]). It is important here to assume that the distribution of $\boldsymbol{Z}_D$ does not depend on the origin (as in Definition 2.9).

Remark 2.13. One can easily extend the definition of equivariant processes to allow the base space to be marked. Therefore, for point-stationary point processes, one can replace condition (i) by invariance under translations only (see Example 2.7). In particular, every stationary stochastic process on $\mathbb{Z}^k$ defines an equivariant process on $\mathbb{Z}^k$.

Definition 2.14. *An **equivariant subset** S is the set of points with mark 1 in some $\{0, 1\}$-valued equivariant process. In addition, if $[\boldsymbol{D}, \boldsymbol{o}]$ is a unimodular discrete space, then the **intensity** of S in $\boldsymbol{D}$ is defined by $\rho_{\boldsymbol{D}}(\boldsymbol{S}) := \mathbb{P}[\boldsymbol{o} \in \boldsymbol{S}_{\boldsymbol{D}}]$.*

For example, $\boldsymbol{S}_D := \{v \in D : \#N_1(v) = 4\}$ defines an equivariant subset. Also, let $D = \mathbb{Z}$ and $\boldsymbol{S}_D$ be the set of even numbers with probability p and the set of odd numbers with probability $1 - p$. Then, $\boldsymbol{S}$ is an equivariant subset of $\mathbb{Z}$ if and only if $p = \frac{1}{2}$ (notice Condition (i) of Definition 2.9).

Lemma 2.15. *Let $[\boldsymbol{D}, \boldsymbol{o}]$ be a unimodular discrete space and S an equivariant subset. Then $\boldsymbol{S_D} \neq \emptyset$ with positive probability if and only if it has positive intensity. Equivalently, $\boldsymbol{S_D} = \boldsymbol{D}$ a.s. if and only if $\rho_{\boldsymbol{D}}(\boldsymbol{S}) = 1$.*

Proof. The claim is implied by the mass transport principle (2.2) for the function $g[D, u, v; S] := 1_{\{v \in S\}}$. The details are left to the reader. $\square$

The above lemma is a generalization of similar results in [5] and [2].

2.6 Notes and Bibliographical Comments

The mass transport principle was introduced in [31]. The concept of unimodular graphs was first defined for deterministic transitive graphs in [12] and generalized to random rooted graphs and networks in [2].

Unimodular graphs have many analogies and connections to (Palm versions of) stationary point processes and point-stationary point processes, as discussed in Example 9.5 of [2] and also in [5] and [35]. As already explained, the framework of unimodular discrete spaces introduced in this section can be regarded as a common generalization of these concepts.

Special cases of the notion of equivariant processes have been considered in the literature. The first formulation in Subsection 2.5 is considered in [2] for unimodular graphs. *Factors of IID* [40] are special cases of equivariant processes where the marks of the points are obtained from i.i.d. marks (Example 2.11) in an equivariant way. *Covariant subsets* and *covariant partitions* of unimodular graphs are defined similarly in [5], but no extra randomness is allowed therein. In the case of stationary (marked) point processes, the first formulation of Subsection 2.5 is used in the literature. However, the authors believe that the general formulation of Definition 2.9 is new even in those special cases.

3 The Unimodular Minkowski and Hausdorff Dimensions

This section presents the new notions of dimension for unimodular discrete spaces. As mentioned in the introduction, the statistical homogeneity of unimodular discrete

spaces is used to define discrete analogues of the Minkowski and Hausdorff dimensions. Also, basic properties of these definitions are discussed.

3.1 The Unimodular Minkowski Dimension

Definition 3.1. *Let $[\boldsymbol{D}, \boldsymbol{o}]$ be a unimodular discrete space and $r \geq 0$. An **equivariant r-covering $\boldsymbol{R}$ of $\boldsymbol{D}$** is an equivariant subset of $\boldsymbol{D}$ such that the set of balls $\{N_r(v) : v \in \boldsymbol{R_D}\}$ cover $\boldsymbol{D}$ almost surely. Here, the same symbol $\boldsymbol{R}$ is used for the following equivariant process (Definition 2.9):*

$$\boldsymbol{R}(v) := \boldsymbol{R_D}(v) := \begin{cases} r, & \text{there is a ball centered at } v \text{ in the covering,} \\ 0, & \text{otherwise,} \end{cases}$$

for $v \in \boldsymbol{D}$. Note that an equivariant covering may use extra randomness and is not necessarily a function of $\boldsymbol{D}$. This is essential in the following definition.

Let $\mathcal{C}_r$ be the set of all equivariant r-coverings. Define

$$\lambda_r := \lambda_r(\boldsymbol{D}) := \inf\{\text{intensity of } \boldsymbol{R} \text{ in } \boldsymbol{D} : \boldsymbol{R} \in \mathcal{C}_r\}, \tag{3.1}$$

where the intensity is defined in Definition 2.14.

Note that λ_r is non-increasing in terms of r. A smaller λ_r heuristically means that a *smaller number of balls per point* is needed to cover $\boldsymbol{D}$. So define

Definition 3.2. *The **upper and lower unimodular Minkowski dimensions** of $\boldsymbol{D}$ are defined by*

$$\begin{aligned} \overline{\operatorname{udim}}_M(\boldsymbol{D}) &:= \overline{\operatorname{decay}}(\lambda_r), \\ \underline{\operatorname{udim}}_M(\boldsymbol{D}) &:= \underline{\operatorname{decay}}(\lambda_r), \end{aligned}$$

*as $r \to \infty$. If the decay rate of λ_r exists, define the **unimodular Minkowski dimension** of $\boldsymbol{D}$ by*

$$\operatorname{udim}_M(\boldsymbol{D}) := \operatorname{decay}(\lambda_r).$$

Here are some first illustrations of the definition.

Example 3.3. The randomly shifted lattice $\boldsymbol{S}_n := (2n+1)\mathbb{Z}^k - \boldsymbol{U}_n$, where $\boldsymbol{U}_n \in \{-n, \dots, n\}^k$ is chosen uniformly, is an equivariant n-covering of $\mathbb{Z}^k$ equipped with the l_∞ metric (other metrics can be treated similarly). This implies that $\lambda_n \leq \mathbb{P}[0 \in \boldsymbol{S}_n] = (2n+1)^{-k}$, and hence, $\underline{\operatorname{udim}}_M(\mathbb{Z}^k) \geq k$.

Example 3.4. If $\boldsymbol{D}$ is finite with positive probability, then it can be seen that any non-empty equivariant subset has intensity at least $\mathbb{E}[1/\#\boldsymbol{D}]$ (use the mass transport principle when sending mass $1/\#\boldsymbol{D}$ from every point of the subset to every point of $\boldsymbol{D}$). This implies that $\operatorname{udim}_M(\boldsymbol{D}) = 0$.

Remark 3.5 (Bounding the Minkowski Dimension)**.** In all examples in this work, lower bounds on the unimodular Minkowski dimension are obtained by constructing explicit examples of r-coverings. Upper bounds can be obtained by constructing *disjoint* or *bounded* coverings, as discussed in Subsection 3.2 below, or by comparison with the unimodular Hausdorff dimension defined in Subsection 3.3 below (see Theorem 3.22).

3.2 Optimal Coverings for the Minkowski Dimension

Definition 3.6. *Let $[\boldsymbol{D}, \boldsymbol{o}]$ be a unimodular discrete space and $r \geq 0$. If the infimum in the definition of λ_r (3.1) is attained by an equivariant r-covering $\boldsymbol{S}$; i.e., $\mathbb{P}[\boldsymbol{o} \in \boldsymbol{S_D}] = \lambda_r$, then $\boldsymbol{S}$ is called an **optimal r-covering** for $\boldsymbol{D}$.*

Theorem 3.7. *Every unimodular discrete space has an optimal r-covering for every $r \geq 0$.*

Sketch of the proof. Let $\boldsymbol{S}_1, \boldsymbol{S}_2, \ldots$ be a sequence of equivariant r-coverings of $\boldsymbol{D}$ such that $\mathbb{P}\left[\boldsymbol{o} \in \boldsymbol{S}_n\right] \to \lambda_r$. By a tightness argument and choosing a subsequence if necessary, one may assume that $[\boldsymbol{D}, \boldsymbol{o}; \boldsymbol{S}_n]$ converges weakly, say to $[\boldsymbol{D}, \boldsymbol{o}; \boldsymbol{S}]$, where $\boldsymbol{S}$ is an equivariant subset $\boldsymbol{S}$ of $\boldsymbol{D}$ (see [7] for more details). Since each $\boldsymbol{S}_n$ is an r-covering, $\mathbb{P}\left[\boldsymbol{S}_n \cap N_r(\boldsymbol{o}) = \emptyset\right] = 0$. It is straightforward to deduce that $\mathbb{P}\left[\boldsymbol{S} \cap N_r(\boldsymbol{o}) = \emptyset\right] = 0$. So by putting balls of radius r on the points of $\boldsymbol{S}$, the root is covered a.s. So Lemma 2.15 implies that every point is covered a.s.; i.e., $\boldsymbol{S}$ is an r-covering. Also, by weak convergence, $\mathbb{P}\left[\boldsymbol{o} \in \boldsymbol{S}\right] = \lim_n \mathbb{P}\left[\boldsymbol{o} \in \boldsymbol{S}_n\right] = \lambda_r$. This implies that $\boldsymbol{S}$ is an optimal r-covering. □

In general, finding an optimal covering is difficult. In some examples, the following is easier to study.

Definition 3.8. *Let $K < \infty$ and $r \geq 0$. An r-covering of D is K-**bounded** if each point of D is covered at most K times a.s. A sequence $(\boldsymbol{R}_n)_n$ of equivariant coverings of D is called **uniformly bounded** if there is $K < \infty$ such that each $\boldsymbol{R}_n$ is K-bounded.*

Lemma 3.9. *If $\boldsymbol{R}$ is a K-bounded equivariant r-covering of $\boldsymbol{D}$, then*

$$\frac{1}{K}\mathbb{P}\left[\boldsymbol{R}(\boldsymbol{o}) \neq 0\right] \leq \lambda_r \leq \mathbb{P}\left[\boldsymbol{R}(\boldsymbol{o}) \neq 0\right]. \tag{3.2}$$

So if $(\boldsymbol{R}_n)_n$ is a sequence of equivariant coverings which is uniformly bounded, with $\boldsymbol{R}_n$ an n-covering for each $n \geq 1$, then

$$\begin{aligned}
\overline{\mathrm{udim}}_M(\boldsymbol{D}) &= \overline{\mathrm{decay}}\left(\mathbb{P}\left[\boldsymbol{R}_n(\boldsymbol{o}) \neq 0\right]\right), \\
\underline{\mathrm{udim}}_M(\boldsymbol{D}) &= \underline{\mathrm{decay}}\left(\mathbb{P}\left[\boldsymbol{R}_n(\boldsymbol{o}) \neq 0\right]\right).
\end{aligned}$$

Proof. The rightmost inequality in (3.2) is immediate from the definition of λ_r. Let $\boldsymbol{R}'$ be another equivariant r-covering. Let $g(u, v) = 1$ if $\boldsymbol{R}'(u) = \boldsymbol{R}(v) = r$ and $d(u, v) \leq r$. Then $g^+(\boldsymbol{o}) \leq K 1_{\{\boldsymbol{R}'(\boldsymbol{o}) \neq 0\}}$ and $g^-(\boldsymbol{o}) \geq 1_{\{\boldsymbol{R}(\boldsymbol{o}) \neq 0\}}$. Hence by the mass transport principle (2.2), $\frac{1}{K}\mathbb{P}\left[\boldsymbol{R}(\boldsymbol{o}) \neq 0\right] \leq \mathbb{P}\left[\boldsymbol{R}'(\boldsymbol{o}) \neq 0\right]$ and the leftmost inequality in (3.2) then follows from the definition of λ_r. The last two equalities follow immediately from (3.2). □

Corollary 3.10. *If $\boldsymbol{R}$ is an equivariant **disjoint** r-**covering** of $\boldsymbol{D}$ (i.e., the balls used in the covering are pairwise disjoint a.s.), then it is an optimal r-covering for $\boldsymbol{D}$.*

Example 3.11. The covering of $\mathbb{Z}^k$ (equipped with the l_∞ metric) constructed in Example 3.3 is a disjoint covering. So it is optimal and hence $\mathrm{udim}_M(\mathbb{Z}^k) = k$. For $\mathbb{Z}^k$ equipped with the Euclidean metric, one can construct a 3^k-bounded covering similarly and deduce the same result.

Example 3.12. Let T_k be the k-regular tree. For $r \geq 1$, consider a deterministic covering of T_k by disjoint balls of radius r. By choosing $\boldsymbol{o}$ in one of these balls uniformly at random, it can be seen that an equivariant disjoint r-covering of $[T_k, \boldsymbol{o}]$ is obtained (the proof is left to the reader). So Corollary 3.10 implies that $\lambda_r = 1/\#N_r(\boldsymbol{o})$ which has exponential decay when $k \geq 3$. Hence, $\mathrm{udim}_M(T_k) = \infty$ for $k \geq 3$.

Proposition 3.13. *For any point-stationary point process Φ on $\mathbb{R}$ endowed with the Euclidean metric, by letting $p(r) := \mathbb{P}\left[\Phi \cap (0, r) = \emptyset\right]$, one has*

$$\begin{aligned}
\overline{\mathrm{udim}}_M(\Phi) &= \overline{\mathrm{decay}}\left(\tfrac{1}{r}\int_0^r p(s)\mathrm{d}s\right) \leq 1 \wedge \overline{\mathrm{decay}}\left(p(r)\right), \\
\underline{\mathrm{udim}}_M(\Phi) &= \underline{\mathrm{decay}}\left(\tfrac{1}{r}\int_0^r p(s)\mathrm{d}s\right) = 1 \wedge \underline{\mathrm{decay}}\left(p(r)\right).
\end{aligned}$$

Proof. Let $r > 0$ and φ be a discrete subset of $\mathbb{R}$. Let $\boldsymbol{U}_r$ be a random number in $[0, r)$ chosen uniformly. For each $n \in \mathbb{Z}$, put a ball of radius r centered at the largest element of $\varphi \cap [nr + \boldsymbol{U}_r, (n+1)r + \boldsymbol{U}_r)$. Denote this random r-covering of φ by $\boldsymbol{R}_\varphi$. One can see that $\boldsymbol{R}$ is equivariant under translations (see Remark 2.13). This implies that $\boldsymbol{R}$ is an

equivariant covering (verifying Condition (ii) of Definition I. 2.9 is skipped here). One has

$$\mathbb{P}\left[0 \in \boldsymbol{R}_\Phi\right] = \mathbb{P}\left[\Phi \cap (0, \boldsymbol{U}_r) = \emptyset\right] = \frac{1}{r} \int_0^r \mathbb{P}\left[\Phi \cap (0, s) = \emptyset\right] \mathrm{d}s =: q(r).$$

Now, since $\boldsymbol{R}$ is a 3-bounded covering, Lemma 3.9 implies the two left-hand-side equalities. For all $\beta < \underline{\operatorname{decay}}\,(p(r))$, one has $p(r) < r^{-\beta}$ for large enough r. So, if in addition, $\beta < 1$, then $\overline{q(r) < cr^{-\beta}}$ for some constant c, so that $\underline{\operatorname{decay}}\,(q(r)) \geq \beta$. Therefore $\underline{\operatorname{decay}}\,(q(r)) \geq 1 \wedge \underline{\operatorname{decay}}\,(p(r))$. Now, the final equality in the claim is deduced from $q(r) \geq p(r)$. Similarly, if $\overline{\operatorname{decay}}\,(p(r)) < 1$, one can deduce $\overline{\operatorname{decay}}\,(q(r)) \leq \overline{\operatorname{decay}}\,(p(r))$. Also, $q(r) \geq \frac{1}{r} \int_0^1 p(s)ds$, and hence $\overline{\operatorname{decay}}\,(q(r)) \leq 1$. This implies the first inequality and completes the proof. $\qquad\square$

3.3 The Unimodular Hausdorff Dimension

The definition of the unimodular Hausdorff dimension is based on coverings of the discrete space by balls of possibly different radii. Such a covering can be represented by an assignment of marks to the points, where the mark of a point v represents the radius of the ball centered at v. As mentioned earlier, it is convenient to assume that the radii are at least 1 (in fact, this condition is technically necessary in what follows). Also, by convention, if there is no ball centered at v, the mark of v is defined to be 0. In relation with this convention, the following notation is used for all discrete spaces D and points $v \in D$:

$$N_r(v) := \left\{ \begin{array}{ll} \{u \in D : d(v, u) \leq r\}, & r \geq 1, \\ \emptyset, & r = 0. \end{array} \right.$$

In words, $N_r(v)$ is the *closed ball* of radius r centered at v, except when $r = 0$.

Definition 3.14. *Let $[\boldsymbol{D}, \boldsymbol{o}]$ be a unimodular discrete space. An **equivariant (ball-) covering $\boldsymbol{R}$ of $\boldsymbol{D}$** is an equivariant process on $\boldsymbol{D}$ (Definition 2.9) with values in $\Xi := \{0\} \cup [1, \infty)$, such that the family of balls $\{N_{\boldsymbol{R}(v)}(v) : v \in \boldsymbol{D}\}$ covers the points of $\boldsymbol{D}$ almost surely. For simplicity, $N_{\boldsymbol{R}(v)}(v)$ will also be denoted by $N_{\boldsymbol{R}}(v)$. Also, for $0 \leq \alpha < \infty$ and $1 \leq M < \infty$, let*

$$\mathcal{H}_M^\alpha(\boldsymbol{D}) := \inf \left\{\mathbb{E}\left[\boldsymbol{R}(\boldsymbol{o})^\alpha\right] : \boldsymbol{R}(v) \in \{0\} \cup [M, \infty),\ \forall v,\ \textit{a.s.}\right\}, \tag{3.3}$$

where the infimum is over all equivariant coverings $\boldsymbol{R}$ such that almost surely, $\forall v \in \boldsymbol{D}$: $\boldsymbol{R}(v) \in \{0\} \cup [M, \infty)$, and, by convention, $0^0 := 0$. Note that $\mathcal{H}_M^\alpha(\boldsymbol{D})$ is a non-decreasing function of both α and M.

In the *ergodic* case, $\mathbb{E}\left[\boldsymbol{R}(\boldsymbol{o})^\alpha\right]$ can be interpreted as the average of $\boldsymbol{R}(\cdot)^\alpha$ over the vertices. Also, $\mathbb{P}\left[\boldsymbol{R}(\boldsymbol{o}) > 0\right]$ (which is used for defining the unimodular Minkowski dimension) can be interpreted as *the number of balls per point*. Ergodicity is however a special case, and there is no need to assume it in what follows; for more on the matter, see Example 3.19 and the discussion after it.

Definition 3.15. *Let $[\boldsymbol{D}, \boldsymbol{o}]$ be a unimodular discrete space. The number $\mathcal{H}_1^\alpha(\boldsymbol{D})$, defined in (3.3), is called the α-**dimensional Hausdorff content** of $\boldsymbol{D}$. The **unimodular Hausdorff dimension** of $\boldsymbol{D}$ is defined by*

$$\operatorname{udim}_H(\boldsymbol{D}) := \sup\{\alpha \geq 0 : \mathcal{H}_1^\alpha(\boldsymbol{D}) = 0\}, \tag{3.4}$$

with the convention that $\sup \emptyset = 0$.

The key point of assuming equivariance in the above definition is that by Lemma 2.12, $[\boldsymbol{D}, \boldsymbol{o}; \boldsymbol{R}]$ is a unimodular marked discrete space. Note also that extra randomness is allowed in the definition of equivariant coverings. Note also that

$$0 \leq \mathcal{H}_1^\alpha(\boldsymbol{D}) \leq 1,$$

since for the covering by balls of radius 1, one has $\mathbb{E}\left[\boldsymbol{R}(\boldsymbol{o})^\alpha\right] = 1$.

Examples 3.16 and 3.19 below provide basic illustrations of the unimodular Hausdorff dimension.

Example 3.16. If $\boldsymbol{D}$ is finite with positive probability, then one can show similarly to Example 3.4 that $\mathbb{E}\left[\boldsymbol{R}(\boldsymbol{o})^\alpha\right] \geq \mathbb{E}\left[1/\#\boldsymbol{D}\right]$ for every $\boldsymbol{R}$, and hence, $\operatorname{udim}_H(\boldsymbol{D}) = 0$. Also, for the covering $\boldsymbol{S}_n$ of $\mathbb{Z}^k$ constructed in Example 3.3, one has $\mathbb{E}\left[\boldsymbol{S}_n(\boldsymbol{o})^\alpha\right] = (2n+1)^{\alpha-k}$. If $\alpha < k$, this implies that $\mathcal{H}_1^\alpha(\mathbb{Z}^k) = 0$, and hence, $\operatorname{udim}_H(\mathbb{Z}^k) \geq k$. The upper bound $\operatorname{udim}_H(\mathbb{Z}^k) \leq k$ is implied by Lemma 3.17 below. So $\operatorname{udim}_H(\mathbb{Z}^k) = k$.

Lemma 3.17. *Let $[\boldsymbol{D}, \boldsymbol{o}]$ be a unimodular discrete space and $\alpha \geq 0$. If there exists $c \geq 0$ such that $\forall r \geq 1 : \#N_r(\boldsymbol{o}) \leq cr^\alpha$ a.s., then $\operatorname{udim}_H(\boldsymbol{D}) \leq \alpha$.*

Proof. Let $\boldsymbol{R}$ be an arbitrary equivariant covering. For all discrete spaces D and $u, v \in D$, let $g_D(u, v)$ be 1 if $d(u, v) \leq \boldsymbol{R}_D(u)$ and 0 otherwise. One has $g^+(u) = \#N_{\boldsymbol{R}}(u)$ and $g^-(u) \geq 1$ a.s. (since $\boldsymbol{R}$ is a covering). By the assumption and the mass transport principle (2.2), one gets

$$\mathbb{E}\left[\boldsymbol{R}(\boldsymbol{o})^\alpha\right] \geq \frac{1}{c}\mathbb{E}\left[\#N_{\boldsymbol{R}}(\boldsymbol{o})\right] = \frac{1}{c}\mathbb{E}\left[g^+(\boldsymbol{o})\right] = \frac{1}{c}\mathbb{E}\left[g^-(\boldsymbol{o})\right] \geq \frac{1}{c}.$$

Since $\boldsymbol{R}$ is arbitrary, one gets $\mathcal{H}_1^\alpha(\boldsymbol{D}) \geq \frac{1}{c} > 0$, and hence, $\operatorname{udim}_H(\boldsymbol{D}) \leq \alpha$. □

Remark 3.18 (Bounding the Hausdorff Dimension). In most examples in this work, a lower bound on the unimodular Hausdorff dimension is provided, either by comparison with the Minkowski dimension (see Subsection 3.4 below), or by explicit construction of a sequence of equivariant coverings $\boldsymbol{R}_1, \boldsymbol{R}_2, \ldots$ such that $\mathbb{E}\left[\boldsymbol{R}_n(\boldsymbol{o})^\alpha\right] \to 0$ as $n \to \infty$. Note that this gives $\mathcal{H}_1^\alpha(\boldsymbol{D}) = 0$, which implies that $\operatorname{udim}_H(\boldsymbol{D}) \geq \alpha$. Constructing coverings does not help to find upper bounds for the Hausdorff dimension. The derivation of upper bounds is mainly discussed in Section 5. The main tools are *the mass distribution principle* (Theorem 5.2), which is a stronger form of Lemma 3.17 above, and the *unimodular Billingsley's lemma* (Theorem 5.6).

Example 3.19. Let $[\boldsymbol{D}, \boldsymbol{o}]$ be $[\mathbb{Z}, 0]$ with probability $\frac{1}{2}$ and $[\mathbb{Z}^2, 0]$ with probability $\frac{1}{2}$. It is shown below that $\operatorname{udim}_M(\boldsymbol{D}) = \operatorname{udim}_H(\boldsymbol{D}) = 1$.

For $n \in \mathbb{N}$, the equivariant n-covering of Example 3.3 makes sense for $\boldsymbol{D}$ and is uniformly bounded. One has $\mathbb{P}\left[\boldsymbol{R}(0) > 0\right] = \frac{1}{2}(n^{-1} + n^{-2})$. This implies that $\operatorname{udim}_M(\boldsymbol{D}) = \operatorname{decay}\left(\frac{1}{2}(n^{-1} + n^{-2})\right) = 1$. Also, for $\alpha < 1$, one has $\mathbb{E}\left[\boldsymbol{R}(\boldsymbol{o})^\alpha\right] = \frac{1}{2}(n^{\alpha-1} + n^{\alpha-2}) \to 0$ as $n \to \infty$. This implies that $\mathcal{H}_1^\alpha(\boldsymbol{D}) = 0$ for all $\alpha < 1$ and hence $\operatorname{udim}_H(\boldsymbol{D}) \geq 1$. On the other hand, for any equivariant covering $\boldsymbol{S}$, one has

$$\mathbb{E}\left[\boldsymbol{S}(\boldsymbol{o})\right] \geq \mathbb{E}\left[\boldsymbol{S}(\boldsymbol{o})\,|\,\boldsymbol{D} = \mathbb{Z}\right]\mathbb{P}\left[\boldsymbol{D} = \mathbb{Z}\right] = \frac{1}{2}\mathbb{E}\left[\boldsymbol{S}(\boldsymbol{o})\,|\,\boldsymbol{D} = \mathbb{Z}\right].$$

Let $c > 2$. The proof of Lemma 3.17 for $[\mathbb{Z}, 0]$ implies that $\mathbb{E}\left[\boldsymbol{S}(\boldsymbol{o})\,|\,\boldsymbol{D} = \mathbb{Z}\right] \geq \frac{1}{c}$. This implies that $\mathcal{H}_1^1(\boldsymbol{D}) \geq \frac{1}{2c} > 0$. So $\operatorname{udim}_H(\boldsymbol{D}) \leq 1$.

Remark 3.20. The result of this example might seem counterintuitive at first glance as the union of a filled square and a segment is two dimensional. The number of balls of radius ϵ required to cover the square dominates the number of balls required to cover the segment, but in Example 3.19, the situation is reversed: a larger *fraction* of points is needed to cover $\mathbb{Z}$ than $\mathbb{Z}^2$. This is a consequence of considering large balls and also counting the number of balls *per point*. See also Subsection 8.3.

In fact, the following example justifies more clearly why Example 3.19 is one dimensional: Let G_n be the union of a $n \times n$ square grid (regarded as a graph) and a path of length n^2 sharing a vertex with the grid. To cover G_n by balls of radius r, a fraction of order $1/r$ of the vertices of G_n are needed (as r is fixed and $n \to \infty$). So it is not counterintuitive to

say that G_n is one dimensional asymptotically. Indeed, G_n tends to the random graph of Example 3.19 in the local weak convergence [2] as $n \to \infty$ (if one chooses the root of G_n randomly and uniformly).

Remark 3.21. In Example 3.19 above, different samples of $\boldsymbol{D}$ have different natures heuristically. This is formalized by saying that $[\boldsymbol{D}, \boldsymbol{o}]$ is *non-ergodic*; i.e., there is an event $A \subseteq \mathcal{D}_*$ such that the proposition $[D, o] \in A$ does not depend on the origin of D and $0 < \mathbb{P}[[\boldsymbol{D}, \boldsymbol{o}] \in A] < 1$. In such cases, it is desirable to assign a dimension to every sample of $\boldsymbol{D}$. In easy examples like Example 3.19, this might be achieved by conditioning. For instance, in some examples, it is convenient to condition on having infinite cardinality (which is common, e.g., in branching processes). However, in general, it doesn't seem easier to define the dimension of samples separately in a way that is compatible with the definitions of this paper. In the future work [6], the notion of *sample dimension* is defined by combining the definitions in this paper with either *ergodic decomposition* or conditional expectation. In this work, the reader may focus mainly on the ergodic case, but it should be noted that the definitions and results do not require ergodicity.

3.4 Comparison of Hausdorff and Minkowski Dimensions

Theorem 3.22 (Minkowski vs. Hausdorff). *One has*

$$\underline{\mathrm{udim}}_M(\boldsymbol{D}) \leq \overline{\mathrm{udim}}_M(\boldsymbol{D}) \leq \mathrm{udim}_H(\boldsymbol{D}).$$

Proof. The first inequality holds by the definition. For the second one, the definition of λ_r (3.1) implies that for every $\alpha \geq 0$ and $r \geq 1$,

$$\inf\{\mathbb{E}[\boldsymbol{R}(\boldsymbol{o})^\alpha] : \boldsymbol{R} \text{ is an equivariant } r\text{-covering}\} = r^\alpha \lambda_r.$$

This readily implies that $\mathcal{H}_1^\alpha(\boldsymbol{D}) \leq r^\alpha \lambda_r$ for every $r \geq 1$. So, if $\alpha < \overline{\mathrm{decay}}(\lambda_r)$, one gets $\mathcal{H}_1^\alpha(\boldsymbol{D}) = 0$, and hence, $\mathrm{udim}_H(\boldsymbol{D}) \geq \alpha$. This implies the claim. $\qquad\square$

Remark 3.23. There exist examples in which the inequalities in Theorem 3.22 are strict (see e.g., Subsections 4.2.2 and 4.4). However, equality holds in most examples. In what follows, the equality $\mathrm{udim}_M(\boldsymbol{D}) = \mathrm{udim}_H(\boldsymbol{D})$ will be referred to as *regularity* for the unimodular discrete space $\boldsymbol{D}$, regarded as a fractal object.

3.5 The Unimodular Hausdorff Size

Consider the setting of Subsection 3.3. For $0 \leq \alpha < \infty$, let

$$\mathcal{H}_\infty^\alpha(\boldsymbol{D}) := \lim_{M \to \infty} \mathcal{H}_M^\alpha(\boldsymbol{D}) \in [0, \infty], \tag{3.5}$$

where $\mathcal{H}_M^\alpha(\boldsymbol{D})$ is defined in (3.3). Note that the limit exists because of monotonicity.

Definition 3.24. *The **unimodular α-dimensional Hausdorff size** of D (in short, unimodular α-dim H-size of $\boldsymbol{D}$) is*

$$\mathcal{M}^\alpha(\boldsymbol{D}) := (\mathcal{H}_\infty^\alpha(\boldsymbol{D}))^{-1}. \tag{3.6}$$

This definition resembles the Hausdorff measure of compact sets. But since $\mathcal{M}^\alpha$ is not a measure, the term *size* is used instead. It can be used to compare unimodular spaces with equal dimension. The following results gather some elementary properties of the function $\mathcal{H}_M^\alpha$ and the Hausdorff size.

Lemma 3.25. *One has*

(i) $\mathcal{H}_1^\alpha(\boldsymbol{D}) \leq \mathcal{H}_M^\alpha(\boldsymbol{D}) \leq M^\alpha \mathcal{H}_1^\alpha(\boldsymbol{D})$.

(ii) $\mathcal{H}_1^\alpha(\boldsymbol{D}) = 0 \Leftrightarrow \mathcal{H}_\infty^\alpha(\boldsymbol{D}) = 0 \Leftrightarrow \mathcal{M}^\alpha(\boldsymbol{D}) = \infty$.

(iii) If $\alpha \geq \beta$, then $\mathcal{H}_M^\alpha(\boldsymbol{D}) \geq M^{\alpha-\beta}\mathcal{H}_M^\beta(\boldsymbol{D})$.

Proof. (i). If $\boldsymbol{R}$ is an equivariant covering, them $M\boldsymbol{R}$ is also an equivariant covering and satisfies $\forall v \in \boldsymbol{D} : MR(v) \in \{0\} \cup [M, \infty)$ a.s.

(ii). The claim is implied by part (i).

(iii). If $\boldsymbol{R}$ is an equivariant covering such that $\forall v \in \boldsymbol{D} : \boldsymbol{R}(v) \in \{0\} \cup [M, \infty)$ a.s., then $\boldsymbol{R}(\boldsymbol{o})^\alpha \geq M^{\alpha-\beta}\boldsymbol{R}(\boldsymbol{o})^\beta$ a.s. $\qquad\square$

Lemma 3.26. *If $\alpha < \mathrm{udim}_H(\boldsymbol{D})$, then $\mathcal{H}_\infty^\alpha(\boldsymbol{D}) = 0$ and $\mathcal{M}^\alpha(\boldsymbol{D}) = \infty$. Also, if $\alpha > \mathrm{udim}_H(\boldsymbol{D})$, then $\mathcal{H}_\infty^\alpha(\boldsymbol{D}) = \infty$ and $\mathcal{M}^\alpha(\boldsymbol{D}) = 0$.*

Proof. For $\alpha < \mathrm{udim}_H(\boldsymbol{D})$, one has $\mathcal{H}_1^\alpha(\boldsymbol{D}) = 0$. So part (ii) of Lemma 3.25 implies that $\mathcal{M}^\alpha(\boldsymbol{D}) = \infty$. For $\alpha > \mathrm{udim}_H(\boldsymbol{D})$, there exists β such that $\alpha > \beta > \mathrm{udim}_H(\boldsymbol{D})$. For this β, one has $\mathcal{H}_1^\beta(\boldsymbol{D}) > 0$ and part (iii) of the same lemma implies that $\mathcal{H}_M^\alpha(\boldsymbol{D}) \geq M^{\alpha-\beta}\mathcal{H}_M^\beta(\boldsymbol{D}) \geq M^{\alpha-\beta}\mathcal{H}_1^\beta(\boldsymbol{D})$. This implies that $\mathcal{H}_\infty^\alpha(\boldsymbol{D}) = \infty$, which proves the claim. $\qquad\square$

Remark 3.27. For $\alpha := \mathrm{udim}_H(\boldsymbol{D})$, the α-dim H-size of $\boldsymbol{D}$ can be zero, finite or infinite. The lattice $\mathbb{Z}^k$ provides a case where $\mathcal{M}^\alpha(\boldsymbol{D})$ is positive and finite (Proposition 3.29 below). Examples 6.1 and 6.2 provide examples of the infinite and zero cases respectively.

The following propositions provide basic examples of the computation of the Hausdorff size.

Proposition 3.28 (0-dim H-size). *One has $\mathcal{M}^0(\boldsymbol{D}) = (\mathbb{E}\left[1/\#\boldsymbol{D}\right])^{-1}$.*

Proof. As in Example 3.16, one gets $\mathcal{H}_M^0(\boldsymbol{D}) \geq \mathbb{E}\left[1/\#\boldsymbol{D}\right]$. It is enough to show that equality holds. If $\boldsymbol{D}$ is finite a.s., this can be proved by putting a single ball of radius $M \vee \mathrm{diam}(\boldsymbol{D})$ centered at a point of $\boldsymbol{D}$ chosen uniformly at random. Second, assume $\boldsymbol{D}$ is infinite a.s. It is enough to construct an equivariant covering $\boldsymbol{R}$ such that $\mathbb{P}\left[\boldsymbol{R}(\boldsymbol{o}) > 0\right]$ is arbitrarily small. Let $p > 0$ be arbitrary and $\boldsymbol{S}$ be the *Bernoulli equivariant subset* obtained by selecting each point with probability p in an i.i.d. manner. For all infinite discrete spaces D and $v \in D$, let $\boldsymbol{\tau}_D(v)$ be the closest point of $\boldsymbol{S}_D$ to v (if there is a tie, choose one of them uniformly at random independently). It can be seen that $\boldsymbol{\tau}_D^{-1}(u)$ is finite almost surely (use the mass transport principle for $g(v, u) := 1_{\{u=\boldsymbol{\tau}_D(v)\}}$). For $u \in \boldsymbol{S}_D$, let $\boldsymbol{R}(u) := 1 \vee \mathrm{diam}(\boldsymbol{\tau}^{-1}(u))$ be the diameter of the *Voronoi cell* of u. For $u \in D \setminus \boldsymbol{S}_D$, let $\boldsymbol{R}(u) := 0$. It is clear that $\boldsymbol{R}$ is a covering, and in fact, an equivariant covering. One has $\mathbb{P}\left[\boldsymbol{R}(\boldsymbol{o}) > 0\right] = \mathbb{P}\left[\boldsymbol{o} \in \boldsymbol{S}_D\right] = p$, which is arbitrarily small. So the claim is proved in this case.

