



Local behavior of the Eden model on graphs and tessellations of manifolds

Dongming Merrick Hua¹ · Fedor Manin¹ · Tahda Queer² · Tianyi Wang¹

Received: 28 December 2022 / Revised: 17 October 2023 / Accepted: 10 November 2023

© The Author(s) 2023

Abstract

The Eden Model in \mathbb{R}^n constructs a blob as follows: initially a single unit hypercube is infected, and each second a hypercube adjacent to the infected ones is selected randomly and infected. Manin, Roldán, and Schweinhart investigated the topology of the Eden model in \mathbb{R}^n by considering the possible shapes which can appear on the boundary. In particular, they give probabilistic lower bounds on the Betti numbers of the Eden model. In this paper, we prove analogous results for the Eden model on any infinite, vertex-transitive, locally finite graph: with high probability as time goes to infinity, every “possible” subgraph (with mild conditions on what “possible” means) occurs on the boundary of the Eden model at least a number of times proportional to an isoperimetric profile of the graph. Using this, we can extend the results about the topology of the Eden model to non-Euclidean spaces, such as hyperbolic n -space and universal covers of certain Riemannian manifolds.

Keywords Eden model · First-passage percolation · Stochastic topology

Mathematics Subject Classification 60K35 · 82B43

✉ Tianyi Wang
tianyiwang@ucsb.edu

Dongming Merrick Hua
dongming@ucsb.edu

Fedor Manin
manin@math.ucsb.edu

Tahda Queer
taqueer@proton.me

¹ University of California Santa Barbara, Santa Barbara, USA

² Hunter College, City University of New York, New York, USA

1 Introduction

In 1961, Murray Eden proposed a stochastic model to simulate the growth of a bacterial colony or tumor on flat surfaces, known as the Eden growth model, as follows. Tessellate \mathbb{R}^n with unit cubes. Choose a starting cube to “infect” as the origin. Then, at each time step, choose a new cube at random to infect out of those adjacent to the infected cubes.

By the Cox–Durrett shape theorem Cox and Durrett (1981), the growth of the Eden model is ball-like (rather than fractal, as one might guess). That is, over time, the model (after rescaling by $t^{-1/n}$) converges to a specific convex set in \mathbb{R}^n . However, this leaves a region near the boundary of the convex set, of thickness $\sim \sqrt{t^{-1/n}}$, in which the model can behave in very complicated ways. Understanding the topology of the Eden model is one way of measuring the amount of complexity in this region, along with e.g. the area of the boundary studied by Damron et al. (2018).

In previous work by Manin et al. (2023), the homology of the Eden model near the boundary in \mathbb{R}^n was studied in detail. For every $1 \leq k \leq n - 1$, they proved an asymptotic lower bound for the k th Betti number $\beta_k(t)$ of the Eden model at time t . More specifically, they showed that there is a constant $C = C(n, k) > 0$ such that

$$\beta_k(t) \geq Ct^{\frac{n-1}{n}}$$

with high probability as $t \rightarrow \infty$. This work has recently been strengthened in some cases and extended to other first-passage percolation models by Damron et al. (2022) (Fig. 1).

The definition of the Eden model naturally extends from the cubical lattice in \mathbb{R}^n to other graphs. Cayley graphs of groups are a natural setting because like the cubical lattice, which is a Cayley graph of the group \mathbb{Z}^n , they are homogeneous: they look the same from the point of view of every vertex. Such generalized models have been




Fig. 1 A demonstration of the Eden growth model after 100, 500, and 1500 time steps respectively. This demonstration is written in the Wolfram Language, and is available here <https://demonstrations.wolfram.com/TumorGrowthModel/>

studied in the context of first-passage percolation on Cayley graphs, mostly in the past decade, although interest in percolation on Cayley graphs more generally dates back earlier Benjamini and Schramm (1996). This work has uncovered a contrast based on the growth of the group. Many of the phenomena first discovered for first-passage percolation in \mathbb{Z}^n generalize to Cayley graphs of virtually nilpotent groups (equivalently, by Gromov's theorem, groups of polynomial growth), but often not to those of groups of higher growth (Benjamini and Tessera 2015; Auffinger and Gorski 2023; Gorski 2022). For example, the shape theorem has an analogue for Cayley graphs of virtually nilpotent groups, but in groups of exponential growth no such convergence result can hold (Benjamini and Tessera 2015). Moreover, first-passage percolation in hyperbolic groups exhibits bi-infinite geodesics, a feature conjectured not to occur in \mathbb{Z}^n (Benjamini and Tessera 2017).

In this paper, we generalize the results of Manin et al. (2023) beyond \mathbb{R}^n . We first study the Eden model $A(t)$ on an infinite, locally finite, connected, vertex-transitive graph T . This class includes Cayley graphs, but also other examples such as Diestel–Leader graphs which are not Cayley graphs and not even quasi-isometric to a Cayley graph (Diestel and Leader 2001; Eskin et al. 2012). Subsets of a graph do not have higher-dimensional topology; instead, we prove that certain local patterns occur many times on the perimeter of the Eden model with high probability. To be precise:

Theorem 1.1 *Let T be an infinite, locally finite, connected, vertex-transitive graph where each vertex has degree d . Let S be any connected subgraph of T which satisfies the following:*

- S is contained in the R -ball centered at a vertex x of T ,
- S contains the inner boundary of this R -ball (all vertices at distance R from x).

Let $A(t)$ denote the Eden model on T at time t . There is a constant $C(R, d) > 0$ so that with high probability as $t \rightarrow \infty$, there are at least $CF(t)$ disjoint R -balls in T whose intersection with $A(t)$ is isomorphic to S .

Here $F(N)$ is the *isoperimetric profile* of T : the minimal boundary of a set of N vertices in T . This is a quasi-isometry invariant (up to a multiplicative constant) of T closely related to the Følner function studied e.g. in Erschler (2003). Notably, the actual boundary of $A(t)$ may grow much faster than the isoperimetric profile. For example, for amenable groups of exponential growth, such as solvable groups which are not virtually nilpotent, even the perimeter of a ball grows faster. Specifically, if T is a lattice in the three-dimensional Lie group SOL, then the perimeter of a ball is linear in its volume since both are exponential in the radius, but $F(N) \sim N/\log N$. It would be interesting to strengthen the bound for such groups.

The proof of Theorem 1.1 uses the reformulation of the Eden model as a first-passage percolation model with passage times distributed according to the exponential distribution. Our proof uses the memorylessness of the exponential distribution, and another direction for further study is to prove similar results for other distributions, perhaps using the methods of Damron et al. (2022).

We apply Theorem 1.1 to obtain topological results similar to those of Manin et al. (2023). Those results consider the Eden model as defining a random domain in \mathbb{R}^n consisting of a union of cubical tiles. Therefore, to generalize them, one looks for other

spaces that admit natural tessellations. The best-known (and, in a precise sense, most symmetric) such space is hyperbolic space: for example, for any pair (r, s) , the regular r -gon tessellates either the Euclidean plane, the sphere, or the hyperbolic plane with s meeting at every vertex. Although they are harder to classify, all higher-dimensional hyperbolic spaces also admit infinite families of inequivalent tessellations by convex polyhedra, making this setting considerably richer than the Euclidean one, where there are only finitely many types of tessellation. The Eden model on such a tessellation is defined similarly to the usual Euclidean Eden model: one starts with a single infected cell, and at each time step, one chooses a random adjacent tile to infect.

Theorem 1.2 *Consider a regular tessellation of the hyperbolic n -space \mathbb{H}^n by compact, convex tiles. (Here, **regular** means that the symmetry group of the tessellation acts transitively on the cells.) Let $A(t)$ be the Eden model at time t on this tessellation, and let $\beta_k(t)$ denote the k th Betti number of $A(t)$. Then for every $1 \leq k \leq n - 1$, there are constants $C > c > 0$ depending on k and the tessellation such that*

$$ct \leq \beta_k(t) \leq Ct$$

with high probability as $t \rightarrow \infty$.

Here the growth is linear since hyperbolic spaces have a linear isoperimetric inequality: the perimeter of every set of tiles is linear in its volume.

We also discuss a generalization that encompasses both the hyperbolic and Euclidean settings. We can think of the unit cubes of the Eden model on \mathbb{R}^n as fundamental domains of the universal cover of T^n . Similarly, the tiles of many hyperbolic tessellations are fundamental domains of the universal cover of a hyperbolic manifold. For any compact manifold M , we define a similar Eden model which builds random subsets—unions of adjacent fundamental domains—of the universal cover of M . We prove a result similar to Theorem 1.2 for the class of manifolds with non-positive sectional curvature. This class includes the Euclidean (curvature zero) and hyperbolic (constant curvature -1) cases, but also a large class of tessellations of manifolds which do not admit any constant curvature metric, such as symmetric spaces of noncompact type, as well as countless examples with less local symmetry.

Theorem 1.3 *Let M be a closed Riemannian manifold of non-positive sectional curvature, and let $A(t)$ be the Eden model at time t on the universal cover \tilde{M} . Let $\beta_k(t)$ denote the k th Betti number of $A(t)$ with coefficients in $\mathbb{Z}/2\mathbb{Z}$, and let $F(N)$ denote the isoperimetric profile of the Cayley graph of $\pi_1(M)$.¹ Then for every $1 \leq k \leq \dim M - 1$ there is a constant $C = C(M, k) > 0$ such that*

$$\beta_k(t) \geq CF(t)$$

with high probability as $t \rightarrow \infty$.

In particular, if M is negatively curved, then $\beta_k(t)$ grows linearly in t (since growth of the lower bound matches that of the trivial upper bound).

¹ We show in Appendix A that the isoperimetric profiles of two Cayley graphs of the same group differ only by a multiplicative constant.

In this case, unlike in Theorem 1.2, we don't give a matching upper bound.

Moreover, if the fundamental domains we choose (or their intersections) themselves have topology, then the lower bound in Theorem 1.3 may in fact be trivial, in the sense that it holds for every subset of T consisting of t fundamental domains. One might even worry that this is true for all possible choices of fundamental domain, and therefore our results are completely uninteresting. To forestall this objection, we show that if M has constant curvature, one can choose a tessellation (by Voronoi cells) for which any nonempty intersection of cells is contractible. It would be interesting to construct tessellations with this property in a more general setting, perhaps even for any aspherical manifold:

Conjecture 1.4 *If M is an aspherical manifold, there is a compact set K in its universal cover \tilde{M} such that translates of K under the deck action of $\pi_1(M)$ cover \tilde{M} and any nonempty intersection of these translates is contractible.*

If this condition is satisfied, then the union of a set of translates of K is homotopy equivalent to the nerve of the covering by tiles; in other words, the topology is determined by the lattice of intersections and not any local information about K .

2 Preliminaries

In this section we collect some background knowledge required to read the rest of the paper.

Graph theory

One of the main settings in which we will study the Eden model is graphs. More specifically, we will consider the Eden model on infinite, connected, locally finite, vertex-transitive graphs. Here, *locally finite* means that each vertex in the graph has finite degree, and *vertex-transitive* means that for any vertices v, w in our graph, there is a graph automorphism sending v to w . Informally, vertex-transitive graphs “look the same” at each vertex. One large class of infinite, connected, locally finite, vertex-transitive graphs is given by the Cayley graph of any infinite, finitely-generated group with respect to a finite generating set.

The Eden model on an infinite, connected, locally finite, vertex-transitive graph is defined as follows: at time $t = 0$, an arbitrary vertex of the graph is infected (it does not matter which one, since the graph is vertex-transitive). Then, at each integer time, a vertex adjacent to at least one of the previously infected vertices is randomly chosen (with equal probability) and infected. For a given graph, we will use $A(t)$ to denote the Eden model on the graph at time t .

Note that every infinite, connected, locally finite, vertex-transitive graph has a well-defined degree d , such that each vertex in the graph has degree d . This follows from the fact that the graph is locally finite and vertex-transitive, hence each vertex has the same finite degree. This gives us an analogue of the dimension of Euclidean space for graphs.

Furthermore, we can consider each infinite, connected, locally finite, vertex-transitive graph as a metric space, by endowing it with the **graph metric**: the distance between any two vertices is the length of the shortest edge path between them. Since our graph is always connected, such a path always exists. The closed balls in the graph metric will be important to our results.

Geometric group theory

The language and techniques of geometric group theory occur frequently in this paper, so we briefly collect the most important ones here. The proofs of the results in this section can be found in standard books on geometric group theory, such as Clay and Margalit (2017), Druţu and Kapovich (2018), or Löh (2017).

Definition 2.1 Given a group G with a finite generating set S , the **Cayley graph** $\text{Cay}(G, S)$ has vertex set G and an edge between g and h iff $gh^{-1} \in S \cup S^{-1}$.

We quickly prove a recognition theorem for Cayley graphs:

Lemma 2.2 *Let T be a connected, locally-finite graph with $|T| > 1$, and suppose G is a group acting on T by graph isomorphisms (for each $g, x \mapsto gx$ is a graph isomorphism of T). If the action of G on T is free and transitive, then T is a Cayley graph for G .*

Proof Pick a vertex $x \in T$. Since G acts transitively and freely on T , each vertex of T can be written uniquely as gx for some $g \in G$.