Finally, assume $\boldsymbol{D}$ is finite with probability q. For all deterministic discrete spaces D, let $\boldsymbol{R}_D$ be one of the above coverings depending on whether D is finite or infinite. It satisfies $\mathbb{P}\left[\boldsymbol{R}(\boldsymbol{o}) > 0\right] = \mathbb{E}\left[1/\#\boldsymbol{D}\right] + p(1-q)$. Since p is arbitrary, the claim is proved. $\quad\square$

Proposition 3.29. *For all $\delta > 0$, the k-dim H-size of the scaled lattice $[\delta\mathbb{Z}^k, 0]$, equipped with the l_∞ metric, is equal to $(2/\delta)^k$.*

Proof. Let $\boldsymbol{S}_n$ be the covering in Example 3.3 scaled by factor δ. One has $\mathbb{E}\left[\boldsymbol{S}_n(\boldsymbol{o})^k\right] = (n\delta)^k/(2n+1)^k$. This easily implies that $\mathcal{H}_\infty^k(\delta\mathbb{Z}^k) \leq (\delta/2)^k$. On the other hand, the proof of Lemma 3.17 shows that $\mathcal{H}_\infty^k(\delta\mathbb{Z}^k) \geq c\delta^k$, where c is any constant such that $r^k \geq c\#N_r(0)$ for large enough r. It follows that $\mathcal{H}_\infty^k(\delta\mathbb{Z}^k) \geq (\delta/2)^k$, and the claim is proved. $\qquad\square$

3.6 The Effect of a Change of Metric

To avoid confusion when considering two metrics, a pointed discrete space is denoted by $((D, d), o)$ here, where d is the metric on D and o is the origin. Note that if d' is another metric on D, then $d' \in \mathbb{R}^{D \times D}$. So d' can be considered as a marking of D in the sense of Definition 2.8 and $((D, d), o; d')$ is a pointed marked discrete space.

Definition 3.30. *An **equivariant (boundedly finite) metric** is an $\mathbb{R}$-valued equivariant process $\boldsymbol{d}'$ such that, for all discrete spaces (D, d), $\boldsymbol{d}'_{(D,d)}$ is almost surely (w.r.t. the extra randomness) a metric on D and $(D, \boldsymbol{d}'_{(D,d)})$ is a boundedly finite metric space.*

If in addition, $[(\boldsymbol{D}, \boldsymbol{d}), \boldsymbol{o}]$ is a unimodular discrete space, then $[(\boldsymbol{D}, \boldsymbol{d}), \boldsymbol{o}; \boldsymbol{d}']$ is a unimodular marked discrete space by Lemma 2.12. It can be seen that $[(\boldsymbol{D}, \boldsymbol{d}'), \boldsymbol{o}; \boldsymbol{d}]$, obtained by swapping the metrics, makes sense as a random pointed marked discrete space (see [7] for the measurability requirements). By verifying the mass transport principle (2.2) directly, it is easy to show that $[(\boldsymbol{D}, \boldsymbol{d}'), \boldsymbol{o}; \boldsymbol{d}]$ is unimodular.

The following result is valid for both the Hausdorff and the (upper and lower) Minkowski dimensions.

Theorem 3.31 (Change of Metric)**.** *Let $[(\boldsymbol{D}, \boldsymbol{d}), \boldsymbol{o}]$ be a unimodular discrete space and $\boldsymbol{d}'$ be an equivariant metric. If $\boldsymbol{d}' \leq c\boldsymbol{d} + a$ a.s., with c and a constants, then the dimension of $(\boldsymbol{D}, \boldsymbol{d}')$ is larger than or equal to that of $(\boldsymbol{D}, \boldsymbol{d})$. Moreover, for every $\alpha \geq 0$, $\mathcal{M}^{\alpha}(\boldsymbol{D}, \boldsymbol{d}') \geq c^{-\alpha} \mathcal{M}^{\alpha}(\boldsymbol{D}, \boldsymbol{d})$.*

Proof. The claim is implied by the fact that the ball $N_{cr+a}((\boldsymbol{D}, \boldsymbol{d}'), v)$ contains the ball $N_r((\boldsymbol{D}, \boldsymbol{d}), v)$ and is left to the reader. $\qquad\square$

As a corollary, if $\frac{1}{c}\boldsymbol{d} - a \leq \boldsymbol{d}' \leq c\boldsymbol{d} + a$ a.s., then $(\boldsymbol{D}, \boldsymbol{d}')$ has the same unimodular dimensions as $(\boldsymbol{D}, \boldsymbol{d})$. Also, $c\boldsymbol{D}$ has the same dimension as $\boldsymbol{D}$ and $\mathcal{M}^{\alpha}(c\boldsymbol{D}) = c^{-\alpha}\mathcal{M}^{\alpha}(\boldsymbol{D})$.

For instance, this result can be applied to Cayley graphs, which are an important class of unimodular graphs [2]. It follows that the unimodular dimensions of a Cayley graph do not depend on the generating set. In fact, it will be proved in Subsection 6.6 that these dimensions are equal to the *polynomial growth degree* of H.

Example 3.32. Let $[G, o]$ be a unimodular graph. Examples of equivariant metrics on G are the graph-distance metric corresponding to an equivariant spanning subgraph (e.g., the drainage network model of Subsection 4.5 below) and metrics *generated by equivariant edge lengths*. More precisely, if l is an equivariant process which assigns a positive *weight* to the edges of every deterministic graph, then one can let $\boldsymbol{d}'(u, v)$ be the minimum weight of the paths that connect u to v. If $\boldsymbol{d}'$ is a metric for almost every realization of G and is boundedly-finite a.s., then it is an equivariant metric.

3.7 Dimension of Subspaces

Let $[\boldsymbol{D}, \boldsymbol{o}]$ be a unimodular discrete space and $\boldsymbol{S}$ be an equivariant subset which is almost surely nonempty. Lemma 2.15 implies that $\mathbb{P}[\boldsymbol{o} \in \boldsymbol{S}_{\boldsymbol{D}}] > 0$. So one can consider $[\boldsymbol{S}_{\boldsymbol{D}}, \boldsymbol{o}]$ conditioned on $\boldsymbol{o} \in \boldsymbol{S}_{\boldsymbol{D}}$. By directly verifying the mass transport principle (2.1), it is easy to see that $[\boldsymbol{S}_{\boldsymbol{D}}, \boldsymbol{o}]$ conditioned on $\boldsymbol{o} \in \boldsymbol{S}_{\boldsymbol{D}}$ is unimodular (see the similar claim for unimodular graphs in [5]).

Convention 3.33. For an equivariant subset $\boldsymbol{S}$ as above, the unimodular Hausdorff dimension of $[\boldsymbol{S}_{\boldsymbol{D}}, \boldsymbol{o}]$ (conditioned on $\boldsymbol{o} \in S_{\boldsymbol{D}}$) is denoted by $\mathrm{udim}_H(\boldsymbol{S}_{\boldsymbol{D}})$. The same convention is used for the Minkowski dimension, the Hausdorff size, etc.

Theorem 3.34. *Let $[\boldsymbol{D}, \boldsymbol{o}]$ be a unimodular discrete space and $\boldsymbol{S}$ an equivariant subset such that $\boldsymbol{S}_{\boldsymbol{D}}$ is nonempty a.s. Then,*

(i) One has

$$\begin{aligned}
\operatorname{udim}_H(\boldsymbol{S_D}) &= \operatorname{udim}_H(\boldsymbol{D}), \\
\overline{\operatorname{udim}}_M(\boldsymbol{S_D}) &\geq \overline{\operatorname{udim}}_M(\boldsymbol{D}), \\
\underline{\operatorname{udim}}_M(\boldsymbol{S_D}) &\geq \underline{\operatorname{udim}}_M(\boldsymbol{D}).
\end{aligned}$$

(ii) If ρ is the intensity of $\boldsymbol{S}$ in $\boldsymbol{D}$, then for every $\alpha \geq 0$, the α-dim H-size of $\boldsymbol{S_D}$ satisfies

$$2^{-\alpha} \rho \, \mathcal{M}^\alpha(\boldsymbol{D}) \leq \mathcal{M}^\alpha(\boldsymbol{S_D}) \leq \rho \, \mathcal{M}^\alpha(\boldsymbol{D}).$$

Theorem 3.34 is proved below by using the fact that every covering of the larger set induces a covering of the subset by deleting some balls and then re-centering and enlarging the remaining balls. This matches the analogous idea in the continuum setting. The apparently surprising direction of the inequalities is due to the definition of dimension which implies that having less balls means having larger or equal dimension. For more on the matter, see the discussion on negative dimension in Subsection 8.3.

Remark 3.35. Subsection 3.8 below defines a modification $\mathcal{M}'_\alpha(\boldsymbol{D})$ of the unimodular Hausdorff size by considering coverings by arbitrary sets. With this definition, one has $\mathcal{M}'_\alpha(\boldsymbol{S_d}) = \rho \, \mathcal{M}'_\alpha(\boldsymbol{D})$. This can be proved similarly to Theorem 3.34. with the modification that there is no need to double the radii.

Remark 3.36. In the setting of Theorem 3.34, $\operatorname{udim}_M(\boldsymbol{S_D})$ can be strictly larger than $\operatorname{udim}_M(\boldsymbol{D})$ (see, e.g., Subsection 4.4). However, equality holds when $\boldsymbol{D}$ is regular (see Remark 3.23), which immediately follows from Theorems 3.22 and 3.34. Also, equality is guaranteed if $\boldsymbol{S_D}$ is a r-covering of $\boldsymbol{D}$ for some constant r. In other words, roughly speaking, the unimodular dimensions are *quasi-isometry invariant* (see e.g., [28]) and do not depend on the fine details of the discrete space.

Proof of Theorem 3.34. The first claim of (i) is implied by (ii) and Lemma 3.26, and hence, is skipped. Let $\boldsymbol{R}$ be an arbitrary equivariant r-covering of $\boldsymbol{D}$. For every $v \in \boldsymbol{R}$, let $\boldsymbol{\tau}(v)$ be an element picked uniformly at random in $N_r(v) \cap \boldsymbol{S_D}$, which is defined only when $N_r(v) \cap \boldsymbol{S_D} \neq \emptyset$. Let $\boldsymbol{R}' := \{\boldsymbol{\tau}(v) : v \in \boldsymbol{R}, N_r(v) \cap \boldsymbol{S_D} \neq \emptyset\}$ and note that $\boldsymbol{R}'$ is a $2r$-covering of $\boldsymbol{S_D}$. One has

$$\begin{aligned}
\mathbb{P}\left[\boldsymbol{o} \in \boldsymbol{R}'\right] &\leq \mathbb{E}\left[\sum_v \mathbf{1}_{\{v \in \boldsymbol{R}\}} \mathbf{1}_{\{\boldsymbol{\tau}(v) = \boldsymbol{o}\}}\right] \\
&= \mathbb{E}\left[\sum_v \mathbf{1}_{\{\boldsymbol{o} \in \boldsymbol{R}\}} \mathbf{1}_{\{\boldsymbol{\tau}(\boldsymbol{o}) = v\}}\right] \leq \mathbb{P}\left[\boldsymbol{o} \in \boldsymbol{R}\right],
\end{aligned}$$

where the equality is by the mass transport principle. This gives $\rho \lambda_{2r}(\boldsymbol{S_D}) \leq \lambda_r(\boldsymbol{D})$, which implies the claims regarding the Minkowski dimension.

Now, part (ii) is proved. The definition of $\mathcal{H}^\alpha_\infty(\boldsymbol{S_D})$ implies that there exists a sequence $\boldsymbol{R}_n$ of equivariant coverings of $\boldsymbol{S_D}$ such that $\boldsymbol{R}_n(\cdot) \in \{0\} \cup [n, \infty)$ for all $n = 1, 2, \dots$ and $\mathbb{E}\left[\boldsymbol{R}_n(\boldsymbol{o})^\alpha \mid \boldsymbol{o} \in \boldsymbol{S_D}\right] \to \mathcal{H}^\alpha_\infty(\boldsymbol{S_D})$. One may extend $\boldsymbol{R}_n$ to be defined on $\boldsymbol{D}$ by letting $\boldsymbol{R}_n(v) := 0$ for $v \in \boldsymbol{D} \setminus \boldsymbol{S_D}$. Let $\epsilon > 0$ be arbitrary and $\boldsymbol{B}_n \subseteq \boldsymbol{D}$ be the union of $N_{(1+\epsilon)\boldsymbol{R}_n}(v)$ for all $v \in \boldsymbol{D}$. Define $\boldsymbol{R}'_n(u) := (1+\epsilon)\boldsymbol{R}_n(u)$ for $u \in \boldsymbol{B}_n$ and $\boldsymbol{R}'_n(u) := 1/\epsilon$ for $u \notin \boldsymbol{B}_n$. It is clear that $\boldsymbol{R}'_n$ is an equivariant covering of $\boldsymbol{D}$. Also,

$$\begin{aligned}
\mathbb{E}\left[\boldsymbol{R}'_n(\boldsymbol{o})^\alpha\right] &= (1+\epsilon)^\alpha \mathbb{E}\left[\boldsymbol{R}_n(\boldsymbol{o})^\alpha\right] + \frac{1}{\epsilon^\alpha} \mathbb{P}\left[\boldsymbol{o} \notin \boldsymbol{B}_n\right] \\
&= \rho(1+\epsilon)^\alpha \mathbb{E}\left[\boldsymbol{R}_n(\boldsymbol{o})^\alpha \mid \boldsymbol{o} \in \boldsymbol{S_D}\right] + \frac{1}{\epsilon^\alpha} \mathbb{P}\left[\boldsymbol{o} \notin \boldsymbol{B}_n\right]. \quad (3.7)
\end{aligned}$$

Since the radii of the balls in $\boldsymbol{R}_n$ are at least n, one gets that $\boldsymbol{B}_n$ includes the ϵn-neighborhood of $\boldsymbol{S_D}$. Therefore, $\mathbb{P}\left[\boldsymbol{o} \notin \boldsymbol{B}_n\right] \leq \mathbb{P}\left[N_{\epsilon n}(\boldsymbol{o}) \cap \boldsymbol{S_D} = \emptyset\right]$. Since $\boldsymbol{S_D}$ is non-empty a.s., this in turn implies that $\mathbb{P}\left[\boldsymbol{o} \notin \boldsymbol{B}_n\right] \to 0$ as $n \to \infty$ (note that the events $N_{\epsilon n}(\boldsymbol{o}) \cap \boldsymbol{S_D} = \emptyset$ are nested and converge to the event $\boldsymbol{S_D} = \emptyset$). So (3.7) implies that

$$\liminf_{n \to \infty} \mathbb{E}\left[\boldsymbol{R}_n'(\boldsymbol{o})^\alpha\right] = \rho(1+\epsilon)^\alpha \liminf_{n \to \infty} \mathbb{E}\left[\boldsymbol{R}_n(\boldsymbol{o})^\alpha \,|\, \boldsymbol{o} \in \boldsymbol{S_D}\right] = \rho(1+\epsilon)^\alpha \mathcal{H}_\infty^\alpha(\boldsymbol{S_D}).$$

Note that the radii of the balls in $\boldsymbol{R}_n'$ are at least $n \wedge (1/\epsilon)$. Therefore, one obtains $\mathcal{H}_{1/\epsilon}^\alpha(\boldsymbol{D}) \leq \rho(1+\epsilon)^\alpha \mathcal{H}_\infty^\alpha(\boldsymbol{S_D})$. By letting $\epsilon \to 0$, one gets $\mathcal{H}_\infty^\alpha(\boldsymbol{D}) \leq \rho \mathcal{H}_\infty^\alpha(\boldsymbol{S_D})$; i.e., $\mathcal{M}^\alpha(\boldsymbol{S_D}) \leq \rho \mathcal{M}^\alpha(\boldsymbol{D})$.

Conversely, let $\boldsymbol{R}_n$ be a sequence of equivariant coverings of $\boldsymbol{D}$ for $n = 1, 2, \ldots$ such that $\boldsymbol{R}_n(\cdot) \in \{0\} \cup [n, \infty)$ a.s. and $\mathbb{E}\left[\boldsymbol{R}_n(\boldsymbol{o})^\alpha\right] \to \mathcal{H}_\infty^\alpha(\boldsymbol{D})$. Fix n in the following. Let $\boldsymbol{B} := \boldsymbol{B_D} := \{v : N_{\boldsymbol{R}_n}(v) \cap \boldsymbol{S_D} \neq \emptyset\}$. For each $v \in \boldsymbol{B}$, let $\boldsymbol{\tau}_n(v)$ be an element chosen uniformly at random in $N_{\boldsymbol{R}_n}(v) \cap \boldsymbol{S_D}$. For $v \notin \boldsymbol{B}$, let $\boldsymbol{\tau}_n(v)$ be undefined. For $w \in \boldsymbol{S_D}$, let $\boldsymbol{R}_n'(w) := 2\max\{\boldsymbol{R}_n(v) : v \in \boldsymbol{\tau}_n^{-1}(w)\}$. It can be seen that $\boldsymbol{R}_n'$ is an equivariant covering of $\boldsymbol{S_D}$. One has

$$
\begin{aligned}
\mathbb{E}\left[\boldsymbol{R}_n'(\boldsymbol{o})^\alpha\right] &\leq 2^\alpha \mathbb{E}\left[\sum_v \boldsymbol{R}_n(v)^\alpha 1_{\{v \in \boldsymbol{\tau}_n^{-1}(\boldsymbol{o})\}}\right] \\
&= 2^\alpha \mathbb{E}\left[\sum_v \boldsymbol{R}_n(\boldsymbol{o})^\alpha 1_{\{\boldsymbol{o} \in \boldsymbol{\tau}_n^{-1}(v)\}}\right] \leq 2^\alpha \mathbb{E}\left[\boldsymbol{R}_n(\boldsymbol{o})^\alpha\right],
\end{aligned}
$$

where the equality is by the mass transport principle. It follows that

$$\rho \liminf_{n \to \infty} \mathbb{E}\left[\boldsymbol{R}_n'(\boldsymbol{o})^\alpha \,|\, \boldsymbol{o} \in \boldsymbol{S_D}\right] \leq 2^\alpha \mathcal{H}_\infty^\alpha(\boldsymbol{D}).$$

So $\rho \mathcal{H}_\infty^\alpha(\boldsymbol{S_D}) \leq 2^\alpha \mathcal{H}_\infty^\alpha(\boldsymbol{D})$. Hence, $\mathcal{M}^\alpha(\boldsymbol{S_D}) \geq 2^{-\alpha} \rho \, \mathcal{M}^\alpha(\boldsymbol{D})$ and the claim is proved. $\quad\square$

3.8 Covering By Arbitrary Sets

According to Remark 3.35, it is more natural to redefine the Hausdorff size by considering coverings by finite subsets which are not necessarily balls (as in the continuum setting). A technical challenge is to define such coverings in an equivariant way. This will be done at the end of this subsection using the notion of equivariant processes of Subsection 2.5. Once an equivariant covering $\boldsymbol{C}$ is defined (which is an equivariant collection of finite subsets), one can define the *average diameter of sets $U \in \boldsymbol{C}$ per point* by

$$\mathbb{E}\left[\sum_{U \in \boldsymbol{C}} \frac{1}{\#U} 1_{\{\boldsymbol{o} \in U\}} \operatorname{diam}(U)\right].$$

The same idea is used to redefine $\mathcal{H}_M^\alpha(\boldsymbol{D})$ as follows:

$$\mathcal{H}_{\alpha, M}'(\boldsymbol{D}) := \inf_{\boldsymbol{C}} \left\{ \mathbb{E}\left[\sum_{U \in \boldsymbol{C}} \frac{1}{\#U} 1_{\{\boldsymbol{o} \in U\}} \left(M \vee \frac{1}{2}\operatorname{diam}(U)\right)^\alpha\right] \right\},$$

where the infimum is over all equivariant coverings $\boldsymbol{C}$. Here, taking the maximum with M is *similar* to the condition that the subsets have diameter at least $2M$ (note that a ball of radius M might have diameter strictly less than $2M$). Finally, define the **modified unimodular Hausdorff size** $\mathcal{M}_\alpha'(\boldsymbol{D})$ similarly to (3.6). Remark 3.35 shows an advantage of this definition. Also, the reader can verify that $2^{-\alpha}\mathcal{H}_{2M}^\alpha(\boldsymbol{D}) \leq \mathcal{H}_{\alpha, M}'(\boldsymbol{D}) \leq \mathcal{H}_M^\alpha(\boldsymbol{D})$. Therefore,

$$\mathcal{M}^\alpha(\boldsymbol{D}) \leq \mathcal{M}_\alpha'(\boldsymbol{D}) \leq 2^\alpha \mathcal{M}^\alpha(\boldsymbol{D}).$$

This implies that the notion of unimodular Hausdorff dimension is not changed by this modification. One can also obtain a similar equivalent form of the unimodular Minkowski dimension. This is done by redefining λ_r by considering equivariant coverings by sets of diameter at most $2r$. The details are left to the reader. A similar idea will be used in Subsection 4.1.2 to calculate the Minkowski dimension of one-ended trees.

Finally, here is the promised representation of the above coverings as equivariant processes (it should be noted that it is not always possible to number the subsets in an equivariant way and the collection should be necessarily unordered). To show the idea, consider a covering $C = \{U_1, U_2, \ldots\}$ of a deterministic discrete space D, where each U_i is bounded. For each U_i, assign the mark $(\boldsymbol{X}_i, \mathrm{diam}(U_i))$ to every point of U_i, where $\boldsymbol{X}_i \in [0, 1]$ is chosen i.i.d. and uniformly. Note that multiple marks are assigned to every point and the covering can be reconstructed from the marks. With this idea, let the mark space Ξ be the set of discrete subsets of $\mathbb{R}^2$ (regard every discrete set as a counting measure and equip Ξ with a metrization of the vague topology). This mark space can be used to represent equivariant coverings by equivariant processes (for having a complete mark space, one can extend Ξ to the set of discrete multi-sets in $\mathbb{R}^2$).

3.9 Notes and Bibliographical Comments

Several definitions and basic results of this section have analogues in the continuum setting. A list of such analogies is given below. Note however that there is no systematic way of translating the results in the continuum setting to that of unimodular discrete spaces. In particular, inequalities are most often, but not always, in the other direction. The comparison of the unimodular Minkowski and Hausdorff dimensions (Theorem 3.22) is analogous to the similar comparison in the continuum setting (see e.g., (1.2.3) of [14]), but in the reverse direction. Theorem 3.31, regarding changing the metric, is analogous to the fact that the ordinary Minkowski and Hausdorff dimensions are not increased by applying a Lipschitz function. Theorem 3.34 regarding the dimension of subsets is analogous to the fact that the ordinary dimensions do not increase by passing to subsets. Note however that equality holds in Theorem 3.34 for the unimodular Hausdorff dimension (and also for the unimodular Minkowski dimension in most usual examples), in contrast to the continuum setting.

For point processes (Example 2.7), one can redefine the unimodular Hausdorff dimension by using *dyadic cubes* instead of balls. This changes the value of the Hausdorff size up to a constant factor, and hence, the value of Hausdorff dimension is not changed. Since dyadic cubes are nested, this simplifies some of the arguments. This approach will be used in Subsection 7.3.

4 Examples

This section presents a set of examples of unimodular discrete spaces together with discussions about their dimensions. Recall that the tools for bounding the dimensions are summarized in Remarks 3.5 and 3.18. As mentioned in Remark 3.18, bounding the Hausdorff dimension from above usually requires the unimodular mass transport principle or the unimodular Billingsley lemma, which will be stated in Section 5. So the upper bounds for some of the following examples are completed later in Subsection 6.1.

4.1 General Unimodular Trees

In this subsection, general results are presented regarding the dimension of unimodular trees with the graph-distance metric. Specific instances are presented later in the section. It turns out that the number of *ends* of the tree plays a key role (an **end** in a tree is an equivalence class of simple paths in the tree, where two such paths are equivalent

if their symmetric difference is finite).

It is well known that the number of ends in a unimodular tree belongs to $\{0, 1, 2, \infty\}$ [2]. Unimodular trees without end are finite, and hence, are zero dimensional (Example 3.16). The only thing to mention is that there exists an algorithm to construct an optimal n-covering for such trees. This algorithm is similar to the algorithm for one-ended trees, discussed below, and is skipped for brevity. The case of infinitely many ends will be studied in Subsection 6.2. The remaining two cases are discussed below.

4.1.1 Unimodular Two-Ended Trees

If T is a tree with two ends, then there is a unique bi-infinite path in T called its **trunk**. Moreover, each connected component of the complement of the trunk is finite.

Theorem 4.1. *For all unimodular two-ended trees $[\boldsymbol{T}, \boldsymbol{o}]$ endowed with the graph-distance metric, one has $\mathrm{udim}_M(\boldsymbol{T}) = \mathrm{udim}_H(\boldsymbol{T}) = 1$. Moreover, if ρ is the intensity of the trunk of $\boldsymbol{T}$, then the modified 1-dim H-size of $\boldsymbol{T}$ is $\mathcal{M}'_1(\boldsymbol{T}) = 2\rho^{-1}$.*

Proof. For all two-ended trees T, let $\boldsymbol{S}_T$ be the trunk of T. Then, $\boldsymbol{S}$ is an equivariant subset. Therefore, Theorem 3.34 implies that $\mathrm{udim}_H(\boldsymbol{T}) = \mathrm{udim}_H(\boldsymbol{S_T})$. Since the trunk is isometric to $\mathbb{Z}$ as a metric space, Example 3.16 implies that $\mathrm{udim}_H(\boldsymbol{T}) = 1$. In addition, Remark 3.35 and Proposition 3.29 imply that $\mathcal{M}'_1(\boldsymbol{T}) = \rho^{-1}\mathcal{M}'_1(\mathbb{Z}) = 2\rho^{-1}$.

The claim concerning the unimodular Minkowski dimension is implied by Corollary 5.10 of the next section, which shows that any unimodular infinite graph satisfies $\underline{\mathrm{udim}}_M(\boldsymbol{G}) \geq 1$ (this theorem will not be used throughout). $\square$

4.1.2 Unimodular One-Ended Trees

Unimodular one-ended trees arise naturally in many examples (see [2]). In particular, the (local weak) limit of many interesting sequences of finite trees/graphs are one-ended ([3, 2]). In terms of unimodular dimensions, it will be shown that unimodular one-ended trees are the richest class of unimodular trees.

First, the following notation is borrowed from [5]. Every one-ended tree T can be regarded as a family tree as follows. For every vertex $v \in T$, there is a unique infinite simple path starting from v. Denote by $F(v)$ the next vertex in this path and call it the **parent** of v. By deleting $F(v)$, the connected component containing v is finite. This set is denoted by $D(v)$ and its elements are called the **descendants** of v. The maximum distance of v to its descendants will be called the **height** of v and be denoted by $h(v)$.

Theorem 4.2. *If $[\boldsymbol{T}, \boldsymbol{o}]$ is a unimodular one-ended tree endowed with the graph-distance metric, then*

$$\overline{\mathrm{udim}}_M(\boldsymbol{T}) = 1 + \overline{\mathrm{decay}}\left(\mathbb{P}\left[h(\boldsymbol{o}) \geq n\right]\right), \tag{4.1}$$

$$\underline{\mathrm{udim}}_M(\boldsymbol{T}) = 1 + \underline{\mathrm{decay}}\left(\mathbb{P}\left[h(\boldsymbol{o}) \geq n\right]\right). \tag{4.2}$$

In addition,

$$\mathrm{udim}_H(\boldsymbol{T}) \geq \overline{\mathrm{decay}}\left(\mathbb{P}\left[h(\boldsymbol{o}) = n\right]\right) \geq \overline{\mathrm{udim}}_M(\boldsymbol{T}). \tag{4.3}$$

It is not known whether the first inequality in (4.3) is always an equality. It should also be noted that $\overline{\mathrm{decay}}\left(\mathbb{P}\left[h(\boldsymbol{o}) = n\right]\right)$ can be strictly larger than $1 + \overline{\mathrm{decay}}\left(\mathbb{P}\left[h(\boldsymbol{o}) \geq n\right]\right)$ (see e.g., Subsection 4.2.2), however, they are equal in most *usual* examples.

The proof of Theorem 4.2 is based on a recursive construction of an optimal covering by *cones*, defined below, rather than balls. It is shown below that considering cones instead of balls does not change the Minkowski dimension. An optimal ball-covering is also discussed in [7].

The **cone** with height n at $v \in \boldsymbol{T}$ is defined by $C_n(v) := N_n(v) \cap D(v)$; i.e., the first n generations of the descendants of v, including v itself. Let λ'_n be the infimum intensity of equivariant coverings by cones of height n. The claim is that

$$\lambda'_{2n} \leq \lambda_n \leq \lambda'_n. \tag{4.4}$$

This immediately implies that

$$\underline{\mathrm{udim}}_M(\boldsymbol{T}) = \underline{\mathrm{decay}}\left(\lambda'_n\right), \quad \overline{\mathrm{udim}}_M(\boldsymbol{T}) = \overline{\mathrm{decay}}\left(\lambda'_n\right). \tag{4.5}$$

To prove (4.4), note that any covering by cones of height n is also a covering by balls of radius n. This implies that $\lambda_n \leq \lambda'_n$. Also, if $\boldsymbol{S}$ is a covering by balls of radius n, then $\{F^n(v) : v \in \boldsymbol{S}\}$ is a covering by cones of height $2n$. By the mass transport principle (2.2), one can show that the intensity of the latter is not greater than the intensity of $\boldsymbol{S}$. This implies that $\lambda'_{2n} \leq \lambda_n$. So (4.4) is proved.

Lemma 4.3. *For every unimodular one-ended tree $[\boldsymbol{T}, \boldsymbol{o}]$, the output $\boldsymbol{S}$ of the following greedy algorithm is an optimal equivariant covering of $\boldsymbol{T}$ by cones of height n.*

> $\boldsymbol{S} := \emptyset$;
> **while** *true* **do**
> > *Add all vertices of height n in $\boldsymbol{T}$ to $\boldsymbol{S}$;*
> > $\boldsymbol{T} := \boldsymbol{T} \setminus \bigcup_{v \in \boldsymbol{S}} D(v)$;
> **end**

Note that the algorithm does not finish in finite time, but for each vertex v of $\boldsymbol{T}$, it is determined in finite time whether a cone is put at v or not. So the output of the algorithm is well defined.

Proof. Let $\boldsymbol{A}$ be any equivariant covering of $\boldsymbol{T}$ by cones of height n. Consider a realization $(T; A)$ of $[\boldsymbol{T}; \boldsymbol{A}]$. Let v be a vertex such that $h(v) = n$. Since A is a covering by cones of height n, A should have at least one vertex in $D(v)$ (to see this, consider the farthest leaf from v in $D(v)$). Now, for all such vertices v, delete the vertices in $A \cap D(v)$ from A and then add v to A. Let A_1 be the subset of T obtained by doing this operation for all vertices v of height n. So A_1 is also a covering of T by cones of height n. Now, remove all vertices $\{v : h(v) = n\}$ and their descendants from T to obtain a new one-ended tree. Consider the same procedure for the remaining tree and its intersection with A. Inductively, one obtains a sequence of subsets $A = A_0, A_1, \ldots$ of T such that, for each i, A_i is a covering of T by cones of height n which agrees with $\boldsymbol{S}_T$ on the set of vertices that are removed from the tree up to step i.

By letting $[\boldsymbol{T}; \boldsymbol{A}]$ be random, the above induction gives a sequence of equivariant subsets $\boldsymbol{A} = \boldsymbol{A}_0, \boldsymbol{A}_1, \ldots$ on $\boldsymbol{T}$. It can be seen that the intensity of $\boldsymbol{A}_1$ is at most that of $\boldsymbol{A}$ (this can be verified by the mass transport principle (2.1)) and more generally, the intensity of $\boldsymbol{A}_{i+1}$ is at most that of $\boldsymbol{A}_i$; i.e., $\mathbb{P}\left[\boldsymbol{o} \in \boldsymbol{A}_{i+1}\right] \leq \mathbb{P}\left[\boldsymbol{o} \in \boldsymbol{A}_i\right]$. Also, $\lim_{i \to \infty} \boldsymbol{A}_i = \boldsymbol{S}$ as equivariant subsets of $\boldsymbol{T}$. This implies that $\mathbb{P}\left[\boldsymbol{o} \in \boldsymbol{A}\right] \geq \mathbb{P}\left[\boldsymbol{o} \in \boldsymbol{S}\right]$, hence, $\boldsymbol{S}$ is an optimal covering by cones of height n. $\square$

Lemma 4.4. *Under the above setting, one has*

$$\mathbb{P}\left[h(\boldsymbol{o}) \bmod (n+1) = -1\right] \leq \lambda'_n \leq \mathbb{P}\left[h(\boldsymbol{o}) \bmod \left\lfloor \frac{n}{2}+1 \right\rfloor = -1\right]. \tag{4.6}$$

Proof. The proof of the second inequality in (4.6) is based on the construction of the following equivariant covering. Let $B_n := \{v \in \boldsymbol{T} : h(v) \bmod n = -1\}$ and $B'_n := \{F^{n-1}(v) : v \in B_n\}$. The claim is that B'_n is a covering of $\boldsymbol{T}$ by cones of height $2n - 2$. Let $v \in \boldsymbol{T}$ be an arbitrary vertex. Let k be the unique integer such that $(k-1)n - 1 <$

$h(v) \leq kn - 1$. Let j be the first nonnegative integer such that $h(F^j(v)) \geq kn - 1$ and let $w := F^j(v)$. One has $0 \leq j \leq n-1$. By considering the longest path in $D(w)$ from w to the leaves, one finds $z \in D(w)$ such that $h(z) \bmod n = -1$ and $0 \leq d(w, z) \leq n - 1$. Therefore w (and hence v) is a descendant of $F^{n-1}(z)$. Also, $d(w, F^{n-1}(z)) \leq n - 1$. It follows that $d(v, F^{n-1}(z)) \leq 2n - 2$. So v is covered by the cone of height $2n - 2$ at $F^{n-1}(z)$. Since $F^{n-1}(z) \in B'_n$, it is proved that B'_n gives a $(2n - 2)$-covering by cones. It follows that $\lambda'_{2n-2} \leq \mathbb{P}\left[o \in B'_n\right] \leq \mathbb{P}\left[o \in B_n\right]$ (where the last inequality can be verified by the mass transport principle (2.1)). This implies the second inequality in (4.6).

To prove the first inequality in (4.6), let S be the optimal covering by cones of height n given by the algorithm of Lemma 4.3. Send unit mass from each vertex $v \in S$ to the first vertex in $v, F(v), \ldots, F^n(v)$ which belongs to B_{n+1} (if there is any). So the outgoing mass from v is at most $1_{\{v \in S\}}$. In the next paragraph, it is proved that the incoming mass to each $w \in B_{n+1}$ is at least 1. This in turn (by the mass transport principle) implies that $\mathbb{P}\left[o \in S\right] \geq \mathbb{P}\left[o \in B_{n+1}\right]$, which proves the first inequality in (4.6).

The final step consists in proving that the incoming mass to each $w \in B_{n+1}$ is at least 1. If $h(w) = n$, then $w \in S$ and the claim is proved. So assume $h(w) > n$. By considering the longest path from w in $D(w)$, one can find a vertex z such that $w = F^{n+1}(z)$ and $h(z) = h(w) - (n + 1)$. This implies that no vertex in $\{F(z), \ldots, F^n(z)\}$ is in B_{n+1}. So to prove the claim, it suffices to show that at least one of these vertices or w itself lies in S. Note that in the algorithm in Lemma 4.3, at each step, the height of w decreases by a value at least 1 and at most $n + 1$ until w is removed from the tree. So in the last step before w is removed, the height of w is in $\{0, 1, \ldots, n\}$. This is possible only if in the same step of the algorithm, an element of $\{F(z), \ldots, F^n(z), w\}$ is added to S. This implies the claim and the lemma is proved. $\qquad\square$

Now, the tools needed to prove the main results are all available.

Proof of Theorem 4.2. Lemma 4.4 and (4.5) imply that the upper and lower Minkowski dimensions of T are exactly the upper and lower decay rates of $\mathbb{P}\left[h(o) \bmod n = -1\right]$ respectively. So one should prove that these rates are equal to the upper and lower decay rates of $\mathbb{P}\left[h(o) \geq n\right]$ plus 1.

The first step consists in showing that $\mathbb{P}\left[h(o) = n\right]$ is non-increasing in n. To see this, send unit mass from each vertex v to $F(v)$ if $h(v) = n$ and $h(F(v)) = n + 1$. Then the outgoing mass is at most $1_{\{h(v)=n\}}$ and the incoming mass is at least $1_{\{h(v)=n+1\}}$. The result then follows by the mass transport principle. This monotonicity implies that $n \cdot \mathbb{P}\left[h(o) \bmod n = -1\right] \geq \mathbb{P}\left[h(o) \geq n - 1\right]$. Similarly, by monotonicity,

$$
\begin{aligned}
\frac{n}{2} \mathbb{P}\left[h(o) \bmod n = -1\right] &\leq \mathbb{P}\left[h(o) \bmod n \in \left\{-1, -2, \ldots, -\left\lceil \frac{n}{2} \right\rceil\right\}\right] \\
&\leq \mathbb{P}\left[h(o) \geq \left\lfloor \frac{n}{2} \right\rfloor\right].
\end{aligned}
$$

These inequalities conclude the proof of (4.1) and (4.2).

It remains to prove (4.3). The second inequality follows from (4.1) and the fact that $\overline{\operatorname{decay}}\left(\mathbb{P}\left[h(o) = n\right]\right) \geq \overline{\operatorname{decay}}\left(\mathbb{P}\left[h(o) \geq n\right]\right) + 1$, which is not hard to see. We now prove the first inequality. Fix $0 < \epsilon < \alpha < \overline{\operatorname{decay}}\left(\mathbb{P}\left[h(o) = n\right]\right)$. So there is a sequence $n_1 < n_2 < \cdots$ such that $\mathbb{P}\left[h(o) = n_i\right] < n_i^{-\alpha}$ for each i. One may assume the sequence is such that $n_i \geq 2^i$ for each i. Now, for each $k \in \mathbb{N}$, consider the following covering of T:

$$
\boldsymbol{R}_k(v) := \begin{cases} 2(n_i - n_{i-1}), & \text{if } h(v) = n_i \text{ and } i > k, \\ 2n_k, & \text{if } h(v) = n_k, \\ 0, & \text{otherwise.} \end{cases}
$$

By arguments similar to Lemma 4.4, it can be seen that $\boldsymbol{R}_k$ is indeed a (ball-)covering. It is claimed that $\mathbb{E}\left[\boldsymbol{R}_k(\boldsymbol{o})^{\alpha-\epsilon}\right] \to 0$ as $k \to \infty$ If the claim is proved, then $\mathrm{udim}_H(\boldsymbol{T}) \geq \alpha - \epsilon$ and the proof of (4.3) is concluded. Let $c := 2^{\alpha-\epsilon}$. One has

$$
\begin{aligned}
\mathbb{E}\left[\boldsymbol{R}_k(\boldsymbol{o})^{\alpha-\epsilon}\right] &= c n_k^{\alpha-\epsilon} \mathbb{P}\left[h(\boldsymbol{o}) = n_k\right] + c \sum_{i=k+1}^{\infty} (n_i - n_{i-1})^{\alpha-\epsilon} \mathbb{P}\left[h(\boldsymbol{o}) = n_i\right] \\
&\leq c n_k^{-\epsilon} + c \sum_{i=k+1}^{\infty} (n_i - n_{i-1})^{\alpha-\epsilon} n_i^{-\alpha}.
\end{aligned}
$$

Therefore, it is enough to prove that

$$
\sum_{i=1}^{\infty} (n_i - n_{i-1})^{\alpha-\epsilon} n_i^{-\alpha} < \infty. \tag{4.7}
$$

It is easy to see that the maximum of the function $(x - n_{i-1})^{\alpha-\epsilon} x^{-\alpha}$ over $x \geq n_{i-1}$ happens at $\frac{\alpha}{\epsilon} n_{i-1}$ and the maximum value is $c' n_{i-1}^{-\epsilon}$, where $c' = (\frac{\alpha}{\epsilon} - 1)^{\alpha-\epsilon}$ is a constant. So the left hand side of (4.7) is at most $c' \sum_{i=0}^{\infty} n_i^{-\epsilon}$, which is finite by the assumption $n_i \geq 2^i$. So (4.7) is proved and the proof is completed. $\qquad\square$

4.2 Instances of Unimodular Trees

This subsection discusses the dimension of some explicit unimodular trees. More examples are given in Subsection 4.5, in Section 6, and also in the ongoing work [6] (e.g., uniform spanning forests).

4.2.1 The Canopy Tree

The canopy tree C_k with offspring cardinality k [1] is constructed as follows. Its vertex set is partitioned in levels $L_0, L_1, \ldots$. Each vertex in level n is connected to k vertices in level $n-1$ (if $n \neq 0$) and one vertex (its parent) in level $n+1$. Let $\boldsymbol{o}$ be a random vertex of C_k such that $\mathbb{P}\left[\boldsymbol{o} \in L_n\right]$ is proportional to k^{-n}. Then, $[C_k, \boldsymbol{o}]$ is a unimodular random tree.

Below, three types of metrics are considered on C_k. First, consider the graph-distance metric. Given $n \in \mathbb{N}$, let $S := \{v \in C_k : h(v) \geq n\}$, where $h(v)$ is the height of v defined in Subsection 4.1.2. The set S gives an equivariant n-covering and $\mathbb{P}\left[\boldsymbol{o} \in S\right]$ is exponentially small as $n \to \infty$. So $\mathrm{udim}_M(C_k) = \mathrm{udim}_H(C_k) = \infty$.

Second, for each n, let the length of each edge between L_n and L_{n+1} be a^n, where $a > 1$ is constant. Let d_1 be the resulting metric on C_k. Given $r > 0$, let S_1 be the set of vertices having distance at least r/a to L_0 (under d_1). One can show that S_1 is an r-covering of (C_k, d_1) and $\mathrm{decay}\left(\mathbb{P}\left[\boldsymbol{o} \in S_1\right]\right) = \log k / \log a$. Therefore, $\underline{\mathrm{udim}}_M(C_k, d_1) \geq \log k / \log a$. On the other hand, one can see that the ball of radius a^n centered at $\boldsymbol{o}$ (under d_1) has cardinality of order k^n. One can then use Lemma 3.17 to show that $\mathrm{udim}_H(C_k, d_1) \leq \log k / \log a$. So $\mathrm{udim}_M(C_k, d_1) = \mathrm{udim}_H(C_k, d_1) = \log k / \log a$.

Third, replace a^n by $n!$ in the second case and let d_2 be the resulting metric. Then, the cardinality of the ball of radius r centered at $\boldsymbol{o}$ has order less than r^{α} for every $\alpha > 0$. One can use Lemma 3.17 again to show that $\mathrm{udim}_H(C_k, d_2) \leq \alpha$. This implies that $\mathrm{udim}_M(C_k, d_2) = \mathrm{udim}_H(C_k, d_2) = 0$.

4.2.2 The Generalized Canopy Tree

This example generalizes the canopy tree of Subsection 4.2.1. The goal is to provide an example where the lower Minkowski dimension, the upper Minkowski dimension and the Hausdorff dimension are all different when suitable parameters are chosen.

Fix $p_0, p_1, \ldots > 0$ such that $\sum p_i = 1$. Let $\boldsymbol{U}_0, \boldsymbol{U}_1, \ldots$ be an i.i.d. sequence of random number in $[0, 1]$ with the uniform distribution. For each $n \geq 0$, let $\Phi_n := \left(\frac{1}{p_n} (\mathbb{Z} + \boldsymbol{U}_n) \right) \times \{n\}$, which is a point process on the horizontal line $y = n$ in the plane. Let $\boldsymbol{o}_n := \left(\frac{1}{p_n} \boldsymbol{U}_n, n \right) \in \Phi_n$ and $\Phi := \cup_i \Phi_i$. Then, Φ is a point process in the plane which is stationary under horizontal translations. Choose $\boldsymbol{m}$ independent of the sequence $(\boldsymbol{U}_i)_i$ such that $\mathbb{P}[\boldsymbol{m} = n] = p_n$ for each n. Then, let $\boldsymbol{o} := \boldsymbol{o_m}$.

Construct a graph $\boldsymbol{T}$ on Φ as follows: For each n, connect each $x \in \Phi_n$ to its closest point (or closest point on its right) in Φ_{n+1}. Note that $\boldsymbol{T}$ is a forest by definition. However, the next lemma shows that $[\boldsymbol{T}, \boldsymbol{o}]$ is a unimodular tree.

Definition 4.5. *The **generalized canopy tree** with parameters $p_0, p_1, \ldots$ is the unimodular tree $[\boldsymbol{T}, \boldsymbol{o}]$ constructed above.*

Note that in the case where p_n is proportional to k^{-n} for k fixed, $[\boldsymbol{T}, \boldsymbol{o}]$ is just the ordinary canopy tree C_k of Subsection 4.2.1. Also, one can generalize the above construction by letting Φ_n be a sequence of point processes which are (jointly) stationary under horizontal translations.

Lemma 4.6. *One has*

(i) $[\Phi, \boldsymbol{o}]$, endowed with the Euclidean metric, is a unimodular discrete space.

(ii) $\boldsymbol{T}$ is a tree a.s. and $[\boldsymbol{T}, \boldsymbol{o}]$ is unimodular.

Proof. For part (i), it is enough to show that $\Phi - \boldsymbol{o}$ is a point-stationary point process in the plane (see Example 2.7). This is skipped for brevity (see [7]). The main ingredients are using stationarity of Φ under horizontal translations and the fact that $\Phi_n - \boldsymbol{o}_n$ is point-stationary (the proof is similar to that of the formula for the Palm version of the superposition of stationary point processes, e.g., in [47].)

To prove (ii), note that $\boldsymbol{T}$ can be realized as an equivariant process on Φ (see Definition 2.9 and Remark 2.13). Therefore, by Lemma 2.12 and Theorem 3.31, it is enough to prove that $\boldsymbol{T}$ is connected a.s. Nevertheless, the same lemma implies that the connected component $\boldsymbol{T}'$ of $\boldsymbol{T}$ containing $\boldsymbol{o}$ is a unimodular tree. Since it is one-ended, Theorem 3.9 of [5] implies that the *foils* $\boldsymbol{T}' \cap \Phi_i$ are infinite a.s. By noting that the edges do not cross (as segments in the plane), one obtains that $\boldsymbol{T}' \cap \Phi_i$ should be the whole Φ_i; hence, $\boldsymbol{T}' = \boldsymbol{T}$. Therefore, $\boldsymbol{T}$ is connected a.s. and the claim is proved. $\square$

Proposition 4.7. *The sequence $(p_n)_n$ can be chosen such that*

$$\underline{\mathrm{udim}}_M(\boldsymbol{T}) < \overline{\mathrm{udim}}_M(\boldsymbol{T}) < \mathrm{udim}_H(\boldsymbol{T}),$$

where $\boldsymbol{T}$ is endowed with the graph-distance metric. Moreover, for any $0 \leq \alpha \leq \beta \leq \gamma \leq \infty$, the sequence $(p_n)_n$ can be chosen such that

$$\underline{\mathrm{udim}}_M(\boldsymbol{T}) \leq \alpha, \quad \overline{\mathrm{udim}}_M(\boldsymbol{T}) = \beta, \quad \mathrm{udim}_H(\boldsymbol{T}) \geq \gamma.$$

For example, it is possible to have $\mathrm{udim}_M(\boldsymbol{T}) = 0$ and $\mathrm{udim}_H(\boldsymbol{T}) = \infty$ simultaneously.

Proof. $\boldsymbol{T}$ is a one-ended tree (see Subsection 4.1.2). Assume the sequence $(p_n)_n$ is non-increasing. So the construction implies that there is no leaf of the tree in Φ_n for all $n > 0$. Therefore, for all $n \geq 0$, the height of every vertex in Φ_n is precisely n. So by letting $q_n := \sum_{i \geq n} p_i$, Theorem 4.2 implies that

$$
\begin{aligned}
\mathrm{udim}_H(\boldsymbol{T}) &\geq \overline{\mathrm{decay}}(p_n), \\
\overline{\mathrm{udim}}_M(\boldsymbol{T}) &= 1 + \overline{\mathrm{decay}}(q_n), \\
\underline{\mathrm{udim}}_M(\boldsymbol{T}) &= 1 + \underline{\mathrm{decay}}(q_n).
\end{aligned}
$$

For simplicity, assume $0 < \alpha$ and $\gamma < \infty$ (the other cases can be treated similarly). Define $n_0, n_1, \dots$ recursively as follows. Let $n_0 := 0$. Given that n_i is defined, let n_{i+1} be large enough such that the line connecting points $(n_i, n_i^{-\beta})$ and $(n_{i+1}, n_{i+1}^{-\beta})$ intersects the graph of the function $x^{-\alpha}$ and has slope larger than $-n^{-\gamma}$. Now, let $q_{n_i} := n_i^{-\beta}$ for each i and define q_n linearly in the interval $[n_i, n_{i+1}]$. Let $p_n := q_n - q_{n+1}$. It can be seen that p_n is non-increasing, $\underline{\text{decay}}(q_n) \le \alpha$, $\overline{\text{decay}}(q_n) = \beta$ and $\overline{\text{decay}}(p_n) \ge \gamma$. $\qquad\square$

4.2.3 Unimodular Eternal Galton-Watson Trees

Eternal Galton-Watson (EGW) trees are defined in [5]. Unimodular EGW trees (in the nontrivial case) can be characterized as unimodular one-ended trees in which the descendants of the root constitute a Galton-Watson tree. Also, unimodularity implies that the latter Galton-Watson tree is necessarily critical (use the mass transport principle when sending a unit mass from each vertex to its *parent*). Here, the trivial case that each vertex has exactly one offspring is excluded (where the corresponding EGW tree is a bi-infinite path). In particular, the *Poisson skeleton tree* [3] is an eternal Galton-Watson tree.

Recall that the offspring distribution of a Galton-Watson tree is the probability measure $(p_0, p_1, \dots)$ on $\mathbb{Z}^{\ge 0}$ where p_n is the probability that the root has n offsprings.

Proposition 4.8. *Let $[\boldsymbol{T}, \boldsymbol{o}]$ be a unimodular eternal Galton-Watson tree. If the offspring distribution has finite variance, then $\text{udim}_M(\boldsymbol{T}) = \text{udim}_H(\boldsymbol{T}) = 2$.*

Proof (first part). Here, it is only proved that $\text{udim}_M(\boldsymbol{T}) = 2$. The other equality will be proved in Subsection 6.1. By Kesten's theorem [33] for the Galton-Watson tree formed by the descendants of the root, $\lim_n n\mathbb{P}[h(\boldsymbol{o}) \ge n]$ exists and is positive. It follows that $\text{decay}(\mathbb{P}[h(\boldsymbol{o}) \ge n]) = 1$. So the claim is implied by Theorem 4.2. $\qquad\square$

4.3 Examples Associated with Random Walks

Let μ be a probability measure on $\mathbb{R}^k$. Consider the (double-sided) simple random walk $(S_n)_{n\in\mathbb{Z}}$ in $\mathbb{R}^k$ starting from $S_0 := 0$ such that and the jumps $S_n - S_{n-1}$ are i.i.d. with distribution μ. In this subsection, unimodular discrete spaces are constructed based on the image and the zero set of this random walk and their dimensions are studied in some special cases. The graph of the simple random walk will be studied in Subsection 6.4.

4.3.1 The Image of the Simple Random Walk

Assume the random walk is transient; i.e., it visits every given ball only finitely many times. It follows that the image $\Phi = \{S_n\}_{n\in\mathbb{Z}}$ is a random discrete subset of $\mathbb{R}^k$. If no point of $\mathbb{R}^k$ is visited more than once (e.g., when S_n is in the positive cone a.s.), then it can be seen that Φ is a point-stationary point process, hence, $[\Phi, 0]$ is a unimodular discrete space. Hence, $[\Phi, 0]$ is a unimodular discrete space. In the general case, by similar arguments, one should bias the distribution of $[\Phi, 0]$ by the inverse of the *multiplicity* of the origin; i.e., by $1/\#\{n : S_n = 0\}$, to obtain a unimodular discrete space. This claim can be proved by direct verification of the mass transport principle.

Below, the focus is on the case where the jumps are real-valued and strictly positive. In this case, Φ is actually a point stationary *renewal process* [24].

Proposition 4.9. *Let $\Phi := \{S_n\}_{n\in\mathbb{Z}}$ be the image of a simple random walk S in $\mathbb{R}$ starting from $S_0 := 0$. Assume the jumps $S_n - S_{n-1}$ are positive a.s. Then*

(i) $\underline{\text{udim}}_M(\Phi) = \underline{\text{decay}}\left(\frac{1}{r}\mathbb{E}[S_1 \wedge r]\right) = 1 \wedge \underline{\text{decay}}(\mathbb{P}[S_1 > r])$.

(ii) $\overline{\text{udim}}_M(\Phi) = \overline{\text{decay}}\left(\frac{1}{r}\mathbb{E}[S_1 \wedge r]\right) \le 1 \wedge \overline{\text{decay}}(\mathbb{P}[S_1 > r])$.

(iii) $\mathrm{udim}_H(\Phi) \leq 1 \wedge \overline{\mathrm{decay}}\left(\mathbb{P}\left[S_1 > r\right]\right)$.

(iv) If $\beta := \mathrm{decay}\left(\mathbb{P}\left[S_1 > r\right]\right)$ *exists, then* $\mathrm{udim}_M(\Phi) = \mathrm{udim}_H(\Phi) = 1 \wedge \beta$.

Proof (first part). For every $r > 0$, one has $\mathbb{P}\left[\Phi \cap (0, r) = \emptyset\right] = \mathbb{P}\left[S_1 \geq r\right]$. So the claims regarding the Minkowski dimension are direct consequences of Proposition 3.13 and do not require the i.i.d. assumption. The proofs of the last two claims, will be given in Subsection 6.1. □

The image of the nearest-neighbor simple random walk in $\mathbb{Z}^k$ will be studied in [6]. It will be shown that it has dimension 2 when $k \geq 2$. Furthermore, a *doubling property* will be proved in this case.

As another example, if $[\boldsymbol{T}, \boldsymbol{o}]$ is any unimodular tree such that the simple random walk on $\boldsymbol{T}$ is transient a.s., then the image of the (two sided) simple random walk on $\boldsymbol{T}$ is another unimodular tree (after biasing by the inverse of the multiplicity of the root). The new tree is two-ended a.s., and hence, is 1-dimensional by Theorem 4.1.

4.3.2 Zeros of the Simple Random Walk

Proposition 4.10. *Let* Ψ *be the zero set of the symmetric simple random walk on* $\mathbb{Z}$ *with uniform jumps in* $\{\pm 1\}$. *Then,* $\mathrm{udim}_M(\Psi) = \mathrm{udim}_H(\Psi) = \frac{1}{2}$.

Proof. Represent Ψ uniquely as $\Psi := \{S_n : n \in \mathbb{Z}\}$ such that $S_0 := 0$ and $S_n < S_{n+1}$ for each n. Then, $(S_n)_n$ is another simple random walk and Ψ is its image. The distribution of the jump S_1 is explicitly computed in the classical literature on random walks (using the reflection principle). In particular, there exist $c_1, c_2 > 0$ such that $c_1 r^{-\frac{1}{2}} < \mathbb{P}\left[S_1 > r\right] < c_2 r^{-\frac{1}{2}}$ for every $r \geq 1$. Therefore, the claim is implied by part (iv) of Proposition 4.9 (recall that this part of Proposition 4.9 will be proved later). □

4.4 A Subspace with Larger Minkowski Dimension

Let $\Phi \subseteq \mathbb{R}$ be an arbitrary point-stationary point process. Let S_1 be the first point of Φ on the right of the origin. Assume $\beta := \mathrm{decay}\left(\mathbb{P}\left[S_1 > r\right]\right)$ exists with $\beta < 1$. Then Proposition 4.9 gives that $\mathrm{udim}_M(\Phi) = \beta$.

Let $\alpha < \beta < 1$. Consider the intervals defined by consecutive points of Φ. In each such interval, say (a, b), add $\lceil(b-a)^\alpha\rceil - 1$ points so as to split the interval into $\lceil(b-a)^\alpha\rceil$ equal parts. Let Φ' denote the resulting point process (with the points of Φ and the additional points). The assumption $\alpha < \beta$ implies that $\mathbb{E}\left[S_1^\alpha\right] < \infty$. Now, by biasing the distribution of Φ' by $\lceil S_1^\alpha\rceil$ and changing the origin to a point of $\Phi' \cap [0, S_1)$ chosen uniformly at random, one obtains a point-stationary point process Ψ (see Theorem 5 in [35] and also the examples in [2]), (it is not a renewal process). The distribution of Ψ is determined by the following equation, where h is any measurable nonnegative function:

$$\mathbb{E}\left[h(\Psi)\right] = \frac{1}{\mathbb{E}\left[\lceil S_1^\alpha\rceil\right]}\mathbb{E}\left[\sum_{x \in \Phi' \cap [0, S_1)} h(\Phi' - x)\right]. \tag{4.8}$$

Proposition 4.11. *Let* Φ *and* Ψ *be as above. Then,* Φ *has the same distribution as an equivariant subspace of* Ψ *(conditioned on having the root) and*

$$\mathrm{udim}_M(\Phi) = \beta > \frac{\beta - \alpha}{1 - \alpha} = \mathrm{udim}_M(\Psi).$$

Note that Theorem 3.34 implies that $\mathrm{udim}_H(\Phi) = \mathrm{udim}_H(\Psi)$. Therefore, the proposition implies $\mathrm{udim}_M(\Psi) < \mathrm{udim}_H(\Psi)$.

Proof. Let A be the set of newly-added points in Ψ, which can be defined by adding marks from the beginning and is an equivariant subset of Ψ. By (4.8), one can verify that $\Psi \setminus A$ conditioned on $0 \notin A$ has the same distribution as Φ (see also Proposition 6 in [35]). Also, by letting $c := \mathbb{E}\left[\lceil S_1^\alpha \rceil\right]$, (4.8) gives

$$
\begin{aligned}
\mathbb{P}\left[\Psi \cap (0, r) = 0\right] &= \frac{1}{c}\mathbb{E}\left[\sum_{x \in \Phi' \cap [0, S_1)} 1_{\{(\Phi' - x) \cap (0, r) = \emptyset\}}\right] \\
&= \frac{1}{c}\mathbb{E}\left[\lceil S_1^\alpha \rceil 1_{\{\Phi' \cap (0, r) = \emptyset\}}\right] \\
&= \frac{1}{c}\mathbb{E}\left[\lceil S_1^\alpha \rceil 1_{\{S_1 / \lceil S_1^\alpha \rceil > r\}}\right].
\end{aligned}
$$

Now, by the assumption $\mathrm{decay}\left(\mathbb{P}\left[S_1 > r\right]\right) = \beta$ and integration by parts, it is straightforward to deduce that $\mathrm{decay}\left(\mathbb{P}\left[\Psi \cap (0, r) = 0\right]\right) = (\beta - \alpha)/(1 - \alpha)$. Therefore, Proposition 3.13 gives the claim. $\square$

Remark 4.12. The fact that Ψ has a smaller Minkowski dimension than Φ means that the tail of the distribution of the jumps (or inter-arrivals) of Ψ is heavier than that of the inter-arrivals of Φ. This may look surprising as the inter-arrival times of Ψ are obtained by subdividing those of Φ into smaller sub-intervals. The explanation of this apparent contradiction is of the same nature as that of Feller's paradox (Section I.4 of [24]). It comes from the renormalization of *size-biased sampling*: the typical inter-arrival of Ψ has more chance to be found in a larger inter-arrival of Φ, and this length-biasing dominates the effect of the subdivision.

4.5 A Drainage Network Model

Practical observations show that large river basins have a fractal structure. For example, [30] discovered a power law relating the area and the height of river basins. There are various ways to model river basins and their fractal properties in the literature. In particular, [45] formalizes and proves a power law with exponent $3/2$ for a specific model called *Howard's model*. Below, the simpler model of [43] is studied. One can ask similar questions for Howard's model or other drainage network models.

Connect each (x, y) in the even lattice $\{(x, y) \in \mathbb{Z}^2 : x + y \bmod 2 = 0\}$ to either $(x - 1, y - 1)$ or $(x + 1, y - 1)$ with equal probability in an i.i.d. manner to obtain a directed graph T. Note that the downward path starting at a given vertex is the rotated graph of a simple random walk. It is known that T is connected and is a one-ended tree (see e.g., [45]). Also, by Lemma 2.12, $[T, 0]$ is unimodular.

Note that by considering the Euclidean metric on T, the Hausdorff dimension of T is 2. In the following, the graph-distance metric is considered on T.

Proposition 4.13. *One has* $\mathrm{udim}_M(T) = \mathrm{udim}_H(T) = \frac{3}{2}$ *under the graph-distance metric.*

Proof (first part). Here, it will be proved that $\mathrm{udim}_M(T) = \frac{3}{2}$. The rest of the proof is postponed to Subsection 6.1. The idea is to use Theorem 4.2. Following [45], there are two *backward paths* (going upward) in the odd lattice that surround the descendants $D(o)$ of the origin. These two paths have exactly the same distribution as (rotated) graphs of independent simple random walks starting at $(-1, 0)$ and $(1, 0)$, respectively, until they hit for the first time. In this setting, $h(o)$ is exactly the hitting time of these random walks. So classical results on random walks imply that $\mathbb{P}\left[h(o) \geq n\right]$ is bounded between two constant multiples of $n^{-\frac{1}{2}}$ for all n. So Theorem 4.2 implies that $\mathrm{udim}_M(T) = \frac{3}{2}$. $\square$

4.6 Self Similar Unimodular Discrete Spaces

This section provides a class of examples of unimodular discrete spaces obtained by discretizing self-similar sets. Let $l \geq 1$ and $f_1, \ldots, f_l$ be similitudes of $\mathbb{R}^k$ with similarity ratios $r_1, \ldots, r_l$ respectively (i.e., $\forall x, y \in \mathbb{R}^k : |f_i(x) - f_i(y)| = r_i |x - y|$). For every $n \geq 0$ and every string $\sigma = (j_1, \ldots, j_n) \in \{1, \ldots, l\}^n$, let $f_\sigma := f_{j_1} \cdots f_{j_n}$. Also let $|\sigma| := n$. Fix a point $o \in \mathbb{R}^k$ (one can similarly start with a finite subset of $\mathbb{R}^k$ instead of a single point). Let $K_0 := \{o\}$ and $K_{n+1} := \bigcup_j f_j(K_n)$ for each $n \geq 0$. Equivalently,

$$K_n = \{f_\sigma(o) : |\sigma| = n\}. \tag{4.9}$$

Recall that if $r_i < 1$ for all i, then by contraction arguments, K_n converges in the Hausdorff metric to *the attractor* of $f_1, \ldots, f_l$ (see e.g., Section 2.1 of [14]). The attractor is the unique compact set $K \subseteq \mathbb{R}^k$ such that $K = \bigcup_i f_i(K)$. In addition, if the f_i's satisfy the *open set condition*; i.e., there is a bounded open set $V \subseteq \mathbb{R}^k$ such that $f_i(V) \subseteq V$ and $f_i(V) \cap f_j(V) = \emptyset$ for each i, j, then the Minkowski and Hausdorff dimensions of K are equal to the *similarity dimension*, which is the unique $\alpha \geq 0$ such that $\sum r_i^\alpha = 1$.

The following is the main result of this section. It introduces a discrete analogue of self-similar sets by scaling the sets K_n and taking local weak limits.

Theorem 4.14. *Let $\boldsymbol{o}_n$ be a point of K_n chosen uniformly at random, where K_n is defined in (4.9). Assume that $r_i = r < 1$ for all i and the open set condition is satisfied. Then,*

(i) $[r^{-n} K_n, \boldsymbol{o}_n]$ converges weakly to some unimodular discrete space.

(ii) The unimodular Minkowski and Hausdorff dimension of the limiting space are equal to $\alpha := \log l / |\log r|$. Moreover, it has positive and finite α-dim H-size.

The proof is given at the end of this subsection. In fact, a point process Ψ in $\mathbb{R}^k$ will be constructed such that $[r^{-n} K_n, \boldsymbol{o}_n]$ converges weakly to $[\Psi, o]$. In addition, $\Psi - o$ is point-stationary. It can also be constructed directly by the algorithm in Remark 4.21 below.

Definition 4.15. *The unimodular discrete space in Theorem 4.14 is called a **self similar unimodular discrete space**.*

It should be noted that self similar unimodular discrete spaces depend on the choice of the initial point o in general.

The following are examples of unimodular self similar discrete spaces. The reader is also invited to construct a unimodular discrete version of the Sierpinski carpet similarly.

Example 4.16. If $f_1(x) := x/2$ and $f_2(x) := (1+x)/2$, then the limiting space is just $\mathbb{Z}$. Similarly, the lattice $\mathbb{Z}^k$ and the triangular lattice in the plane are self similar unimodular discrete spaces.

Example 4.17 (Unimodular Discrete Cantor Set). Start with two points $K_0 := \{0, 1\}$. Let $f_1(x) := x/3$ and $f_2(x) := (2+x)/3$. Then, K_n is the set of the interval ends in the n-th step of the definition of the Cantor set. Here, it is easy to see that the random set $\Psi_n := 3^n (K_n - \boldsymbol{o}_n) \subseteq \mathbb{Z}$ converges weakly to the random set $\Psi \subseteq \mathbb{Z}$ defined as follows: $\Psi := \cup_n \boldsymbol{T}_n$, where $\boldsymbol{T}_n$ is defined by letting $\boldsymbol{T}_0 := \{0, \pm 1\}$ and $\boldsymbol{T}_{n+1} := \boldsymbol{T}_n \cup (\boldsymbol{T}_n \pm 2 \times 3^n)$, where the sign is chosen i.i.d., each sign with probability $1/2$. Note that each $\boldsymbol{T}_n$ has the same distribution as Ψ_n, but the sequence $\boldsymbol{T}_n$ is nested. In addition, since $\boldsymbol{o}_n$ is chosen uniformly, Ψ_n and Ψ are point-stationary point processes, and hence $[\Psi, 0]$ is unimodular (a deterministic discrete Cantor set exists in the literature which is not unimodular). Theorem 4.14 implies that $\mathrm{udim}_M(\Psi) = \mathrm{udim}_H(\Psi) = \log 2 / \log 3$.

Example 4.18 (Unimodular Discrete Koch Snowflake). Let C_n be the set of points in the n-th step of the construction of the Koch snowflake. Let $\boldsymbol{x}_n$ be a random point of
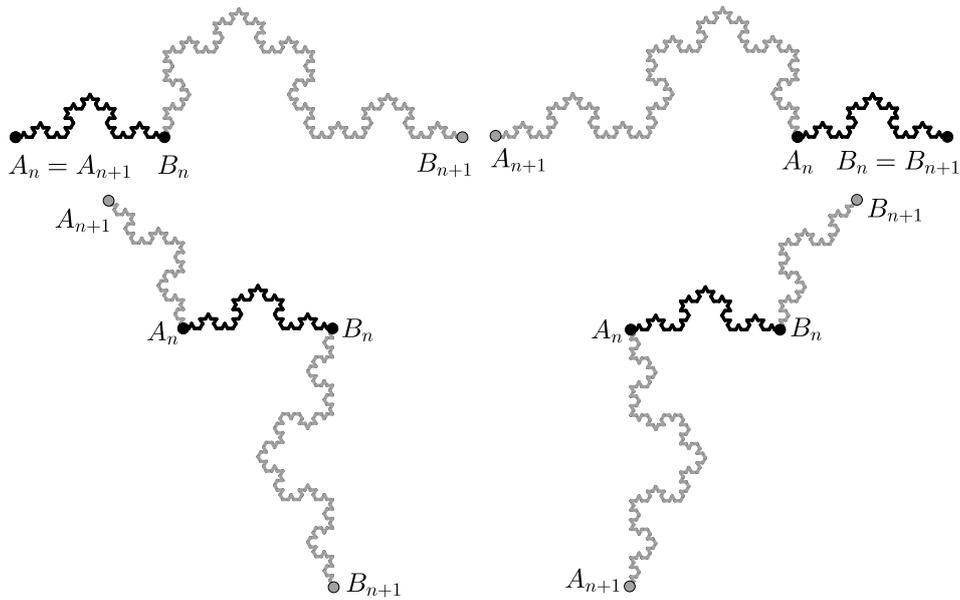
Figure 1: Four ways to attach 3 isometric copies to T_n in the construction of the unimodular discrete Koch snowflake, where each copy is a rotated/translated version of T_n (relative to A_n and B_n). Here, T_n is shown in black.

C_n chosen uniformly and $\Phi_n := 3^n(C_n - x_n)$. It can be seen that Φ_n tends weakly to a random discrete subset Φ of the triangular lattice which is almost surely a bi-infinite path (note that the cycle disappears in the limit). It can be seen that Φ can be obtained by Theorem 4.14. In this paper, Φ is called the **unimodular discrete Koch snowflake**. Also, Theorem 4.14 implies that $\mathrm{udim}_M(\Phi) = \mathrm{udim}_H(\Phi) = \log 4/\log 3$.

In addition, Φ can be constructed explicitly as $\Phi := \cup_n T_n$, where T_n is a random finite path in the triangular lattice with distinguished end points A_n and B_n defined inductively as follows: Let $T_1 := \{A_1, B_1\}$, where A_1 is the origin and B_1 is a neighbor of the origin in the triangular lattice chosen uniformly at random. For each $n \geq 1$, given (T_n, A_n, B_n), let $(T_{n+1}, A_{n+1}, B_{n+1})$ be obtained by attaching to T_n three isometric copies of itself as shown in Figure 1. There are 4 ways to attach the copies and one of them should be chosen at random with equal probability (the copies should be attached to T_n relative to the position of A_n and B_n). It can be seen that no points overlap.

Remark 4.19. If the r_i's are not all equal, the guess is that there is no scaling of the sequence $[K_n, o_n]$ that converges to a nontrivial unimodular discrete space (which is not a single point). This has been verified by the authors in the case $o \in V$. In this case, by letting a_n be the distance of o_n to its closest point in K_n, it is shown that for any $\epsilon > 0$, $\mathbb{P}\left[a_n/(\bar{r})^n < \epsilon\right] \to \frac{1}{2}$ and $\mathbb{P}\left[a_n/(\bar{r})^n > \frac{1}{\epsilon}\right] \to \frac{1}{2}$, where $\bar{r}$ is the geometric mean of $r_1, \ldots, r_l$. This implies the claim (note that the counting measure matters for convergence; e.g., $\{0, \frac{1}{n}\}$ does not converge to $\{0\}$).

To prove Theorem 4.14, it is useful to consider the following nested version of the sets K_n (note that K_n is not necessarily contained in K_{n+1}, unless o is a fixed point of some f_i). Let $u_1, u_2, \ldots$ be i.i.d. uniform random numbers in $\{1, \ldots, l\}$ and $\delta_n := (u_n, \ldots, u_1)$. Let $o'_n := f_{\delta_n}(o)$. Let $\hat{K}_n := f_{\delta_n}^{-1} K_n = f_{u_1}^{-1} \cdots f_{u_n}^{-1} K_n$. The chosen order of the indices in δ_n ensures that $\hat{K}_n \subseteq \hat{K}_{n+1}$ for all n. It is easy to see that $[\hat{K}_n, o]$ has the same

distribution as $[r^{-n}K_n, \boldsymbol{o}'_n]$. For $v \in \hat{\boldsymbol{K}}_n$, let

$$\boldsymbol{w}_n(v) := \#\{\sigma : |\sigma| = n, f_\sigma(o) = f_{\boldsymbol{\delta}_n}(v)\}.$$

One has $\boldsymbol{w}_n(v) \leq \boldsymbol{w}_{n+1}(v)$. Note that in the case $o \in V$, $\boldsymbol{w}_n(\cdot) = 1$ and the arguments are much simpler. The reader can assume this at first reading.

In the following, for $x \in \mathbb{R}^k$, $B_r(x)$ represents the closed ball of radius r centered at x in $\mathbb{R}^k$.