Let

$$S = \{g \in G : x, gx \text{ are adjacent}\}.$$

Since T is connected and locally-finite, S is nonempty and finite. We claim that S generates G .

First, note that since $x \mapsto g^{-1}x$ is a graph isomorphism, gx and hx are adjacent if and only if x and $g^{-1}hx$ are adjacent. In other words, $gx, hx \in T$ are adjacent if and only if $g^{-1}h$ is in S . (Note that this also shows $S = S^{-1}$.)

Let $h \in G$ be an arbitrary nonidentity element. Since T is connected, there exists a path in T from x to hx . In other words, there exist vertices $g_1x, g_2x, \dots, g_nx \in T$ such that x and g_1x are adjacent, $g_ix, g_{i+1}x$ are adjacent for all $i \in \{1, \dots, n-1\}$, and g_nx is adjacent to hx . By the previous paragraph, we have

$$g_1, g_i^{-1}g_{i+1}, g_n^{-1}h \in S$$

for all $i \in \{1, \dots, n-1\}$. We have

$$h = g_1(g_1^{-1}g_2)(g_2^{-1}g_3) \cdots (g_n^{-1}h),$$

which proves any element in G can be written as the product of elements in S , hence S generates G .

So, we can define the Cayley graph $\text{Cay}(G, S)$. Let $f: \text{Cay}(G, S) \rightarrow T$ be the map sending $g \in \text{Cay}(G, S)$ to gx . This is a bijection, since each vertex of T can be written uniquely as gx for some $g \in G$.

Furthermore, since $g, h \in \text{Cay}(G, S)$ are adjacent iff $g^{-1}h \in S \cup S^{-1} = S$, and $gx, hx \in T$ are adjacent iff $g^{-1}h \in S$, we conclude that g, h are adjacent iff $f(g) = gx$ and $f(h) = hx$ are adjacent, hence f is a graph isomorphism. \square

Equipping $\text{Cay}(G, S)$ with the graph metric, we obtain a metric space associated with G . While the Cayley graphs of a group G with respect to different generating sets S, S' are usually not isometric, they are metrically equivalent in a weaker sense:

Definition 2.3 Let X, Y be metric spaces and let $C \geq 1$ and $K \geq 0$ be constants. A function $f: X \rightarrow Y$ is a (C, K) -*quasi-isometry* if

- (Quasi-isometric embedding) For any $x, x' \in X$,

$$\frac{1}{C}d_Y(f(x), f(x')) - K \leq d_X(x, x') \leq Cd_Y(f(x), f(x')) + K.$$

- (Coarse surjectivity) For any $y \in Y$, there exists some $x \in X$ such that $d_Y(f(x), y) \leq K$.

If for some C and K there is a (C, K) -quasi-isometry $f: X \rightarrow Y$, then X and Y are *quasi-isometric*. This is an equivalence relation.

Proposition 2.4 Given a group G and two finite generating sets S, S' , the map $f: \text{Cay}(G, S) \rightarrow \text{Cay}(G, S')$ sending g to g is a quasi-isometry.

As a consequence of this proposition, a finitely generated group G is well-defined as a metric space up to quasi-isometry, and so we can talk about groups being quasi-isometric to one another or to other spaces.

Definition 2.5 A metric space X is *proper* if closed balls in X are compact.

Definition 2.6 A metric space X is a *geodesic space* if for any $x, y \in X$, there exist $a, b \in \mathbb{R}$ with $a \leq b$ and an isometric embedding $\gamma: [a, b] \rightarrow X$ such that $\gamma(a) = x$ and $\gamma(b) = y$. Such an isometric embedding is called a *geodesic*.

Note that a geodesic in this metric sense is a shortest path between any of its points. This does not always coincide with the variational definition of a geodesic in Riemannian geometry. However, in a complete Riemannian manifold, a metric geodesic is always a Riemannian geodesic; conversely, there is always a shortest Riemannian geodesic between two points and this shortest path is always a metric geodesic. In particular, every complete Riemannian manifold is a geodesic space.

Definition 2.7 Let G be a group and X a metric space. A group action $G \curvearrowright X$ is *geometric* if it satisfies the following properties:

- It *acts by isometries*: for each $g \in G$, the map $X \rightarrow X$ sending x to $g \cdot x$ is an isometry.

- It is **cocompact**: the quotient X/G is a compact space.
- It is **properly discontinuous**: for any compact $K \subseteq X$, there are only finitely many $g \in G$ such that $K \cap g \cdot K \neq \emptyset$.

A central result in geometric group theory is the Milnor–Schwarz Lemma, which states the following:

Theorem 2.8 (Milnor–Schwarz) *If G acts geometrically on a proper geodesic space X , then G is finitely generated, and for any finite generating set S of G and any $x \in X$ the map f_x sending $g \in G$ to $g \cdot x \in X$ is a quasi-isometry from $\text{Cay}(G, S)$ to X .*

One important source of geometric group actions come from fundamental groups of manifolds. Let M be a compact, connected Riemannian manifold, and equip the universal cover \tilde{M} with the unique Riemannian metric such that the covering map $p: \tilde{M} \rightarrow M$ is a local isometry. Then, the action of $\pi_1(M)$ on \tilde{M} by deck transformations is geometric (Löh (2017), Corollary 5.4.10).

Another source of geometric group actions come from lattices of isometries of hyperbolic n -space, which will be discussed in more detail in the next subsection.

Hyperbolic geometry

Hyperbolic n -space \mathbb{H}^n is the unique (up to isometry) simply-connected n -dimensional Riemannian manifold of constant sectional curvature -1 . There are multiple models of n -dimensional hyperbolic space, but the one we use in this paper is the **Poincaré disc model**: we view \mathbb{H}^n as the open unit ball in \mathbb{R}^n with Riemannian metric

$$ds^2 = \frac{4(dx_1^2 + \cdots + dx_n^2)}{(1 - (x_1^2 + \cdots + x_n^2))^2}.$$

Much of the material in this subsection can be found in (Druţu and Kapovich (2018), Chapter 4).

Hyperbolic spaces are particularly symmetric: for every $x, y \in \mathbb{H}^n$ and every orthogonal transformation $L: T_x \mathbb{H}^n \rightarrow T_y \mathbb{H}^n$, there is an isometry $f: \mathbb{H}^n \rightarrow \mathbb{H}^n$ such that $f(x) = y$ and $Df_x = L$. Since isometries take geodesics to geodesics, each such isometry is unique. Therefore, there is a bundle structure

$$O(n) \rightarrow \text{Isom}(\mathbb{H}^n) \xrightarrow{p} \mathbb{H}^n,$$

where the map p sends an isometry f to $f(x)$, for an arbitrary basepoint x .

In particular, the group of isometries fixing x is isomorphic to $O(n)$. Moreover, these subgroups are conjugate for every choice of x . Therefore, to understand them, it is enough to consider the isometries fixing the origin in the Poincaré disk model. These are exactly the orthogonal transformations of \mathbb{R}^n , since the Poincaré disk metric is rotationally invariant. The fixed set of each such transformation is a linear subspace of \mathbb{R}^n ; its intersection with the Poincaré disk is isometric to a lower-dimensional hyperbolic space inside \mathbb{H}^n . We state this as a lemma:

Lemma 2.9 *For every nontrivial isometry of hyperbolic space, the fixed set, if nonempty, is a hyperbolic subspace of lower dimension.*

As a Lie group, the isometry group is isomorphic to $O^+(n, 1)$, an index 2 subgroup of the group of linear isometries of the quadratic form

$$x_1^2 + \cdots + x_n^2 - x_{n+1}^2.$$

Definition 2.10 A *lattice* Γ in $\text{Isom}(\mathbb{H}^n)$ is a discrete subgroup $\Gamma < \text{Isom}(\mathbb{H}^n)$, where $\text{Isom}(\mathbb{H}^n)$ is given the compact-open topology. We say Γ is a *cocompact* or *uniform lattice* if the quotient $\text{Isom}(\mathbb{H}^n)/\Gamma$ is compact.

If Γ is a cocompact lattice of isometries of \mathbb{H}^n , then the action of Γ on \mathbb{H}^n is geometric. This is a special case of a more general fact: if M is a homogeneous Riemannian manifold, and Γ is a cocompact lattice in $\text{Isom}(M)$, then Γ acts geometrically on M (Druţu and Kapovich (2018), p. 142).

Cocompact lattices in \mathbb{H}^n are abundant, with infinitely many distinct types in every dimension $n \geq 2$, as demonstrated by work of Borel, Mostow, and Gromov and Piatetski-Shapiro among others; see (Druţu and Kapovich (2018), Chapter 12) for an overview.

In this paper we use such lattices to obtain tilings of \mathbb{H}^n by convex sets called Voronoi cells, on which we can then define a version of the Eden model. The full construction is given in Sect. 4.

Homology and cohomology

We use standard tools from homology theory which can be found in any algebraic topology textbook such as Hatcher (2002). Our results hold for homology both with integer coefficients and with coefficients in any field.

First-passage percolation and the Eden model

First-passage percolation (FPP) is a well-studied family of stochastic growth models; see Auffinger et al. (2017) for an extensive survey of FPP on \mathbb{Z}^n , and Benjamini and Tessera (2015), Benjamini and Tessera (2017) for some results on other Cayley graphs. In our proof in the following sections, we will make use of a kind of equivalence, first noticed by Richardson (1973), between the Eden model and a specific FPP model: site FPP with exponentially distributed passage times. While our results can be interpreted equally as applying to the Eden model and to this FPP model, the proof requires working with the FPP model.

We first define site FPP on a graph $T = (V, E)$. Every vertex of T is called a *site*. Every site p is assigned an independent and identically distributed (i.i.d.) number ρ_p called the *passage time*. The passage time determines the following: if the first site adjacent to p gets infected at time t , then at time $t + \rho_p$, the site p is infected. A graph with passage times assigned to each vertex, together with an initial vertex which is infected at time 0, determines the infected region at every time $t > 0$. We usually call

this region the **FPP ball** at time t to emphasize its similarity to a ball around the initial vertex with respect to a random metric. A probability distribution on ρ_p determines an **FPP model**, which is a distribution on the set of functions $V \rightarrow [0, \infty)$ determining the passage times of each site.

We choose our i.i.d. passage times from an exponential distribution with mean 1. The exponential distribution is **memoryless**, that is,

$$\mathbb{P}(X > t + s | X > s) = \mathbb{P}(X > t).$$

In particular, if at time t a site p is located at distance 1 from the infected region, the additional time required to infect p is again distributed exponentially with mean 1, and does not depend on the passage times of any other site.

Moreover, the probability of two sites adjacent to A having the same passage time is zero (that is, two sites are never filled in simultaneously). Therefore, for a creature that can only tell the order in which events happen, but not the duration between them, the site FPP model with exponentially distributed passage times will be statistically indistinguishable from the Eden growth model. In this sense, an Eden model is just an FPP model with time rescaled. This correspondence was first observed by Richardson (1973). Therefore, in the remainder of this paper, we will consider the FPP model with exponentially distributed ($\mu = 1$) passage times, as it is easier to work with than the Eden model itself.

One possible issue is this. We will prove several results that state that an event happens **with high probability**, that is, that the probability approaches 1 as time increases. We will prove this with respect to FPP time, but would like to state it with respect to Eden time. In Sect. 7, we justify this discrepancy by showing that in any infinite, connected, locally finite, vertex-transitive graph, as Eden time grows, FPP time also grows without bound with high probability. (In a Euclidean grid, this follows from the shape theorem, but in general, the relationship may be considerably looser.)

3 The Eden model on graphs

While the homology of the Eden model on graphs is not very interesting, we can instead generalize one of the main results shown in Manin et al. (2023), a consequence of which is the lower bound on homology. In Manin et al. (2023), it is proved that for a given “nice” subset of \mathbb{R}^n built up from unit hypercubes, with high probability we can find at least $Ct^{\frac{n-1}{n}}$ copies of this subset on the boundary of the Eden model in \mathbb{R}^n . In this section, we will show an analogous result for graphs.

First, we introduce some definitions.

Definition 3.1 (Boundary). Let $A = (V_A, E_A)$ be a subgraph of a graph $T = (V, E)$. Then the **boundary** ∂A of A is

$$\partial A = \{w \in V \setminus V_A : \exists v \in A, (v, w) \in E.\}$$

This definition, although it is convenient for our results, is non-standard: the boundary $\partial_{\text{edge}}A$ of A is usually defined to be the set of edges connecting vertices in A to vertices outside A . However, it is not hard to see that

$$|\partial A| \leq |\partial_{\text{edge}}(A)| \leq \deg T |\partial A|.$$

That is, the two notions of perimeter (size of the boundary) differ by at most a multiplicative constant, which means that many ideas are unaffected by substituting one for the other, such as the definition of non-amenability below.

Definition 3.2 Let $T = (V, E)$ be a connected, locally finite graph with the graph metric d . Given a vertex $x \in T$ and a subgraph A of T , we let

$$\begin{aligned} d(A, x) &= \min\{d(x, v) : v \in A\}. \\ B_R(A) &= \{v \in T : d(A, v) \leq R\} \\ D_R(A) &= \{v \in T : d(A, v) = R\}. \end{aligned}$$

When the subscript is too long, we will sometimes use $B(A, R)$ and $D(A, R)$ to denote the same sets. Note that $\partial A = D_1(A)$, and that in general we have $\partial B_R(A) = D_{R+1}(A)$.

The isoperimetric inequality in \mathbb{R}^n , states that given a fixed amount of volume, the shape with the least perimeter that encloses at least that much volume is a sphere, hence any shape has perimeter greater than or equal to that of a sphere of the same volume. The term $t^{\frac{n-1}{n}}$ in the lower bound on the homology of the Euclidean Eden model derives from this isoperimetric inequality. Our generalization to graphs involves defining a corresponding notion of isoperimetry for graphs.

Definition 3.3 The *Følner isoperimetric profile* $F : \mathbb{N} \rightarrow \mathbb{N}$ of a graph $T = (V, E)$ is defined by

$$F(n) = \min \{|\partial A| : |A| \geq n\}$$

where A is any subgraph of T with size at least n .

This definition is nonstandard in two ways. Firstly, it is not defined in terms of the more standard edge boundary. However, one can verify that if F_{edge} is the Følner isoperimetric profile defined in terms of the edge boundary, then by the fact that

$$|\partial A| \leq |\partial_{\text{edge}}(A)| \leq \deg T |\partial A|$$

we also have

$$F(n) \leq F_{\text{edge}}(n) \leq \deg T \cdot F(n).$$

The first inequality follows from choosing A such that $|\partial_{\text{edge}}(A)| = F_{\text{edge}}(n)$, and the second follows from choosing A such that $|\partial A| = F(n)$. Since ultimately we are only

concerned with the isoperimetric profile up to a multiplicative constant, this means the two are interchangeable for our purposes.

Secondly, F is defined ranging over all subgraphs of size at least n , whereas it is more common to define F ranging over all subgraphs of size exactly n . The reason we choose to define it this way is so that F becomes a nondecreasing function, which is necessary for certain arguments to go through.

Proposition 3.4 *Let T be an infinite vertex-transitive graph. Then*

$$F(m+n) \leq F(m) + F(n).$$

In particular, if k and n are positive integers, we have $F(kn) \leq kF(n)$.

Proof From the definition of F , there exist subgraphs A_m, A_n of T with $|A_m| = m$, $|A_n| = n$ which attain the minimum for F , i.e. $|\partial A_m| = F(m)$ and $|\partial A_n| = F(n)$. Furthermore, since T is infinite and vertex-transitive, we can choose $A_m, A_n \subseteq T$ so that $d(A_m, A_n) \geq 3$. In particular, $A_m \cap A_n = \emptyset$ and $|\partial A_m| \cap |\partial A_n| = \emptyset$.

Now take $A = A_m \cup A_n$. Then $|\partial A| = |\partial A_m| + |\partial A_n| = F(m) + F(n)$, and $|A| = m + n$. Therefore, by the definition of F ,

$$F(m+n) \leq |\partial A| = F(m) + F(n).$$

The second statement follows by induction on k . □

Examples 3.5 (a) The classical isoperimetric inequality implies that the isoperimetric profile of the Cayley graph of \mathbb{Z}^d is $Cn^{\frac{d-1}{d}}$ for some constant C .

(b) In any regular tree, the isoperimetric profile is linear in n .

(c) In general, a graph whose isoperimetric profile is bounded below by a linear function is called **non-amenable**. This is one of many equivalent definitions of non-amenability, due to Følner (1955). Let \tilde{M} be a complete manifold of negative sectional curvature and let Γ be a group acting on \tilde{M} geometrically. Then \tilde{M} is a hyperbolic metric space (Löh 2017, Example 7.2.3), and Γ is a Gromov hyperbolic group. In particular, this means that Γ is non-amenable (Löh (2017), Corollary 9.1.11) and therefore the isoperimetric profile of Γ is bounded below by a linear function. This is related to the fact that \tilde{M} also has a linear isoperimetric profile (in the Riemannian sense).

The main result of this section is Theorem 1.1, which we restate here:

Theorem 3.6 *Let T be an infinite, locally finite, connected, vertex-transitive graph where each vertex has degree d . Let S be any connected subgraph of T which satisfies the following:*

- S is contained in an R -ball $B_R(x) = \{v \in T : d(v, x) \leq R\}$ whose center x is a vertex of T ,
- S contains $D_R(x) \subset B_R(x)$.

Let $A(t)$ denote the Eden model on T at time t . There is a constant $C(R, d) > 0$ so that with high probability as $t \rightarrow \infty$, there are at least $CF(|A(t)|)$ disjoint R -balls in T whose intersection with $A(t)$ is isomorphic to S .

Note that since T is connected, we can assume $d \geq 2$.

To prove Theorem 3.6, we proceed as follows. We first show that at any time t , we can put $\Omega(F(|A(t)|))$ disjoint R -balls adjacent to $A(t)$, in the sense that these balls touch $A(t)$ at the boundary, but they do not intersect with $A(t)$. This is proven in Lemma 3.9. In particular, we can find $\Omega(F(|A(t-2)|))$ such balls adjacent to $A(t-2)$. Lemma 3.10 shows that this implies we can find $\Omega(F(|A(t)|))$ many such balls adjacent to $A(t-2)$. Finally, Lemma 3.14 shows that at time t , a positive fraction of these balls will intersect with $A(t)$ in a translate of S .

Remark 3.7 Further in the paper, we actually use a slight strengthening of Theorem 3.6. Fix a base vertex v_0 of T and suppose that for every $v \in T$ we have fixed a graph automorphism $\phi_v : T \rightarrow T$ that takes v_0 to v . Then we can choose $CF(|A(t)|)$ disjoint balls such that whenever v is the center of one of them, then $B_R(v) \cap A(t) = \phi_v(S)$. One can see that this is true by tracing through the proof of the theorem.

We first give a rough estimate of the volume of R -balls in our graph:

Lemma 3.8 *Let T be an infinite vertex-transitive graph of degree $d \geq 2$, and let $B_R(x)$ denote the ball of radius R centered at $x \in T$. Then $|B_R(x)| \leq d^{R+1}$.*

Proof Recall that $D_R(x)$ denotes the set of points in T at distance R from x . Note that for each $R \geq 1$, every point in $D_R(x)$ is distance one from some point in $D_{R-1}(x)$, and therefore

$$D_R(x) \subseteq \bigcup_{y \in D_{R-1}(x)} B_1(y).$$

Since each vertex has degree d , we have $|B_1(y)| = d$ for all y , and hence

$$|D_R(x)| \leq \left| \bigcup_{y \in D_{R-1}(x)} B_1(y) \right| \leq |D_{R-1}(x)| \cdot d.$$

Since $|D_0(x)| = 1$, by induction we have $|D_R(x)| \leq d^R$ for all R . It follows that

$$|B_R(x)| = \left| \bigcup_{i=0}^R D_i(x) \right| \leq 1 + d + d^2 + \dots + d^R = \frac{d^{R+1} - 1}{d - 1} \leq d^{R+1}.$$

□

Lemma 3.9 *There is a constant $C_1(R, d) > 0$ such that we can find $C_1 F(|A(t)|)$ disjoint balls of radius R which are adjacent to but do not intersect $A(t)$.*

Proof We use A as a shorthand notation for $A(t)$. As before, $D_R(A)$ denotes the set of points in T whose distance from A is exactly R .

As noted above, we have $D_{R+1}(A) = \partial B_R(A)$, and we have $|A| \leq |B_R(A)|$ (because $A \subseteq B_R(A)$). So, by definition of the Følner isoperimetric profile,

$$|D_{R+1}(A)| = |\partial B_R(A)| \geq F(|B_R(A)|) \geq F(|A|).$$

Now consider the set of all possible centers of R -balls which are adjacent to but do not intersect $A(t)$. These are exactly the points of $D_{R+1}(A)$. We now pick the centers of our balls by a greedy procedure. First, pick any $x_1 \in D_{R+1}(A)$ as the center of our first R -ball. Then any vertex v with $d(v, x_1) \leq 2R$ in the graph metric is no longer an available candidate to be the center of another such R -ball (since we need the R -balls to be disjoint). At worst, we have removed $\text{Vol}(B_{2R}) \leq d^{2R+1}$ points of $D_{R+1}(A)$ from the candidate pool. Repeatedly applying this argument, we see that we can put at least

$$\left\lceil \frac{|D_{R+1}(A)|}{d^{2R+1}} \right\rceil \geq \left\lceil \frac{F(|A|)}{d^{2R+1}} \right\rceil \geq d^{-(2R+1)} F(|A|)$$

disjoint balls of radius R which are adjacent to but do not intersect $A(t)$. This completes the proof with $C_1 = d^{-(2R+1)}$. \square

Lemma 3.10 *Let G (we use G to denote a graph here since we use T to denote a tree in the proof) be an infinite, connected, locally finite, vertex-transitive graph of degree d , and let $A(t)$ be the ball at time t in the mean 1 exponentially distributed FPP model on G . There is a constant $C_2 > 0$ depending on d such that*

$$F(|A(t)|) \leq C_2 F(|A(t-2)|)$$

with high probability as $t \rightarrow \infty$.

Proof We will show that there is a constant C_2 such that

$$|A(t)| \leq C_2 |A(t-2)| \tag{3.11}$$

with high probability as $t \rightarrow \infty$. Without loss of generality, we can assume C_2 is an integer (otherwise take its ceiling). Since F is nondecreasing by definition and subadditive by Proposition 3.4, (3.11) implies that

$$F(|A(t)|) \leq F(C_2 |A(t-2)|) \leq C_2 F(|A(t-2)|),$$

completing the proof.

To prove (3.11), we will show that there is a constant $\varepsilon > 0$ such that with high probability as $t \rightarrow \infty$ we have

$$|A(t+\varepsilon)| \leq |A(t)| + |\partial A(t)|. \tag{3.12}$$

In particular we have $|A(t+\varepsilon)| \leq (d+1)|A(t)|$. Then taking $C_2 = (d+1)^{\lceil 2/\varepsilon \rceil}$ does the trick.

To prove (3.12), we first consider a rooted tree T of degree $d-1$ whose nodes are equipped with i.i.d., exponentially distributed passage times with mean 1. Choose $\varepsilon > 0$ such that for any node $p' \in T$, we have

$$\mathbb{P}(\rho_{p'} < \varepsilon) = \frac{1}{d+1}.$$

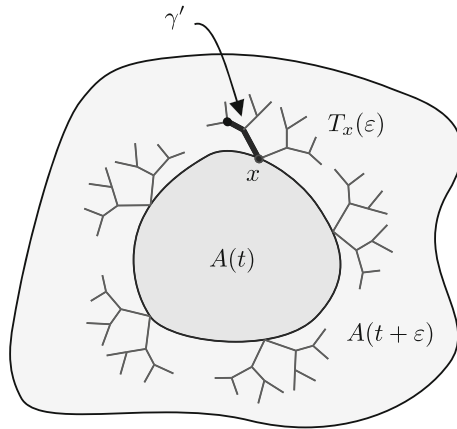


Fig. 2 An illustration of how we couple sites in $A(t + \varepsilon) - A(t)$ with nodes in $\coprod_{x \in \partial A(t)} T_x$. For every $x \in \partial A(t)$, a tree T_x rooted at x is mapped into $A(t + \varepsilon)$. A path $\gamma' : [0, 1] \rightarrow \coprod_{x \in \partial A(t)} T_x$ is illustrated in the figure, identified with $f(\gamma') : [0, 1] \rightarrow G$. The same two sites in G can be connected by multiple paths in different trees, since $f : \coprod_{x \in \partial A(t)} T_x \rightarrow G$ is not injective. However, each path from some $x \in \partial A(t)$ to another vertex y is associated with a unique path in T_x rooted at x

Let $T(\varepsilon) \subseteq T$ be the maximal subtree of T containing the root whose nodes all have passage time less than ε . Then the expected size E of $T(\varepsilon)$ satisfies the recurrence relation

$$E = \frac{1}{d + 1}(1 + (d - 1)E).$$

Solving this we get $E = \frac{1}{2}$.

Now we start to prove (3.12). Fix $t > 0$. We give a coupling between the distribution of passage times on a collection of $(d - 1)$ -ary rooted trees and the distribution of passage times for sites in G outside $A(t)$. We describe this coupling by assigning passage times to sites in G using the passage times on the collection of trees. With every $x \in \partial A(t)$ we associate a tree T_x and a graph homomorphism $f_x : T_x \rightarrow G$. This map sends the root o_x of the tree to x , and it sends the $d - 1$ vertices adjacent to o_x to $d - 1$ different vertices adjacent to x , avoiding the one vertex contained in $A(t)$; and from there the map is defined inductively so that the d vertices adjacent to a given v are mapped to all d vertices adjacent to $f_x(v)$. Thus we obtain a map

$$f : \coprod_{x \in \partial A(t)} T_x \rightarrow G$$

such that the root o_x of T_x is mapped via f to $x \in \partial A(t)$, and every path in $G \setminus A(t)$ connecting a vertex in $\partial A(t)$ to another vertex in $G \setminus A(t)$ corresponds to a unique rooted path in $\coprod_x T_x$.

See Fig. 2 for an intuitive picture for this construction.