Lemma 4.20. *Let $\hat{\boldsymbol{K}} := \cup_n \hat{\boldsymbol{K}}_n$ and $\boldsymbol{w}(v) := \lim_n \boldsymbol{w}_n(v)$ for $v \in \hat{\boldsymbol{K}}$.*

  (i) *$\boldsymbol{w}(\cdot)$ is uniformly bounded.*

 (ii) *Almost surely, $\hat{\boldsymbol{K}}$ is a discrete set.*

(iii) *The distribution of $[\hat{\boldsymbol{K}}, o]$, biased by $1/\boldsymbol{w}(o)$, is the limiting distribution alluded to in Theorem 4.14.*

Proof. (i). Assume $f_{\sigma_1(o)} = \cdots = f_{\sigma_l}(o)$ and $|\sigma_j| = n$ for each $j \leq l$. Let D be a fixed number such that V intersects $B_D(o)$. Now, the sets $f_{\sigma_j}(V)$ for $1 \leq j \leq l$ are disjoint and intersect a common ball of radius Dr^n. Moreover, each of them contains a ball of radius ar^n and each is contained in a ball of radius br^n (for some fixed $a, b > 0$). Therefore, Lemma 2.2.5 of [14] implies that $l \leq (\frac{D+2b}{a})^k =: C$. This implies that $\boldsymbol{w}_n(\cdot) \leq C$ a.s., hence $\boldsymbol{w}(\cdot) \leq C$ a.s.

(ii). Let D be arbitrary as in the previous part. Assume $f_{\boldsymbol{\delta}_n}^{-1} f_{\sigma_j(o)} \in B_D(o)$ for $j = 1, \ldots, l$. Now, for $j = 1, \ldots, l$, the sets $f_{\sigma_j}(V)$ are disjoint and intersect a common ball of radius $2Dr^n$. As in the previous part, one obtains $l \leq (\frac{2D+2b}{a})^k$. Therefore, $\#N_D(o) \leq (\frac{2D+2b}{a})^k$ a.s. Since this holds for all large enough D, one obtains that $\hat{\boldsymbol{K}}$ is a discrete set a.s.

(iii). Note that the distribution of $\boldsymbol{o}'_n$ is just the distribution of $\boldsymbol{o}_n$ biased by the *multiplicities* of the points in K_n. It follows that biasing the distribution of $[\hat{\boldsymbol{K}}_n, o]$ by $1/\boldsymbol{w}_n(o)$ gives just the distribution of $[r^{-n}K_n, \boldsymbol{o}_n]$. The latter is unimodular since $\boldsymbol{o}_n$ is uniform in K_n. So the distribution of $[\hat{\boldsymbol{K}}, o]$ biased by $1/\boldsymbol{w}(o)$ is also unimodular and satisfies the claim of Theorem 4.14. $\square$

Proof of Theorem 4.14. Convergence is proved in Lemma 4.20. The rest of the proof is base on the construction of a sequence of equivariant coverings of $\hat{\boldsymbol{K}}$. In this proof, with an abuse of notation, the dimension of $\hat{\boldsymbol{K}}$ means the dimension of the unimodular space obtained by biasing the distribution of $\hat{\boldsymbol{K}}$ by $1/\boldsymbol{w}(o)$ (see Lemma 4.20). Let $D > \mathrm{diam}(K)$ be given, where K is the attractor of $f_1, \ldots, f_l$. Let $m > 0$ be large enough so that $\mathrm{diam}(K_m) < D$. Note that each element in $\hat{\boldsymbol{K}}$ can be written as $f_{\boldsymbol{\delta}_n}^{-1} f_\sigma(o)$ for some n and some string σ of length n. Let $\boldsymbol{\gamma}_m$ be a string of length m chosen uniformly at random and independently of other variables. For an arbitrary n and a string σ of length n, let

$$\begin{aligned}
\boldsymbol{U}_\sigma &:= f_{\boldsymbol{\delta}_{n+m}}^{-1} f_\sigma(K_m), \\
\boldsymbol{z}_\sigma &:= f_{\boldsymbol{\delta}_{n+m}}^{-1} f_\sigma f_{\boldsymbol{\gamma}_m}(o).
\end{aligned}$$

Note that $\boldsymbol{U}_\sigma \subseteq \hat{\boldsymbol{K}}$ is always a scaling of K_m with ratio r^{-m} and $\boldsymbol{z}_\sigma \in \boldsymbol{U}_\sigma$. Now, define the following covering of $\hat{\boldsymbol{K}}$:

$$\boldsymbol{R}_m(v) := \begin{cases} Dr^{-m}, & \text{if } v = \boldsymbol{z}_\sigma \text{ for some } \sigma, \\ 0, & \text{otherwise.} \end{cases}$$

It can be seen that $\boldsymbol{R}_m$ gives an equivariant covering. Also, note that $\boldsymbol{R}_m(o) > 0$ if and only if $f_\sigma f_{\boldsymbol{\gamma}_m}(o) = f_{\boldsymbol{\delta}_{n+m}}(o)$ for some n and some string σ of length n. Let $A_{n,m}(o)$ be

the set of possible outcomes for $\boldsymbol{\gamma}_m$ such that there exists a string σ of length n such that the last equation holds. One can see that this set is increasing with n and deduce that $\boldsymbol{w}_m(o) \leq \#A_{n,m}(o) \leq \boldsymbol{w}_{n+m}(o)$. By letting $\boldsymbol{w}'_m(o) := \# \cup_n A_{n,m}(o)$, it follows that $\boldsymbol{w}_m(o) \leq \boldsymbol{w}'_m(o) \leq \boldsymbol{w}(o)$. According to the above discussion, $\boldsymbol{R}_m(o) > 0$ if and only if $\boldsymbol{\gamma}_m \in \cup_n A_{n,m}(o)$. So

$$\mathbb{P}\left[\boldsymbol{R}_m(o) > 0 \,|\, \boldsymbol{u}_0, \boldsymbol{u}_1, \ldots\right] = \boldsymbol{w}'_m(o) r^{m\alpha}.$$

Therefore, by considering the biasing that makes $\hat{\boldsymbol{K}}$ unimodular, one gets

$$\mathbb{E}\left[\frac{1}{\boldsymbol{w}(o)} 1_{\{\boldsymbol{R}_m(o) > 0\}}\right] = \mathbb{E}\left[\frac{\boldsymbol{w}'_m(o) r^{m\alpha}}{\boldsymbol{w}(o)}\right] \leq r^{m\alpha}. \tag{4.10}$$

Since the balls in the covering have radius Dr^{-m}, one gets $\underline{\dim}_M(\hat{\boldsymbol{K}}) \geq \alpha$.

On the other hand, by (4.10) and monotone convergence, one finds that

$$\mathbb{E}\left[\frac{1}{\boldsymbol{w}(o)} 1_{\{\boldsymbol{R}_m(o) > 0\}}\right] \geq \frac{1}{2} r^{m\alpha},$$

for large enough m. Similar to the proof of part (i) of Lemma 4.20, one can show that the sequence of coverings $\boldsymbol{R}_m$ (for $m = 1, 2, \ldots$) is uniformly bounded. Therefore, Lemma 3.9 implies that $\mathrm{udim}_M(\hat{\boldsymbol{K}}) = \alpha$. Moreover, since $\mathbb{E}\left[\boldsymbol{R}_m(o)^\alpha / \boldsymbol{w}(o)\right]$ is bounded (by D^α), one can get that $\mathcal{M}^\alpha(\hat{\boldsymbol{K}}) > 0$.

Lemma 3.17 will be used to bound the Hausdorff dimension. Let $D > 1$ be arbitrary. Choose m such that $r^{-m} \leq D < r^{-m-1}$. By Lemma 4.20, there are finitely many points in $\hat{\boldsymbol{K}} \cap B_D(o)$. Therefore, one finds n such that $\hat{\boldsymbol{K}} \cap B_D(o) = \hat{\boldsymbol{K}}_{n+m} \cap B_D(o)$. It follows that the sets $\{\boldsymbol{U}_\sigma : |\sigma| = n\}$ cover $\hat{\boldsymbol{K}}_{n+m}$. Now, assume $\sigma_1, \ldots, \sigma_k$ are strings of length n such that $\boldsymbol{U}_{\sigma_i}$ are distinct and intersects $B_D(o)$. One obtains that

$$\#B_D(o) \cap \hat{\boldsymbol{K}} \leq \sum_{j=1}^k \#B_D(o) \cap \boldsymbol{U}_{\sigma_j} \leq kl^m = kr^{-\alpha m} \leq kD^\alpha. \tag{4.11}$$

Consider the sets $\boldsymbol{V}_{\sigma_j} := f_{\boldsymbol{\delta}_{n+m}}^{-1} f_{\sigma_j}(V)$ which are disjoint (since σ_j's have the same length). Note that if $\epsilon > \mathrm{diam}(V \cup \{o\})$ is fixed, then the ϵ-neighborhood of V contains K_m. Therefore, all $\boldsymbol{V}_{\sigma_j}$'s intersect a common ball of radius $D + \epsilon r^{-m} \leq (1+\epsilon)D$. Moreover, each of them contains a ball of radius $ar^{-m} \geq arD$ and is contained in a ball of radius $br^{-m} \leq bD$ (for some $a, b > 0$ not depending on D). Therefore, Lemma 2.2.5 of [14] implies that $k \leq (\frac{(1+\epsilon)+2b}{ar})^k$. Therefore, (4.11) implies that

$$\#B_D(o) \cap \hat{\boldsymbol{K}} \leq CD^\alpha, \quad \text{a.s.}$$

Therefore, Lemma 3.17 implies that $\mathrm{udim}_H(\hat{\boldsymbol{K}}) \leq \alpha$. Moreover, the proof of the lemma shows that $\mathcal{M}^\alpha(\hat{\boldsymbol{K}}) < \infty$. This completes the proof. $\qquad \square$

Remark 4.21. Motivated by Examples 4.17 and 4.18, it can be seen that every unimodular self similar discrete space can be constructed by successively attaching copies of a set to itself. This is expressed in the following algorithm.

4.7 Notes and Bibliographical Comments

Some of the examples in this section, listed below, are motivated by analogous examples in the continuum setting. In fact, the unimodular dimensions of these examples are equal to the ordinary dimensions of the analogous continuum examples. This connection will be discussed further in [6] via scaling limits.

Proposition 4.8 is inspired by the dimension of the Brownian continuum random tree (see [29] or Theorem 5.5 of [20]), which is the scaling limit of Galton-Watson trees

$\hat{\boldsymbol{K}}_0 := \{o\}$;

Let g_0 be the identity map;

Choose i.i.d. random numbers $i_1, i_2, \ldots$ uniformly in $\{1, \ldots, l\}$;

for $n = 1, 2, \ldots$ **do**

    let $\hat{\boldsymbol{K}}_n$ consist of l isometric copies of $\hat{\boldsymbol{K}}_{n-1}$ as follows

$$\hat{\boldsymbol{K}}_n := \bigcup_{j=1}^{l} g_{n-1} f_{i_n}^{-1} f_j g_{n-1}^{-1}(\hat{\boldsymbol{K}}_{n-1});$$

    Let $g_n := g_{n-1} f_{i_n}^{-1}$;

end

conditioned to be large. The zero set of the simple random walk (Proposition 4.10) is analogous to the zero set of Brownian motion. Self-similar unimodular discrete spaces are inspired by continuum self-similar sets (see e.g., Section 2.1 of [14]) as discussed in Subsection 4.6.

5 The Unimodular Mass Distribution Principle and Billingsley Lemma

Let D be a discrete space and $o \in D$. The **upper and lower (polynomial) volume growth rates** of D are

$$\overline{\text{growth}}\,(\#N_r(o)) = \limsup_{r \to \infty} \log \#N_r(o)/\log r,$$
$$\underline{\text{growth}}\,(\#N_r(o)) = \liminf_{r \to \infty} \log \#N_r(o)/\log r.$$

D has **polynomial growth** if $\overline{\text{growth}}\,(\#N_r(o)) < \infty$. These definitions have various other names in the literature (e.g., mass dimension [8], fractal dimension, or growth degree); volume growth will be used in the present paper since it is common in the context of graphs and discrete groups.

If the upper and lower volume growth rates are equal, the common value is called the **volume growth rate** of D. Note that for $v \in D$, one has $N_r(o) \subseteq N_{r+c}(v)$ and $N_r(v) \subseteq N_{r+c}(o)$, where $c := d(o, v)$. This implies that $\overline{\text{growth}}\,(\#N_r(o))$ and $\underline{\text{growth}}\,(\#N_r(o))$ do not depend on the choice of the point o.

In various situations in this paper, some *weight* in $\mathbb{R}^{\geq 0}$ will be assigned to each point of D. In these cases, it is natural to redefine the volume growth rate by considering the weights; i.e., by replacing $\#N_r(o)$ with the sum of the weights of the points in $N_r(o)$. This will be formalized below using the notion of *equivariant processes*. Recall that an equivariant process should be defined for all discrete spaces D. However, if a random pointed discrete space $[\boldsymbol{D}, \boldsymbol{o}]$ is considered, it is enough to define weights in almost every realization (see Subsection 2.5 for more on the matter). Also, given D, the weights are allowed to be random.

Definition 5.1. *An **equivariant weight function** w is an equivariant process (Definition 2.5) with values in $\mathbb{R}^{\geq 0}$. For all discrete spaces D and $v \in D$, the (random) value $\boldsymbol{w}(v) := \boldsymbol{w}_D(v)$ is called the **weight** of v. Also, for $S \subseteq D$, let $\boldsymbol{w}(S) := \boldsymbol{w}_D(S) := \sum_{v \in S} \boldsymbol{w}(v)$.*

The last equation shows that one could also call w an *equivariant measure*.

Assume $[\boldsymbol{D}, \boldsymbol{o}]$ is a unimodular discrete space (Subsection 2.4). Lemma 2.12 shows that $[\boldsymbol{D}, \boldsymbol{o}; \boldsymbol{w}_D]$ is a random pointed marked discrete space and is unimodular.

5.1 Unimodular Mass Distribution Principle

Theorem 5.2 (Mass Distribution Principle)**.** *Let $[\boldsymbol{D}, \boldsymbol{o}]$ be a unimodular discrete space.*

(i) *Let $\alpha, c, M > 0$ and assume there exists an equivariant weight function $\boldsymbol{w}$ such that $\forall r \geq M : \boldsymbol{w}(N_r(\boldsymbol{o})) \leq cr^\alpha$, a.s. Then, $\mathcal{H}^\alpha_M(\boldsymbol{D})$ defined in (3.3) satisfies $\mathcal{H}^\alpha_M(\boldsymbol{D}) \geq \frac{1}{c}\mathbb{E}\left[\boldsymbol{w}(\boldsymbol{o})\right]$.*

(ii) *If in addition, $\boldsymbol{w}_{\boldsymbol{D}}(\boldsymbol{o}) > 0$ with positive probability, then $\mathrm{udim}_H(\boldsymbol{D}) \leq \alpha$.*

Proof. Let $\boldsymbol{R}$ be an arbitrary equivariant covering such that $\boldsymbol{R}(\cdot) \in \{0\} \cup [M, \infty)$ a.s. By the assumption on $\boldsymbol{w}$, $\boldsymbol{R}(\boldsymbol{o})^\alpha \geq \frac{1}{c}\boldsymbol{w}(N_{\boldsymbol{R}}(\boldsymbol{o}))$ a.s. Therefore,

$$\mathbb{E}\left[\boldsymbol{R}(\boldsymbol{o})^\alpha\right] \geq \frac{1}{c}\mathbb{E}\left[\boldsymbol{w}(N_{\boldsymbol{R}}(\boldsymbol{o}))\right]. \tag{5.1}$$

Consider the *independent coupling* of $\boldsymbol{w}$ and $\boldsymbol{R}$; i.e., for each deterministic discrete space G, choose $\boldsymbol{w}_G$ and $\boldsymbol{R}_G$ independently (see Definition 2.9). Then, it can be seen that the pair $(\boldsymbol{w}, \boldsymbol{R})$ is an equivariant process. So by Lemma 2.12, $[\boldsymbol{G}, \boldsymbol{o}; (\boldsymbol{w}, \boldsymbol{R})]$ is unimodular. Now, the mass transport principle (2.2) can be used for $[\boldsymbol{G}, \boldsymbol{o}; (\boldsymbol{w}, \boldsymbol{R})]$. By letting $g(u, v) := \boldsymbol{w}(v)1_{\{v \in N_{\boldsymbol{R}}(u)\}}$, one gets $g^+(\boldsymbol{o}) = \boldsymbol{w}(N_{\boldsymbol{R}}(\boldsymbol{o}))$. Also, $g^-(\boldsymbol{o}) = \boldsymbol{w}(\boldsymbol{o})\sum_{u \in \boldsymbol{D}} 1_{\{\boldsymbol{o} \in N_{\boldsymbol{R}}(u)\}} \geq \boldsymbol{w}(\boldsymbol{o})$ a.s., where the last inequality follows from the fact that $\boldsymbol{R}$ is a covering. Therefore, the mass transport principle implies that $\mathbb{E}\left[\boldsymbol{w}(N_{\boldsymbol{R}}(\boldsymbol{o}))\right] \geq \mathbb{E}\left[\boldsymbol{w}(\boldsymbol{o})\right]$ (recall that by convention, $N_{\boldsymbol{R}}(\boldsymbol{o})$ is the empty set when $\boldsymbol{R}(\boldsymbol{o}) = 0$). So by (5.1), one gets $\mathbb{E}\left[\boldsymbol{R}(\boldsymbol{o})^\alpha\right] \geq \frac{1}{c}\mathbb{E}\left[\boldsymbol{w}(\boldsymbol{o})\right]$. Since this holds for any $\boldsymbol{R}$, one gets that $\mathcal{H}^\alpha_M(\boldsymbol{D}) \geq \frac{1}{c}\mathbb{E}\left[\boldsymbol{w}(\boldsymbol{o})\right]$ and the first claim is proved. If in addition, $\boldsymbol{w}(\boldsymbol{o}) > 0$ with positive probability, then $\mathbb{E}\left[\boldsymbol{w}(\boldsymbol{o})\right] > 0$. Therefore, $\mathcal{H}^\alpha_1(\boldsymbol{D}) > 0$ and the second claim is proved. $\qquad\square$

5.2 Unimodular Billingsley Lemma

The main result of this subsection is Theorem 5.6. It is based on Lemmas 5.3 and 5.4 below. Lemma 5.3 is a stronger version of the mass distribution principle (Theorem 5.2).

Lemma 5.3 (An Upper Bound)**.** *Let $[\boldsymbol{D}, \boldsymbol{o}]$ be a unimodular discrete space and $\alpha \geq 0$.*

(i) *If there exist $c \geq 0$ and $\boldsymbol{w}$ is an equivariant weight function such that $\limsup_{r \to \infty} \boldsymbol{w}(N_r(\boldsymbol{o}))/r^\alpha \leq c$, a.s., then $\mathcal{H}^\alpha_\infty(\boldsymbol{D}) \geq \frac{1}{2^\alpha c}\mathbb{E}\left[\boldsymbol{w}(\boldsymbol{o})\right]$.*

(ii) *In addition, if $\boldsymbol{w}_{\boldsymbol{D}}(\boldsymbol{o}) > 0$ with positive probability, then $\mathrm{udim}_H(\boldsymbol{D}) \leq \alpha$.*

Proof. Let $c' > c$ be arbitrary. The assumption implies that $\sup\{r \geq 0 : \boldsymbol{w}(N_r(\boldsymbol{o})) > c'r^\alpha\} < \infty$ a.s. For $m \geq 1$, let $A_m := \{v \in \boldsymbol{D} : \forall r \geq m : \boldsymbol{w}(N_r(v)) \leq c'r^\alpha\}$, which is an increasing sequence of equivariant subsets. So

$$\lim_{m \to \infty} \mathbb{P}\left[\boldsymbol{o} \in A_m\right] = 1. \tag{5.2}$$

Let $\boldsymbol{R}$ be an equivariant covering such that $\boldsymbol{R}(\cdot) \in \{0\} \cup [m, \infty)$ a.s. One has

$$\mathbb{E}\left[\boldsymbol{R}(\boldsymbol{o})^\alpha\right] \geq \mathbb{E}\left[\boldsymbol{R}(\boldsymbol{o})^\alpha 1_{\{N_{\boldsymbol{R}}(\boldsymbol{o}) \cap A_m \neq \emptyset\}}\right]. \tag{5.3}$$

If $N_{\boldsymbol{R}}(\boldsymbol{o}) \cap A_m \neq \emptyset$, then $\boldsymbol{R}(\boldsymbol{o}) \neq 0$ and hence $\boldsymbol{R}(\boldsymbol{o}) \geq m$. In the next step, assume that this is the case. Let v be an arbitrary point in $N_{\boldsymbol{R}}(\boldsymbol{o}) \cap A_m$. By the definition of A_m, one gets that for all $r \geq m$, $\boldsymbol{w}(N_r(v)) \leq c'r^\alpha$. Since $N_{\boldsymbol{R}(\boldsymbol{o})}(\boldsymbol{o}) \subseteq N_{2\boldsymbol{R}(\boldsymbol{o})}(v)$, it follows that $\boldsymbol{w}(N_{\boldsymbol{R}}(\boldsymbol{o})) \leq \boldsymbol{w}(N_{2\boldsymbol{R}(\boldsymbol{o})}(v)) \leq 2^\alpha c'\boldsymbol{R}(\boldsymbol{o})^\alpha$. Therefore, (5.3) gives

$$\mathbb{E}\left[\boldsymbol{R}(\boldsymbol{o})^\alpha\right] \geq \frac{1}{2^\alpha c'}\mathbb{E}\left[\boldsymbol{w}(N_{\boldsymbol{R}}(\boldsymbol{o}))1_{\{N_{\boldsymbol{R}}(\boldsymbol{o}) \cap A_m \neq \emptyset\}}\right]. \tag{5.4}$$

Let $g(u,v) := \boldsymbol{w}(v)1_{\{v \in N_{\boldsymbol{R}}(u)\}}1_{\{N_{\boldsymbol{R}}(u) \cap A_m \neq \emptyset\}}$. The outgoing mass from the root is $g^+(\boldsymbol{o}) = \boldsymbol{w}(N_{\boldsymbol{R}}(\boldsymbol{o}))1_{\{N_{\boldsymbol{R}}(\boldsymbol{o}) \cap A_m \neq \emptyset\}}$. Also, since there is a ball $N_{\boldsymbol{R}}(u)$ that covers $\boldsymbol{o}$ a.s., one has $g^-(\boldsymbol{o}) \geq \boldsymbol{w}(\boldsymbol{o})1_{\{\boldsymbol{o} \in A_m\}}$ a.s. Therefore, the mass transport principle (2.2) and (5.4) imply that $\mathbb{E}[\boldsymbol{R}(\boldsymbol{o})^\alpha] \geq \frac{1}{2^\alpha c'}\mathbb{E}[\boldsymbol{w}(\boldsymbol{o})1_{\{\boldsymbol{o} \in A_m\}}]$. This implies that $\mathcal{H}_m^\alpha(\boldsymbol{D}) \geq \frac{1}{2^\alpha c'}\mathbb{E}[\boldsymbol{w}(\boldsymbol{o})1_{\{\boldsymbol{o} \in A_m\}}]$. By using (5.2) and letting m tend to infinity, one gets $\mathcal{H}_\infty^\alpha(\boldsymbol{D}) \geq \frac{1}{2^\alpha c'}\mathbb{E}[\boldsymbol{w}(\boldsymbol{o})]$. Since $c' > c$ is arbitrary, the first claim is proved.

Part (ii) is proved by the same argument as the corresponding statement in Theorem 5.2. The proof leverages Lemma 3.26. □

Lemma 5.4 (Lower Bounds). *Let $[\boldsymbol{D}, \boldsymbol{o}]$ be a unimodular discrete space, $\alpha \geq 0$ and $c > 0$. Let $\boldsymbol{w}$ be an arbitrary equivariant weight function such that $\mathbb{E}[\boldsymbol{w}(\boldsymbol{o})] < \infty$.*

  (i) If $\exists r_0 : \forall r \geq r_0 : \boldsymbol{w}(N_r(\boldsymbol{o})) \geq cr^\alpha$ a.s., then $\underline{\mathrm{udim}}_M(\boldsymbol{D}) \geq \alpha$.

  (ii) If $\underline{\mathrm{growth}}(\boldsymbol{w}(N_r(\boldsymbol{o}))) \geq \alpha$ a.s., then $\mathrm{udim}_H(\boldsymbol{D}) \geq \alpha$.

  (iii) If $\lim_{\delta \downarrow 0} \liminf_{r \to \infty} \mathbb{P}[\boldsymbol{w}(N_r(\boldsymbol{o})) \leq \delta r^\alpha] = 0$, then $\mathrm{udim}_H(\boldsymbol{D}) \geq \alpha$.

  (iv) If $\underline{\mathrm{decay}}\left(\mathbb{E}\left[\exp\left(-\frac{\boldsymbol{w}(N_n(\boldsymbol{o}))}{n^\alpha}\right)\right]\right) \geq \alpha$, then $\underline{\mathrm{udim}}_M(\boldsymbol{D}) \geq \alpha$.

Proof. The proofs of the first two parts are very similar. The second part is proved first.

(ii). Let β, γ and κ be such that $\gamma < \beta < \kappa < \alpha$. Fix $n \in \mathbb{N}$. Let $\boldsymbol{S} = \boldsymbol{S}_{\boldsymbol{D}}$ be the equivariant subset obtained by selecting each point $v \in \boldsymbol{D}$ with probability $1 \wedge (n^{-\beta}\boldsymbol{w}(v))$ (the selection variables are assumed to be conditionally independent given $[\boldsymbol{D}, \boldsymbol{o}; \boldsymbol{w}]$). Let $\boldsymbol{R}_n(v) = n$ if $v \in \boldsymbol{S}_{\boldsymbol{D}}$, $\boldsymbol{R}_n(v) = 1$ if $N_n(v) \cap \boldsymbol{S}_{\boldsymbol{D}} = \emptyset$, and $\boldsymbol{R}_n(v) = 0$ otherwise. Then $\boldsymbol{R}_n$ is an equivariant covering. It is shown below that $\mathbb{E}[\boldsymbol{R}_n(\boldsymbol{o})^\gamma] \to 0$. Let $M := \sup\{r \geq 0 : \boldsymbol{w}(N_r(\boldsymbol{o})) < r^\kappa\}$. By the assumption, $M < \infty$ a.s. One has

$$
\begin{aligned}
\mathbb{E}[\boldsymbol{R}_n(\boldsymbol{o})^\gamma] &= n^\gamma \mathbb{P}[\boldsymbol{o} \in \boldsymbol{S}_{\boldsymbol{D}}] + \mathbb{P}[N_n(\boldsymbol{o}) \cap \boldsymbol{S}_{\boldsymbol{D}} = \emptyset] \\
&= n^\gamma \mathbb{E}[1 \wedge n^{-\beta}\boldsymbol{w}(\boldsymbol{o})] + \mathbb{E}\left[\prod_{v \in N_n(\boldsymbol{o})}\left(1 - (1 \wedge n^{-\beta}\boldsymbol{w}(v))\right)\right] \\
&\leq n^{\gamma-\beta}\mathbb{E}[\boldsymbol{w}(\boldsymbol{o})] + \mathbb{E}\left[\exp\left(-n^{-\beta}\boldsymbol{w}(N_n(\boldsymbol{o}))\right)\right] \\
&= n^{\gamma-\beta}\mathbb{E}[\boldsymbol{w}(\boldsymbol{o})] + \mathbb{E}\left[\exp\left(-n^{-\beta}\boldsymbol{w}(N_n(\boldsymbol{o}))\right)|M < n\right]\mathbb{P}[M < n] \\
&\qquad + \mathbb{E}\left[\exp\left(-n^{-\beta}\boldsymbol{w}(N_n(\boldsymbol{o}))\right)|M \geq n\right]\mathbb{P}[M \geq n] \\
&\leq n^{\gamma-\beta}\mathbb{E}[\boldsymbol{w}(\boldsymbol{o})] + \exp\left(-n^{\kappa-\beta}\right) + \mathbb{P}[M \geq n],
\end{aligned}
$$

where the first inequality holds because $1 - (1 \wedge x) \leq e^{-x}$ for all $x \geq 0$. Therefore, $\mathbb{E}[\boldsymbol{R}_n(\boldsymbol{o})^\gamma] \to 0$ when $n \to \infty$. It follows that $\mathrm{udim}_H(\boldsymbol{D}) \geq \gamma$. Since γ is arbitrary, this implies $\mathrm{udim}_H(\boldsymbol{D}) \geq \alpha$.

(i). Only a small change is needed in the above proof. For $n \geq r_0$, let $\boldsymbol{R}_n(v) = n$ if either $v \in \boldsymbol{S}_{\boldsymbol{D}}$ or $N_n(v) \cap \boldsymbol{S}_{\boldsymbol{D}} = \emptyset$, and let $\boldsymbol{R}_n(v) = 0$ otherwise. Note that $\boldsymbol{R}_n$ is a covering by balls of equal radii. By the same computations and the assumption $M \leq r_0$, one gets

$$\mathbb{P}[\boldsymbol{R}_n(\boldsymbol{o}) \neq 0] \leq n^{-\beta}\mathbb{E}[\boldsymbol{w}(\boldsymbol{o})] + \exp\left(-n^{\kappa-\beta}\right),$$

which is of order $n^{-\beta}$ for large n. This implies that $\underline{\mathrm{udim}}_M(\boldsymbol{D}) \geq \beta$. Since β is arbitrary, one gets $\underline{\mathrm{udim}}_M(\boldsymbol{D}) \geq \alpha$ and the claim is proved.

(iii). Let $\beta < \alpha$. It will be proved below that under the assumption of (iii), there is a sequence $r_1, r_2, \dots$ such that $\mathbb{E}\left[\exp\left(-r_n^{-\beta}\boldsymbol{w}(N_{r_n}(\boldsymbol{o}))\right)\right] \to 0$. If so, by a slight modification of the proof of part (ii), one can find a sequence of equivariant coverings $\boldsymbol{R}_n$ such that $\mathbb{E}[\boldsymbol{R}_n(\boldsymbol{o})^\beta] < \infty$ and (iii) is proved.

Let $\epsilon > 0$ be arbitrary. By the assumption, there is $\delta > 0$ and $r \geq 1$ such that $\mathbb{P}\left[\boldsymbol{w}(N_r(\boldsymbol{o})) \leq \delta r^\alpha\right] < \epsilon$. So

$$
\begin{aligned}
\mathbb{E}\left[\exp\left(-r^{-\beta}\boldsymbol{w}(N_r(\boldsymbol{o}))\right)\right] &\leq \mathbb{E}\left[\exp\left(-r^{-\beta}\boldsymbol{w}(N_r(\boldsymbol{o}))\right)|\boldsymbol{w}(N_r(\boldsymbol{o})) > \delta r^\alpha\right] \\
&\quad + \mathbb{P}\left[\boldsymbol{w}(N_r(\boldsymbol{o})) \leq \delta r^\alpha\right] \\
&\leq \exp(-\delta r^{\alpha-\beta}) + \epsilon.
\end{aligned}
$$

Note that for fixed ϵ and δ as above, r can be arbitrarily large. Now, choose r large enough for the right hand side to be at most 2ϵ. This shows that $\mathbb{E}\left[\exp\left(-r^{-\beta}\boldsymbol{w}(N_r(\boldsymbol{o}))\right)\right]$ can be arbitrarily small and the claim is proved.

(iv). As before, let $\boldsymbol{R}_n(v) = n$ if either $v \in \boldsymbol{S_D}$ or $N_n(v) \cap \boldsymbol{S_D} = \emptyset$, and let $\boldsymbol{R}_n(v) = 0$ otherwise. The calculations in the proof of part (ii) show that

$$
\mathbb{P}\left[\boldsymbol{R}_n(\boldsymbol{o}) \neq 0\right] \leq n^{-\beta}\mathbb{E}\left[\boldsymbol{w}(\boldsymbol{o})\right] + \mathbb{E}\left[\exp\left(-n^{-\beta}\boldsymbol{w}(N_n(\boldsymbol{o}))\right)\right].
$$

Now, the assumption implies the claim. $\qquad\square$

Remark 5.5. The assumption in part (iii) of Lemma 5.4 is equivalent to the condition that there exists a sequence $r_n \to \infty$ such that the family of random variables $r_n^\alpha/\boldsymbol{w}(N_{r_n}(\boldsymbol{o}))$ is tight. Also, from the proof of the lemma, one can see that this assumption is equivalent to

$$
\liminf_{n\to\infty} \mathbb{E}\left[\exp\left(-\frac{\boldsymbol{w}(N_n(\boldsymbol{o}))}{n^\alpha}\right)\right] = 0.
$$

Theorem 5.6 (Unimodular Billingsley Lemma). *Let $[\boldsymbol{D}, \boldsymbol{o}]$ be a unimodular discrete metric space. Then, for all equivariant weight functions $\boldsymbol{w}$ such that $0 < \mathbb{E}\left[\boldsymbol{w}(\boldsymbol{o})\right] < \infty$, one has*

$$
\begin{aligned}
\text{ess inf}\left(\underline{\text{growth}}\left(\boldsymbol{w}(N_r(\boldsymbol{o}))\right)\right) &\leq \text{udim}_H(\boldsymbol{D}) \\
&\leq \text{ess inf}\left(\overline{\text{growth}}\left(\boldsymbol{w}(N_r(\boldsymbol{o}))\right)\right) \\
&\leq \overline{\text{growth}}\left(\mathbb{E}\left[\boldsymbol{w}(N_r(\boldsymbol{o}))\right]\right).
\end{aligned}
$$

Proof. The first inequality is implied by part (ii) of Lemma 5.4. For the second inequality, assume that $\overline{\text{growth}}\left(\boldsymbol{w}(N_r(\boldsymbol{o}))\right) < \alpha$ with positive probability. On this event, one has $\boldsymbol{w}(N_r(\boldsymbol{o})) \leq r^\alpha$ for large r; i.e., $\limsup_r \boldsymbol{w}(N_r(\boldsymbol{o}))/r^\alpha \leq 1$. Now, Lemma 5.3 implies that $\text{udim}_H(\boldsymbol{D}) \leq \alpha$. This proves the second inequality. The last claim follows because $\overline{\text{growth}}(X_n) \leq \overline{\text{growth}}(\mathbb{E}[X_n])$ for any monotone sequence of nonnegative random variables X_n (see Lemma C.3 of [7]). $\qquad\square$

Corollary 5.7. *Under the assumptions of Theorem 5.6, if the upper and lower growth rates of $\boldsymbol{w}(N_r(\boldsymbol{o}))$ are almost surely constant (e.g., when $[\boldsymbol{D}, \boldsymbol{o}; \boldsymbol{w}]$ is ergodic), then,*

$$
\underline{\text{growth}}\left(\boldsymbol{w}(N_r(\boldsymbol{o}))\right) \leq \text{udim}_H(\boldsymbol{D}) \leq \overline{\text{growth}}\left(\boldsymbol{w}(N_r(\boldsymbol{o}))\right) \quad a.s. \tag{5.5}
$$

In particular, if $\text{growth}\left(\boldsymbol{w}(N_r(\boldsymbol{o}))\right)$ exists and is constant a.s., then

$$
\text{udim}_H(\boldsymbol{D}) = \text{growth}\left(\boldsymbol{w}(N_r(\boldsymbol{o}))\right).
$$

In fact, without the assumption of this corollary, an inequality similar to (5.5) is valid for the *sample Hausdorff dimension* of $\boldsymbol{D}$, which will be studied in [6].

Remark 5.8. In many examples, in the unimodular Billingsley lemma, it is enough to take $\boldsymbol{w}$ equal to the counting measure, i.e., $\forall v : \boldsymbol{w}(v) = 1$ and $\boldsymbol{w}(N_r(\boldsymbol{o})) = \#N_r(\boldsymbol{o})$. Analogously, for many natural fractals in the continuum setting, there is a natural *mass measure* that can be used in Billingsley's lemma.