Now we assign passage times to sites in $G \setminus A(t)$. To each $x \in \partial A(t)$ we assign the passage time of o_x . To define the passage time for other vertices, we first give a couple of definitions:

- The **passage time along a path** in a graph is the sum of the passage times of all the vertices of the path, including endpoints.
- Given a vertex $y \in G \setminus A(t)$, the **passage time from $A(t)$ to y** is the minimal passage time along any path from any vertex of $\partial A(t)$ to y .
- Given a vertex $z \in \coprod_x T_x$, the **passage time to just before z** is the sum of the passage times of the vertices along the unique simple path from x to z , **not** including z .

Then to every remaining $y \in G \setminus A(t)$, we assign the passage time of \tilde{y} , where \tilde{y} minimizes the passage time to just before \tilde{y} among all vertices in $f^{-1}(y)$.

Since the choice of \tilde{y} does not depend on the passage time of \tilde{y} , the resulting passage time of y is exponentially distributed with mean 1 and independent of all other passage times. Therefore this indeed defines a coupling.

Lemma 3.13 *For every $y \in G \setminus A(t)$, if the passage time from $A(t)$ to y is ρ , then there is a $\tilde{y} \in f^{-1}(y)$ such that the passage time from \tilde{y} to the root of the tree containing it is at most ρ .*

Proof We prove this for every $y \in G \setminus A(t)$ by induction on the passage time from $A(t)$ to y . For every site $x \in \partial A(t)$, its passage time is the passage time of o_x , so the statement is true. Now let $y \in G \setminus A(t)$ be another site, let $\tilde{y} \in T_x$ be the site that it is coupled to and consider the path which realizes the passage time from $A(t)$ to y (the “shortest path”). By induction, the passage time from $A(t)$ to the second-to-last vertex of this path is greater than the passage time to just before \tilde{y} . Therefore, by the definition of the coupling, the passage time from $A(t)$ to y is greater than the passage time from o_x to \tilde{y} . □

In particular, if $y \in A(t + \varepsilon) \setminus A(t)$, then there is some T_x and some $\tilde{y} \in f^{-1}(y) \cap T_x$ such that the passage time from o_x to \tilde{y} is at most ε . Therefore \tilde{y} is contained in the subtree $T_x(\varepsilon)$ of T_x , which is again the maximal subtree of T_x containing the root whose nodes all have passage time less than ε . It follows that

$$A(t + \varepsilon) \setminus A(t) \subseteq f\left(\coprod_{x \in \partial A(t)} T_x(\varepsilon)\right).$$

We determined above that $\mathbb{E}(|T_x(\varepsilon)|) = \frac{1}{2}$. Since the $|T_x(\varepsilon)|$ are i.i.d. random variables, the law of large numbers implies that with high probability as $t \rightarrow \infty$,

$$\left| \coprod_{x \in \partial A(t)} T_x(\varepsilon) \right| \leq |\partial A(t)|.$$

This proves (3.12). □

By Lemma 3.9, at time $t - 2$, we can find $N \geq C_1 F(|A(t - 2)|)$ disjoint R -balls B_1, \dots, B_N that are adjacent to but do not intersect $A(t - 2)$. By Lemma 3.10,

$$N \geq C_1 F(|A(t - 2)|) \geq C_1 C_2^{-1} F(|A(t)|).$$

Lemma 3.14 *Let S be as in the statement of Theorem 3.6. There is a $C_3 = C_3(R, d) > 0$ such that with high probability as $t \rightarrow \infty$, the number of values of j , for $1 \leq j \leq N$, such that $B_j \cap A(t)$ is isomorphic to S is at least $C_3(R, d)C_1C_2^{-1}F(|A(t)|)$. Here, the B_j are disjoint R -balls B_1, \dots, B_N that are adjacent to but do not intersect $A(t - 2)$, whose existence is guaranteed by Lemma 3.9.*

Proof Let p be a site that is not contained in $A(t - 2)$. Denote its passage time by ρ_p . By the memorylessness of the exponential distribution, these passage times ρ_p are all i.i.d. Let S_j be a translate of S inside B_j . Now let X_j be the event that for all $p \in B_j$,

$$\begin{aligned} \rho_p &\leq d^{-R-1} \equiv \varepsilon && \text{if } p \in S_j \\ \rho_p &\geq 3 && \text{if } p \in B_j - S. \end{aligned}$$

Then all the X_j are also i.i.d. and the probability that X_j occurs is positive since

$$\begin{aligned} \mathbb{P}(X_j) &= \mathbb{P}(\rho_p \leq \varepsilon)^{|S|} \cdot \mathbb{P}(\rho_p \geq 3)^{|B_j - S_j|} \\ &\geq \min \{ \mathbb{P}(\rho_p \leq \varepsilon), \mathbb{P}(\rho_p \geq 3) \}^{|B_j|} \\ &\geq \min \{ \mathbb{P}(\rho_p \leq \varepsilon), \mathbb{P}(\rho_p \geq 3) \}^{d^{R+1}} > 0. \end{aligned}$$

Taking $C_3(R, d) = \frac{1}{2} \min \{ \mathbb{P}(\rho_p \leq \varepsilon), \mathbb{P}(\rho_p \geq 3) \}^{d^{R+1}} > 0$ as our constant, by the law of large numbers, with high probability we know as least $C_3(R, d)C_1C_2^{-1}F(|A(t)|)$ of the X_j will occur.

Suppose X_j occurs. Let p be a vertex of $A(t - 2)$ which is adjacent to B_j . Since S_j is connected and (since it contains the whole outer layer of B_j) contains a vertex adjacent to p , every point of S_j is connected to p by a path γ of length at most $d^{R+1} - 1$. So the time it takes for all points in S to get infected is at most

$$\sum_{q \in \gamma} \rho_q \leq \ell(\gamma)\varepsilon \leq (d^{R+1} - 1)\varepsilon = \frac{d^{R+1} - 1}{d^{R+1}} < 1.$$

Therefore, for $t - 1 < s < t + 1$, we know $A(s)$ will contain all points of S_j and no points in $B_j - S_j$. In particular, when $s = t$, we know $A(t)$ will contain all points of S_j and no points of $B_j - S_j$. Since we have shown that at least $C_3C_1C_2^{-1}F(|A(t)|)$ of such events will occur with high probability, this concludes the proof. \square

Taking $C = C_3C_1C_2^{-1}$ completes the proof of Theorem 3.6.

4 The Eden model on hyperbolic tessellations

We have established our key technical result Theorem 3.6. In the following two sections we demonstrate how Theorem 3.6 can be applied to obtain topological information about the Eden model on hyperbolic tessellations, parallel the results of Manin et al. (2023). Later, we will give a further generalization of these results.

General setup

Denote the n -dimensional hyperbolic space by \mathbb{H}^n . A *tessellation* of \mathbb{H}^n is a covering by a finite set of isometry classes of *tiles*: compact sets K_α with nonempty interior such that a point is in the boundary of K_α if and only if it is shared with some other K_β .

We construct a large family of regular tessellations (ones using one isometry type of cell) using Voronoi cells of orbits of lattices. Let $\Gamma \subset \text{Isom}(\mathbb{H}^n)$ be a cocompact lattice.

Lemma 4.1 *There is a point $x \in \mathbb{H}^n$ with trivial stabilizer with respect to the action of Γ .*

Proof For every non-identity element $g \in \Gamma$, let X_g be its fixed set. By Lemma 2.9, each X_g is either empty or a hyperbolic subspace of lower dimension. Since Γ is countable, the union of all the X_g is a strict subset of \mathbb{H}^n . \square

Now pick a point $x \in \mathbb{H}^n$ with trivial stabilizer, and consider the orbit \mathcal{O}_x of x under the action of Γ . The *Voronoi cell* of an element $x_\alpha \in \mathcal{O}_x$ is the set

$$V_\alpha = \{y \in \mathbb{H}^n : \forall x_\beta \in \mathcal{O}_x \ d(y, x_\alpha) \leq d(y, x_\beta)\}.$$

Clearly, the Voronoi cells associated to \mathcal{O}_x are closed and isometric, since they are translates of each other under the action of Γ . We will show that they are compact, convex, and the tiles of a tessellation.

Lemma 4.2 *The Voronoi cells associated to \mathcal{O}_x have bounded diameter, hence they are compact.*

Proof Since Γ is a cocompact lattice, by definition there is a compact set $K \subset \mathbb{H}^n$ such that $\bigcup_{g \in \Gamma} g \cdot K = \mathbb{H}^n$. In particular, $h \cdot x \in K$ for some $h \in \Gamma$. Then K has bounded diameter D . Since every point $y \in \mathbb{H}^n$ is contained in $g \cdot K$ for some $g \in \Gamma$, y is at distance at most D from $gh \cdot x$. Therefore every point in the Voronoi cell V_α is at distance at most D from x_α . By the triangle inequality, $\text{diam}(V_\alpha) \leq 2D$. \square

Lemma 4.3 *Voronoi cells associated with \mathcal{O}_x are the tiles of a tessellation.*

Proof Let $\{V_\alpha\}$ be the set of Voronoi cells associated with \mathcal{O}_x . It suffices to show two things:

- (1) The V_α cover \mathbb{H}^n .
- (2) If a point $y \in \mathbb{H}^n$ is contained in $V_\alpha \cap V_\beta$ for some $\alpha \neq \beta$, then $y \in \partial V_\beta$.

To see (1), notice that \mathcal{O}_x is a closed set. Therefore for any $y \in \tilde{X}$ there is a closest point of \mathcal{O}_x .

To prove (2), consider a point $y \in V_\alpha \cap V_\beta$ and consider a minimal geodesic $\gamma : [0, 1] \rightarrow \tilde{X}$ from y to x_α . We claim that $\gamma((0, 1])$ does not intersect V_β . Indeed, suppose that $z = \gamma(t)$ were contained in V_β . Then $d(z, x_\beta) \leq d(z, x_\alpha)$.

Suppose first that $d(z, x_\beta) < d(z, x_\alpha)$. But this would then mean that

$$d(y, x_\beta) \leq d(y, z) + d(z, x_\beta) < d(y, z) + d(z, x_\alpha) = d(y, x_\alpha),$$

which contradicts the fact that $y \in V_\alpha$.

Now suppose that $d(z, x_\beta) = d(z, x_\alpha)$. But then concatenating $\gamma|_{[0,t]}$ with a minimal geodesic from z to x_β would give a shortest path from y to x_β . Since \mathbb{H}^n is a complete Riemannian manifold, this shortest path is a geodesic. But since $\gamma|_{[0,t]}$ is an initial segment of this geodesic, the geodesic is in fact γ and $x_\alpha = x_\beta$, again a contradiction.

Since the sequence $\{\gamma(1/n)\}_{n \in \mathbb{N}}$ converges to y and consists of points outside V_β , y is in the boundary of V_β . □

Definition 4.4 A set $\Omega \subseteq \mathbb{H}^n$ is said to be **convex** if for any pair $x, y \in \Omega$, the unique geodesic connecting x to y is contained in Ω .

Definition 4.5 Let $p, q \in \mathbb{H}^n$. Then their **bisector** $B(p, q)$ is the set

$$B(p, q) = \{x \in \mathbb{H}^n : d(p, x) = d(q, x)\}.$$

Lemma 4.6 Let $p, q \in \mathbb{H}^n$. Then $B(p, q)$ is a totally geodesic subspace.

Proof We use the upper half-plane model of \mathbb{H}^n . We can always apply an elliptic isometry to rotate p around q so that they differ only in the x_1 -coordinate. Then $B(p, q)$ is an $n - 1$ -dimensional hyperplane with constant x_1 -coordinate, which is a totally geodesic subspace in \mathbb{H}^n . Since isometries preserve geodesics, this concludes the proof. □

In Beem (1975), it is shown that the converse of this statement is also essentially true: if the bisector of any pair of points in a Riemannian (or pseudo-Riemannian) manifold M is totally geodesic, then M has constant curvature.

The next result tells us that the Voronoi cells constructed above are convex when $\tilde{X} = \mathbb{H}^n$:

Lemma 4.7 Suppose that $\{V_\alpha\}$ is the set of Voronoi cells generated by a discrete set of points $\{x_\alpha\}$ in \mathbb{H}^n . Then any nonempty intersection of the V_α is a convex set, and therefore contractible.


Proof Since the intersection of convex sets is always convex, it suffices to prove that a single V_α is convex. Now notice that

$$V_\alpha = \bigcap_{\beta \neq \alpha} \{y \in \mathbb{H}^n : d(y, x_\alpha) \leq d(y, x_\beta)\}.$$

Write $H_{\alpha,\beta} = \{y \in \mathbb{H}^n : d(y, x_\alpha) \leq d(y, x_\beta)\}$. It suffices to prove that each $H_{\alpha,\beta}$ is convex.

To do this, we use the convexity of $B(x_\alpha, x_\beta) = \partial H_{\alpha,\beta}$. Suppose that the shortest path $\gamma : [0, 1] \rightarrow \mathbb{H}^n$ between $y, z \in H_{\alpha,\beta}$ has a point outside of $H_{\alpha,\beta}$, say $\gamma(t)$.