Remark 5.9. In fact, the assumption $\mathbb{E}\left[\boldsymbol{w}(\boldsymbol{o})\right] < \infty$ in Theorem 5.6 is only needed for the lower bound while the assumption $\mathbb{E}\left[\boldsymbol{w}(\boldsymbol{o})\right] > 0$ is only needed for the upper bound. These assumptions are also necessary as shown below.

For example, assume Φ is a point-stationary point process in $\mathbb{R}$ (see Example 2.7). For $v \in \Phi$, let $\boldsymbol{w}(v)$ be the sum of the distances of v to its next and previous points in Φ. This equivariant weight function satisfies $\boldsymbol{w}(N_r(v)) \geq 2r$ for all r, and hence $\underline{\text{growth}}\left(\boldsymbol{w}(N_r(\boldsymbol{o}))\right) \geq 1$. But $\text{udim}_H(\Phi)$ can be strictly less than 1 as shown in Subsection 4.3.1.

Also, the condition that $\mathbb{E}\left[\boldsymbol{w_D}(\boldsymbol{o})\right] > 0$ is trivially necessary for the upper bound.

Corollary 5.10. *Let $[G, \boldsymbol{o}]$ be a unimodular random graph equipped with the graph-distance metric. If G is infinite almost surely, then $\underline{\text{udim}}_M(G) \geq 1$ and else, $\text{udim}_M(G) = \text{udim}_H(G) = 0$.*

Proof. If G is infinite a.s., then for $\boldsymbol{w_G} \equiv 1$, one has $\boldsymbol{w}(N_r(\boldsymbol{o})) \geq r$ for all r. So part (i) of Lemma 5.4 implies the first claim. The second claim is implied by Example 3.16 (this can be deduced from the unimodular Billingsley lemma as well). □

Corollary 5.11. *The unimodular Minkowski and Hausdorff dimensions of any unimodular two-ended tree are equal to one.*

This result has already been shown in Theorem 4.1, but can also be deduced from the unimodular Billingsley lemma directly. For this, let $\boldsymbol{w}(v)$ be 1 if v belongs to the trunk of the tree and 0 otherwise.

Problem 5.12. In the setting of Corollary 5.7, is it always the case that $\text{udim}_H(\boldsymbol{D}) = \overline{\text{growth}}\left(\boldsymbol{w}(N_r(\boldsymbol{o}))\right)$?

The claim of this problem holds in all of the examples in which both quantities are computed in this work. This problem is a corollary of Problem 5.21 and the unimodular Frostman lemma (Theorem 7.2) below. Note that there are examples where $\underline{\text{growth}}\left(\cdot\right) \neq \overline{\text{growth}}\left(\cdot\right)$ as shown in Subsections 6.3.1 and 6.5.1.

5.3 Bounds for Point Processes

Example 2.7, explains that for point processes containing the origin, unimodularity is, roughly speaking, equivalent to point-stationarity. To study the dimension of such processes, the following covering is used in the next results. Let φ be a discrete subset of $\mathbb{R}^k$ equipped with the l_∞ metric and $r \geq 1$. Let $C := C_r := [0, r)^k$, $\boldsymbol{U} := \boldsymbol{U}_r$ be a point chosen uniformly at random in $-C$, and consider the partition $\{C + \boldsymbol{U} + z : z \in r\mathbb{Z}^k\}$ of $\mathbb{R}^k$ by cubes. Then, for each $z \in r\mathbb{Z}^k$, choose a random element in $(C + \boldsymbol{U} + z) \cap \varphi$ independently (if the intersection is nonempty). The distribution of this random element should depend on the set $(C + \boldsymbol{U} + z) \cap \varphi$ in a translation-invariant way (e.g., choose with the uniform distribution or choose the least point in the lexicographic order). Let $\boldsymbol{R} = \boldsymbol{R}_\varphi$ assign the value r to the selected points and zero to the other points of φ. Then, $\boldsymbol{R}$ is an equivariant covering. Also, each point is covered at most 3^k times. So $\boldsymbol{R}$ is 3^k-bounded (Definition 3.8).

Theorem 5.13 (Minkowski Dimension in the Euclidean Case). *Let Φ be a point-stationary point process in $\mathbb{R}^k$ and assume the metric on Φ is equivalent to the Euclidean metric. Then, for all equivariant weight functions $\boldsymbol{w}$ such that $\boldsymbol{w}_\Phi(0) > 0$ a.s., one has*

$$\overline{\text{udim}}_M(\Phi) = \overline{\text{decay}}\left(\mathbb{E}\left[\boldsymbol{w}(0)/\boldsymbol{w}(C_r + \boldsymbol{U}_r)\right]\right) \leq \overline{\text{decay}}\left(\mathbb{E}\left[\boldsymbol{w}(0)/\boldsymbol{w}(N_r(0))\right]\right)$$
$$\leq \overline{\text{growth}}\left(\mathbb{E}\left[\boldsymbol{w}(N_r(0))\right]\right),$$
$$\underline{\text{udim}}_M(\Phi) = \underline{\text{decay}}\left(\mathbb{E}\left[\boldsymbol{w}(0)/\boldsymbol{w}(C_r + \boldsymbol{U}_r)\right]\right) \leq \underline{\text{decay}}\left(\mathbb{E}\left[\boldsymbol{w}(0)/\boldsymbol{w}(N_r(0))\right]\right)$$
$$\leq \underline{\text{growth}}\left(\mathbb{E}\left[\boldsymbol{w}(N_r(0))\right]\right),$$

where $\boldsymbol{U}_r$ is a uniformly at random point in $-C_r$ independent of Φ and $\boldsymbol{w}$.

Proof. By Theorem 3.31, one may assume the metric on Φ is the l_∞ metric without loss of generality. Given any $r > 0$, consider the equivariant covering $\boldsymbol{R}$ described above, but when choosing a random element of $(C_r + \boldsymbol{U}_r + z) \cap \varphi$, choose point v with probability $\boldsymbol{w}_\varphi(v)/\boldsymbol{w}_\varphi(C_r + \boldsymbol{U}_r + z)$ (conditioned on $\boldsymbol{w}_\varphi$). One gets $\mathbb{P}\left[0 \in \boldsymbol{R}\right] = \mathbb{E}\left[\boldsymbol{w}(0)/\boldsymbol{w}(C_r + \boldsymbol{U}_r)\right]$. As mentioned above, $\boldsymbol{R}$ is equivariant and uniformly bounded (for all $r > 0$). So Lemma 3.9 implies both equalities in the claim. The inequalities are implied by the facts that $\boldsymbol{w}(C_r + \boldsymbol{U}_r) \leq \boldsymbol{w}(N_r(0))$ and

$$\mathbb{E}\left[\frac{\boldsymbol{w}(0)}{\boldsymbol{w}(N_r(0))}\right] \mathbb{E}\left[\boldsymbol{w}(N_r(0))\right] \geq \mathbb{E}\left[\sqrt{\boldsymbol{w}(0)}\right]^2 > 0,$$

which is implied by the Cauchy-Schwartz inequality. $\qquad\square$

Example 5.14. The right-most inequalities in the above theorem can be strict. For example, let $T > 0$ be a random number and let $\Phi := \frac{1}{T}\mathbb{Z}$. Then $\#N_r(0) \sim 1 + T/r$. So, $\mathrm{decay}\,(1/\#N_r(0)) = 1$, but it might be the case that $\mathbb{E}\left[\#N_r(0)\right] = \infty$. For an ergodic example, let $1 \leq T_i \in \mathbb{Z}$ be i.i.d. with finite mean but infinite variance (for $i \in \mathbb{Z}$). In each interval $[i, i+1]$, put $T_i - 1$ equidistant points and let Φ_0 be the union of these points together with $\mathbb{Z}$. Bias the distribution of Φ_0 by T_0 (Definition 2.1) and then translate Φ_0 by moving a random point in $\Phi_0 \cap [0, 1)$ to the origin. Let Φ be the resulting point process. It can be seen that Φ is unimodular and point-stationary. Since $\mathbb{E}\left[T_0^2\right] = \infty$, one gets $\mathbb{E}\left[N_r(0)\right] \geq \mathbb{E}\left[N_1(0)\right] = \infty$. But one can show that $\mathrm{decay}\,(\mathbb{E}\left[1/N_r(0)\right]) = 1$.

Proposition 5.15. *If Φ is a point-stationary point process in $\mathbb{R}^k$ and the metric on Φ is equivalent to the Euclidean metric, then $\mathrm{udim}_H(\Phi) \leq k$.*

Proof. One may assume the metric on Φ is the l_∞ metric without loss of generality. Let $C := [0, 1)^k$ and $\boldsymbol{U}$ be a random point in $-C$ chosen uniformly. For all discrete subsets $\varphi \subseteq \mathbb{R}^k$ and $v \in \varphi$, let $\boldsymbol{C}(v)$ be the cube containing v of the form $C + \boldsymbol{U} + z$ (for $z \in \mathbb{Z}^k$) and $\boldsymbol{w}_\varphi(v) := 1/\#(\varphi \cap \boldsymbol{C}(v))$. Now, $\boldsymbol{w}$ is an equivariant weight function. The construction readily implies that $\boldsymbol{w}(N_r(\boldsymbol{o})) \leq (2r + 1)^k$. Moreover, by $\boldsymbol{w} \leq 1$, one has $\mathbb{E}\left[\boldsymbol{w}(0)\right] < \infty$. Therefore, the unimodular Billingsley lemma (Theorem 5.6) implies that $\mathrm{udim}_H(\Phi) \leq k$. $\qquad\square$

Proposition 5.16. *If Ψ is a stationary point process in $\mathbb{R}^k$ with finite intensity and Ψ_0 is its Palm version, then $\mathrm{udim}_M(\Psi_0) = \mathrm{udim}_H(\Psi_0) = k$. Moreover, the modified unimodular Hausdorff size of Ψ_0, defined in Section 3.8, satisfies $\mathcal{M}'_k(\Psi_0) = 2^k\rho(\Psi)$, where $\rho(\Psi)$ is the intensity of Ψ.*

Notice that if $\Psi_0 \subseteq \mathbb{Z}^k$, then the claim is directly implied by Theorem 3.34. The general case is treated below.

Proof. For the first claim, by Proposition 5.15 and Theorem 3.22, it is enough to prove that $\underline{\mathrm{udim}}_M(\Psi_0) \geq k$. Let Ψ' be a shifted square lattice independent of Ψ (i.e., $\Psi' = \mathbb{Z}^k + \boldsymbol{U}$, where $\boldsymbol{U} \in [0, 1)^k$ is chosen uniformly, independently of Ψ). Let $\Psi'' := \Psi \cup \Psi'$. Since Ψ'' is a superposition of two independent stationary point processes, it is a stationary point process itself. By letting $p := \rho(\Psi)/(\rho(\Psi) + 1)$, the Palm version Ψ''_0 of Ψ'' is obtained by the superposition of Ψ_0 and an independent stationary lattice with probability p (heads), and the superposition of $\mathbb{Z}^k$ and Ψ with probability $1 - p$ (tails). So part (i) of Lemma 5.4 implies that $\underline{\mathrm{udim}}_M(\Psi''_0) \geq k$. Note that Ψ''_0 has two natural equivariant subsets which, after conditioning to contain the origin, have the same distributions as Ψ_0 and $\mathbb{Z}^k$ respectively. Therefore, one can use Theorem 3.34 to deduce that $\underline{\mathrm{udim}}_M(\Psi_0) \geq \underline{\mathrm{udim}}_M(\Psi''_0) = k$. Therefore, Proposition 5.15 implies that $\mathrm{udim}_H(\Psi_0) = \mathrm{udim}_M(\Psi_0) = k$.

Also, by using Theorem 3.34 twice, one gets $\mathcal{M}'_k(\Psi_0) = p\mathcal{M}'_k(\Psi''_0)$ and $\mathcal{M}'_k(\mathbb{Z}^k) = (1-p)\mathcal{M}'_k(\Psi''_0)$. Therefore, $\mathcal{M}'_k(\Psi_0) = p/(1-p)\mathcal{M}'_k(\mathbb{Z}^k)$. By the definition of $\mathcal{M}'_k$, one can directly show that $\mathcal{M}'_k(\mathbb{Z}^k) = 2^k$ (see also Proposition 3.29). This implies the claim. $\quad\square$

The last claim of Proposition 5.16 suggests the following, which is verified when $k = 1$ in the next proposition.

Conjecture 5.17. *If Φ is a point-stationary point process in $\mathbb{R}^k$ which is not the Palm version of any stationary point process, then $\mathcal{M}^k(\Phi) = 0$.*

Proposition 5.18. *Conjecture 5.17 is true when $k = 1$.*

Proof. Denote Φ as $\Phi = \{S_n : n \in \mathbb{Z}\}$ such that $S_0 = 0$ and $S_n < S_{n+1}$ for each n. Then, the sequence $T_n := S_{n+1} - S_n$ is stationary under shifting the indices (see Example 2.7). The assumption that Φ is not the Palm version of a stationary point process is equivalent to $\mathbb{E}[S_1] = \infty$ (see [18] or Proposition 6 of [35]). Indeed, if $\mathbb{E}[S_1] < \infty$, then one could bias the probability measure by S_1 (Definition 2.1) and then shift the whole process by $-\boldsymbol{U}$, where $\boldsymbol{U} \in [0, S_1]$ is chosen uniformly and independently.

Since $\mathbb{E}[S_1] = \infty$, Birkhoff's pointwise ergodic theorem [44] implies that $\lim_n (T_1 + \cdots + T_n)/n = \infty$. This in turn implies that $\lim_r \#N_r(0)/r = 0$. Therefore, Lemma 5.3 gives that $\mathcal{H}^1_\infty(\Phi) = \infty$; i.e., $\mathcal{M}^1(\Phi) = 0$. $\quad\square$

5.4 Connections to Birkhoff's Pointwise Ergodic Theorem

The following corollary of the unimodular Billingsley lemma is of independent interest. Note that the statement does not involve dimension.

Theorem 5.19. *Let $[\boldsymbol{D}, \boldsymbol{o}]$ be a unimodular discrete space. For any two equivariant weight functions $\boldsymbol{w}_1$ and $\boldsymbol{w}_2$ such that $\mathbb{P}[\exists v \in \boldsymbol{D} : \boldsymbol{w}_2(v) \neq 0] = 1$ and $\mathbb{E}[\boldsymbol{w}_1(\boldsymbol{o})] < \infty$, one has*

$$\underline{\text{growth}}(\boldsymbol{w}_1(N_r(\boldsymbol{o}))) \leq \overline{\text{growth}}(\boldsymbol{w}_2(N_r(\boldsymbol{o}))), \quad a.s.$$

In particular, if $\boldsymbol{w}_1(N_r(\boldsymbol{o}))$ and $\boldsymbol{w}_2(N_r(\boldsymbol{o}))$ have well defined growth rates, then their growth rates are equal.

Note that the condition $\mathbb{E}[\boldsymbol{w}_1(\boldsymbol{o})] < \infty$ is necessary as shown in Remark 5.9.

Proof. Let $\epsilon > 0$ be arbitrary and

$$A := \{[D, o] \in \mathcal{D}_* : \underline{\text{growth}}(\boldsymbol{w}_1(N_r(o))) > \overline{\text{growth}}(\boldsymbol{w}_2(N_r(o))) + \epsilon\}.$$

It can be seen that A is a measurable subset of $\mathcal{D}_*$. Assume $\mathbb{P}[[\boldsymbol{D}, \boldsymbol{o}] \in A] > 0$. Denote by $[\boldsymbol{D}', \boldsymbol{o}']$ the random pointed discrete space obtained by conditioning $[\boldsymbol{D}, \boldsymbol{o}]$ on A. Since A does not depend on the root (i.e., if $[D, o] \in A$, then $\forall v \in D : [D, v] \in A$), by a direct verification of the mass transport principle (2.1), one can show that $[\boldsymbol{D}', \boldsymbol{o}']$ is unimodular. So by using the unimodular Billingsley lemma (Theorem 5.6) twice, one gets

$$\text{ess inf}\left(\underline{\text{growth}}(\boldsymbol{w}_1(N_r(\boldsymbol{o}')))\right) \leq \text{udim}_H(\boldsymbol{D}') \leq \text{ess inf}\left(\overline{\text{growth}}(\boldsymbol{w}_2(N_r(\boldsymbol{o}')))\right).$$

By the definition of A, this contradicts the fact that $[\boldsymbol{D}', \boldsymbol{o}'] \in A$ a.s. So $\mathbb{P}[[\boldsymbol{D}, \boldsymbol{o}] \in A] = 0$ and the claim is proved. $\quad\square$

Remark 5.20. Theorem 5.19 is a generalization of a weaker form of Birkhoff's pointwise ergodic theorem as explained below. In the cases where $\boldsymbol{D}$ is either $\mathbb{Z}$, the Palm version of a stationary point process in $\mathbb{R}^k$ or a point-stationary point process in $\mathbb{R}$, Birkhoff's pointwise ergodic theorem (or its generalizations) implies that $\lim \boldsymbol{w}_1(N_r(\boldsymbol{o}))/\boldsymbol{w}_2(N_r(\boldsymbol{o})) = \mathbb{E}[\boldsymbol{w}_1(0)]/\mathbb{E}[\boldsymbol{w}_2(0)]$ a.s. This is stronger than the claim of Theorem 5.19. Note that Theorem 5.19 implies nothing about $\lim \boldsymbol{w}_1(N_r(\boldsymbol{o}))/\boldsymbol{w}_2(N_r(\boldsymbol{o}))$. On the other side, note

that *amenability* is not assumed in this Theorem, which is a general requirement in the study of ergodic theorems. However, it will be proved in [6] that, roughly speaking, non-amenability implies $\overline{\text{growth}}\,(w_2(N_r(o))) = \infty$, which makes the claim of Theorem 5.19 trivial in this case. In this case, using exponential gauge functions seems more interesting.

Problem 5.21. Is it true that for every unimodular discrete space $[\boldsymbol{D}, \boldsymbol{o}]$, the growth rates $\overline{\text{growth}}\,(\boldsymbol{w}(N_r(\boldsymbol{o})))$ and $\underline{\text{growth}}\,(\boldsymbol{w}(N_r(\boldsymbol{o})))$ do not depend on $\boldsymbol{w}$ as long as $0 < \mathbb{E}\,[\boldsymbol{w}(\boldsymbol{o})] < \infty$?

5.5 Notes and Bibliographical Comments

As already mentioned, the unimodular mass distribution principle and the unimodular Billingsley lemma have analogues in the continuum setting (see e.g., [14]) and are named accordingly. Note however that there is no direct or systematic reduction to these continuum results. For instance, in the continuum setting, one should assume that the space under study is a subset of the Euclidean space, or more generally, satisfies the *bounded subcover property* (see e.g., [14]). Theorem 5.6 does not require such assumptions. Note also that the term $\overline{\text{growth}}\,(\boldsymbol{w}(N_r(\boldsymbol{o})))$ in Theorem 5.6 does not depend on the origin in contrast to the analogous term in the continuum version. Similar observations can be made on Theorem 5.2.

6 Examples Continued

This section presents further examples for illustrating the results of the previous section.

6.1 Remaining Proofs from Section 4

The unimodular Billingsley lemma can be used to complete the computation of the unimodular Hausdorff dimension in the examples of Section 4. These examples include Eternal GW trees, the image of a random walk and the drainage network model.

Proof of Proposition 4.8 (second part). The equality $\text{udim}_M(\boldsymbol{T}) = 2$ is proved in Subsection 4.2.3. So it remains to prove $\text{udim}_H(\boldsymbol{T}) \leq 2$. By the unimodular Billingsley lemma, it is enough to show that $\mathbb{E}\,[\#N_n(\boldsymbol{o})] \leq cn^2$ for a constant c. Recall from Subsection 4.1.2 that $F(v)$ represents the parent of vertex v and $D(v)$ denotes the subtree of descendants of v. Write $N_n(\boldsymbol{o}) = Y_0 \cup Y_1 \cup \cdots \cup Y_n$, where $Y_n := N_n(\boldsymbol{o}) \cap D(\boldsymbol{o})$ and $Y_i := N_n(\boldsymbol{o}) \cap D(F^{n-i}(\boldsymbol{o})) \setminus D(F^{n-i-1}(\boldsymbol{o}))$ for $0 \leq i < n$. By the explicit construction of EGW trees in [5], Y_n is a critical Galton-Watson tree up to generation n. Also, for $0 \leq i < n$, Y_i has the same structure up to generation i, except that the distribution of the first generation is *size-biased minus one* (i.e., $(np_{n+1})_n$ with the notation of Subsection 6.3.2). So the assumption of finite variance implies that the first generation in each Y_i has finite mean, namely m'. Now, one can inductively show that $\mathbb{E}\,[\#Y_n] = n$ and $\mathbb{E}\,[\#Y_i] = im'$, for $0 \leq i < n$. It follows that $\mathbb{E}\,[\#N_n(\boldsymbol{o})] \leq (1 + m')n^2$ and the claim is proved. $\qquad\square$

Proof of Proposition 4.9 (second part). In Subsection 4.3.1, it is proved that $\underline{\text{udim}}_M(\Phi) \geq 1 \wedge \underline{\text{decay}}\,(\mathbb{P}\,[S_1 > r])$. So part (iv) is implied by part (iii), which is proved below. Since Φ is a point-stationary point process in $\mathbb{R}$ (see Subsection 4.3.1), Proposition 5.15 implies that $\text{udim}_H(\Phi) \leq 1$. Now, assume $\overline{\text{decay}}\,(\mathbb{P}\,[S_1 > r]) < \beta$. Then, there exists $c > 0$ such that $\mathbb{P}\,[S_1 > r] > cr^{-\beta}$ for all $r \geq 1$. This implies that there exists $C < \infty$ and a random number $r_0 > 0$ such that for all $r \geq r_0$, one has $\#N_r(\boldsymbol{o}) \leq Cr^\beta \log \log r$ a.s. (for the proof, see [7] or Theorem 4 of [26]). Therefore, the unimodular Billingsley lemma

(Theorem 5.6) implies that $\mathrm{udim}_H(\Phi) \leq \beta + \epsilon$ for every $\epsilon > 0$, which in turn implies that $\mathrm{udim}_H(\Phi) \leq \beta$. $\qquad\square$

Example 6.1 (Infinite H-Size). In Proposition 4.9, assume that $\mathbb{P}[S_1 > r] = 1/\log r$ for large enough r. Then, part (iii) of the proposition implies that $\mathrm{udim}_H(\Phi) = 0$. However, since Φ is infinite a.s., it has infinite 0-dim H-size (Proposition 3.28).

Example 6.2 (Zero H-Size). In Proposition 4.9, assume $\mathbb{P}[S_1 > r] = 1/r$ for large enough r. Then, part (iii) of the proposition implies that $\mathrm{udim}_H(\Phi) = 1$. Since $\mathbb{E}[S_1] = \infty$, Φ is not the Palm version of any stationary point process (see Proposition 5.18). Therefore, Proposition 5.18 implies that $\mathcal{M}^1(\Phi) = 0$.

Proof of Proposition 4.13 (second part). The equality $\mathrm{udim}_M(\boldsymbol{T}) = \frac{3}{2}$ is proved in Subsection 4.5. So it remains to prove $\mathrm{udim}_H(\boldsymbol{T}) \leq \frac{3}{2}$. To use Billingsley's lemma, an upper bound on $\mathbb{E}[\#N_n(\boldsymbol{o})]$ is derived. Let $e_{k,l} := \#\left(F^{-k}(F^l(\boldsymbol{o})) \setminus F^{-(k-1)}(F^{l-1}(\boldsymbol{o}))\right)$ be the number of descendants of order k of $F^l(\boldsymbol{o})$ which are not a descendant of $F^{l-1}(\boldsymbol{o})$ (for $l = 0$, let it be just $\#F^{-k}(\boldsymbol{o})$). One has $\#N_n(\boldsymbol{o}) = \sum_{k,l} e_{k,l} 1_{\{k+l \leq n\}}$. It can be seen that $\mathbb{E}[e_{k,l}]$ is equal to the probability that two independent paths of length k and l starting both at $\boldsymbol{o}$ do not collide at another point. Therefore, $\mathbb{E}[e_{k,l}] \leq c(k \wedge l)^{-\frac{1}{2}}$ for some c and all k, l. This implies that (in the following, c is updated at each step to a new constant without changing the notation)

$$\mathbb{E}\left[\sum_{k,l \geq 0} e_{k,l} 1_{\{k+l \leq n\}}\right] \leq \sum_{k=0}^{\lfloor \frac{n}{2} \rfloor} ck^{-\frac{1}{2}}(n-k) \leq cn \sum_{k=0}^{\lfloor \frac{n}{2} \rfloor} k^{-\frac{1}{2}} \leq cn^{\frac{3}{2}}.$$

The above inequalities imply that $\mathbb{E}[\#N_n(\boldsymbol{o})] \leq cn^{\frac{3}{2}}$ for some c and all n. Therefore, the unimodular Billingsley lemma (Theorem 5.6) implies that $\mathrm{udim}_H(\boldsymbol{T}) \leq \frac{3}{2}$. So the claim is proved. $\qquad\square$

6.2 General Unimodular Trees Continued

The following is a direct corollary of Theorem 4.2 and the unimodular Billingsley lemma. Since the statement does not involve dimension, it is of independent interest and believed to be new.

Corollary 6.3. *For every unimodular one-ended tree $[\boldsymbol{T}, \boldsymbol{o}]$ and every equivariant weight function $\boldsymbol{w}$, almost surely,*

$$\overline{\mathrm{decay}}\left(\mathbb{P}[h(\boldsymbol{o}) = n]\right) \leq \overline{\mathrm{growth}}\left(\boldsymbol{w}(N_r(\boldsymbol{o}))\right) \leq \overline{\mathrm{growth}}\left(\mathbb{E}[\boldsymbol{w}(N_r(\boldsymbol{o}))]\right).$$

The rest of this subsection is focused on unimodular trees with infinitely many ends.

Proposition 6.4. *Let $[\boldsymbol{T}, \boldsymbol{o}]$ be a unimodular tree with infinitely many ends such that $\mathbb{E}[\deg(\boldsymbol{o})] < \infty$. Then $\boldsymbol{T}$ has exponential volume growth a.s. and $\mathrm{udim}_H(\boldsymbol{T}) = \infty$.*

In fact, the graph-distance metric on $\boldsymbol{T}$ can be replaced by an arbitrary equivariant metric. This will be proved in [6].

The following proof uses the definitions and results of [2], but they are not recalled for brevity.

Proof of Proposition 6.4. By Corollary 8.10 of [2], $[\boldsymbol{T}, \boldsymbol{o}]$ is *non-amenable* (this will be discussed further in [6]). So Theorem 8.9 of [2] implies that the critical probability p_c of percolation on $\boldsymbol{T}$ is less than one with positive probability. In fact, it can be shown that $p_c < 1$ a.s. (if not, condition on the event $p_c = 1$ to get a contradiction). For any tree, p_c is equal to the inverse of the *branching number*. So the branching number is more than one, which implies that the tree has exponential volume growth. Finally, the unimodular Billingsley lemma (Theorem 5.6) implies that $\mathrm{udim}_H(\boldsymbol{T}) = \infty$. $\qquad\square$

The following example shows that the Minkowski dimension can be finite.

Example 6.5. Let T be the 3-regular tree. Split each edge e by adding a random number $\boldsymbol{l}_e$ of new vertices and let $\boldsymbol{T}_0$ be the resulting tree. Let v_e be the middle vertex in this edge (assuming $\boldsymbol{l}_e$ is always odd) and assign marks by $\boldsymbol{m}_0(v_e) := \boldsymbol{l}_e$. Assume that the random variables $\boldsymbol{l}_e$ are i.i.d. If $\mathbb{E}\left[\boldsymbol{l}_e\right] < \infty$, then one can bias the probability measure and choose a new root to obtain a unimodular marked tree, namely $[\boldsymbol{T}, \boldsymbol{o}; \boldsymbol{m}]$ (see Example 9.8 of [2] or [35]). It will be shown below that $\overline{\mathrm{udim}}_M(\boldsymbol{T})$ may be finite.
Let $\boldsymbol{R}$ be an arbitrary equivariant r-covering of $\boldsymbol{T}$. Consider the set of middle vertices $\boldsymbol{A}_r := \{v \in \boldsymbol{T} : \boldsymbol{m}(v) \geq r\}$. Since these vertices have pairwise distance at least r, they belong to different balls in the covering. So, by the mass transport principle, one can show that $\rho(\boldsymbol{R}) \geq \rho(\boldsymbol{A}_r)$, where $\rho(\cdot) = \mathbb{P}\left[\boldsymbol{o} \in \cdot\right]$ denotes the intensity. On the other hand, let $\boldsymbol{S}$ be the equivariant subset of vertices with degree 3. Send unit mass from every point of $\boldsymbol{A}_r$ to its two closest points in $\boldsymbol{S}$. Then the mass transport principle implies that $2\rho(\boldsymbol{A}_r) = 3\rho(\boldsymbol{S})\mathbb{P}\left[\boldsymbol{l}_e \geq r\right]$. Hence, $\rho(\boldsymbol{R}) \geq \frac{3}{2}\rho(\boldsymbol{S})\mathbb{P}\left[\boldsymbol{l}_e \geq r\right]$. This gives that $\overline{\mathrm{udim}}_M(\boldsymbol{T}) \leq \overline{\mathrm{decay}}\left(\mathbb{P}\left[\boldsymbol{l}_e \geq r\right]\right)$, which can be finite. In fact, if $\mathrm{decay}\left(\mathbb{P}\left[\boldsymbol{l}_e \geq r\right]\right)$ exists, Proposition 6.6 below implies that $\mathrm{udim}_M(\boldsymbol{T}) = \mathrm{decay}\left(\mathbb{P}\left[\boldsymbol{l}_e \geq r\right]\right)$.

The following proposition gives a lower bound on the Minkowski dimension.

Proposition 6.6. *Let $[\boldsymbol{T}, \boldsymbol{o}]$ be a unimodular tree with infinitely many ends and without leaves. Let $\boldsymbol{S}$ be the equivariant subset of vertices of degree at least 3. For every $v \in \boldsymbol{S}$, let $\boldsymbol{w}(v)$ be the sum of the distances of v to its neighbors in $\boldsymbol{S}$. If $\mathbb{E}\left[\boldsymbol{w}(\boldsymbol{o})^\alpha\right] < \infty$, then $\underline{\mathrm{udim}}_M(\boldsymbol{T}) \geq \alpha$.*

The proof is based on the following simpler result. This will be used in Subsection 6.3.3 as well.

Proposition 6.7. *Let $[\boldsymbol{T}, \boldsymbol{o}]$ be a unimodular tree such that the degree of every vertex is at least 3. Let $\boldsymbol{d}'$ be an equivariant metric on $\boldsymbol{T}$. Let $\boldsymbol{w}(v) := \sum_u \boldsymbol{d}'(v, u)$, where the sum is over the 3 neighbors of v which are closest to v under the metric $\boldsymbol{d}'$. If $\mathbb{E}\left[\boldsymbol{w}(\boldsymbol{o})^\alpha\right] < \infty$, then $\underline{\mathrm{udim}}_M(\boldsymbol{T}, \boldsymbol{d}') \geq \alpha$.*

Proof. Define $\boldsymbol{w}'(v) := \sum_u \boldsymbol{d}'(u, v)^\alpha$, where the sum is over the three closest neighbors of v. It is enough to assume that $\boldsymbol{d}'$ is generated by equivariant edge lengths since increasing the edge lengths does not increase the dimension (by Theorem 3.31). By the same argument, it is enough to assume $\boldsymbol{d}'(u, v) \geq 1$ for all $u \sim v$. Then, it can be seen that there exists a constant c, that depends only on α, such that $\boldsymbol{w}'(N_r(v)) \geq cr^\alpha$ for all $v \in \boldsymbol{T}$ and $r \geq 0$ (see Lemma C.5 in [7]). Also, the assumption implies that $\mathbb{E}\left[\boldsymbol{w}'(\boldsymbol{o})\right] < \infty$. So Lemma 5.4 implies that $\underline{\mathrm{udim}}_M(T_3, \boldsymbol{d}') \geq \alpha$ and the claim is proved. $\square$

Proof of Proposition 6.6. For $v \in \boldsymbol{S}$, let $\boldsymbol{w}'(v) := \sum_u d(u, v)^\alpha$, where the sum is over the neighbors of v in $\boldsymbol{S}$. For $v \in \boldsymbol{T} \setminus \boldsymbol{S}$, if u_1 and u_2 are the two closest points of $\boldsymbol{S}$ to v, let $g(v, u_i) := d(u_i, v)^{\alpha-1}$ and $\boldsymbol{w}'(v) := g(v, u_1) + g(v, u_2)$. The assumption implies that $\mathbb{E}\left[\boldsymbol{w}'(\boldsymbol{o})\right] < \infty$ (use the mass transport principle for g defined above). Similarly to Proposition 6.7, there exists $c = c(\alpha)$, such that $\boldsymbol{w}'(N_r(v)) \geq cr^\alpha$ for all $v \in \boldsymbol{T}$ and $r \geq 0$ (see Lemma C.5 in [7]) and the claim is proved. $\square$

6.3 Instances of Unimodular Trees Continued

6.3.1 A Unimodular Tree With No Volume Growth Rate

Recall the generalized canopy tree $[\boldsymbol{T}, \boldsymbol{o}]$ from Subsection 4.2.2. Here, it is shown that $\underline{\mathrm{growth}}(\boldsymbol{T}) \neq \overline{\mathrm{growth}}(\boldsymbol{T})$ if the parameters are suitably chosen. Similarly, it provides an example where the exponential growth rate does not exist. The existence of unimodular trees without exponential growth rate is already proved in [49], but with a more difficult construction.

Choose the sequence $(p_n)_n$ in the definition of $[\boldsymbol{T}, \boldsymbol{o}]$ such that $p_n = c2^{-q_n}$ and $\sum_n p_n = 1$, where c is constant and $q_0 \leq q_1 \leq \cdots$ is a sequence of integers. In this case, $\boldsymbol{T}$ is obtained by splitting the edges of the canopy tree by adding new vertices or concatenating them, depending only on the *level* of the edges. It can be seen that if v is a vertex in the n-th level of $\boldsymbol{T}$, then the number of descendants of v is $(p_0 + \cdots + p_n)/p_n$. It follows that $\underline{\text{growth}}(\boldsymbol{T}) = \underline{\text{decay}}(p_n)$ and $\overline{\text{growth}}(\boldsymbol{T}) = \overline{\text{decay}}(p_n)$. So, by choosing $(p_n)_n$ appropriately, $\boldsymbol{T}$ can have no polynomial (or exponential) volume growth rate. This proves the claim. Note also that the unimodular Billingsley lemma and Theorem 4.2 imply that $\text{udim}_H(\boldsymbol{T}) = \overline{\text{growth}}(\boldsymbol{T})$ here.

6.3.2 Unimodular Galton-Watson Trees

Here, it is shown that the unimodular Galton-Watson tree [2] is infinite dimensional (note that this tree differs from the Eternal Galton-Watson tree of Subsection 4.2.3 which is a directed tree). Consider an ordinary Galton-Watson tree with offspring distribution $\mu = (p_0, p_1, \ldots)$, where μ is a probability measure on $\mathbb{Z}^{\geq 0}$. The unimodular Galton-Watson tree $[\boldsymbol{T}, \boldsymbol{o}]$ has a similar construction with the difference that the offspring distribution of the origin is different from that of the other vertices: It has for distribution the size-biased version $\hat{\mu} = (\frac{n}{m} p_n)_n$, where m is the mean of μ (assumed to be finite).

In what follows, the trivial case $p_1 = 1$ is excluded. If $m \leq 1$, then $\boldsymbol{T}$ is finite a.s.; i.e., there is extinction a.s. Therefore, $\text{udim}_H(\boldsymbol{T}) = 0$. So assume the *supercritical case*, namely $m > 1$. If $p_0 > 0$, then $\boldsymbol{T}$ is finite with positive probability. So $\text{udim}_H(\boldsymbol{T}) = 0$ for the same reason. Nevertheless, one can condition on non-extinction as follows.