Fig. 3 An example of a tessellation of \mathbb{H}^n by Voronoi cells is the regular hyperbolic tiling of type (n, m) , which is a tiling where m regular n -gons share a vertex. This image by DmitryBrant <https://dmitrybrant.com/2007/01/24/hyperbolic-tessellations> shows a $(7, 3)$ regular tiling in blue and its dual graph T in red (color figure online)



Then there must be two points $s < t < s'$ such that $\gamma(s), \gamma(s') \in B(x_\alpha, x_\beta)$. But then the shortest path between $\gamma(s)$ and $\gamma(s')$ lies in $B(x_\alpha, x_\beta)$. Substituting this for the portion of γ between s and s' creates a path between y and z which is shorter than γ , a contradiction. \square

Hyperbolic Eden models

We have shown that Voronoi cells of an orbit \mathcal{O}_x of an action of Γ on \mathbb{H}^n are convex fundamental domains. The generalization of the Eden model on a tessellation of \mathbb{H}^n is then clear: pick a tile to start with; then at every time step, randomly choose a tile adjacent to the infected region to infect next. Here we say two tiles are *adjacent* if they share an $(n - 1)$ -dimensional face. We will use $A(t)$ to denote the Eden model at time t .

We can view each tile V_α as the vertex α of a graph T , and we stipulate that there is an edge from α to β in T if and only if

$$\dim(V_\alpha \cap V_\beta) = n - 1$$

in \mathbb{H}^n (Fig. 3).

Lemma 4.8 *A closed ball in \mathbb{H}^n intersects finitely many compact tiles of any tessellation.*

Proof Let d denote the maximum diameter of all the tiles, since there are finitely many isometry classes. Similarly, let v denote the minimal volume of a tile. If a closed ball $B(y, r)$ intersects a tile V , then the ball $B(y + d, r)$ contains V . Therefore the ball $B(y, r)$ contains at most $v^{-1} \text{Vol } B(y + d, r)$ tiles. \square

We will also use the following well-known result to prove that T is a Cayley graph for Γ :

Lemma 4.9 (*Thom's transversality theorem*). *Let P be a smooth manifold and M be a smooth submanifold of manifold N . Let $C_{\text{th}M}^\infty(P, N)$ denote the set of all smooth*

maps $P \rightarrow N$ that are transversal to M . Then $C_{\text{rh}M}^\infty(P, N)$ is a dense subspace of $C(P, N)$ with the compact-open topology.

This result was originally obtained in Thom (1954). The proof of a slightly weaker version can be found in (Kosinski (1993), chapter IV.2).

Lemma 4.10 *Suppose that Γ acts transitively on the tiles of a tessellation by convex, compact tiles. Then the graph T is a Cayley graph for Γ .*

In particular, this means that T is infinite, locally finite, connected, and vertex-transitive.

Proof By Lemma 2.2, it suffices to prove that T is locally finite and connected and Γ acts freely and transitively on T by graph isomorphisms.

The action of Γ on T is induced by an action by isometries on \mathbb{H}^n , which takes $(n - 1)$ -dimensional intersections to $(n - 1)$ -dimensional intersections. Therefore Γ acts by graph homomorphisms. Moreover, the action is free since we chose x to have trivial stabilizer and transitive since the centers of cells are exactly the orbit of x under Γ .

By Lemma 4.8, there are finitely many tiles that intersect any closed ball, hence T is locally finite.

To show that T is connected, let us pick $\alpha, \beta \in T$. We want to show that there is a path in T that starts at α and ends at β . This amounts to showing that in \mathbb{H}^n , we can find a path starting in V_α and ending in V_β that only travels between adjacent tiles, i.e., it avoids any intersections of dimension $\leq n - 2$. Since \mathbb{H}^n is path-connected, we can certainly find a smooth path $\gamma : [0, 1] \rightarrow \mathbb{H}^n$ starting in V_α and ending in V_β . Take an open ball B that contains V_α, V_β , and the image of γ . By Lemma 4.8, there are finitely many intersections of dimension $\leq n - 2$ inside B , each of which is contained in a hyperplane \mathcal{H}^{n-2} of dimension $n - 2$.

Now we claim that we can perturb the path γ to a new path γ' in B which is transversal to the finitely many hyperplanes of dimension $n - 2$ mentioned above, and whose endpoints still lie in V_α and V_β . Since for every point $u \in \gamma' \cap \mathcal{H}^{n-2}$, we have

$$\dim([0, 1]) + \dim \mathcal{H}^{n-2} = n - 1 < n = \dim \mathbb{H}^n,$$

transversality implies that $\gamma' \cap \mathcal{H}^{n-2} = \emptyset$, so γ' is the path that we want. Let

$$\varepsilon = \min\{d(\gamma(0), \partial V_\alpha), d(\gamma(1), \partial V_\beta), d(\gamma([0, 1]), \partial B)\}.$$

By Lemma 4.9, we may perturb γ by less than $\varepsilon/2$ to a new path $\gamma^{(1)}$ such that it is transversal to the first hyperplane, i.e. $\gamma^{(1)} \pitchfork \mathcal{H}_1^{n-2}$. Then, inductively, at j -th step, let $d_j = \min\{d(\gamma^{(j-1)}, \mathcal{H}_1^{n-2}), \dots, d(\gamma^{(j-1)}, \mathcal{H}_{j-1}^{n-2})\}$. We perturb $\gamma^{(j-1)}$ by less than $\min(d_j, \varepsilon/2^j)$ such that $\gamma^{(j)} \pitchfork \mathcal{H}_j^{n-2}$. Note that we can perturb $\gamma^{(j-1)}$ in such that way that it stays transversal to all the previous hyperplanes, i.e., $\gamma^{(j)} \pitchfork \mathcal{H}_k^{n-2}$ for all $k \leq j$. Since there are only finitely many such hyperplanes, say N of them, this

process must terminate. In the end, we will obtain a path γ' such that $\gamma' \cap \mathcal{H}_k^{n-2}$ for all $1 \leq k \leq N$, and we can control the total perturbation by

$$d(\gamma, \gamma') < \varepsilon \left(\frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \cdots + \frac{1}{2^N} \right) < \varepsilon.$$

Hence the starting point of γ' is in V_α and the ending point of γ' is in V_β , as desired. Moreover, since the image of γ' still lies in B , it does not hit any new $(n-2)$ -dimensional intersections. \square

By the Milnor–Schwarz lemma, it follows that the map $T \rightarrow \mathbb{H}^n$ which sends each $\alpha \mapsto x_\alpha$ is a quasi-isometry. Moreover, since Γ acts on hyperbolic space, it is non-amenable by Example 3.5(c), that is, its isoperimetric profile function is bounded below by a linear function.

The main result of this section is:

Theorem 4.11 *Let $A(t)$ be the Eden model at time t on a Voronoi tessellation of \mathbb{H}^n whose centers are an orbit of the action of a lattice Γ . Let $1 \leq k \leq n-1$ and let $\beta_k(t)$ denote the k th Betti number of $A(t)$, i.e. $\beta_k(t) = \text{rank } H_k(A(t))$. Then there are constants $C > c > 0$ depending on k and the tessellation such that*

$$c|A(t)| \leq \beta_k(t) \leq C|A(t)|$$

with high probability as $t \rightarrow \infty$.

Remark 4.12 In the proof we use only that we have a regular tessellation of \mathbb{H}^n by convex compact tiles. Thus we prove the slightly more general statement in Theorem 1.2.

Here is a rough sketch of the proof. By Theorem 3.6, there is a constant $C(R, d) > 0$ (where d is the degree of the graph T constructed above) so that for any connected subgraph S that satisfies certain hypotheses, the intersection of $A(t)$ with at least $C|A(t)|$ disjoint R -balls is isomorphic to S , with high probability as $t \rightarrow \infty$. Each such subgraph S will correspond to a union of Voronoi cells in \mathbb{H}^n , which we denote by S . This means that we can find at least $C|A(t)|$ copies of S along the periphery of the Eden model $A(t)$ at time t . Then to prove Theorem 4.11, it suffices to construct, for each $1 \leq k \leq n-1$, an S such that each copy of S contributes at least 1 to the k th Betti number β_k . The next section will be devoted to the construction of such an S .

There is one subtlety. Theorem 3.6 lets us find many copies of the subgraph S in T , and moreover (as shown in the proof) there is a graph automorphism of T sending any of these subgraphs to any other. But in \mathbb{H}^n , we will need copies of S (as geometric realization in \mathbb{H}^n) to be homeomorphic, which cannot be guaranteed by the existence of such a graph automorphism unless the automorphism is induced by the group action of G . Therefore we must use the modification of Theorem 3.6 discussed in Remark 3.7.

Upper bound

The upper bound of Theorem 4.11 is essentially trivial, in the sense that it holds for every union of tiles, not just those generated by the Eden model. Consider any union

U of m tiles of our tessellation. The set of tiles forms what is known as a **good cover** of U : each tile as well as any nonempty intersection of tiles is contractible, because our tiles are convex. This means that the homology of U is the same as the homology of the nerve N of the good cover, since by Leray's nerve theorem the nerve of the good cover N is homotopy equivalent to U . Every tile intersects the same finite number of other tiles, say r . Therefore, for every k , the number of k -simplices of N is at most $\binom{r}{k+1}m$. This is an upper bound for the rank of $H_k(U)$.

5 Constructing hyperbolic handles

We now construct S with the properties described above for $1 \leq k \leq n - 2$. The case $k = n - 1$ involves a simpler S , so we address it separately in the proof below.

We use the Poincaré disc model of \mathbb{H}^n . In this model, \mathbb{H}^n is identified with the open unit disc in \mathbb{R}^n , with metric

$$ds^2 = \frac{4(dx_1^2 + \dots + dx_n^2)}{(1 - (x_1^2 + \dots + x_n^2))^2}.$$

We first construct a copy of S at the origin. Then we can use a deck transformation to shift S .

Notation

We consider two different metric spaces with distinct, if related, notions of distance. The first is \mathbb{H}^n with the hyperbolic metric d induced by the metric ds^2 . The second is the graph T induced by the tessellation, equipped with the graph metric d_T . For the rest this section, we establish the following notational conventions:

- Any subgraph of the graph T will be denoted by calligraphic letters. For example, the subgraph of T associated with the hyperbolic handle will be denoted by \mathcal{S} . A closed ball of radius R in the graph metric d_T centered at x will be denoted by $\mathcal{B}_R(x)$.
- For any subgraph \mathcal{A} of the graph T , we denote the union of the corresponding tiles in \mathbb{H}^n by $\langle \mathcal{A} \rangle$.
- A subset of the hyperbolic space \mathbb{H}^n will be denoted by ordinary italic letters. The hyperbolic handle viewed as a subset of \mathbb{H}^n (a union of tiles) will be denoted by S . A closed ball of radius R in the hyperbolic metric (which we denote by d or $d_{\mathbb{H}^n}$) centered at x will be denoted by $B_R(x)$, and its boundary will be denoted by $\partial B_R(x) = S_R(x)$.

Construction at the origin

We assume without loss of generality that the origin is the center of one of our tiles.

Fix $1 \leq k \leq n - 2$. Choose a k -dimensional plane P^k through the origin in \mathbb{R}^n and let P^{n-k} be the plane through the origin perpendicular to P^k . These two planes

intersect the Poincaré disc \mathbb{H}^n . Denote the union of all tiles in \mathbb{H}^n that intersect $\overline{P^k}$ by $\overline{P^k}$. Denote the union of all tiles in \mathbb{H}^n that intersect P^{n-k} but do not intersect $\overline{P^k}$ by $\overline{P^{n-k}}$. We claim:

Lemma 5.1 *There exists $R_1 > 0$ such that for all vertices α of T satisfying $d_T(\alpha, 0) \geq R_1$, if the tile V_α intersects P^{n-k} , then $V_\alpha \in \overline{P^{n-k}}$.*

Proof Let $u, v \in \mathbb{R}^n$ be two vectors in P^k and P^{n-k} respectively with Euclidean norm less than 1, so that they live in the Poincaré disc \mathbb{H}^n . Their distance in \mathbb{H}^n is given by

$$d_{\mathbb{H}^n}(u, v) = 2 \ln \left(\frac{\|u - v\| + \sqrt{\|u\|^2\|v\|^2 - 2\langle u, v \rangle + 1}}{\sqrt{(1 - \|u\|^2)(1 - \|v\|^2)}} \right).$$

Since $\langle u, v \rangle = 0$, as $\|u\|, \|v\| \rightarrow 1$, we have $d(u, v) \rightarrow \infty$. Therefore as $R \rightarrow \infty$, the hyperbolic distance between $P^k \setminus B_R(0)$ and $P^{n-k} \setminus B_R(0)$ also goes to infinity. Since every Voronoi cell has bounded diameter, there is an $R_0 > 0$ such that whenever V_α contains a point with hyperbolic distance at least R_0 from 0 and V_α intersects P^{n-k} , V_α does not intersect P^k . The Milnor–Schwarz lemma gives us a quasi-isometry between T and \mathbb{H}^n , so taking the constants $C > 1, K > 0$ from the quasi-isometry we can set $R_1 = CR_0 + K$. □

Consider the ball $B_{R_1}(0)$ in the graph metric. By the proof of the lemma, the union of tiles $\langle B_{R_1}(0) \rangle \cap \overline{(P^k)^c}$ contains the $(n - k - 1)$ -sphere Ω of radius R_0 (in the hyperbolic metric) inside P^{n-k} .

Lemma 5.2 *Let $\varepsilon > 0$ be given. There are constants $R_2, R_3 > 0$ such that*

$$N_\varepsilon(S_{R_2}(0)) \subset \langle B_{R_3}(0) \setminus B_{R_1}(0) \rangle$$

where $N_\varepsilon(K)$ denotes the closed ε -neighborhood of K .

Proof The lemma says, more or less, that one can nest the R_1 -ball in the graph metric inside an R_2 -ball in the hyperbolic metric, which in turn is nested inside an R_3 -ball in the graph metric, with some wiggle room to spare (see Fig. 4). This follows easily from the fact that the inclusion $T \hookrightarrow \mathbb{H}^n$ is a quasi-isometry.

Denote $\text{diam}(V_\alpha) = \lambda < \infty$. There are C and K such that for every point $x \in \langle B_{R_1}(0) \rangle$,


$$d(x, 0) \leq CR_1 + K + \lambda.$$

So we can choose $R_2 = CR_1 + K + \lambda + \varepsilon + 1$. Similarly, for every α such that V_α intersects $N_\varepsilon(S_{R_2}(0))$,

$$d_T(\alpha, 0) < R_3 = C(R_2 + \lambda + \varepsilon + 1) + K.$$

□

Fig. 4 The hyperbolic handle S



Now we are ready to finish our construction of S . Let J denote the union of all Voronoi cells contained inside the open R_2 -ball around 0, excluding those Voronoi cells which also intersect P^k . Then, we define

$$S = \text{closure}(\langle \mathcal{B}_{R_3}(0) \rangle \setminus J).$$

This set splits into a “spherical shell” outside the hyperbolic R_2 -ball and a “handle” inside. Figure 4 gives an illustration of this construction. The reason we take the closure at the end is because when we remove J , we remove some parts of boundaries of Voronoi cells in $\langle \mathcal{B}_{R_3}(0) \rangle$ as well. This way, S corresponds precisely to $\langle S \rangle$ for some subgraph S of T . This subgraph S satisfies the conditions in Theorem 3.6:

Lemma 5.3 *The subgraph S is connected, is contained in a ball \mathcal{B}_R centered at a vertex of T , and contains the sphere of the same radius and center.*

Proof For the latter two conditions it suffices to show that $J \cap \langle \partial \mathcal{B}_{R_3}(0) \rangle$ is empty, but this is guaranteed by the construction in Lemma 5.2.

For connectedness, we use the same transversality approach from 4.10. First, note that $\langle \mathcal{B}_{R_3}(0) \rangle \setminus J$ is path-connected, because every point x in this set falls into two categories: either it is a point contained in a Voronoi cell intersecting $S_{R_2} \cup P^k$, or it is a point inside $\langle \mathcal{B}_{R_3}(0) \rangle$ outside of B_{R_2} . In the latter case, taking a path in the graph connecting the vertex corresponding to the Voronoi cell containing x and the origin, we see this path must contain some vertex corresponding to a Voronoi cell intersecting S_{R_2} (because this path can be translated to a piecewise geodesic path in hyperbolic space from outside B_{R_2} to inside, which must intersect S_{R_2}).

Therefore, any point in the second category is connected to a point in the first category by a path, hence it suffices to show points in the first category are connected to each other by paths. For two such points x and y , we first connect them to points

on $S_{R_2} \cup P^k$ by paths in the respective Voronoi cells, and then observe that $S_{R_2} \cup P^k$ is path-connected.

Now, given two centers of Voronoi cells in S (which are necessarily also contained in $\langle \mathcal{B}_{R_3}(0) \setminus J \rangle$), we can take a path between them in $\langle \mathcal{B}_{R_3}(0) \setminus J \rangle$. Since J is closed this path has a minimum distance from J . Using the argument in 4.10, we can perturb this path so that it remains in $\langle \mathcal{B}_{R_3}(0) \setminus J \rangle$, starts and ends in the same Voronoi cells, and avoids all intersections of dimension $\leq n - 2$. Translating this to a path in the graph we conclude that \mathcal{S} is connected. \square

Now we note that the action of an element of Γ induces both an automorphism of T and a homeomorphism of \mathbb{H}^n which takes tiles to tiles. In particular, for every point $x_\alpha \in \mathcal{O}_x$, there is an element $g \in \Gamma$ such that $g(0) = x_\alpha$. Then the action of g takes S to a copy S_α centered at x_α , and it takes the subgraph \mathcal{S} of T to an isomorphic subgraph centered at α . Applying Theorem 3.6, or rather the variant in Remark 3.7, we get:

Corollary 5.4 *There is a constant $C = C(n, \Gamma) > 0$ such that with high probability, there are at least $C|A(t)|$ disjoint R -balls \mathcal{B}_j in the graph metric such that $\langle \mathcal{B}_j \rangle \cap A(t)$ is homeomorphic to S .*

Proof Let $\mathcal{A}(t)$ be the subgraph of T corresponding to $A(t)$. By Remark 3.7, there are $C_1(R_3, n)F(|A(t)|)$ disjoint R -balls \mathcal{B}_j such that $\mathcal{A}(t) \cap \mathcal{B}_j = g \cdot S$, where F is the isoperimetric profile of Γ . The statement of the corollary follows from the fact that F is linear. \square

Homology of the hyperbolic Eden model

To finish the proof of Theorem 4.11, it remains to show that each S adds at least 1 to $\beta_k(t)$.

We have shown that at time t , we can find $C|A(t)|$ homeomorphic copies S_α of the hyperbolic handle, centered at points $\{x_\alpha\}_{\alpha \in J}$. For each such handle S_α , denote $N_\alpha^+ = S_\alpha \setminus B_{R_2-\varepsilon}(z_\alpha)$, where R_2 is the radius of the hyperbolic sphere $S_{R_2}(0)$ chosen in the course of constructing S . Similarly, denote $N_\alpha^- = S_\alpha \cap B_{R_2+\varepsilon}(z_\alpha)$. The overlapping subsets N^+ and N^- of S , whose shifted copies are the sets defined here, are marked in Fig. 4. Now define

$$M = A(t) \setminus \coprod_{\alpha \in J} B_{R_2-\varepsilon}(z_\alpha) \quad N = \coprod_{\alpha \in J} N_\alpha^-.$$

Intuitively, M is the part of $A(t)$ that does not include the handles, plus all disjoint copies of N^+ , and N is just disjoint copies of N^- . Then

$$A(t) = \text{int } M \cup \text{int } N \quad \text{and} \quad M \cap N = \coprod_{\alpha} N_\varepsilon(S^{n-1}) \simeq \coprod_{\alpha} S^{n-1}.$$

Therefore, since $1 \leq k \leq n - 2$, we have that $H_k(M \cap N) \cong 0$.

Write A for $A(t)$ to simplify notation. The Mayer–Vietoris theorem gives an exact sequence

$$\dots \rightarrow H_k(M \cap N) \rightarrow H_k(M) \oplus H_k(N) \xrightarrow{\psi} H_k(A) \rightarrow H_{k-1}(M \cap N) \rightarrow \dots,$$

and since $H_k(M \cap N) \cong 0$, the map ψ is injective. Since $N = \coprod_{\alpha} (N_{\alpha}^-)$, we have that $H_k(N) \cong \bigoplus_{\alpha} H_k(N_{\alpha}^-)$. Therefore, if we can show that for each α , the homology group $H_k(N_{\alpha}^-)$ has rank at least 1, this completes the proof of Theorem 4.11. Note that all the N_{α}^- are isometric, so we just need to understand the homology of the set $N^- \subseteq S$.

Note that N^- is contained in $B_{R_2+\varepsilon}(0) \setminus \Omega$ because $\overline{P^k}$ is disjoint from Ω and $\langle \mathcal{B}_{R_1}(0) \rangle$ contains Ω . On the other hand, by definition, N^- contains $S_{R_2+\varepsilon} \cup (B_{R_2+\varepsilon}(0) \cap P^k)$.

Now, note that $B_{R_2+\varepsilon}(0) \setminus \Omega$ deformation retracts to $S_{R_2+\varepsilon} \cup (B_{R_2+\varepsilon}(0) \cap P^k)$. (A low-dimensional case is the deformation retraction of a solid ball minus a circle lying on a plane deformation retracting to a sphere union a diameter perpendicular to that plane. In general as in that case, each point moves out along the line linking it to the closest point on Ω .) This deformation retraction restricts to a retraction from N^- to $S_{R_2+\varepsilon} \cup (B_{R_2+\varepsilon}(0) \cap P^k)$, and so the inclusion map from $S_{R_2+\varepsilon} \cup (B_{R_2+\varepsilon}(0) \cap P^k)$ to N^- induces an injection on homology. Since our aim is to show that $H_k(N_{\alpha}^-)$ has rank at least 1, it suffices to do so for $H_k(S_{R_2+\varepsilon} \cup (B_{R_2+\varepsilon}(0) \cap P^k))$.

We apply Mayer–Vietoris once more. For some sufficiently small $\varepsilon' > 0$, let V be an open ε' -neighborhood of $S_{R_2+\varepsilon}$ in $S_{R_2+\varepsilon} \cup (B_{R_2+\varepsilon}(0) \cap P^k)$, and let W be an open ε' -neighborhood of $B_{R_2+\varepsilon}(0) \cap P^k$ in $S_{R_2+\varepsilon} \cup (B_{R_2+\varepsilon}(0) \cap P^k)$. Then V, W are open in $S_{R_2+\varepsilon} \cup (B_{R_2+\varepsilon}(0) \cap P^k) = V \cup W$, so we may apply Mayer–Vietoris.

We see that V deformation retracts to $S_{R_2+\varepsilon} \cong S^{n-1}$, W deformation retracts to a point, and $V \cap W$ deformation retracts to S^{k-1} . We have the Mayer–Vietoris exact sequence

$$\begin{aligned} \dots \rightarrow H_k(V) \oplus H_k(W) &\rightarrow H_k\left(S_{R_2+\varepsilon} \cup (B_{R_2+\varepsilon}(0) \cap P^k)\right) \\ &\rightarrow H_{k-1}(V \cap W) \rightarrow H_{k-1}(V) \oplus H_{k-1}(W) \rightarrow \dots \end{aligned}$$

We see that $H_k(V), H_{k-1}(V), H_k(W), H_{k-1}(W)$ are all zero, so we have

$$H_k\left(S_{R_2+\varepsilon} \cup (B_{R_2+\varepsilon}(0) \cap P^k)\right) \cong H_{k-1}(V \cap W) \cong H_{k-1}(S^{k-1}) \cong \mathbb{Z},$$

completing the proof of Theorem 4.11 for $1 \leq k \leq n - 2$.

Finally, we address the $k = n - 1$ case. Here we can simply take our S to be the union of tiles touching some sufficiently large closed ball B_R , minus the interior of the tile containing the center of the ball. Then the corresponding subgraph S is connected and satisfies the other conditions of 3.6 by a similar argument as above, hence we can find sufficiently many copies of it. Via the same Mayer–Vietoris argument it suffices to show that each copy of S has $n - 1$ th homology with rank at least 1, but this follows

from the fact that S retracts to S^{n-1} , hence there is a surjection from the $n - 1$ th homology of S to $H^{n-1}(S^{n-1}) \cong \mathbb{Z}$.

Remark 5.5 Note that throughout this proof we have used the fact that balls and spheres in \mathbb{R}^n centered at the origin and contained in the unit disk are exactly the balls and spheres centered at the origin in the disk model for hyperbolic space.

6 Further generalizations

We have shown that in \mathbb{H}^n , the Eden model associated with a Voronoi tessellation satisfies Theorem 4.11. In this section we generalize this theorem to a broader class of manifolds and actions on them: those with non-positive sectional curvature.

Let X be a compact connected Riemannian manifold. We endow the universal cover \tilde{X} of X with the Riemannian metric lifted from X , that is, the Riemannian metric that makes the covering map into a local isometry. The set of homeomorphisms (indeed, isometries) $\phi : \tilde{X} \rightarrow \tilde{X}$ such that the diagram

$$\begin{array}{ccc} \tilde{X} & \xrightarrow{\phi} & \tilde{X} \\ & \searrow p & \swarrow p \\ & X & \end{array}$$

commutes forms a group under composition. This group is called the **deck group** of \tilde{X} and is isomorphic to the fundamental group $\pi_1(X)$. This gives an action of $\pi_1(X)$ by isometries on \tilde{X} .

We want to generalize the Eden model to a tessellation of \tilde{X} whose symmetries are given by the action of $\pi_1(X)$. In the case that X is the flat torus T^n and $\tilde{X} = \mathbb{R}^n$, this reconstructs the usual Eden model; in the case that X is a hyperbolic manifold, this gives some of the hyperbolic examples described above. In other cases, we would like to divide \tilde{X} into similar “grid cubes” or “tiles”.