Proposition 6.8. *Let $[\boldsymbol{T}, \boldsymbol{o}]$ be a supercritical unimodular Galton-Watson tree conditioned on non-extinction. Then, $\text{udim}_M(\boldsymbol{T}) = \text{udim}_H(\boldsymbol{T}) = \infty$.*

Proof. The result for the Hausdorff dimension follows from the unimodular Billingsley lemma (Theorem 5.6) and the Kesten-Stigum theorem [34], which implies that $\lim_n \#N_n(\boldsymbol{o})m^{-n}$ exists and is positive a.s. Computing the Minkowski dimension is more difficult. By part (iv) of Lemma 5.4, it is enough to prove that $\mathbb{E}\left[(1 - n^{-\alpha})^{\#N_n(\boldsymbol{o})}\right]$ has infinite decay rate for every $\alpha \geq 0$. Denote by $[\widetilde{\boldsymbol{T}}, \widetilde{\boldsymbol{o}}]$ the Galton-Watson tree with the same parameters. Using the fact that $\#N_n(\boldsymbol{o})$ is stochastically larger than $\#N_{n-1}(\widetilde{\boldsymbol{o}})$, one gets that it is enough to prove the last claim for $[\widetilde{\boldsymbol{T}}, \widetilde{\boldsymbol{o}}]$.

For simplicity, the proof is given for the case $p_0 = 0$ only. By this assumption, the probability of extinction is zero. The general case can be proved with similar arguments and by using the decomposition theorem of supercritical Galton-Watson trees (see e.g., Theorem 5.28 of [41]). In fact, the following proof implies the general claim by the fact that the *trunk*, conditioned on non-extinction, is another supercritical unimodular Galton-Watson tree. The latter can be proved similarly to the decomposition theorem.

Let $f(s) := \sum_n p_n s^n$ be the generating function of μ. By classical results of the theory of branching processes, for all $s \leq 1$, $\mathbb{E}\left[s^{d_n(\widetilde{\boldsymbol{o}})}\right] = f^{(n)}(s)$, where $d_n(\widetilde{\boldsymbol{o}}) := \#N_n(\widetilde{\boldsymbol{o}}) - \#N_{n-1}(\widetilde{\boldsymbol{o}})$ and $f^{(n)}$ is the n-fold composition of f with itself. Let $a > 0$ be fixed and $g(s) := \frac{as}{-s+a+1}$ (such functions are frequently used in the literature on branching processes; see, e.g., [4]). One has $f(0) = g(0) = 0$, $f(1) = g(1) = 1$, $f'(1) = m > 1$, $g'(1) = (1 + a)/a$, and f is convex. Therefore, a can be chosen large enough such that $f(s) \leq g(s)$ for all $s \in [0, 1]$. So

$$f^{(n)}(s) \leq g^{(n)}(s) = \frac{a^n s}{a^n + (a+1)^n(1-s)},$$

where the last equality can be checked by induction. Therefore,

$$f^{(n)}(1 - n^{-\alpha}) \leq \frac{a^n}{a^n + n^{-\alpha}(a+1)^n}.$$

It follows that $\mathrm{decay}\left(f^{(n)}(1-n^{-\alpha})\right)=\infty$. Therefore, the above discussion gives that $\mathbb{E}\left[(1-n^{-\alpha})^{\#N_n(o)}\right]$ has infinite decay rate and the claim is proved. $\qquad\square$

6.3.3 The Poisson Weighted Infinite Tree

The Poisson Weighted Infinite Tree (PWIT) is defined as follows (see e.g., [3]). It is a rooted tree $[\boldsymbol{T},\boldsymbol{o}]$ such that the degree of every vertex is infinite. Regarding $\boldsymbol{T}$ as a family tree with progenitor $\boldsymbol{o}$, the edge lengths are as follows. For every $u\in\boldsymbol{T}$, the set $\{d(u,v):v$ is an offspring of $u\}$ is a Poisson point process on $\mathbb{R}^{\geq0}$ with intensity function x^k, where $k>0$ is a given integer. Moreover, for different vertices u, the corresponding Poisson point processes are jointly independent. It is known that the PWIT is unimodular (notice that although each vertex has infinite degree, the PWIT is boundedly-finite as a metric space). See for example [3] for more details.

Proposition 6.9. *The PWIT satisfies* $\mathrm{udim}_M(\mathit{PWIT})=\mathrm{udim}_H(\mathit{PWIT})=\infty$.

Proof. Denote the neighbors of $\boldsymbol{o}$ by $v_1,v_2,\ldots$ such that $d(\boldsymbol{o},v_i)$ is increasing in i. It is straightforward that all moments of $d(\boldsymbol{o},v_3)$ are finite. Therefore, Proposition 6.7 implies that $\underline{\mathrm{udim}}_M(\boldsymbol{T})=\infty$ (see also Lemma C.5 in [7]). This proves the claim. $\qquad\square$

6.4 The Graph of the Simple Random Walk

As in Subsection 4.3, consider the simple random walk $(S_n)_{n\in\mathbb{Z}}$ in $\mathbb{R}^k$, where $S_0=0$ and the increments S_n-S_{n-1} are i.i.d. The graph of the random walk $(S_n)_{n\in\mathbb{Z}}$ is $\Psi:=\{(n,S_n):n\in\mathbb{Z}\}\subseteq\mathbb{R}^{k+1}$. It can be seen that Ψ is a point-stationary point process, and hence, $[\Psi,0]$ is unimodular (see Subsection 4.3.1).

Since $\#\Psi\cap[-n,n]^{k+1}\leq2n+1$, the mass distribution principle (Theorem 5.2) implies that $\mathrm{udim}_H(\Psi)\leq1$. In addition, if S_1 has finite first moment, then the strong law of large numbers implies that $\lim_n\frac{1}{n}S_n=\mathbb{E}\left[S_1\right]$. This implies that $\liminf_n\frac{1}{n}\#\left(\Psi\cap[-n,n]^{k+1}\right)>0$. Therefore, the unimodular Billingsley lemma (Theorem 5.6) implies that $\mathrm{udim}_H(\Psi)\geq1$. Hence, $\mathrm{udim}_H(\Psi)=1$. This matches the result of [38] that the *macroscopic dimension* of the graph of the Brownian motion is 1, while its *microscopic dimension* is $3/2$ when $k=1$ (see Subsection 8.1).

Below, the focus is on the case $k=1$ and on the following metric:

$$d((x,y),(x',y')):=\max\{\sqrt{|x-x'|},|y-y'|\}. \tag{6.1}$$

Theorem 3.31 implies that, by considering this metric, unimodularity is preserved and dimension is not decreased. Under this metric, the ball $N_n(0)$ is $\Psi\cap[-n^2,n^2]\times[-n,n]$. It is straightforward that $\mathbb{Z}^2$ has volume growth rate 3 and also Minkowski and Hausdorff dimension 3 under this metric.

Proposition 6.10. *If the jumps are ±1 uniformly, under the metric* (6.1)*, the graph Ψ of the simple random walk satisfies* $\mathrm{udim}_M(\Psi)=\mathrm{udim}_H(\Psi)=2$.

Proof. Let $n\in\mathbb{N}$. The ball $N_n(0)$ has at most $2n^2+1$ elements. So the mass distribution principle (Theorem 5.2) implies that $\mathrm{udim}_H(\Psi)\leq2$. For the other side, let $\mathcal{C}$ be the equivariant disjoint covering of $\mathbb{Z}^2$ by translations of the rectangle $[-n^2,n^2]\times[-n,n]$ (similar to Example 3.11). For each rectangle $\sigma\in\mathcal{C}$, select the right-most point in $\sigma\cap\Psi$ and let $\boldsymbol{S}=\boldsymbol{S}_\Psi$ be the set of selected points. By construction, $\boldsymbol{S}$ gives an n-covering of Ψ and it can be seen that it is an equivariant covering. Let σ_0 be the rectangle containing the origin. By construction, $0\in\boldsymbol{S}$ if and only if it is either on a right-edge of σ_0 or on a horizontal edge of σ_0 and the random walk stays outside σ_0. The first case happens with probability $1/(2n^2+1)$. By classical results concerning the hitting time of random walks, one can obtain that the probability of the second case lies between two constant

multiples of n^{-2}. It follows that $\mathbb{P}\left[0 \in \boldsymbol{S}\right]$ lies between two constant multiples of n^{-2}. Therefore, $\underline{\dim}_M(\Psi) \geq 2$. This proves the claim. $\qquad\square$

6.5 Other Self Similar Unimodular Spaces

In this subsection, two examples are presented which have some kind of self-similarity heuristically, but do not fit into the framework of Subsection 4.6.

6.5.1 Unimodular Discrete Spaces Defined by Digit Restriction

Let $J \subseteq \mathbb{Z}^{\geq 0}$. For $n \geq 0$, consider the set of natural numbers with expansion $(a_n a_{n-1} \ldots a_0)$ in base 2 such that $a_i = 0$ for every $i \notin J$. Similarly to the examples in Subsection 4.6, one can shift this set randomly and take a limit to obtain a unimodular discrete space. This can be constructed in the following way as well: Let $\boldsymbol{T}_0 := \{0\}$. If $n \in J$, let $\boldsymbol{T}_{n+1} := \boldsymbol{T}_n \cup (\boldsymbol{T}_n \pm 2 \times 2^n)$, where the sign is chosen i.i.d., each sign with probability $1/2$. If $n \notin J$, let $\boldsymbol{T}_{n+1} := \boldsymbol{T}_n$. Finally, let $\Psi := \cup_n \boldsymbol{T}_n$.

The upper and lower asymptotic densities of J in $\mathbb{Z}^{\geq 0}$ are defined by $\overline{d}(J) := \limsup_n \frac{1}{n} J_n$ and $\underline{d}(J) := \liminf_n \frac{1}{n} J_n$, where $J_n := \#J \cap \{0, \ldots, n\}$.

Proposition 6.11. *Almost surely,*

$$\dim_H(\Psi) = \overline{\dim}_M(\Psi) = \overline{\text{growth}}\,(\#N_n(\boldsymbol{o})) \;\; = \;\; \overline{d}(J),$$
$$\underline{\dim}_M(\Psi) = \underline{\text{growth}}\,(\#N_n(\boldsymbol{o})) \;\; = \;\; \underline{d}(J).$$

In particular, this provides another example of a unimodular discrete space where the (polynomial) volume growth rate does not exist.

Proof. Let $n \geq 0$ be given. Cover $\boldsymbol{T}_n$ by a ball of radius 2^n centered at the minimal element of $\boldsymbol{T}_n$. By the same recursive definition, one can cover $\boldsymbol{T}_{n+1}$ by either 1 or 2 balls of the same radius. Continuing the recursion, an equivariant 2^n-covering $\boldsymbol{R}_n$ is obtained. It is straightforward to see that $\mathbb{P}\left[\boldsymbol{R}_n(\boldsymbol{o}) > 0\right] = 2^{-J_n}$. Since these coverings are uniformly bounded (Definition 3.8), Lemma 3.9 implies that $\underline{\dim}_M(\Psi) = \underline{d}(J)$ and $\overline{\dim}_M(\Psi) = \overline{d}(J)$. One has

$$\#\boldsymbol{T}_m = 2^{J_m}. \tag{6.2}$$

This implies that $\#N_{2^n}(\boldsymbol{o}) \leq 2^{J_n+1}$. One can deduce that $\overline{\text{growth}}\,(\#N_n(\boldsymbol{o})) \leq \overline{d}(J)$. So the unimodular Billingsley lemma (Theorem 5.6) gives $\dim_H(\Psi) \leq \overline{d}(J)$. This proves the claim. $\qquad\square$

6.5.2 Randomized Discrete Cantor set

This subsection proposes a unimodular discrete analogue of the *random Cantor set*, recalled below. Let $0 \leq p \leq 1$ and $b > 1$. The random Cantor set in $\mathbb{R}^k$ [32] (see also [14]) is defined by $\Lambda_k(b,p) := \cap_n E_n$, where E_n is defined by the following random algorithm: Let $E_0 := [0,1]^k$. For each $n \geq 0$ and each *b-adic* cube of edge length b^{-n} in E_n, divide it into b^k smaller b-adic cubes of edge length b^{-n-1}. Keep each smaller b-adic cube with probability p and delete it otherwise independently from the other cubes. Let E_{n+1} be the union of the kept cubes. It is shown in Section 3.7 of [14] that $\Lambda_k(b,p)$ is empty for $p \leq b^{-k}$ and otherwise, has dimension $k + \log_b p$ conditioned on being non-empty.

For each $n \geq 0$, let $\boldsymbol{K}_n$ be the set of lower left corners of the b-adic cubes forming E_n. It is easy to show that $\boldsymbol{K}_n$ tends to $\Lambda_k(b,p)$ a.s. under the Hausdorff metric.

Proposition 6.12. *Let $\boldsymbol{K}'_n$ denote the random set obtained by biasing the distribution of $\boldsymbol{K}_n$ by $\#\boldsymbol{K}_n$ (Definition 2.1). Let $\boldsymbol{o}'_n$ be a point chosen uniformly at random in $\boldsymbol{K}'_n$.*

(i) $[b^n \boldsymbol{K}'_n, \boldsymbol{o}'_n]$ converges weakly to some unimodular discrete space $[\hat{\boldsymbol{K}}, \hat{\boldsymbol{o}}]$.

(ii) If $p < b^{-k}$, then $\hat{K}$ is finite a.s., hence, $\mathrm{udim}_H(\hat{K}) = 0$ a.s.

(iii) If $p \geq b^{-k}$, then $\hat{K}$ is infinite a.s. and

$$\mathrm{udim}_H(\hat{\boldsymbol{K}}) = \mathrm{udim}_M(\hat{\boldsymbol{K}}) = k + \log_b p, \quad a.s.$$

Note that in contrast to the continuum analogue [32], for $p = b^{-k}$, the set is non-empty and even infinite, though still zero dimensional. Also, for $p < b^{-k}$ the set is non-empty as well.

To prove the above proposition, the following construction of $\hat{K}$ will be used. First, consider the usual nested sequence of partitions Π_n of $\mathbb{Z}^k$ by translations of the cube $\{0, \ldots, b^n - 1\}^k$, where $n \geq 0$. To make it *stationary*, shift each Π_n randomly as follows. Let $a_0, a_1, \ldots \in \{0, 1, \ldots, b-1\}^k$ be i.i.d. uniform numbers and let $\boldsymbol{U}_n = \sum_{i=0}^n a_i b^i \in \mathbb{Z}^k$. Shift the partition Π_n by the vector $\boldsymbol{U}_n$ to form a partition denoted by Π'_n. It is easy to see that Π'_n is a nested sequence of partitions.

Lemma 6.13. *Let $(\Pi'_n)_n$ be the stationary nested sequence of partitions of $\mathbb{Z}^k$ defined above. For each $n \geq 0$ and each cube $C \in \Pi'_n$ that does not contain the origin, with probability $1 - p$ (independently for different choices of C), mark all points in $C \cap \mathbb{Z}^k$ for deletion. Then, the set of the unmarked points of $\mathbb{Z}^k$, pointed at the origin, has the same distribution as $[\hat{K}, \hat{o}]$ defined in Proposition 6.12.*

Proof of Lemma 6.13. Let Φ be the set of unmarked points in the algorithm. For $n \geq 0$, let $\boldsymbol{C}_n$ be the cube in Π'_n that contains the origin. It is proved below that $\boldsymbol{C}_n \cap \Phi$ has the same distribution as $b^n(\boldsymbol{K}'_n - \boldsymbol{o}_n)$. This implies the claim.

Let $A_n \subseteq [0, 1]^k$ be the set of possible outcomes of $\boldsymbol{o}'_n$. One has $\#A_n = b^{kn}$. For $v \in A_n$, it is easy to see that the distribution of $b^n(\boldsymbol{K}'_n - \boldsymbol{o}_n)$, conditioned on $\boldsymbol{o}'_n = v$, coincides with the distribution of $\boldsymbol{C}_n \cap \Phi$ conditioned on $\boldsymbol{C}_n = b^n([0, 1)^k - v)$. So it remains to prove that $\mathbb{P}\left[\boldsymbol{o}'_n = v\right] = \mathbb{P}\left[\boldsymbol{C}_n = b^n([0, 1)^k - v)\right]$, which is left to the reader. $\square$

Here is another description of $\hat{K}$. The nested structure of $\bigcup_n \Pi'_n$ defines a tree as follows. The set of vertices is $\bigcup_n \Pi'_n$. For each $n \geq 0$, connect (the vertex corresponding to) every cube in Π'_n to the unique cube in Π'_{n+1} that contains it. This tree is the canopy tree (Subsection 4.2.1) with offspring cardinality $N := b^k$, except that the root (the cube $\{0\}$) is always a leaf. Now, keep each vertex with probability p and remove it with probability $1 - p$ in an i.i.d. manner. Let $\boldsymbol{T}$ be the connected component of the remaining graph that contains the root. Conditioned on the event that $\boldsymbol{T}$ is infinite, $\hat{K}$ corresponds to the set of leaves in the connected component of the root.

Proof of Proposition 6.12. The unimodular Billingsley lemma is used to get an upper bound on the Hausdorff dimension. For this $\mathbb{E}\left[\#N_{b^n}(\boldsymbol{o})\right]$ is studied. Consider the tree $[\boldsymbol{T}, \boldsymbol{o}]$ defined above and obtained by the percolation process on the canopy tree with offspring cardinality $N := b^k$. Let C be any cube in Π'_i that does not contain the origin. Note that the subtree of descendants of C in the percolation cluster (conditioned on keeping C) is a Galton-Watson tree with binomial offspring distribution with parameters (N, p). Classical results on branching processes say $\mathbb{E}\left[\#C \cap \hat{\boldsymbol{K}} \,|\Pi'_i\right] = pm^i$, where $m := pb^k$. So the construction implies that

$$\mathbb{E}\left[\#\boldsymbol{C}_n \cap \hat{\boldsymbol{K}}\right] = 1 + p(N-1)\left(m^{n-1} + m^{n-2} + \cdots + 1\right).$$

For $m > 1$, the latter is bounded by lm^n for some constant l not depending on n. Note that $N_{b^n}(\boldsymbol{o})$ is contained in the union of $\boldsymbol{C}_n$ and $3^k - 1$ other cubes in Π'_n. It follows that $\mathbb{E}\left[\#N_{b^n}(\boldsymbol{o})\right] \leq l'm^n$, where $l' = l + (3^k - 1)p$. So the unimodular Billingsley lemma

(Theorem 5.6) implies that $\text{udim}_H(\hat{\boldsymbol{K}}) \leq k + \log_b p$. The claim for $m = 1$ and $m < 1$ are similar.

Consider now the Minkowski dimension. As above, we assume $m > 1$ and the proofs for the other cases are similar. Let $n \geq 0$ be given. By considering the partition Π'_n by cubes, one can construct a b^n-covering $\boldsymbol{R}_n$ as in Theorem 5.13. This covering satisfies $\mathbb{P}[\boldsymbol{R}_n(\boldsymbol{o}) \geq 0] = \mathbb{E}\left[1/\#(\boldsymbol{C}_n \cap \hat{\boldsymbol{K}})\right]$. Let $[\boldsymbol{T}', \boldsymbol{o}']$ be the eternal Galton-Watson tree of Subsection 4.2.3 with binomial offspring distribution with parameters (N, p). By regarding $\boldsymbol{T}'$ as a family tree, it is straightforward that $[\boldsymbol{T}, \boldsymbol{o}]$ has the same distribution as the part of $[\boldsymbol{T}', \boldsymbol{o}']$, up to the generation of the root (see [5] for more details on eternal family trees). Therefore, Lemma 5.7 of [5] implies that $\mathbb{E}\left[1/\#(\boldsymbol{C}_n \cap \hat{\boldsymbol{K}})\right] = m^{-n}\mathbb{P}[h(\boldsymbol{o}') \geq n]$. Since $m > 1$, $\mathbb{P}[h(\boldsymbol{o}') \geq n]$ tends to the non-extinction probability of the descendants of the root, which is positive. By noticing the fact that the radii of the balls are b^n and the covering is uniformly bounded, one gets that $\text{udim}_M(\hat{\boldsymbol{K}}) = \log_b m = k + \log_b p$.

Finally, it remains to prove that $\hat{\boldsymbol{K}}$ is infinite a.s. when $p = b^{-k}$. In this case, consider the eternal Galton-Watson tree $[\boldsymbol{T}', \boldsymbol{o}']$ as above. Proposition 6.8 of [5] implies that the generation of the root is infinite a.s. This proves the claim. $\qquad\square$

6.6 Cayley Graphs

As mentioned in Subsection 3.6, the dimension of a Cayley graph depends only on the group and not on the generating set. The following result connects it to the volume growth rate of the group. Note that Gromov's theorem [27] implies that the polynomial growth degree exists and is either an integer or infinity.

Theorem 6.14. *For every finitely generated group H with polynomial growth degree $\alpha \in [0, \infty]$, one has $\text{udim}_M(H) = \text{udim}_H(H) = \alpha$. Also, if $\alpha < \infty$, then $\mathcal{M}^\alpha(H) < \infty$.*

Proof. First, assume $\alpha < \infty$. The result of Bass [10] implies that there are constants $c, C > 0$ such that $\forall r \geq 1 : cr^\alpha < \#N_r(o) \leq Cr^\alpha$, where o is an arbitrary element of H. So the mass distribution principle (Theorem 5.2) and part (i) of Lemma 5.4 imply that $\text{udim}_M(H) = \text{udim}_H(H) = \alpha$. In addition, (5.1) in the proof of Theorem 5.2 implies that $\mathcal{H}^\alpha_M(H) \geq 1/C$ for all $M \geq 1$, which implies that $\mathcal{M}^\alpha(H) \leq C < \infty$.

Second, assume $\alpha = \infty$. The result of [50] shows that for any $\beta < \infty$, $\#N_r(o) > r^\beta$ for sufficiently large r. Therefore, part (i) of Lemma 5.4 implies that $\underline{\text{udim}}_M(H) \geq \beta$. Hence, $\text{udim}_M(H) = \text{udim}_H(H) = \infty$ and the claim is proved. $\qquad\square$

It is natural to expect that $\mathcal{M}^\alpha(H) > 0$ as well, but only a weaker inequality will be proved in Proposition 7.18.

6.7 Notes and Bibliographical Comments

The proof of Proposition 6.4 was suggested by R. Lyons. Bibliographical comments on some of the examples discussed in this section can be found at the end of Section 4. The example defined by digit restriction (Subsection 6.5.1) is inspired by an example in the continuum setting (see e.g., Examples 1.3.2 of [14]). The randomized discrete Cantor set (Subsection 6.5.2) is inspired by the *random cantor set* (see e.g., Section 3.7 of [14]).

7 Frostman's Theory

This section provides a unimodular version of Frostman's lemma and some of its applications. In a sense to be made precise later, this lemma gives converses to the mass distribution principle. It is a powerful tool in the theoretical analysis of the unimodular Hausdorff dimension. For example, it is used in this section to derive inequalities for the

dimension of product spaces and *embedded spaces* (Subsections 7.4.2 and 7.4.3). It is also the basis of many of the results in [6].

7.1 Unimodular Frostman Lemma

The statement of the unimodular Frostman lemma requires the definition of *weighted Hausdorff* content. The latter is based on the notion of *equivariant weighted collections of balls* as follows. For this, the following mark space is needed. Let Ξ be the set of functions $c : \mathbb{R}^{\geq 0} \to \mathbb{R}^{\geq 0}$ which are positive in only finitely many points; i.e., $c^{-1}((0, \infty))$ is a finite set. Remark 7.9 below defines a metric on Ξ, so that the notion of Ξ-valued equivariant processes (Definition 2.9) is well defined. Such a process c is called an **equivariant weighted collection of balls** [2]. Consider a unimodular discrete space $[\boldsymbol{D}, \boldsymbol{o}]$ with distribution μ. For $v \in \boldsymbol{D}$, the reader can think of the value $\boldsymbol{c}_r(v) := \boldsymbol{c}(v)(r)$, if positive, to indicate that there is a ball in the collection, with radius r, centered at v, and with **cost** (or weight) $\boldsymbol{c}_r(v)$. Note that extra randomness is allowed in the definition. A ball-covering $\boldsymbol{R}$ can be regarded a special case of this construction by letting $\boldsymbol{c}_r(v)$ be 1 when $r = \boldsymbol{R}(v)$ and 0 otherwise.

Definition 7.1. *Let $f : \mathcal{D}_* \to \mathbb{R}$ be a measurable function and $M \geq 1$. An equivariant weighted collection of balls c is called a (f, M)-**covering** if*

$$\forall v \in \boldsymbol{D} : f(v) \leq \sum_{u \in \boldsymbol{D}} \sum_{r \geq M} \boldsymbol{c}_r(u) 1_{\{v \in N_r(u)\}}, \quad a.s., \tag{7.1}$$

where $f(v) := f[\boldsymbol{D}, v]$ for $v \in \boldsymbol{D}$. For $\alpha \geq 0$, define

$$
\begin{aligned}
\xi_M^\alpha(f) &:= \inf \left\{ \mathbb{E} \left[\sum_r \boldsymbol{c}_r(\boldsymbol{o}) r^\alpha \right] : \boldsymbol{c} \text{ is a } (f, M)\text{-covering} \right\}, \\
\xi_\infty^\alpha(f) &:= \lim_{M \to \infty} \xi_M^\alpha(f).
\end{aligned}
$$

It is straightforward that every equivariant ball-covering of Definition 3.14 gives a $(1, 1)$-covering, where the first 1 is regarded as the constant function $f \equiv 1$ on $\mathcal{D}_*$. This gives (see also Conjecture 7.4 below) that for all $M \geq 1$,

$$\xi_M^\alpha(1) \leq \mathcal{H}_M^\alpha(\boldsymbol{D}). \tag{7.2}$$

Also, by considering the case $\boldsymbol{c}_M(v) := f(v) \vee 0$, one can see that if $f \in L_1(\mathcal{D}_*, \mu)$, then

$$\xi_M^\alpha(f) \leq M^\alpha \mathbb{E}\left[f(\boldsymbol{o}) \vee 0\right] < \infty.$$

In the next theorem, to be consistent with the setting of the paper, the following notation is used: $w(u) := w([D, u])$ for $u \in D$, and $w(N_r(v)) = \sum_{u \in N_r(v)} w(u)$. Also, recall that a deterministic equivariant weight function is given by a measurable function $w : \mathcal{D}_* \to \mathbb{R}^{\geq 0}$ (see Example 2.11).

Theorem 7.2 (Unimodular Frostman Lemma). *Let $[\boldsymbol{D}, \boldsymbol{o}]$ be a unimodular discrete space, $\alpha \geq 0$ and $M \geq 1$.*

(i) *There exists a bounded measurable weight function $w : \mathcal{D}_* \to \mathbb{R}^{\geq 0}$ such that $\mathbb{E}\left[w(\boldsymbol{o})\right] = \xi_M^\alpha(1)$ and almost surely,*

$$\forall v \in \boldsymbol{D}, \ \forall r \geq M, w(N_r(v)) \leq r^\alpha. \tag{7.3}$$

---

[2]The term 'weighted' refers to the weighted sums in Definition 7.1 and should not be confused with equivariant weight functions of Definition 5.1

(ii) In addition, if either $\mathrm{udim}_H(\boldsymbol{D}) < \alpha$ *or* $\alpha = \mathrm{udim}_H(\boldsymbol{D})$ *and* $\mathcal{M}^\alpha(\boldsymbol{D}) < \infty$, *then* $w[\boldsymbol{D}, \boldsymbol{o}] \neq 0$ *with positive probability.*

The proof is given later in this subsection.

Remark 7.3. One can show that if w is an equivariant weight function satisfying (7.3), then $\mathbb{E}[w(\boldsymbol{o})] \leq \xi_M^\alpha(1)$ and $\mathbb{E}[w(\boldsymbol{o})h(\boldsymbol{o})] \leq \xi_M^\alpha(h)$. Therefore, the (deterministic) weight function w given in the unimodular Frostman lemma is a maximal equivariant weight function satisfying (7.3) (it should be noted that such maximal functions are not unique in general). The proof is similar to that of the mass distribution principle (Theorem 5.2) and is left to the reader.

Conjecture 7.4. *One has* $\mathcal{H}_M^\alpha(\boldsymbol{D}) = \xi_M^\alpha(1)$.

Here are a few comments on this conjecture. An analogous equality holds in the continuum setting (see 2.10.24 in [23]). Lemma 7.8 below proves a weaker inequality. It can be seen that the conjecture holds for $\mathbb{Z}^k$ (with the l_∞ metric) and for the non-ergodic example of Example 3.19. In the former case, this is obtained by considering the constant weight function $w(\cdot) \equiv \left(\frac{M}{2M+1}\right)^k$, which satisfies the claim of the unimodular Frostman lemma. The latter case is similar by letting $w[\mathbb{Z}, 0] := \frac{M}{2M+1}$ and $w[\mathbb{Z}^2, 0] := 0$.

Remark 7.5. The unimodular Frostman lemma implies that, in theory, the mass distribution principle (Theorem 5.2) is enough for bounding the Hausdorff dimension from above. However, there are very few examples in which the function w given by the unimodular Frostman lemma can be explicitly computed (in some of the examples, a function w satisfying only (7.3) can be found; e.g., for two-ended trees). Therefore, in practice, the unimodular Billingsley lemma is more useful than the mass distribution principle.

The following lemma is needed to prove Theorem 7.2.

Lemma 7.6. *The function* $\xi_M^\alpha : L_1(\mathcal{D}_*, \mu) \to \mathbb{R}$ *is continuous. In fact, it is* M^α-*Lipschitz; i.e.,* $|\xi_M^\alpha(f_1) - \xi_M^\alpha(f_2)| \leq M^\alpha \mathbb{E}[|f_1(\boldsymbol{o}) - f_2(\boldsymbol{o})|]$.

Proof. Let $\boldsymbol{c}$ be an equivariant weighted collection of balls satisfying (7.1) for f_1. Intuitively, add a ball of radius M at each point v with cost $|f_2(v) - f_1(v)|$. More precisely, let $\boldsymbol{c}_r'(v) := \boldsymbol{c}_r(v)$ for $r \neq M$ and $\boldsymbol{c}_M'(v) := \boldsymbol{c}_M(v) + |f_2(v) - f_1(v)|$. This definition implies that $\boldsymbol{c}'$ satisfies (7.1) for f_2. Also,

$$\xi_M^\alpha(f_2) \leq \mathbb{E}\left[\sum_i \boldsymbol{c}_i'(\boldsymbol{o}) i^\alpha\right] = \mathbb{E}\left[\sum_r \boldsymbol{c}_r(\boldsymbol{o}) i^\alpha\right] + M^\alpha \mathbb{E}[|f_2(\boldsymbol{o}) - f_1(\boldsymbol{o})|].$$

Since $\boldsymbol{c}$ is arbitrary, one obtains $\xi_M^\alpha(f_2) \leq \xi_M^\alpha(f_1) + M^\alpha \mathbb{E}[|f_2(\boldsymbol{o}) - f_1(\boldsymbol{o})|]$, which implies the claim. $\square$

Proof of Theorem 7.2. The theorem is a special case of Proposition 7.7 which is proved below. $\square$

Proposition 7.7. *In the setting of Theorem 7.2, let* $h \in L_1(\mathcal{D}_*, \mu)$ *be any given function such that* $h > 0$ *a.s. Then, one can replace the condition* $\mathbb{E}[w(\boldsymbol{o})] = \xi_M^\alpha(1)$ *in Theorem 7.2 by* $\mathbb{E}[w(\boldsymbol{o})h(\boldsymbol{o})] = \xi_M^\alpha(h)$, *and the conclusions of the theorem are valid.*

Proof. It is easy to see that $\xi_M^\alpha(tf) = t\xi_M^\alpha(f)$ for all f and $t \geq 0$ and also $\xi_M^\alpha(f_1 + f_2) \leq \xi_M^\alpha(f_1) + \xi_M^\alpha(f_2)$ for all f_1, f_2. Let $h \in L_1(\mathcal{D}_*, \mu)$ be given. By the Hahn-Banach theorem (see Theorem 3.2 of [46]), there is a linear functional $l : L_1(\mathcal{D}_*, \mu) \to \mathbb{R}$ such that $l(h) = \xi_M^\alpha(h)$ and $-\xi_M^\alpha(-f) \leq l(f) \leq \xi_M^\alpha(f)$, for all $f \in L_1$. Since l is sandwiched between two functions which are continuous at 0 and are equal at 0 (since $\xi_M^\alpha(0) = 0$), one gets that l is continuous at 0. Since l is linear, this implies that l is continuous. Since the dual

of $L_1(\mathcal{D}_*, \mu)$ is $L_\infty(\mathcal{D}_*, \mu)$, one obtains that there is a function $w \in L_\infty(\mathcal{D}_*, \mu)$ such that $l(f) = \mathbb{E}\left[f(\boldsymbol{o})w(\boldsymbol{o})\right]$, for all $f \in L_1$. Note that if $f \geq 0$, then $\xi_M^\alpha(-f) = 0$ and so $l(f) \geq 0$. This implies that $w(\boldsymbol{o}) \geq 0$ a.s. (otherwise, let $f(\boldsymbol{o}) := 1_{\{w(\boldsymbol{o})<0\}}$ to get a contradiction). Consider a version of w which is nonnegative everywhere. The claim is that w satisfies the requirements.

Let $r \geq M$ be fixed. For all discrete spaces D, let $\boldsymbol{S} := \boldsymbol{S}_D := \{v \in D : w(N_r(v)) > r^\alpha\}$. By the definition of $\boldsymbol{S}_D$, one has

$$\mathbb{E}\left[w(N_r(\boldsymbol{o}))1_{\{\boldsymbol{o}\in\boldsymbol{S}\}}\right] \geq r^\alpha \mathbb{P}\left[\boldsymbol{o} \in \boldsymbol{S}\right]. \tag{7.4}$$

Moreover, if $\mathbb{P}\left[\boldsymbol{o} \in \boldsymbol{S}\right] > 0$, then the inequality is strict. Let $f_r(v) := \#N_r(v) \cap \boldsymbol{S}$. By the mass transport principle for the function $(v, u) \mapsto w(u)1_{\{v\in\boldsymbol{S}\}}1_{\{u\in N_r(v)\}}$, one gets

$$
\begin{aligned}
\mathbb{E}\left[w(N_r(\boldsymbol{o}))1_{\{\boldsymbol{o}\in\boldsymbol{S}\}}\right] &= \mathbb{E}\left[w(\boldsymbol{o})\#N_r(\boldsymbol{o}) \cap \boldsymbol{S}\right] \\
&= \mathbb{E}\left[w(\boldsymbol{o})f_r(\boldsymbol{o})\right] \\
&= l(f_r) \\
&\leq \xi_M^\alpha(f_r) \\
&\leq r^\alpha \mathbb{P}\left[\boldsymbol{o} \in \boldsymbol{S}\right],
\end{aligned}
$$

where the last inequality is implied by considering the following weighted collection of balls for f_r: put balls of radius r with cost 1 centered at the points in $\boldsymbol{S}$. More precisely, let $\boldsymbol{c}_r(v) := 1_{\{v\in\boldsymbol{S}\}}$ and $\boldsymbol{c}_s(v) := 0$ for $s \neq r$. It is easy to see that this satisfies (7.1) for f_r, which implies the last inequality by the definition of $\xi_M^\alpha(\cdot)$. Thus, equality holds in (7.4). Hence, $\mathbb{P}\left[\boldsymbol{o} \in \boldsymbol{S}\right] = 0$; i.e., $w(N_r(\boldsymbol{o})) \leq r^\alpha$ a.s. Lemma 2.15 implies that almost surely, $\forall v \in \boldsymbol{D} : w(N_r(v)) \leq r^\alpha$. So the same holds for all rational $r \geq M$ simultaneously. By monotonicity of $w(N_r(v))$ w.r.t. r, one gets that the latter almost surely holds for all $r \geq M$ as desired. Also, one has $\mathbb{E}\left[w(\boldsymbol{o})h(\boldsymbol{o})\right] = l(h) = \xi_M^\alpha(h)$. Thus, w satisfies the desired requirements.