Voronoi cells and fundamental domains

Often, one defines a fundamental domain in the universal cover \tilde{X} to be a subset which contains exactly one point from each orbit of $\pi_1(X)$. We modify this slightly: for us, a **(closed) fundamental domain** is a closed set D whose translates under the action of $\pi_1(X)$ cover \tilde{X} and such that for any $\phi \in \pi_1(X)$,

$$D \cap (\phi \cdot D) \subseteq \partial D.$$

The closure of a sufficiently nice fundamental domain in the usual sense fits this definition. For example, in Euclidean space \mathbb{R}^n thought of as the universal cover of a flat torus, a closed unit cube is a fundamental domain in our sense. An example of a fundamental domain in the usual sense would be $[0, 1)^n$.

A fundamental domain on \tilde{X} can be obtained from the following construction. Pick a point $x \in \tilde{X}$, and consider its orbit \mathcal{O}_x under the action of $\pi_1(X)$. An element in \mathcal{O}_x is of the form $\phi \cdot x$ for some $\phi \in \pi_1(X)$. As before, the **Voronoi cell** of an element $x_\alpha \in \mathcal{O}_x$ is the set

$$V_\alpha = \{y \in \tilde{X} : \forall x_\beta \in \mathcal{O}_x \ d(y, x_\alpha) \leq d(y, x_\beta)\}.$$

Clearly, the Voronoi cells associated to \mathcal{O}_x are translates of each other under the action of $\pi_1(X)$. In addition:

Lemma 6.1 *For any $x \in \tilde{X}$, Voronoi cells associated with \mathcal{O}_x are fundamental domains in \tilde{X} .*

This is proved in exactly the same way as Lemma 4.3.

Lemma 6.2 *Voronoi cells in \tilde{X} have bounded diameter: $\text{diam}(V_\alpha) < \infty$. In particular, since all Voronoi cells are isometric, they must all have the same finite diameter.*

Proof Consider Voronoi cells $\{V_\alpha\}$ centered at points of the orbit \mathcal{O}_x . Since X is compact, it has finite diameter. We will show that $\text{diam}(V_\alpha) \leq 2 \text{diam } X$.

Take a point $y \in \tilde{X}$. Let $y' = p(y)$ and $x' = p(x)$, where $p : \tilde{X} \rightarrow X$ is the covering map. Connect y' to x' by a minimal geodesic γ' in X . There is a unique lift $\gamma : [0, 1] \rightarrow \tilde{X}$ with $\gamma(0) = y$ and $\gamma(1) = x_\alpha$ for some $x_\alpha \in \mathcal{O}_x$. Therefore the distance from y to the nearest point of \mathcal{O}_x is bounded above by $d(x', y')$.

It follows that for any y and z located in the same Voronoi cell V_α , there is a path from y to z through x_α of length at most $2 \text{diam } X$. □

The Eden model on \tilde{X}

We generalize the Eden model to this setting the same way we did for hyperbolic tessellations: pick a Voronoi cell to start with; then at every time step, randomly choose a Voronoi cell adjacent to the infected region to infect next. So that our model generalizes the Euclidean and hyperbolic ones, we say that two Voronoi cells are **adjacent** if their intersection is $(n - 1)$ -dimensional in the sense of Lebesgue covering dimension. We will use $A(t)$ to denote the Eden model at time t .

As before, we view each Voronoi cell V_α as the vertex α of a graph T , with an edge from α to β in T if and only if

$$\text{dim}(V_\alpha \cap V_\beta) \geq n - 1.$$

Lemma 6.3 *The graph T is a Cayley graph for $\pi_1(X)$.*

In particular, this means that T is infinite, locally finite, connected, and vertex-transitive.

Proof The proof is similar to that of Lemma 4.10. Since $\pi_1(X)$ acts freely and transitively on T by graph isomorphisms, by 2.2 it suffices to prove that T is locally finite and connected. The proof of local finiteness is the same as for Lemma 4.10.

It remains to show that T is connected. Let \mathcal{U} be a nonempty, non-full subset of the vertices of T . We will show that there is at least one edge between \mathcal{U} and \mathcal{U}^c . Denote the union of fundamental domains corresponding to vertices in a set \mathcal{A} by $\langle \mathcal{A} \rangle$. Then

$$\tilde{X} = \langle \mathcal{U} \rangle \cup \langle \mathcal{U}^c \rangle$$

and $\langle \mathcal{U} \rangle$ and $\langle \mathcal{U}^c \rangle$ both have nonempty interior. Choose points $y \in \text{int}\langle \mathcal{U} \rangle$ and $z \in \text{int}\langle \mathcal{U}^c \rangle$ and let V be a neighborhood homeomorphic to \mathbb{R}^n of a simple curve between y and z . By (Engelking (1978), Theorem 1.8.12),

$$\dim(V \cap \langle \mathcal{U} \rangle \cap \langle \mathcal{U}^c \rangle) \geq n - 1.$$

On the other hand,

$$\langle \mathcal{U} \rangle \cap \langle \mathcal{U}^c \rangle = \bigcup_{\alpha \in \mathcal{U}, \beta \in \mathcal{U}^c} V_\alpha \cap V_\beta.$$

Since this is a countable union, and indeed a finite union if we restrict to any compact subset of \tilde{X} , by (Munkres (2000), Corollary 50.3), for some α and β , $V_\alpha \cap V_\beta$ is at least $(n - 1)$ -dimensional. \square

Now suppose that X has non-positive (but not necessarily constant) sectional curvature, $K \leq 0$. Among other nice properties, this condition guarantees that every pair of points $x, y \in \tilde{X}$ is connected by a unique geodesic. We will show the following generalization of Theorem 4.11:

Theorem 6.4 *Let X be a compact connected manifold of non-positive sectional curvature. Let $A(t)$ be the Eden model at time t on a tessellation of \tilde{X} by Voronoi cells of an orbit $\pi_1(X) \cdot x$. Let $0 \leq k \leq n - 1$ and let $\beta_k(t)$ denote the k th Betti number of $A(t)$, i.e. $\beta_k(t) = \text{rank } H_k(A(t))$. Then there is a constant $C = C(X, k) > 0$ such that*

$$\beta_k(t) \geq CF(|A(t)|)$$

with high probability as $t \rightarrow \infty$, where $F : \mathbb{N} \rightarrow \mathbb{N}$ is the isoperimetric profile of $\pi_1(X)$.

Here the isoperimetric profile of $\pi_1(X)$ is the isoperimetric profile of its Cayley graph for any finite generating set. Since all such Cayley graphs are quasi-isometric, this is well-defined up to a multiplicative constant, as shown in Theorem A.4. In the case that X has not just non-positive but *negative* sectional curvature, the fundamental group $\pi_1(X)$ is Gromov hyperbolic and therefore non-amenable (see Example 3.5(c)). It follows that we again have a linear lower bound.

The proof of this theorem is very similar to that of Theorem 4.11. In fact, the only steps at which we use hyperbolicity are the construction of P^k and P^{n-k} and the




Fig. 5 Illustration of the map \exp_p

proof of Lemma 5.1, which use the Poincaré disc model explicitly. We replace this machinery by using the exponential map at a point $x \in \tilde{X}$.

The Exponential Map

Let M be a complete Riemannian manifold and let $p \in M$. Given a tangent vector $v \in T_p M$, by the Hopf–Rinow theorem, there is a unique geodesic γ_v in M with $\gamma_v(0) = p$ and $\gamma'_v(0) = v$. Then the **exponential map** $\exp_p : T_p M \rightarrow M$ is defined by

$$\exp_p(v) = \gamma_v(1).$$

If M is non-positively curved and simply connected, then the Cartan–Hadamard theorem states that since there is a unique geodesic between p and any other point, the exponential map is a diffeomorphism. Moreover, the derivative of the exponential map at the origin of the tangent space is the identity: $d(\exp_p)_0 = \text{id}$.

Alternatively, we may reparametrize all of our geodesics with unit speed. Let η_v be the unit speed parametrization of γ_v . Since the geodesic equation $\nabla_{\dot{\gamma}_v} \dot{\gamma}_v = 0$ implies that γ_v is a curve of constant speed, we see that $\gamma_v(1) = \eta_v(|v|)$ (Fig. 5). Therefore, we can equivalently think of the exponential map as the map $\exp_p : T_p M \rightarrow M$ defined by

$$\exp_p(v) = \eta_v(|v|).$$

Now take M to be \tilde{X} and the point p to be the center x of one of our Voronoi cells. By construction, \tilde{X} is simply connected, complete, and has non-positive sectional curvature. Choose two orthogonal vector subspaces H^k and $H^{n-k} \subseteq T_x \tilde{X}$ of the corresponding dimensions. Define submanifolds $P^k = \exp_p(H^k)$ and $P^{n-k} = \exp_p(H^{n-k})$ in \tilde{X} . Since the geodesic $\eta_v(t)$ only depends on the direction of the tangent vector v , the geodesic ray from the origin in \mathbb{R}^n in the direction $v/|v|$ will be mapped onto the geodesic η_v in \tilde{X} . As in the hyperbolic construction, let $\overline{P^k}$ be the




Fig. 6 Illustration of the proof of Corollary 6.6

union of all Voronoi cells that intersect P^k and $\overline{P^{n-k}}$ be the union of all Voronoi cells that intersect P^{n-k} but not P^k .

We state the following result (which is a consequence of the Rauch comparison theorem) without proof:

Lemma 6.5 (Law of cosines). *Let (M, g) be a complete simply connected Riemannian manifold with non-positive sectional curvature, and consider a geodesic triangle in M whose side lengths are $a, b,$ and c with opposite angles $A, B,$ and C respectively. Then we have*

$$a^2 + b^2 - 2ab \cos(C) \leq c^2.$$

This gives the following result:

Corollary 6.6 *There exists $R_1 > 0$ such that for all vertices α of T satisfying $d_T(\alpha, 0) \geq R_1,$ if the Voronoi cell V_α intersects $P^{n-k},$ then $V_\alpha \in \overline{P^{n-k}}.$ Here 0 is the Voronoi cell centered at $x.$*

Proof Choose $R_0 = \text{diam}(V_\alpha)/\sqrt{2}.$ Choose $y \in P^k$ and $z \in P^{n-k}$ such that $d(x, y)$ and $d(x, z)$ are both greater than $R_0.$ Then consider the geodesic triangle spanned by $x, y,$ and $z.$ Since the exponential map takes lines through the origin to geodesics, the unique geodesic from x to y is the image of a straight line segment contained in $H^k.$ Similarly, the geodesic from x to z is the image of a straight line segment in $H^{n-k}.$ Moreover, since $d(\exp_p)_0 = \text{id},$ it preserves the angle between geodesics through $x.$ Hence our triangle has a right angle at $x,$ and by the law of cosines,

$$d(x, y)^2 \geq d(x, p)^2 + d(y, p)^2 > 2R^2 = \text{diam}(V_\alpha)^2.$$

Therefore, if $d(x, x_\alpha) > R_0$ and V_α intersects $P^{n-k},$ then $V_\alpha \in \overline{P^{n-k}}.$ Since $T \rightarrow \tilde{X}$ is a quasi-isometric embedding by the Milnor–Schwarz lemma, there is an R_1 such that if $d(0, \alpha) > R_1,$ then $d(x, x_\alpha) > R_0$ (Fig. 6). This completes the proof. \square

We have now constructed the first layer of the shell $\mathcal{B}_{R_1}(p)$ in Fig. 4, a handle centered at $x.$ The rest of the construction and proof works the same way it did in the hyperbolic case.

Remark 6.7 Note that the image of a ray from the origin in the tangent space under \exp_p is the unique geodesic emanating from p in that direction, and furthermore that when restricted to this ray \exp_p is an isometry. So, we see that every ball (set of points distance $\leq R$ from p) or sphere (set of points exactly distance R in the p) centered at p in \tilde{X} is the image under \exp_p of a corresponding ball/sphere centered at the origin in the tangent space. Therefore, the union of a sphere in \tilde{X} and the part of the hyperplane $\exp_p(H^k)$ contained in the sphere is just the image under \exp_p of the union of a sphere and a hyperplane portion in the tangent space. Since \exp_p is a diffeomorphism, the two have the same homology, and so the proof of the hyperbolic case works here as well.

Discussion

There is no obvious obstruction to generalizing the theorem further to manifolds which don't have non-positive curvature. While our particular construction of handles with the required topological properties does not generalize, there doesn't seem to be a good reason why such handles cannot be constructed. We therefore conjecture that Theorem 6.4 holds for any manifold X .

Another notion worth discussing is the nontriviality of our results. For example, if every Voronoi cell V_α had a k -handle of its own, then the inequality $\beta_k(E) \geq CF(|E|)$ would hold for every union of Voronoi cells. In that case, Theorem 6.4 would not distinguish the Eden model from an arbitrary set of cells.

For hyperbolic tessellations, we know that this is never the case because Voronoi cells are convex. In fact, since all intersections of Voronoi cells are convex and therefore contractible, a union of such cells is homotopy equivalent to its Čech complex. In other words, the homology of such a union is determined by the lattice of intersections.