To prove (ii), assume $\mathcal{M}^\alpha(\boldsymbol{D}) < \infty$. By Lemma 3.25, one has $\mathcal{H}_M^\alpha(\boldsymbol{D}) > 0$. So Lemma 7.8 below implies that $0 < \xi_M^\alpha(h) = \mathbb{E}\left[w(\boldsymbol{o})h(\boldsymbol{o})\right]$. This implies that w is not identical to zero. $\qquad\square$

The above proof uses the following lemma.

Lemma 7.8. *Let $[\boldsymbol{D}, \boldsymbol{o}]$ be a unimodular discrete metric space.*

(i) *By letting $b := \xi_1^\alpha(1)$, one has $b \leq \mathcal{H}_1^\alpha(\boldsymbol{D}) \leq b + b|\log b|$.*

(ii) *Let $h \in L_1(\mathcal{D}_*, \mu)$ be a non-negative function. For $M \geq 1$, one has*

$$\mathcal{H}_M^\alpha(\boldsymbol{D}) \leq \inf_{a\geq 0}\left\{M^\alpha\mathbb{E}\left[e^{-ah(\boldsymbol{o})}\right] + a\xi_M^\alpha(h)\right\}. \tag{7.5}$$

(iii) *In addition, if $h > 0$ a.s., then $\xi_M^\alpha(h) = 0$ if and only if $\mathcal{H}_M^\alpha(\boldsymbol{D}) = 0$.*

Proof. (i). The first inequality is easily obtained from the definition of $\xi_1^\alpha(1)$ by considering the cases where $\boldsymbol{c}(\cdot) \in \{0, 1\}$. In particular, this implies that $b \leq 1$. The second inequality is implied by part (ii) by letting $h(\cdot) := 1$ and $a := -\log b \geq 0$.

(ii). Let $b' > \xi_M^\alpha(h)$ be arbitrary. So there exists an equivariant weighted collection of balls $\boldsymbol{c}$ that satisfies (7.1) for h and $\mathbb{E}\left[\sum_{r\geq M}\boldsymbol{c}_r(\boldsymbol{o})r^\alpha\right] \leq b'$. Next, given $a \geq 0$, define an equivariant covering $\boldsymbol{R}$ as follows. For each $v \in \boldsymbol{D}$ and $r \geq M$ such that $\boldsymbol{c}_r(v) > 0$, put a ball of radius r at v with probability $a\boldsymbol{c}_r(v) \wedge 1$. Do this independently for all v and r (one should condition on $\boldsymbol{D}$ first). If more than one ball is put at v, keep only the one with maximum radius. Let $\boldsymbol{S}$ be the union of the chosen balls. For $u \in \boldsymbol{D} \setminus \boldsymbol{S}$, put a ball of

radius M at u. This gives an equivariant covering, namely $\boldsymbol{R}$, by balls of radii at least M. Then, one gets

$$\mathbb{E}\left[\boldsymbol{R}(\boldsymbol{o})\right]^\alpha \leq M^\alpha \mathbb{P}\left[\boldsymbol{o} \notin \boldsymbol{S}\right] + \mathbb{E}\left[\sum_{r \geq M}(a\boldsymbol{c}_r(\boldsymbol{o}) \wedge 1)r^\alpha\right] \leq M^\alpha \mathbb{P}\left[\boldsymbol{o} \notin \boldsymbol{S}\right] + ab'. \qquad (7.6)$$

To bound $\mathbb{P}\left[\boldsymbol{o} \notin \boldsymbol{S}\right]$, consider a realization of $[\boldsymbol{D}, \boldsymbol{o}]$. First, if for some $v \in \boldsymbol{D}$ and $r \geq M$, one has $a\boldsymbol{c}_r(v) > 1$ and $\boldsymbol{o} \in N_r(v)$, then $\boldsymbol{o}$ is definitely in $\boldsymbol{S}$. Second, assume this is not the case. By (7.1), one has $\sum_{u \in \boldsymbol{D}}\sum_{r \geq M} \boldsymbol{c}_r(u)1_{\{\boldsymbol{o} \in N_r(u)\}} \geq h(\boldsymbol{o})$. This implies that the probability that $\boldsymbol{o} \notin \boldsymbol{S}$ in this realization is

$$\prod_{(v,r):\boldsymbol{o} \in N_r(v)}(1 - a\boldsymbol{c}_r(v)) \leq \exp\left(-\sum_{(v,r):\boldsymbol{o} \in N_r(v)}a\boldsymbol{c}_r(v)\right) \leq e^{-ah(\boldsymbol{o})}.$$

In both cases, one gets $\mathbb{P}\left[\boldsymbol{o} \notin \boldsymbol{S}\right] \leq \mathbb{E}\left[e^{-ah(\boldsymbol{o})}\right]$. Thus, (7.6) implies that $\mathbb{E}\left[R(\boldsymbol{o})\right]^\alpha \leq M^\alpha \mathbb{E}\left[e^{-ah(\boldsymbol{o})}\right] + ab'$. Since $a \geq 0$ and $b' > b$ are arbitrary, the claim follows.

(iii). Assume $\xi_M^\alpha(h) = 0$. By letting $a \to \infty$ in (7.5) and using dominated convergence, one obtains that $\mathcal{H}_M^\alpha(\boldsymbol{D}) = 0$. Conversely, assume $\mathcal{H}_M^\alpha(\boldsymbol{D}) = 0$. The first inequality in (i) gives that $\xi_M^\alpha(a) = 0$ for any constant a. Therefore, $\xi_M^\alpha(h) \leq \xi_M^\alpha(a) + \xi_M^\alpha((h-a) \vee 0)) \leq M^\alpha \mathbb{E}\left[(h-a) \vee 0\right]$. By letting a tend to infinity, one gets $\xi_M^\alpha(h) = 0$. $\qquad \square$

Remark 7.9. In this subsection, the following metric is used on the mark space Ξ. Let Ξ' be the set of finite measures on $\mathbb{R}^2$. By identifying $c \in \Xi$ with the counting measure on the finite set $\{(x, c(x)) : x \in \mathbb{R}^{\geq 0}, c(x) > 0\}$, one can identify Ξ with a Borel subset of Ξ'. It is well known that Ξ' is a complete separable metric space under the Prokhorov metric (see e.g., [17]). So one can define the notion of Ξ'-valued equivariant processes as in Definition 2.9. Therefore, Ξ-valued equivariant processes also make sense.

7.2 Max-Flow Min-Cut Theorem for Unimodular One-Ended Trees

The result of this subsection is used in the next subsection for a Euclidean version of the unimodular Frostman lemma, but is of independent interest as well.

The max-flow min-cut theorem is a celebrated result in the field of graph theory (see e.g., [25]). In its simple version, it studies the minimum number of edges in a *cut-set* in a finite graph; i.e., a set of edges the deletion of which disconnects two given subsets of the graph. A generalization of the theorem in the case of trees is obtained by considering cut-sets separating a given finite subset from the set of ends of the tree. This generalization is used to prove a version of Frostman's lemma for compact sets in the Euclidean space (see e.g., [14]).

This subsection presents an analogous result for unimodular one-ended trees. It discusses cut-sets separating the set of leaves from the end of the tree. Since the tree has infinitely many leaves a.s. (see e.g., [5]), infinitely many edges are needed in any such cut-set. Therefore, cardinality cannot be used to study minimum cut-sets. The idea is to use unimodularity for a quantification of the size of a cut-set.

Let $[\boldsymbol{T}, \boldsymbol{o}; \boldsymbol{c}]$ be a unimodular marked one-ended tree with mark space $\mathbb{R}^{\geq 0}$. Assume the mark $\boldsymbol{c}(e)$ of each edge e is well defined and call it the **conductance** of e. Let $\boldsymbol{L}$ be the set of leaves of $\boldsymbol{T}$. As in Subsection 4.1.2, let $F(v)$ be the parent of vertex v and $D(v)$ be the descendants subtree of v.

Definition 7.10. *A **legal equivariant flow** on $[\boldsymbol{T}; \boldsymbol{c}]$ is an equivariant way of assigning extra marks $\boldsymbol{f}(\cdot) \in \mathbb{R}$ to the edges (see Definition 2.9 and Remark 2.13), such that almost surely,*

(i) *for every edge e, one has $0 \le \boldsymbol{f}(e) \le \boldsymbol{c}(e)$,*

(ii) *for every vertex $v \in \boldsymbol{T} \setminus \boldsymbol{L}$, one has*

$$\boldsymbol{f}(v, F(v)) = \sum_{w \in F^{-1}(v)} \boldsymbol{f}(w, v). \tag{7.7}$$

*Also, an **equivariant cut-set** is an equivariant subset Π of the edges of $[\boldsymbol{T}; \boldsymbol{c}]$ that separates the set of leaves $\boldsymbol{L}$ from the end in $\boldsymbol{T}$.*

Note that extra randomness is allowed in the above definition. The reader can think of the value $\boldsymbol{f}(v, F(v))$ as the *flow* from v to $F(v)$. So (7.7) can be interpreted as *conservation of flow* at the vertices except the leaves. Also, the leaves are regarded as the *sources* of the flow.

Since the number of leaves is infinite a.s., the sum of the flows exiting the leaves might be infinite. In fact, it can be seen that unimodularity implies that the sum is always infinite a.s. The idea is to use unimodularity to quantify how *large* is the flow. Similarly, in any equivariant cut-set, the sum of the conductances of the edges is infinite a.s. Unimodularity is also used to quantify the *conductance* of an equivariant cut-set. These are done in Definition 7.12 below.

Below, since each edge of $\boldsymbol{T}$ can be uniquely represented as $(v, F(v))$, the following convention is helpful.

Convention 7.11. For the vertices v of $\boldsymbol{T}$, the symbols $\boldsymbol{f}(v)$ and $\boldsymbol{c}(v)$ are used as abbreviations for $\boldsymbol{f}(v, F(v))$ and $\boldsymbol{c}(v, F(v))$, respectively. Also, by $v \in \Pi$, one means that the edge $(v, F(v))$ is in Π.

Definition 7.12. *The **norm** of the legal equivariant flow $\boldsymbol{f}$ is defined as*

$$|\boldsymbol{f}| := \mathbb{E}\left[\boldsymbol{f}(\boldsymbol{o})1_{\{\boldsymbol{o} \in \boldsymbol{L}\}}\right].$$

Also, for the equivariant cut-set Π, define

$$\boldsymbol{c}(\Pi) := \mathbb{E}\left[\boldsymbol{c}(\boldsymbol{o})1_{\{\boldsymbol{o} \in \Pi\}}\right] = \mathbb{E}\left[\sum_{w \in F^{-1}(\boldsymbol{o})} \boldsymbol{c}(w)1_{\{w \in \Pi\}}\right],$$

where the last equality follows from the mass transport principle (2.2).

An equivariant cut-set Π is called **equivariantly minimal** if there is no other equivariant cut-set which is a subset of Π a.s. If so, it can be seen that it is **almost surely minimal** as well; i.e., in almost every realization, it is a minimal cut set (see Lemma C.6 of [7]).

Lemma 7.13. *If $\boldsymbol{f}$ is a legal equivariant flow and Π is an equivariant cut-set, then $|\boldsymbol{f}| \le \boldsymbol{c}(\Pi)$. Moreover, if the pair $(\boldsymbol{f}, \Pi)$ is equivariant, then*

$$|\boldsymbol{f}| \le \mathbb{E}\left[\boldsymbol{f}(\boldsymbol{o})1_{\{\boldsymbol{o} \in \Pi\}}\right] \le \boldsymbol{c}(\Pi).$$

In addition, if Π is minimal, then equality holds in the left inequality.

Proof. One can always consider an *independent coupling* of $\boldsymbol{f}$ and Π (as in the proof of Theorem 5.2). So assume $(\boldsymbol{f}, \Pi)$ is equivariant from the beginning. Note that the whole construction (with conductances, the flow and the cut-set) is unimodular (Lemma 2.12). For every leaf $v \in \boldsymbol{L}$, let $\boldsymbol{\tau}(v)$ be the first ancestor of v such that $(v, F(v)) \in \Pi$. Then,

send mass $\boldsymbol{f}(v)$ from each leaf v to $\boldsymbol{\tau}(v)$. By the mass transport principle (2.2), one gets

$$
\begin{aligned}
\mathbb{E}\left[\boldsymbol{f}(o)1_{\{o\in L\}}\right] &= \mathbb{E}\left[1_{\{o\in\Pi\}}\sum_{v\in\boldsymbol{\tau}^{-1}(o)}\boldsymbol{f}(v)\right] \\
&\leq \mathbb{E}\left[1_{\{o\in\Pi\}}\sum_{v\in D(o)\cap L}\boldsymbol{f}(v)\right] = \mathbb{E}\left[\boldsymbol{f}(o)1_{\{o\in\Pi\}}\right],
\end{aligned}
$$

where the last equality holds because $\boldsymbol{f}$ is a flow. Moreover, if Π is minimal, then the above inequality becomes an equality and the claim follows. $\qquad\square$

The main result is the following converse to the above lemma.

Theorem 7.14 (Max-Flow Min-Cut for Unimodular One-Ended Trees). *For every unimodular marked one-ended tree $[\boldsymbol{T},o;\boldsymbol{c}]$ equipped with conductances $\boldsymbol{c}$ as above, if $\boldsymbol{c}$ is bounded on the set of leaves, then*

$$
\max_f |\boldsymbol{f}| = \inf_{\Pi} \boldsymbol{c}(\Pi),
$$

where the maximum is over all legal equivariant flows $\boldsymbol{f}$ and the infimum is over all equivariant cut-sets Π.

Remark 7.15. The claim of Theorem 7.14 is still valid if the probability measure (the distribution of $[\boldsymbol{T},o;\boldsymbol{c}]$) is replaced by any (possibly infinite) measure $\mathcal{P}$ on $\mathcal{D}'_*$ supported on one-ended trees, such that $\mathcal{P}(o\in L)<\infty$ and the mass transport principle (2.2) holds. The same proof works for this case as well. This will be used in Subsection 7.3.

Proof of Theorem 7.14. For $n\geq 1$, let $\boldsymbol{T}_n$ be the sub-forest of $\boldsymbol{T}$ obtained by keeping only vertices of height at most n in $\boldsymbol{T}$. Each connected component of $\boldsymbol{T}_n$ is a finite tree which contains some leaves of $\boldsymbol{T}$. For each such component, namely T', do the following: if T' has more than one vertex, consider the maximum flow on T' between the leaves and the *top* vertex (i.e., the vertex with maximum height in T'). If there is more than one maximum flow, choose one of them randomly and uniformly. Also, choose a minimum cut-set in T' randomly and uniformly. Similarly, if T' has a single vertex v, do the same for the subgraph with vertex set $\{v, F(v)\}$ and the single edge adjacent to v. By doing this for all components of $\boldsymbol{T}_n$, a (random) function $\boldsymbol{f}_n$ on the edges and a cut-set Π'_n are obtained (by letting $\boldsymbol{f}_n$ be zero on the other edges). Π'_n is always a cut-set, but $\boldsymbol{f}_n$ is not a flow. However, $\boldsymbol{f}_n$ satisfies (7.7) for vertices of $\boldsymbol{T}_n\setminus L$, except the top vertices of the connected components of $\boldsymbol{T}_n$. Also, it can be seen that $\boldsymbol{f}_n$ and Π'_n are equivariant.

For each component T' of $\boldsymbol{T}_n$, the set of leaves of T', excluding the top vertex, is $L\cap T'$. So the max-flow min-cut theorem of Ford-Fulkerson [25] (see e.g., Theorem 3.1.5 of [14]) gives that, for each component T' of $\boldsymbol{T}_n$, one has

$$
\sum_{v\in L\cap T'}\boldsymbol{f}_n(v) = \sum_{e\in\Pi'_n\cap T'}\boldsymbol{c}(e).
$$

If u is the top vertex of T', let $h(u)$ be the common value in the above equation. By using the mass transport principle (2.2) for each of the two representations of $\mathbb{E}[h(o)]$, one can obtain

$$
\mathbb{E}\left[\boldsymbol{f}_n(o)1_{\{o\in L\}}\right] = \mathbb{E}[h(o)] = \mathbb{E}\left[\boldsymbol{c}(o)1_{\{o\in\Pi'_n\}}\right] = \boldsymbol{c}(\Pi'_n).
$$

Since $0\leq\boldsymbol{f}_n(\cdot)\leq\boldsymbol{c}_n(\cdot)$, one can see that the distributions of $\boldsymbol{f}_n$ are tight (see [7]). Therefore, there is a sequence $n_1,n_2,\ldots$ and an equivariant process $\boldsymbol{f}'$ such that $\boldsymbol{f}_{n_i}\to\boldsymbol{f}'$ (weakly). It is not hard to deduce that $\boldsymbol{f}'$ is a legal equivariant flow. Also, since $\boldsymbol{f}'(o)$

and $1_{\{o\in L\}}$ are continuous functions of $[\boldsymbol{T},\boldsymbol{o};\boldsymbol{f}']$ and their product is bounded (by the assumption on $\boldsymbol{c}$), one gets that

$$\left|\boldsymbol{f}'\right| = \mathbb{E}\left[\boldsymbol{f}'(\boldsymbol{o})1_{\{o\in L\}}\right] = \lim_i \mathbb{E}\left[\boldsymbol{f}_{n_i}(\boldsymbol{o})1_{\{o\in L\}}\right] = \lim_i \boldsymbol{c}(\Pi'_{n_i}).$$

Therefore, $\max_f |\boldsymbol{f}| \geq \inf_\Pi \boldsymbol{c}(\Pi)$. Note that the maximum of $|\boldsymbol{f}|$ is attained by the same tightness argument as above. So Lemma 7.13 implies that equality holds and the claim is proved. □

7.3 A Unimodular Frostman Lemma for Point Processes

In the Euclidean case, another form of the unimodular Frostman lemma is given below. Its proof is based on the max-flow min-cut theorem of Subsection 7.2. As will be seen, the claim implies that in this case, Conjecture 7.4 holds up to a constant factor (Corollary 7.17). However, the weight function obtained in the theorem needs extra randomness.

Theorem 7.16. *Let Φ be a point-stationary point process in $\mathbb{R}^k$ endowed with the l_∞ metric, and let $\alpha \geq 0$. Then, there exists an equivariant weight function $\boldsymbol{w}$ on Φ such that, almost surely,*

$$\forall v \in \Phi, \ \forall r \geq 1 : \boldsymbol{w}(N_r(v)) \leq r^\alpha \tag{7.8}$$

and

$$\mathbb{E}\left[\boldsymbol{w}(0)\right] \geq 3^{-k}\mathcal{H}_1^\alpha(\Phi). \tag{7.9}$$

In particular, if $\mathcal{H}_1^\alpha(\Phi) > 0$, then $\boldsymbol{w}(0)$ is not identical to zero.

A similar result holds for the Euclidean metric or other equivalent metrics by just changing the constant 3^{-k} in (7.9).

In the following proof, Φ is regarded as a counting measures; i.e., for all $A \subseteq \mathbb{R}^d$, $\Phi(A) := \#(\Phi \cap A)$.

Proof. Let $b > 1$ be an arbitrary integer (e.g., $b = 2$). x For every integer $n \geq 0$, let $\boldsymbol{Q}_n$ be the stationary partition of $\mathbb{R}^k$ by translations of the cube $[0, b^n)^k$ as in Subsection 5.3. Consider the nested coupling of these partitions for $n \geq 0$ (i.e., every cube of $\boldsymbol{Q}_n$ is contained in some cube of $\boldsymbol{Q}_{n+1}$ for every $n \geq 0$) independent of Φ. Let $\boldsymbol{T}_0$ be the tree whose vertices are the cubes in $\cup_n \boldsymbol{Q}_n$ and the edges are between all pairs of nested cubes in $\boldsymbol{Q}_n$ and $\boldsymbol{Q}_{n+1}$ for all n. Let $\boldsymbol{T} \subseteq \boldsymbol{T}_0$ be the subtree consisting of the cubes $\boldsymbol{q}_n(v)$ for all $v \in \Phi$ and $n \geq 0$. The set $\boldsymbol{L}$ of the leaves of $\boldsymbol{T}$ consists of the cubes $\boldsymbol{q}_0(v)$ for all $v \in \Phi$. Let $\boldsymbol{\sigma} := \boldsymbol{q}_0(0) \in \boldsymbol{L}$. Note that in the correspondence $v \mapsto \boldsymbol{q}_0(v)$, each cube $\sigma \in \boldsymbol{L}$ corresponds to $\Phi(\sigma) \geq 1$ points of Φ. Therefore, by verifying the mass transport principle, it can be seen that the distribution of $[\boldsymbol{L}, \boldsymbol{\sigma}]$, biased by $1/\Phi(\boldsymbol{\sigma})$, is unimodular; i.e.,

$$\mathbb{E}\left[\frac{1}{\Phi(\boldsymbol{\sigma})}\sum_{\sigma'\in\boldsymbol{L}} g(\boldsymbol{L},\boldsymbol{\sigma},\sigma')\right] = \mathbb{E}\left[\frac{1}{\Phi(\boldsymbol{\sigma})}\sum_{\sigma'\in\boldsymbol{L}} g(\boldsymbol{L},\sigma',\boldsymbol{\sigma})\right],$$

for every measurable $g \geq 0$. In addition, g can be allowed to depend on $\boldsymbol{T}$ in this equation (but the sum is still on $\sigma' \in \boldsymbol{L}$). Therefore, one can assume the metric on $\boldsymbol{L}$ is the graph-distance metric induced from $\boldsymbol{T}$ (see Theorem 3.31). Moreover, Theorem 5 of [35] implies that by a further biasing and choosing a new root for $\boldsymbol{T}$, one can make $\boldsymbol{T}$ unimodular. More precisely, the following (possibly infinite) measure on $\mathcal{D}_*$ is unimodular:

$$\mathcal{P}[A] := \mathbb{E}\left[\sum_{n\geq 0} \frac{1}{e_n}1_A[\boldsymbol{T},\boldsymbol{q}_n(0)]\right], \tag{7.10}$$

where $e_n := \Phi(\boldsymbol{q}_n(0))$. Let $\mathcal{E}$ denote the integral operator w.r.t. the measure $\mathcal{P}$. For any equivariant flow $\boldsymbol{f}$ on $\boldsymbol{T}$, the norm of $\boldsymbol{f}$ w.r.t. the measure $\mathcal{P}$ (see Remark 7.15) satisfies

$$|\boldsymbol{f}| = \mathcal{E}[\boldsymbol{f} \cdot 1_{\boldsymbol{L}}] = \mathbb{E}\left[\sum_{n \geq 0} \frac{1}{e_n} \boldsymbol{f}(\boldsymbol{q}_n(0)) 1_{\{\boldsymbol{q}_n(0) \in \boldsymbol{L}\}}\right] = \mathbb{E}\left[\frac{1}{\Phi(\boldsymbol{\sigma})} \boldsymbol{f}(\boldsymbol{\sigma})\right],$$

where the second equality is by (7.10). Consider the conductance function $\boldsymbol{c}(\tau) := b^{n\alpha}$ for all cubes τ of edge length b^n in $\boldsymbol{T}$ and all n. Therefore, Theorem 7.14 and Remark 7.15 imply that the maximum of $\mathbb{E}[\boldsymbol{f}(\boldsymbol{\sigma})]$ over all equivariant legal flows $\boldsymbol{f}$ on $[\boldsymbol{T}, \boldsymbol{\sigma}]$ is attained (note that $[\boldsymbol{T}, \boldsymbol{\sigma}]$ is not unimodular, but the theorem can be used for $\mathcal{P}$). Denote by $\boldsymbol{f}_0$ the maximum flow. Let $\boldsymbol{w}$ be the weight function on Φ defined by $\boldsymbol{w}(v) = \delta \boldsymbol{f}_0(\boldsymbol{q}_0(v))/\Phi(\boldsymbol{q}_0(v))$, for all $v \in \Phi$, where $\delta := (b+1)^{-k}$. The claim is that $\boldsymbol{w}$ satisfies the requirements (7.8) and (7.9). Since $\boldsymbol{f}_0$ is a legal flow, it follows that for every cube $\sigma \in \boldsymbol{T}$, one has

$$\boldsymbol{w}(\sigma) = \delta \boldsymbol{f}_0(\sigma) \leq \delta \boldsymbol{c}(\sigma) = \delta b^{n\alpha}.$$

Each cube σ of edge length $r \in [b^n, b^{n+1})$ in $\mathbb{R}^k$ can be covered with at most $(b+1)^k$ cubes of edge length b^n in $\boldsymbol{T}_0$. If $n \geq 0$, the latter are either in $\boldsymbol{T}$ or do not intersect Φ. So the above inequality implies that $\boldsymbol{w}(\sigma) \leq r^\alpha$. So (7.8) is proved for $\boldsymbol{w}$.

To prove (7.9), given any equivariant cut-set Π of $\boldsymbol{T}$, a covering of Φ can be constructed as follows: For each cube $\sigma \in \Pi$ of edge length say b^n, let $\boldsymbol{\tau}(\sigma)$ be one of the points in $\sigma \cap \Phi$ chosen uniformly at random and put a ball of radius b^n centered at $\boldsymbol{\tau}(\sigma)$. Note that this ball contains σ. Do this independently for all cubes in $\boldsymbol{T}$. If a point in Φ is chosen more than once, consider only the largest radius assigned to it. It can be seen that this gives an equivariant covering of Φ, namely $\boldsymbol{R}$. One has

$$\begin{aligned}
\mathbb{E}[\boldsymbol{R}(0)^\alpha] &\leq& \mathbb{E}\left[\sum_{n \geq 0} b^{n\alpha} 1_{\{\boldsymbol{q}_n(0) \in \Pi\}} 1_{\{0 = \boldsymbol{\tau}(\boldsymbol{q}_n(0))\}}\right] \\
&=& \mathbb{E}\left[\sum_{n \geq 0} \frac{b^{n\alpha}}{e_n} 1_{\{\boldsymbol{q}_n(0) \in \Pi\}}\right].
\end{aligned}$$

On the other hand, by (7.10), one can see that

$$\boldsymbol{c}(\Pi) = \mathbb{E}\left[\sum_{n \geq 0} \frac{1}{e_n} \boldsymbol{c}(\boldsymbol{q}_n(0)) 1_{\{\boldsymbol{q}_n(0) \in \Pi\}}\right] = \mathbb{E}\left[\sum_{n \geq 0} \frac{b^{n\alpha}}{e_n} 1_{\{\boldsymbol{q}_n(0) \in \Pi\}}\right].$$

Therefore, $\mathbb{E}[\boldsymbol{R}(0)^\alpha] \leq \boldsymbol{c}(\Pi)$. So $\mathcal{H}_1^\alpha(\Phi) \leq \boldsymbol{c}(\Pi)$. Since Π is an arbitrary equivariant cut-set, by the unimodular max-flow min-cut theorem established above (Theorem 7.14) and the maximality of the flow $\boldsymbol{f}_0$, one gets that $\mathcal{H}_1^\alpha(\Phi) \leq |\boldsymbol{f}_0| = \mathbb{E}[\boldsymbol{f}_0(\boldsymbol{\sigma})/\Phi(\boldsymbol{\sigma})] = \delta^{-1}\mathbb{E}[\boldsymbol{w}(0)]$. So the claim is proved. $\qquad\square$

The following corollary shows that in the setting of Theorem 7.16, the claim of Conjecture 7.4 holds up to a constant factor (compare this with Lemma 7.8).

Corollary 7.17. *For all point-stationary point processes Φ in $\mathbb{R}^k$ endowed with the l_∞ metric and all $\alpha \geq 0$, $3^{-k}\mathcal{H}_1^\alpha(\Phi) \leq \xi_1^\alpha(\Phi) \leq \mathcal{H}_1^\alpha(\Phi)$.*

Proof. The claim is directly implied by (7.2), Theorem 7.16 and Remark 7.3. $\qquad\square$

7.4 Applications

The following subsections give some basic applications of the unimodular Frostman lemma. This lemma is also the basis of many results of [6].

7.4.1 Cayley Graphs

Proposition 7.18. *For every finitely generated group H with polynomial growth degree $\alpha \in [0, \infty]$, one has $\xi_\infty^\alpha(H) < \infty$.*

Note that if Conjecture 7.4 holds, then this result implies $\mathcal{M}^\alpha(H) > 0$, as conjectured in Subsection 6.6.

Proof. By Theorem 6.14, $\mathcal{M}^\alpha(H) < \infty$. So the unimodular Frostman lemma (Theorem 7.2) implies that for every $M \geq 1$, there exists $w : \mathcal{D}_* \to \mathbb{R}^{\geq 0}$ such that $w(N_M(e)) \leq M^\alpha$ and $\mathbb{E}[w(e)] = \xi_M^\alpha(H)$, where e is the neutral element of H. Since the Cayley graph of H is transitive and w is defined up to rooted isomorphisms, $w(H, \cdot)$ is constant. Hence, $w(H, v) = \xi_M^\alpha(H)$ for all $v \in H$. Therefore, $\xi_M^\alpha(H) \# N_M(e) \leq M^\alpha$. Thus, $\xi_M^\alpha(H) \leq 1/c$, where c is as in the proof of Theorem 6.14. By letting $M \to \infty$, one gets $\xi_\infty^\alpha(H) \leq 1/c < \infty$. $\qquad\square$

7.4.2 Dimension of Product Spaces

Let $[\boldsymbol{D}_1, \boldsymbol{o}_1]$ and $[\boldsymbol{D}_2, \boldsymbol{o}_2]$ be independent unimodular discrete metric spaces. By considering any of the usual product metrics; e.g., the sup metric or the p product metric, the **independent product** $[\boldsymbol{D}_1 \times \boldsymbol{D}_2, (\boldsymbol{o}_1, \boldsymbol{o}_2)]$ makes sense as a random pointed discrete space. It is not hard to see that the latter is also unimodular (see also Proposition 4.11 of [2]).

Proposition 7.19. *Let $[\boldsymbol{D}_1 \times \boldsymbol{D}_2, (\boldsymbol{o}_1, \boldsymbol{o}_2)]$ represent the independent product of $[\boldsymbol{D}_1, \boldsymbol{o}_1]$ and $[\boldsymbol{D}_2, \boldsymbol{o}_2]$ defined above. Then,*

$$\mathrm{udim}_H(\boldsymbol{D}_1) + \underline{\mathrm{udim}}_M(\boldsymbol{D}_2) \leq \mathrm{udim}_H(\boldsymbol{D}_1 \times \boldsymbol{D}_2) \leq \mathrm{udim}_H(\boldsymbol{D}_1) + \mathrm{udim}_H(\boldsymbol{D}_2). \quad (7.11)$$

Proof. By Theorem 3.31, one can assume the metric on $\boldsymbol{D}_1 \times \boldsymbol{D}_2$ is the sup metric without loss of generality. So $N_r(v_1, v_2) = N_r(v_1) \times N_r(v_2)$.

The upper bound is proved first. For $i = 1, 2$, let $\alpha_i > \mathrm{udim}_H(\boldsymbol{D}_i)$ be arbitrary. By the unimodular Frostman lemma (Theorem 7.2), there is a nonnegative measurable functions w_i on $\mathcal{D}_*$ such that $\forall v \in \boldsymbol{D}_i : \forall r \geq 1 : w_i(N_r(v)) \leq r^\alpha$, a.s. In addition, $w_i(\boldsymbol{o}_i) \neq 0$ with positive probability. Consider the equivariant weight function $\boldsymbol{w}$ on $\boldsymbol{D}_1 \times \boldsymbol{D}_2$ defined by $\boldsymbol{w}(v_1, v_2) := w_1[\boldsymbol{D}_1, v_1] \times w_2[\boldsymbol{D}_2, v_2]$. It is left to the reader to show that $\boldsymbol{w}$ is an equivariant weight function. One has $\boldsymbol{w}(N_r(v_1, v_2)) = w_1(N_r(v_1))w_2(N_r(v_2)) \leq r^{\alpha_1 + \alpha_2}$. Also, by the independence assumption, $\boldsymbol{w}(\boldsymbol{o}_1, \boldsymbol{o}_2) \neq 0$ with positive probability. Therefore, the mass distribution principle (Theorem 5.2) implies that $\mathrm{udim}_H(\boldsymbol{D}_1 \times \boldsymbol{D}_2) \leq \alpha_1 + \alpha_2$. This proves the upper bound.

For the lower bound in the claim, let $\alpha < \mathrm{udim}_H(\boldsymbol{D}_1)$, $\beta < \underline{\mathrm{udim}}_M(\boldsymbol{D}_2)$ and $\epsilon > 0$ be arbitrary. It is enough to find an equivariant covering $\boldsymbol{R}$ of $\boldsymbol{D}_1 \times \boldsymbol{D}_2$ such that $\mathbb{E}[\boldsymbol{R}(\boldsymbol{o}_1, \boldsymbol{o}_2)^{\alpha + \beta}] < \epsilon$. One has $\underline{\mathrm{decay}}(\lambda_r(\boldsymbol{D}_2)) > \beta$, where λ_r is defined in (3.1). So there is $M > 0$ such that $\forall r \geq M : \lambda_r(\boldsymbol{D}_2) < r^{-\beta}$. So for every $r \geq M$, there is an equivariant r-covering of $\boldsymbol{D}_2$ with intensity less than $r^{-\beta}$. On the other hand, since $\alpha < \mathrm{udim}_H(\boldsymbol{D}_1)$, one has $\mathcal{H}_M^\alpha(\boldsymbol{D}_1) = 0$ (by Lemma 3.25). Therefore there is an equivariant covering $\boldsymbol{R}_1$ of $\boldsymbol{D}_1$ such that $\mathbb{E}[\boldsymbol{R}_1(\boldsymbol{o}_1)^\beta] < \epsilon$ and $\forall v \in \boldsymbol{D}_1 : \boldsymbol{R}_1(v) \in \{0\} \cup [M, \infty)$ a.s. Choose the extra randomness in $\boldsymbol{R}_1$ independently from $[\boldsymbol{D}_2, \boldsymbol{o}_2]$. Given a realization of $[\boldsymbol{D}_1, \boldsymbol{o}_1]$ and $\boldsymbol{R}_1$, do the following: Let $v_1 \in \boldsymbol{D}_1$ such that $\boldsymbol{R}_1(v_1) \neq 0$ (and hence, $\boldsymbol{R}_1(v_1) \geq M$). One can find an equivariant subset $\boldsymbol{S}_{v_1}$ of $\boldsymbol{D}_2$ that gives a covering of $\boldsymbol{D}_2$ by balls of radius $\boldsymbol{R}_1(v_1)$ and has intensity less than $\boldsymbol{R}_1(v_1)^{-\beta}$. Do this independently for all $v_1 \in \boldsymbol{D}_1$. Now, for all $(v_1, v_2) \in \boldsymbol{D}_1 \times \boldsymbol{D}_2$, define

$$\boldsymbol{R}(v_1, v_2) := \begin{cases} \boldsymbol{R}_1(v_1) & \text{if } \boldsymbol{R}_1(v_1) \neq 0 \text{ and } v_2 \in \boldsymbol{S}_{v_1}, \\ 0 & \text{otherwise.} \end{cases}$$

Now, $\boldsymbol{R}$ is a covering of $\boldsymbol{D}_1 \times \boldsymbol{D}_2$ and it can be seen that it is an equivariant covering. Also, given $[\boldsymbol{D}_1, \boldsymbol{o}_1]$ and $\boldsymbol{R}_1$, the probability that $\boldsymbol{o}_2 \in \boldsymbol{S}_{\boldsymbol{o}_1}$ is less than $\boldsymbol{R}_1(\boldsymbol{o}_1)^{-\beta}$. So one gets

$$
\begin{aligned}
\mathbb{E}\left[\boldsymbol{R}(\boldsymbol{o}_1, \boldsymbol{o}_2)^{\alpha+\beta}\right] &= \mathbb{E}\left[\mathbb{E}\left[\boldsymbol{R}(\boldsymbol{o}_1, \boldsymbol{o}_2)^{\alpha+\beta} | [\boldsymbol{D}_1, \boldsymbol{o}_1], \boldsymbol{R}_1\right]\right] \\
&< \mathbb{E}\left[\boldsymbol{R}_1(\boldsymbol{o}_1)^{\alpha+\beta} \boldsymbol{R}_1(\boldsymbol{o}_1)^{-\beta}\right] = \mathbb{E}\left[\boldsymbol{R}_1(\boldsymbol{o}_1)^{\alpha}\right] < \epsilon.
\end{aligned}
$$

So the claim is proved. $\qquad\square$

The following examples provide instances where the inequalities in (7.11) are strict.

Example 7.20. Assume $[\boldsymbol{D}_1, \boldsymbol{o}_1]$ and $[\boldsymbol{D}_2, \boldsymbol{o}_2]$ are unimodular discrete spaces such that $\underline{\dim}_M(\boldsymbol{G}_1) < \dim_H(\boldsymbol{G}_1)$ and $\underline{\dim}_M(\boldsymbol{G}_2) = \dim_H(\boldsymbol{G}_2)$. By Proposition 7.19, one gets

$$
\dim_H(\boldsymbol{G}_1 \times \boldsymbol{G}_2) \geq \dim_H(\boldsymbol{G}_1) + \underline{\dim}_M(\boldsymbol{G}_2) > \dim_H(\boldsymbol{G}_2) + \underline{\dim}_M(\boldsymbol{G}_1).
$$

So by swapping the roles of the two spaces, an example of strict inequality in the left hand side of (7.11) is obtained.

Example 7.21. Let J be a subset of $\mathbb{Z}^{\geq 0}$ such that $\overline{d}(J) = 1$ and $\underline{d}(J) = 0$ simultaneously (see Subsection 6.5.1 for the definitions). Let Ψ_1 and Ψ_2 be defined as in Subsection 6.5.1 corresponding to J and $\mathbb{Z}^{\geq 0} \setminus J$ respectively. Proposition 6.11 implies that $\dim_H(\Psi_1) = \dim_H(\Psi_2) = 1$. On the other hand, (6.2) implies that

$$
\#N_{2^n}(\boldsymbol{o}_1 \times \boldsymbol{o}_2) \leq 2^{J_n + 1} \times 2^{(n+1-J_n)+1} = 2^{n+3}.
$$

This implies that $\overline{\mathrm{growth}}\left(N_r(\boldsymbol{o})\right) \leq 1$. So the unimodular Billingsley lemma (Theorem 5.6) implies that $\dim_H(\Psi_1 \times \Psi_2) \leq 1$ (in fact, equality holds by Proposition 7.24 below). So the rightmost inequality in (7.11) is strict here.

7.4.3 Dimension of Embedded Spaces

It is natural to think of $\mathbb{Z}$ as a subset of $\mathbb{Z}^2$. However, $[\mathbb{Z}, 0]$ is not an equivariant subspace of $[\mathbb{Z}^2, 0]$. By the following definition, $[\mathbb{Z}, 0]$ is called *embeddable in* $[\mathbb{Z}^2, 0]$. The dimension of embedded subspaces is studied in this subsection. The analysis requires the unimodular Frostman lemma.

Definition 7.22. *Let $[\boldsymbol{D}_0, \boldsymbol{o}_0]$ and $[\boldsymbol{D}, \boldsymbol{o}]$ be random pointed discrete spaces. An **embedding** of $[\boldsymbol{D}_0, \boldsymbol{o}_0]$ in $[\boldsymbol{D}, \boldsymbol{o}]$ is a (not necessarily unimodular) random pointed marked discrete space $[\boldsymbol{D}', \boldsymbol{o}'; \boldsymbol{m}]$ with mark space $\{0, 1\}$ such that*

(i) *$[\boldsymbol{D}', \boldsymbol{o}']$ has the same distribution as $[\boldsymbol{D}, \boldsymbol{o}]$.*

(ii) *$\boldsymbol{m}(\boldsymbol{o}') = 1$ a.s. and by letting $\boldsymbol{S} := \{v \in \boldsymbol{D}' : \boldsymbol{m}(v) = 1\}$ equipped with the induced metric from $\boldsymbol{D}'$, $[\boldsymbol{S}, \boldsymbol{o}']$ has the same distribution as $[\boldsymbol{D}_0, \boldsymbol{o}_0]$.*

*If in addition, $[\boldsymbol{D}_0, \boldsymbol{o}_0]$ is unimodular, then $[\boldsymbol{D}', \boldsymbol{o}'; \boldsymbol{m}]$ is called an **equivariant embedding** if*

(iii) *The mass transport principle holds on $\boldsymbol{S}$; i.e., (2.2) holds for functions $g(u, v) := g(\boldsymbol{D}', u, v; \boldsymbol{m})$ such that $g(u, v)$ is zero when $\boldsymbol{m}(u) = 0$ or $\boldsymbol{m}(v) = 0$.*

*If an embedding (resp. an equivariant embedding) exists, $[\boldsymbol{D}_0, \boldsymbol{o}_0]$ is called **embeddable** (resp. **equivariantly embeddable**) in $[\boldsymbol{D}, \boldsymbol{o}]$.*

It should be noted that $[\boldsymbol{D}', \boldsymbol{o}'; \boldsymbol{m}]$ is not an equivariant process on $\boldsymbol{D}$ except in the trivial case where $\boldsymbol{m}(\cdot) = 1$ a.s.

Example 7.23. Here are instances of Definition 7.22.

(i) Let $[\boldsymbol{D}_0, \boldsymbol{o}_0] := [\mathbb{Z}, 0]$ and $[\boldsymbol{D}, \boldsymbol{o}] := [\mathbb{Z}^2, 0]$ equipped with the sup metric. Consider $m : \mathbb{Z}^2 \to \{0, 1\}$ which is equal to one on the boundary of the positive cone. Then, $[\mathbb{Z}^2, 0; m]$ is an embedding of $[\mathbb{Z}, 0]$ in $[\mathbb{Z}^2, 0]$, but is not an equivariant embedding since it does not satisfy (iii).

(ii) A point-stationary point process in $\mathbb{Z}^k$ (pointed at 0) is equivariantly embeddable in $[\mathbb{Z}^k, 0]$.

(iii) Let H be a finitely generated group equipped with the graph-distance metric of an arbitrary Cayley graph over H. Then, any subgroup of H (equipped with the induced metric) is equivariantly embeddable in H.

Proposition 7.24. *If $[\boldsymbol{D}_0, \boldsymbol{o}_0]$ and $[\boldsymbol{D}, \boldsymbol{o}]$ are unimodular discrete spaces and the former is equivariantly embeddable in the latter, then*

$$\mathrm{udim}_H(\boldsymbol{D}) \geq \mathrm{udim}_H(\boldsymbol{D}_0), \tag{7.12}$$

and for all $\alpha \geq 0$ and $M \geq 1$, with ξ_M^α defined in Definition 7.1,

$$\xi_M^\alpha(\boldsymbol{D}, 1) \quad \leq \quad \mathcal{H}_M^\alpha(\boldsymbol{D}_0). \tag{7.13}$$

Proof. First, assume (7.13) holds. For $\alpha > \mathrm{udim}_H(\boldsymbol{D})$, one has $\mathcal{H}_M^\alpha(\boldsymbol{D}) > 0$ (Lemma 3.25). Therefore, Lemma 7.8 implies that $\xi_M^\alpha(\boldsymbol{D}, 1) > 0$. Hence, (7.13) implies that $\mathcal{H}_M^\alpha(\boldsymbol{D}_0) > 0$, which implies that $\mathrm{udim}_H(\boldsymbol{D}_0) \leq \alpha$. So it is enough to prove (7.13).

By the unimodular Frostman lemma (Theorem 7.2), there is a bounded function $w : \mathcal{D}_* \to \mathbb{R}^{\geq 0}$ such that $\mathbb{E}[w(\boldsymbol{o})] = \xi_M^\alpha(\boldsymbol{D}, 1)$, and almost surely, $w(N_r(\boldsymbol{o})) \leq r^\alpha$ for all $r \geq M$. Assume $[\boldsymbol{D}', \boldsymbol{o}'; \boldsymbol{m}]$ is an equivariant embedding as in Definition 7.22. For $x \in \boldsymbol{D}'$, let $w'(x) := w'_{\boldsymbol{D}'}(x) := w[\boldsymbol{D}', x]$. Consider the random pointed marked discrete space $[\boldsymbol{S}, \boldsymbol{o}'; w']$ obtained by restricting w' to $\boldsymbol{S}$. By the definition of equivariant embeddings and by directly verifying the mass transport principle, the reader can obtain that $[\boldsymbol{S}, \boldsymbol{o}'; w']$ is unimodular. Since $[\boldsymbol{S}, \boldsymbol{o}']$ has the same distribution as $[\boldsymbol{D}_0, \boldsymbol{o}_0]$, there exists an equivariant process w_0 on $\boldsymbol{D}_0$ such that $[\boldsymbol{S}, \boldsymbol{o}'; w']$ has the same distribution as $[\boldsymbol{D}_0, \boldsymbol{o}_0; w_0]$ (see the converse of Lemma 2.12 in Subsection 2.5). According to the above discussion, one has

$$\forall r \geq M : w'(N_r(\boldsymbol{S}, \boldsymbol{o}')) \leq w'(N_r(\boldsymbol{D}', \boldsymbol{o}')) \leq r^\alpha, \quad a.s.$$

This implies that $w_0(N_r(\boldsymbol{o}_0)) \leq r^\alpha$ a.s. Therefore, the mass distribution principle (Theorem 5.2) implies that $\mathbb{E}[w_0(\boldsymbol{o}_0)] \leq \mathcal{H}_M^\alpha(\boldsymbol{D}_0)$. One the other hand,

$$\mathbb{E}[w_0(\boldsymbol{o}_0)] = \mathbb{E}[w'(\boldsymbol{o}')] = \mathbb{E}[w(\boldsymbol{o})] = \xi_M^\alpha(\boldsymbol{D}, 1),$$

where the last equality is by the assumption on w. This implies that $\mathcal{H}_M^\alpha(\boldsymbol{D}_0) \geq \xi_M^\alpha(\boldsymbol{D}, 1)$ and the claim is proved. $\square$

It is natural to expect that an embedded space has a smaller Hausdorff size. This is stated in Conjecture 8.4.

Remark 7.25. Another possible way to prove Proposition 7.24 and Conjecture 8.4 is to consider an arbitrary equivariant covering of $\boldsymbol{D}_0$ and try to extend it to an equivariant covering of $\boldsymbol{D}$ by adding some balls (without adding a ball centered at the root). More generally, given an equivariant processes $\boldsymbol{Z}_0$ on $\boldsymbol{D}_0$, one might try to extend it to an equivariant process on $\boldsymbol{D}$ without changing the mark of the root. But at least the latter is not always possible. A counter example is when $[\boldsymbol{D}_0, \boldsymbol{o}_0]$ is K_2 (the complete graph with two vertices), $[\boldsymbol{D}, \boldsymbol{o}]$ is K_3, $\boldsymbol{Z}_0(\boldsymbol{o}_0) = \pm 1$ chosen uniformly at random, and the mark of the other vertex of $\boldsymbol{D}_0$ is $-\boldsymbol{Z}_0(\boldsymbol{o}_0)$.

7.5 Notes and Bibliographical Comments

The unimodular Frostman lemma (Theorem 7.2) is analogous to Frostman's lemma in the continuum setting (see e.g., Thm 8.17 of [42]). The proof of Theorem 7.2 is also inspired by that of [42], but there are substantial differences. For instance, the proof of Lemma 7.8 and also the use of the duality of L_1 and L_∞ in the proof of Theorem 7.2 are new. The Euclidean version of the unimodular Frostman lemma (Theorem 7.16) and its proof are inspired by the continuum analogue (see e.g., [14]).

As already explained, the unimodular max-flow min-cut theorem (Theorem 7.14) is inspired by the max-flow min-cut theorem for finite trees. Also, the results and examples of Subsection 7.4.2 on product spaces are inspired by analogous in the continuum setting; e.g., Theorem 3.2.1 of [14].

8 Miscellaneous Topics

8.1 Connections with Other Notions of Dimension

Several notions of dimension are already defined in the literature for discrete spaces in special cases. A few of them are listed in this subsection together with their connections to unimodular dimensions.

For subsets of $\mathbb{Z}^d$, the notions of *upper and lower mass dimension* are defined in [8], which are just the volume growth rates defined in Section 5. The paper [38] extends the upper mass dimension to general subsets $A \subseteq \mathbb{R}^d$ and calls it the *macroscopic Minkowski dimension* of A (one may define lower macroscopic Minkowski dimension similarly). This extension is obtained by *pixelizing* A to get a subset of $\mathbb{Z}^d$. The unimodular Billingsley lemma states that for unimodular (i.e., point-stationary) and ergodic subsets of $\mathbb{Z}^d$, the unimodular Hausdorff dimension is between the upper and lower mass dimension. A similar result holds in the non-integer case as well:

Corollary 8.1. *For ergodic point-stationary point processes in $\mathbb{R}^d$, the unimodular Hausdorff dimension is between the upper and lower macroscopic Minkowski dimensions a.s.*

This is a direct corollary of Billingsley's lemma applied to the pixelization by a randomly-shifted lattice. It can also be proved by using weights in Billingsley's lemma similarly to the proof of Proposition 5.15.

Another notion is that of *discrete (Hausdorff) dimension* [8], which uses the idea behind the definition of the classical Hausdorff dimension by considering coverings of $\Phi \subseteq \mathbb{Z}^d$ by large balls and considering the cost $(\frac{r}{r+|x|})^\alpha$ for each ball in the covering, where r and x are the radius and the center of the ball and α is a constant (in fact, this is a modified version of the definition of [8] mentioned in [14]). In the future work [6], it is shown that the discrete dimension is an upper bound for the unimodular Hausdorff dimension, when both notions are defined (i.e., for point-stationary point processes).

The unimodular Hausdorff dimension can be connected to the classical Hausdorff dimension via scaling limits. Such limits are random continuum metric spaces and can be defined by weak convergence w.r.t. the Gromov-Hausdorff-Prokhorov metric [37]. It is shown in the preprint [6] that if the unimodular discrete space admits a scaling limit, then the ordinary Hausdorff dimension of the limit is an upper bound for the unimodular Hausdorff dimension.

The above inequalities are expected to be equalities in most examples. The preprint [6] provides more discussion on the matter. Note that these comparison results imply relations between the volume growth rate, scaling limits and discrete dimension, which are of independent interest and which are new to the best of the authors' knowledge.

A problem of potential interest is the connection of unimodular dimensions to other

notions of dimension. This includes Gromov's notion of *asymptotic dimension* [28], the *spectral dimension* of a graph (defined in terms of the return probabilities of the simple random walk), the *typical displacement exponent* of a graph (see [16] for both notions), the *isoperimetric dimension* of a graph [15], the *resistance growth exponent* of a graph, the *stochastic dimension* of a partition of $\mathbb{Z}^d$ [11], etc. In statistical physics, one also assigns dimension and various exponents to finite models. Famous examples are self-avoiding walks and the boundaries of large percolation clusters.

8.2 Gauge Functions and the Unimodular Dimension Function

There exist unimodular discrete spaces $\boldsymbol{D}$ in which the $\mathrm{udim}_H(\boldsymbol{D})$-dimensional Hausdorff size is either zero or infinity (e.g., Examples 6.1 and 6.2). For such spaces, it is convenient to generalize the unimodular Hausdorff size as follows. Consider an increasing function $\varphi : \{0\} \cup [1, \infty) \to [0, \infty)$; e.g., $\varphi(r) = r^\alpha$, called a *gauge function*. Define $\mathcal{H}_M^\varphi(\boldsymbol{D})$ by $\inf_{\boldsymbol{R}} \{\mathbb{E}\left[\varphi(\boldsymbol{R}(o))\right]\}$ similarly to (3.3). Then, define $\mathcal{M}^\varphi(\boldsymbol{D})$ similarly to (3.6). If $0 < \mathcal{M}^\varphi(\boldsymbol{D}) < \infty$, then φ is called a **unimodular dimension function** for $\boldsymbol{D}$.

In addition, given a family of gauge functions $(\varphi_\alpha)_{\alpha \geq 0}$ that is increasing in α and such that $\forall \alpha > \beta : \lim_{r \to \infty} \varphi_\alpha(r)/\varphi_\beta(r) = \infty$, one can redefine the unimodular Hausdorff dimension by $\sup\{\alpha : \mathcal{M}^{\varphi_\alpha}(\boldsymbol{D}) = 0\}$ (see e.g., the next paragraph). One can redefine the unimodular Minkowski dimension similarly. The authors have verified that the results of the paper can be extended to this setting except that Theorem 3.34 and the results of Subsection 3.8 require the *doubling condition* $\sup_{r \geq 1} \varphi(2r)/\varphi(r) < \infty$. The general result of Subsection 4.1.2 can also be extended under the doubling condition. Also, the upper bounds in the unimodular mass distribution principle, the unimodular Billingsley lemma and the unimodular Frostman lemma hold in this more general setting (some other results require the doubling condition). However, for the ease of reading, the results are presented in the original setting of this paper.

As an example of the above framework, one can define the **exponential dimension** by considering $\varphi_\alpha(r) := e^{\alpha r}$. It might be useful for studying unimodular spaces with super-polynomial volume growth, which are more interesting in group theory (see Subsection 6.6). Other gauge functions may also be useful for groups of *intermediate growth*. Note that exponential gauge functions do not satisfy the doubling condition, and hence, the reader should be careful about using the results of this work for such gauge functions.

8.3 Negative Dimensions

If a compact metric space X is the union of k disjoint copies of $\frac{1}{r}X$, then the *similarity dimension* of X is $\log k / \log r$ (see e.g., [14]). This definition can also be used for some infinite discrete sets as well. For instance, $\mathbb{Z}^d$ is a union of 2^d copies of $2\mathbb{Z}^d$. So, it can be said that the *similarity dimension* of $\mathbb{Z}^d$ is negative. The (deterministic) discrete Cantor set (see e.g., [14]) is also $(-\log 2 / \log 3)$-dimensional. There are several further arguments, listed below, suggesting that one should actually assign negative dimensions to unimodular discrete spaces.

First, this would be natural in terms of definition. The unimodular Minkowski dimensions should be redefined by $\overline{\mathrm{udim}}_M(\boldsymbol{D}) = \overline{\mathrm{growth}}(\lambda_r)$ and $\underline{\mathrm{udim}}_M(\boldsymbol{D}) = \underline{\mathrm{growth}}(\lambda_r)$. Using growth instead of decay would then unify the definition of the ordinary Minkowski dimension of compact sets and the unimodular Minkowski dimension. The former is microscopic (i.e., when r tends to 0), whereas the latter is macroscopic ($r \to \infty$). One may also replace the unimodular Hausdorff by the negative of the definitions given so far.

Secondly, this unification of the definitions would also take care of the puzzling

direction of certain inequalities discussed in the paper: when adopting these negative unimodular dimensions, (i) the classical and unimodular Minkowski and Hausdorff dimensions would be ordered in the same way, i.e., $\mathrm{udim}_H(\boldsymbol{D}) \leq \underline{\mathrm{udim}}_M(\boldsymbol{D}) \leq \overline{\mathrm{udim}}_M(\boldsymbol{D})$; (ii) an equivariant subset of a unimodular set would have a unimodular Minkowski dimension smaller than or equal to that of the set, and possibly strictly smaller (see Subsections 3.7 and 4.4), a property that is expected to hold for any notion of *dimension*; (iii) the mass distribution principle and the Billingsley lemma would provide lower bounds on $\mathrm{udim}_H(\boldsymbol{D})$, while upper bounds would be obtained by constructing explicit coverings; (iv) the dimension of non-ergodic examples (e.g., Example 3.19) would also be the supremum dimension of the components, as one might expect (see also Remark 3.20).

It should however be noted that by assigning such negative dimension, the dimension of a non-equivariant subset (see Subsection 7.4.3) would be larger than or equal to that of the whole space (just as the similarity dimension of $\mathbb{Z}$ is larger than that of $\mathbb{Z}^2$). One should not expect that non-equivariant subsets behave like equivariant subsets, since they do not satisfy the main assumption of *statistical homogeneity*, which is the basis of all of the definitions in this work.

8.4   Problems and Conjectures

This subsection gathers some problems and conjectures pertaining to the theory of unimodular dimensions and to specific examples. Those are already stated in the paper and are also briefly listed.

8.4.1   Further conjectures and problems

1. Connections to other notions . The following conjectures connect unimodular dimensions to the properties of the simple random walk.

Conjecture 8.2. *If $[\boldsymbol{G}, \boldsymbol{o}]$ is a unimodular graph such that $\mathrm{udim}_H(\boldsymbol{G}) < 2$, then the simple random walk in $\boldsymbol{G}$ is recurrent a.s.*

Note that the converse of this conjecture does not hold. For instance, it is not hard to show that any one-ended tree (e.g., the canopy tree) is recurrent.

In [8], the discrete dimension of $A \subseteq \mathbb{Z}^d$ is connected to whether the simple random walk in $\mathbb{Z}^d$ hits A infinitely often or not. Analogously, one can generalized the above conjecture as follows:

Conjecture 8.3. *If $[\boldsymbol{D}, \boldsymbol{o}]$ is equivariantly embedded in the unimodular graph $[\boldsymbol{G}, \boldsymbol{o}]$ and $\mathrm{udim}_H(\boldsymbol{D}) > \mathrm{udim}_H(\boldsymbol{G}) - 2$, then the simple random walk in $\boldsymbol{G}$ hits $\boldsymbol{D}$ infinitely often a.s.*

2. Dimension functions.   Does there exist a unimodular discrete space without any dimension function? The answer is not known yet. [22] gives a positive answer to the analogous question in the continuum setting, but the proof ideas don't seem to work in the unimodular discrete setting. Also, by analogy with stable trees [19], a possible candidate is unimodular EGW trees with infinite offspring variance.

3. Simple random walk.   By analogy with the image of subordinator processes (see e.g., [26]), one may guess the exact unimodular dimension function for the image of the random walk under the assumptions of Proposition 4.9. For example, consider the zero set Ψ of the simple random walk (Proposition 4.10). By analogy with the zero set of Brownian motion [48], it is natural to guess that $\mathcal{M}^{1/2}(\Psi) = \infty$ and $\sqrt{r \log \log r}$ is a dimension function for Ψ. To prove this, one should strengthen the bound $Cr^\beta \log \log r$ in the proof of Proposition 4.9 and also construct a covering of the set which is better than that of Proposition 3.13. For the former, one may use Theorem 4 of [26] (it seems

that the assumption of [26] on the tail of the jumps is not necessary for having only an upper bound). For the latter, one might try to get ideas from [48] (it is necessary to use intervals with different lengths).

Another guess is that the image of the symmetric nearest-neighbor simple random walk in $\mathbb{Z}^d$ is 2-dimensional when $d \geq 3$. More generally, if the jumps are in the domain of attraction of a symmetric α-stable process, then the image is α-dimensional. These might be proved similarly to the analogous results in [9].

For the graph of the simple random walk equipped with the Euclidean metric (Subsection 6.4), the guess is that if the increments are in the domain of attraction of an α-stable distribution, where $0 < \alpha \leq 2$, then $\operatorname{udim}_M(\Psi) = \operatorname{udim}_H(\Psi) = \min\{1, \max(0, 2\alpha - 1)/\alpha\}$ (see Theorem 3.13 of [38]). Also, the guess is that the zero set of the symmetric nearest-neighbor simple random walk in $\mathbb{Z}^2$ is $\frac{1}{4}$-dimensional.

4. Eternal Galton-Watson trees. For unimodular eternal Galton-Watson trees (Subsection 4.2.3), a conjecture is that if the offspring distribution is in the domain of attraction of an α-stable distribution, where $\alpha \in [1, 2]$, then $\operatorname{udim}_M(\boldsymbol{T}) = \operatorname{udim}_H(\boldsymbol{T}) = \frac{\alpha}{\alpha-1}$ (see [29] or Theorem 5.5 of [20]). The guess is that there is no *regularly varying* dimension function (see [19]), except in the finite-variance case ($\alpha = 2$), where one may guess that the dimension function is $r^2 \log \log r$ (see [21]).

5. Drainage networks. One can ask about the dimension of other drainage network models. In particular, the simple model of Subsection 4.5 can be extended to a model in $\mathbb{Z}^k$ for $k > 2$ and the connected component containing the origin is unimodular.

6. Embedded spaces.

Conjecture 8.4. *Under the setting of Proposition 7.24, for all $\alpha > 0$, one has $\mathcal{M}^\alpha(\boldsymbol{D}) \geq \mathcal{M}^\alpha(\boldsymbol{D}_0)$.*

Note that in the case $\alpha = 0$, the conjecture is implied by Proposition 3.28. Also, in the general case, the conjecture is implied by (7.13) and Conjecture 7.4. Another problem is the validity of Proposition 7.24 under the weaker assumption of being non-equivariantly embeddable. As a partial answer, if $\operatorname{growth}(\#N_r(\boldsymbol{o}))$ exists, then (7.12) holds. This is proved as follows:

$$\operatorname{udim}_H(\boldsymbol{D}_0) \leq \operatorname{ess\,inf} \overline{\operatorname{growth}}(\#N_r(\boldsymbol{o}_0)) \leq \operatorname{ess\,inf} \overline{\operatorname{growth}}(\#N_r(\boldsymbol{o}))$$
$$= \operatorname{ess\,inf} \operatorname{growth}(\#N_r(\boldsymbol{o})) = \operatorname{udim}_H(\boldsymbol{D}),$$

where the first inequality and the last equality are implied by the unimodular Billingsley lemma.

8.4.2 List of conjectures and problems mentioned in the previous sections

It is not known whether the lower bound (4.3) for the Hausdorff dimension of unimodular one-ended trees is always an equality or not. Problem 5.12 asks whether the equality $\operatorname{udim}_H(\boldsymbol{D}) = \overline{\operatorname{growth}}(\boldsymbol{w}(N_r(\boldsymbol{o})))$ always holds. This is implied by Problem 5.21, which states that the upper and lower growth rates of $\boldsymbol{w}(N_r(\boldsymbol{o}))$ (used in Billingsley's lemma) do not depend on $\boldsymbol{w}$.

Conjecture 5.17 states that every point-stationary point process has zero Hausdorff size unless when it is the Palm version of some stationary point process. Conjecture 7.4 states that $\mathcal{H}_M^\alpha(\boldsymbol{D}) = \xi_M^\alpha(1)$, where $\xi_M^\alpha(1)$ is define in Subsection 7.1. This implies the conjecture that Cayley graphs have positive Hausdorff size (see Subsection 6.6).

It would be interesting to find connections between unimodular dimensions and other notions of dimension, some of which are discussed in Subsection 8.1. Also, as mentioned in the introduction and Subsection 8.2, the setting of this paper might be useful in the study of examples pertaining to statistical physics or group theory.

References

[1] M. Aizenman and S. Warzel, *The canopy graph and level statistics for random operators on trees*, Math. Phys. Anal. Geom. **9** (2006), no. 4, 291–333 (2007). MR2329431

[2] D. Aldous and R. Lyons, *Processes on unimodular random networks*, Electron. J. Probab. **12** (2007), no. 54, 1454–1508. MR2354165

[3] D. Aldous and J. M. Steele, *The objective method: probabilistic combinatorial optimization and local weak convergence*, Probability on discrete structures, Encyclopaedia Math. Sci., vol. 110, Springer, Berlin, 2004, pp. 1–72. MR2023650

[4] S. Asmussen and H. Hering, *Branching processes*, Progress in Probability and Statistics, vol. 3, Birkhäuser Boston, Inc., Boston, MA, 1983. MR0701538

[5] F. Baccelli, M. O. Haji-Mirsadeghi, and A. Khezeli, *Eternal Family Trees and dynamics on unimodular random graphs*, Unimodularity in randomly generated graphs, Contemp. Math., vol. 719, Amer. Math. Soc., 2018, pp. 85–127. MR3880014

[6] F. Baccelli, M.-O. Haji-Mirsadeghi, and A. Khezeli, *On the dimension of unimodular discrete spaces*, preprint, https://www.researchgate.net/publication/328137256_On_the_Dimension_of_Unimodular_Discrete_Spaces (2020).

[7] F. Baccelli, M.-O. Haji-Mirsadeghi, and A. Khezeli, *Unimodular Hausdorff and Minkowski Dimensions*, Arxiv preprint, https://arxiv.org/abs/1807.02980 (2020).

[8] M. T. Barlow and S. J. Taylor, *Fractional dimension of sets in discrete spaces*, J. Phys. A **22** (1989), no. 13, 2621–2628, With a reply by J. Naudts. MR1003752

[9] M. T. Barlow and S. J. Taylor, *Defining fractal subsets of $\mathbf{Z}^d$*, Proc. London Math. Soc. (3) **64** (1992), no. 1, 125–152. MR1132857

[10] H. Bass, *The degree of polynomial growth of finitely generated nilpotent groups*, Proc. London Math. Soc. (3) **25** (1972), 603–614. MR0379672

[11] I. Benjamini, H. Kesten, Y. Peres, and O. Schramm, *Geometry of the uniform spanning forest: transitions in dimensions $4, 8, 12, \ldots$*, Ann. of Math. (2) **160** (2004), no. 2, 465–491. MR2123930

[12] I. Benjamini, R. Lyons, Y. Peres, and O. Schramm, *Group-invariant percolation on graphs*, Geom. Funct. Anal. **9** (1999), no. 1, 29–66. MR1675890

[13] I. Benjamini and O. Schramm, *Recurrence of distributional limits of finite planar graphs*, Electron. J. Probab. **6** (2001), no. 23, 13. MR1873300

[14] C. J. Bishop and Y. Peres, *Fractals in probability and analysis*, Cambridge Studies in Advanced Mathematics, vol. 162, Cambridge University Press, 2017. MR3616046

[15] F. R. K. Chung and S.-T. Yau, *Eigenvalues of graphs and Sobolev inequalities*, Combin. Probab. Comput. **4** (1995), no. 1, 11–25. MR1336652

[16] N. Curien, T. Hutchcroft, and A. Nachmias, *Geometric and spectral properties of causal maps*, arXiv preprint arXiv:1710.03137 (2017). MR4176785

[17] D. J. Daley and D. Vere-Jones, *An introduction to the theory of point processes. Vol. I*, second ed., Probability and its Applications (New York), Springer-Verlag, New York, 2003, Elementary theory and methods. MR1950431

[18] D. J. Daley and D. Vere-Jones, *An introduction to the theory of point processes. Vol. II*, second ed., Probability and its Applications (New York), Springer, New York, 2008, General theory and structure. MR2371524

[19] T. Duquesne, *Packing and Hausdorff measures of stable trees*, Lévy matters I, Lecture Notes in Math., vol. 2001, Springer, Berlin, 2010, pp. 93–136. MR2731897

[20] T. Duquesne and J. F. Le Gall, *Probabilistic and fractal aspects of Lévy trees*, Probab. Theory Related Fields **131** (2005), no. 4, 553–603. MR2147221

[21] Thomas Duquesne and Jean-François Le Gall, *The Hausdorff measure of stable trees*, ALEA Lat. Am. J. Probab. Math. Stat. **1** (2006), 393–415. MR2291942

[22] M. Elekes and T. Keleti, *Borel sets which are null or non-σ-finite for every translation invariant measure*, Adv. Math. **201** (2006), no. 1, 102–115. MR2204751

[23] Herbert Federer, *Geometric measure theory*, Die Grundlehren der mathematischen Wissenschaften, Band 153, Springer-Verlag New York Inc., New York, 1969. MR0257325

[24] W. Feller, *An introduction to probability theory and its applications. Vol. II*, John Wiley & Sons, Inc., New York-London-Sydney, 1966. MR0210154

[25] L. R. Ford, Jr. and D. R. Fulkerson, *Flows in networks*, Princeton University Press, Princeton, N.J., 1962. MR0159700

[26] B. E. Fristedt and W. E. Pruitt, *Lower functions for increasing random walks and subordinators*, Z. Wahrscheinlichkeitstheorie und Verw. Gebiete **18** (1971), 167–182. MR0292163

[27] M. Gromov, *Groups of polynomial growth and expanding maps*, Inst. Hautes Études Sci. Publ. Math. (1981), no. 53, 53–73. MR0623534

[28] M. Gromov, *Asymptotic invariants of infinite groups*, Geometric group theory, Vol. 2 (Sussex, 1991), London Math. Soc. Lecture Note Ser., vol. 182, Cambridge Univ. Press, Cambridge, 1993, pp. 1–295. MR1253544

[29] Bénédicte Haas and Grégory Miermont, *The genealogy of self-similar fragmentations with negative index as a continuum random tree*, Electron. J. Probab. **9** (2004), no. 4, 57–97. MR2041829

[30] J. T. Hack, *Studies of longitudinal stream profiles in virginia and maryland*, vol. 294, US Government Printing Office, 1957.

[31] O. Häggström, *Infinite clusters in dependent automorphism invariant percolation on trees*, Ann. Probab. **25** (1997), no. 3, 1423–1436. MR1457624

[32] J. Hawkes, *Trees generated by a simple branching process*, J. London Math. Soc. (2) **24** (1981), no. 2, 373–384. MR0631950

[33] H. Kesten, P. Ney, and F. Spitzer, *The Galton-Watson process with mean one and finite variance*, Teor. Verojatnost. i Primenen. **11** (1966), 579–611. MR0207052

[34] H. Kesten and B. P. Stigum, *A limit theorem for multidimensional Galton-Watson processes*, Ann. Math. Statist. **37** (1966), 1211–1223. MR0198552

[35] A. Khezeli, *Shift-coupling of random rooted graphs and networks*, Unimodularity in randomly generated graphs, Contemp. Math., vol. 719, Amer. Math. Soc., Providence, RI, 2018, pp. 175–211. MR3880017

[36] A. Khezeli, *A unified framework for generalizing the Gromov-Hausdorff metric*, Arxive preprint, https://arxiv.org/abs/1812.03760 (2018).

[37] A. Khezeli, *Metrization of the Gromov-Hausdorff (-Prokhorov) topology for boundedly-compact metric spaces*, Stochastic Process. Appl. (2019). MR4092421

[38] D. Khoshnevisan and Y. Xiao, *On the macroscopic fractal geometry of some random sets*, Stochastic analysis and related topics, Progr. Probab., vol. 72, Birkhäuser/Springer, Cham, 2017, pp. 179–206. MR3737630

[39] G. Last and H. Thorisson, *Invariant transports of stationary random measures and mass-stationarity*, Ann. Probab. **37** (2009), no. 2, 790–813. MR2510024

[40] R. Lyons, *Factors of IID on trees*, Combin. Probab. Comput. **26** (2017), no. 2, 285–300. MR3603969

[41] R. Lyons and Y. Peres, *Probability on trees and networks*, Cambridge Series in Statistical and Probabilistic Mathematics, vol. 42, Cambridge University Press, New York, 2016. MR3616205

[42] P. Mattila, *Geometry of sets and measures in Euclidean spaces*, Cambridge Studies in Advanced Mathematics, vol. 44, Cambridge University Press, Cambridge, 1995, Fractals and rectifiability. MR1333890

[43] B. G. Nguyen, *Percolation of coalescing random walks*, J. Appl. Probab. **27** (1990), no. 2, 269–277. MR1052299

[44] K. Petersen, *Ergodic theory*, Cambridge Studies in Advanced Mathematics, vol. 2, Cambridge University Press, Cambridge, 1989, Corrected reprint of the 1983 original. MR1073173

[45] R. Roy, K. Saha, and A. Sarkar, *Hack's law in a drainage network model: a Brownian web approach*, Ann. Appl. Probab. **26** (2016), no. 3, 1807–1836. MR3513607

[46] W. Rudin, *Functional analysis*, McGraw-Hill Book Co., New York-Düsseldorf-Johannesburg, 1973, McGraw-Hill Series in Higher Mathematics. MR0365062

[47] R. Schneider and W. Weil, *Stochastic and integral geometry*, Probability and its Applications (New York), Springer-Verlag, Berlin, 2008. MR2455326

[48] S. J. Taylor and J. G. Wendel, *The exact Hausdorff measure of the zero set of a stable process*, Z. Wahrscheinlichkeitstheorie und Verw. Gebiete **6** (1966), 170–180. MR0210196

[49] Á. Timár, *A stationary random graph of no growth rate*, Ann. Inst. Henri Poincaré Probab. Stat. **50** (2014), no. 4, 1161–1164. MR3269989

[50] L. van den Dries and A. J. Wilkie, *Gromov's theorem on groups of polynomial growth and elementary logic*, J. Algebra **89** (1984), no. 2, 349–374. MR0751150