Nevertheless, it seems reasonable to suggest that for any manifold X , one can find fundamental domains in the universal cover that have this property:

Conjecture 6.8 *If X is an aspherical closed manifold, then the action of $\pi_1(X)$ on \tilde{X} admits a closed fundamental domain D such that any nontrivial intersection of translates of D under the $\pi_1(X)$ -action is contractible.*

7 The relationship between FPP and Eden time

Consider the exponential-time FPP model $A(t)$. As previously discussed, it is equivalent to the Eden model up to time rescaling. In this section, we show (informally speaking) that as the Eden time $|A(t)|$ approaches infinity, so does the FPP time t . Therefore, when we show a statement holds with high probability for the geometric and topological behavior of the FPP model, this is also the case for the Eden model.

Lemma 7.1 *The size $|A(t)|$ of the Eden Model on a regular tree of vertex degree n at (FPP) time t satisfies the recurrence relation*

$$\mathbb{E}(|A(t)|) = 2 - e^{-t} + (n - 1) \int_0^t e^{-s} \mathbb{E}(|A(t - s)|) ds.$$

Proof We can view an n -ary tree as a root vertex with n disjoint subtrees coming out of it. Each of these subtrees is an n -ary tree minus one of its subtrees.

In order for a subtree to contribute to $|A(t)|$, the root vertex of that subtree must be infected within t seconds. Therefore, the expected contribution from one such subtree is

$$\int_0^t e^{-s} \left(1 + \frac{n-1}{n} \mathbb{E}(|A(t-s)| - 1) \right) ds.$$

Here, e^{-s} is the probability density function for the passage time of the root of the subtree, and $1 + \frac{n-1}{n} \mathbb{E}(|A(t-s)| - 1)$ is the expected size of the subtree once the root of the subtree has been infected: it consists of the root plus $n - 1$ additional subtrees.

Since there are n such subtrees, we conclude that

$$\begin{aligned} \mathbb{E}(|A(t)|) &= 1 + n \int_0^t e^{-s} \left(1 + \frac{n-1}{n} \mathbb{E}(|A(t-s)| - 1) \right) ds \\ &= 1 + n \int_0^t e^{-s} \left(\frac{1}{n} + \frac{n-1}{n} \mathbb{E}(|A(t-s)|) \right) ds \\ &= 2 - e^{-t} + (n-1) \int_0^t e^{-s} \mathbb{E}(|A(t-s)|) ds. \end{aligned}$$

□

Lemma 7.2 *Let $A(t)$ be the Eden model at time t growing on an n -ary regular tree. There exists a constant $C > 0$ such that*

$$\mathbb{E}(|A(t)|) < C e^{(n-1)t}$$

for all t .

Proof Let C be a constant such that

$$\mathbb{E}(|A(t)|) < C e^{(n-1)t}$$

for all $t \in [0, n]$. Suppose, for contradiction, that there is a t such that $\mathbb{E}(|A(t)|) \geq C e^{(n-1)t}$. Since both sides are continuous in t , there exists some t_0 such that $\mathbb{E}(|A(t_0)|) = C e^{(n-1)t_0}$, and $\mathbb{E}(|A(t)|) < C e^{(n-1)t}$ for all $t < t_0$. Note that we must have $t_0 > n$.

Then by Lemma 7.1,

$$\begin{aligned} \mathbb{E}(|A(t_0)|) &= 2 - e^{-t} + (n - 1) \int_0^{t_0} e^{-s} \mathbb{E}(|A(t_0 - s)|) \, ds \\ &< 2 + (n - 1) \int_0^{t_0} e^{-s} C e^{(n-1)(t_0-s)} \, ds \\ &= 2 + C(n - 1)e^{(n-1)t_0} \int_0^{t_0} e^{-ns} \, ds \\ &= 2 + C(n - 1)e^{(n-1)t_0} \left(\frac{1}{n} - \frac{1}{n} e^{-nt_0} \right) \\ &< 2 + C(n - 1)e^{(n-1)t_0} \cdot \frac{1}{n} \\ &< C e^{(n-1)t_0} \\ &= \mathbb{E}(|A(t_0)|) \end{aligned}$$

where the inequality in the second-to-last line follows from the fact that

$$2 = \frac{1}{n} \cdot 2n < \frac{1}{n} e^n < \frac{1}{n} e^{t_0} \leq \frac{1}{n} e^{(n-1)t_0}.$$

This is a contradiction, therefore such a t_0 cannot exist. Therefore, we have

$$\mathbb{E}(|A(t)|) < C e^{(n-1)t}$$

for all t . □

Lemma 7.3 *Let $A(t)$ be the Eden model on an n -ary regular tree T . Let t denote the FPP time and let $|A(t)|$ denote the size of the Eden model at (FPP) time t . Then there is a constant $D' > 0$ such that*

$$|A(t)| \leq (D')^t$$

with high probability as $t \rightarrow \infty$.

Proof The perimeter of a ball of radius k in T has size n^k . To obtain an upper bound on the size of the subtree infected at time t , we compare our instance of the Eden model to one in which everything within a ball of radius $\lfloor \log_n t \rfloor$ immediately gets infected (everything within this ball has passage time 0) and all other passage times are unchanged. Clearly, this modified model has a strictly larger infected region at time t .

On the perimeter of the $\lfloor \log_n t \rfloor$ -ball, there are $n^{\lfloor \log_n t \rfloor} < t$ disjoint subtrees coming out, each with its root on the perimeter. Each of these subtrees is a regular n -ary tree minus one of its n primary subtrees (subtrees adjacent to the central vertex). By Lemma 7.2, the expected size of an n -ary tree at time t is bounded above by CD^t , where $D = e^{(n-1)} > 0$. Since each of the subtrees is a n -ary tree minus one of its main subtrees, its expected size at time t is also certainly bounded above by CD^t .

Then as $t \rightarrow \infty$, by the law of large numbers, the total size of the $n^{\lfloor \log_n t \rfloor}$ subtrees will converge to $n^{\lfloor \log_n t \rfloor} E < tE < tD^t$, where E is the expected size of each subtree. Thus as $t \rightarrow \infty$, with high probability we have

$$|A(t)| \leq \text{Vol}(B_{\lfloor \log_n t \rfloor}) + 2n^{\lfloor \log_n t \rfloor} E < \text{Vol}(B_{\lfloor \log_n t \rfloor}) + 2CtD^t.$$

Now

$$\text{Vol}(B_{\lfloor \log_n t \rfloor}) = \sum_{k=0}^{\lfloor \log_n t \rfloor} n^k = \frac{n^{\lfloor \log_n t \rfloor} - 1}{n - 1} < \frac{t - 1}{n - 1} < t.$$

Therefore, we get that

$$|A(t)| < t + 2CtD^t \leq t + (2CD)^t < 2^t + (2CD)^t \leq (2 + 2CD)^t.$$

Setting $D' = 2 + 2CD$ concludes the proof. \square

Proposition 7.4 *Let t denote the FPP time and let $|A(t)|$ denote the size of the Eden model at (FPP) time t on an infinite, locally finite, vertex transitive graph G . Then there exists some constant $D > 0$ such that*

$$|A(t)| \leq D^t$$

with high probability as $t \rightarrow \infty$.

Proof Let $d = \deg G$. We use basically the same coupling trick we used in the proof of Lemma 3.10. But this time we couple G with just one $(d - 1)$ -ary tree T , equipped with exponential i.i.d. passage times with mean 1 on its nodes. The origin $o \in G$ of the infection is assigned passage time 0 since it gets infected instantly. We consider a rooted d -ary tree T with its root mapped to O via the natural map $f : T \rightarrow G$. For each site $p \in G$, its passage time ρ_p is defined to be $\rho_{\tilde{p}}$, where \tilde{p} minimizes the passage time to just before \tilde{p} among all vertices in $f^{-1}(p)$.

As before, this defines a coupling between the Eden models on T and G , since the only thing that depends on other vertices is *which* i.i.d. exponentially distributed weight corresponds to a given site of G . By the same argument as in the proof of Lemma 3.13, for every site $p \in G$ infected at time t , there exists a $\tilde{p} \in f^{-1}(p)$ such that the passage time from \tilde{p} to the root of the tree T is at most t .

Therefore, the Eden model $A_G(t)$ on the graph G at time t is in bijection with a subset of the Eden model $A_T(t)$ on the tree at time t . Then Proposition 7.3 implies that with high probability as $t \rightarrow \infty$, we have

$$|A_G(t)| \leq |A_T(t)| \leq D^t$$

for some constant $D > 0$. \square

Corollary 7.5 *As a function of the Eden time s , the FPP time t satisfies $t \geq \log_D s$ with high probability.*

Proof Suppose that $t \leq \log_D s$. Then at FPP time $\log_D s$, we have

$$|A(\log_D s)| \geq t = D^{\log_D s}.$$

But the probability of this converges to 0 as $\log_D s \rightarrow \infty$. Therefore the probability that $t \leq \log_D s$ converges to 0 as $s \rightarrow \infty$. \square

This is the result we needed.

Acknowledgements This paper was written as part of a project for the UCSB REU. We thank Katherine Merkl for many inspiring conversations, Maribel Bueno Cachadina for organizing the REU, and everyone else in the program for being extremely friendly and supportive. The authors were supported in part by NSF DMS-1850663 and a research fellowship from the College of Creative Studies at UCSB.

Declarations

Conflict of interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A: Quasi-isometry invariance of the isoperimetric profile

In this section, we prove that (up to a natural equivalence relation) the isoperimetric profile is an quasi-isometry invariant of infinite, locally finite, vertex-transitive graphs. First, we give some definitions.

Definition A.1 Let T be an infinite, locally finite, vertex transitive graph equipped with the graph metric. We define the **growth function** $N_T: \mathbb{N} \rightarrow \mathbb{N}$ of G as

$$N_T(n) = |B_n(x)|.$$

In other words, the growth function returns the size of a (closed) ball of radius n centered at an arbitrary point in T ; the choice of center x does not matter since G is vertex-transitive.

Lemma 3.8 shows that for any infinite, locally finite, vertex-transitive graph T where each vertex has degree d ,

$$N_T(n) \leq d^{n+1}.$$

We can give some bounds on how much the size of a finite set of vertices of a graph can change under quasi-isometry.

Lemma A.2 *Let T, W be infinite, locally finite, vertex-transitive graphs equipped with the graph metric, and $f: T \rightarrow W$ be a (C, K) -quasi-isometry. Let $S \subset T$ be a finite subset. Then*

$$\frac{1}{N_T(CK)}|S| \leq |f(S)| \leq |S|.$$

Proof The inequality $|f(S)| \leq |S|$ is true for any function. For the other inequality, we prove a bound on the size of the preimage $f^{-1}(y)$ of any point $y \in W$.

Suppose $x, x' \in T$ satisfy $f(x) = f(x') = y$. Since f is a quasi-isometry,

$$\frac{1}{C}d_T(x, x') - K \leq d_W(f(x), f(x')) = 0,$$

hence

$$d_T(x, x') \leq CK.$$

It follows that for any point $x \in f^{-1}(y)$, all of $f^{-1}(y)$ is contained in the CK -ball around x . Such a ball contains $N_T(CK)$ points, so we have

$$|f^{-1}(y)| \leq N_T(CK).$$

This means that f sends at most $N_T(CK)$ points in S to the same point in $f(S)$, hence $\frac{1}{N_T(CK)}|S| \leq |f(S)|$, as desired. \square

Now, we define an equivalence relation on functions from the naturals to the positive reals:

Definition A.3 Let f, g be functions from \mathbb{N} to \mathbb{R}^+ . We write $g \lesssim f$ if there is a constant $a > 0$ such that

$$g(n) \leq af(n)$$

for all $n \geq 1$. If we also have $f \lesssim g$, we write $f \sim g$.

One can check that \sim is indeed an equivalence relation.

To show that the isoperimetric profile is a quasi-isometry invariant, we want to be able to take an isoperimetric set (one that has minimal boundary given its size) in one space and produce something close to isoperimetric in another quasi-isometric space. While taking a quasi-isometry can “scatter” the points of the isoperimetric set, the fact that a quasi-isometry is coarsely surjective allows us to overcome this obstacle: we can take the K -ball around the image of an isoperimetric set to obtain something close to being isoperimetric. The following theorem formalizes this:

Theorem A.4 *Let T, W be infinite, locally finite, vertex-transitive graphs equipped with the graph metric. If T, W are quasi-isometric, then we have $F_T \sim F_W$.*

Proof To differentiate between balls in the two graphs, we will write $B_T(x, R)$ and $B_W(y, R)$ for the closed balls of radius R around $x \in T$ and $y \in W$, respectively.

Fix some $n \geq 1$, and let $S_T \subseteq T$ be a subgraph of size n satisfying $|\partial S_T| = F_T(n)$. Let $f: T \rightarrow W$ be a (C, K) -quasi-isometry. We define

$$S_W = \bigcup_{y \in f(S_T)} B_W(y, K).$$

Consider the edge boundary ∂S_W . For each $y \in \partial S_W$, there exists an $x \in T$ such that $d_W(f(x), y) \leq K$ (due to f being a quasi-isometry). For each $y \in \partial S_W$, we choose such an x , and let $X \subseteq T$ denote the collection of these x 's.

It is possible that for distinct points $y, y' \in \partial S_W$, we choose the same corresponding $x \in X$. If this is the case, then both y, y' are contained in the closed K -ball in W centered at $f(x)$. Therefore, the maximum number of points $y \in \partial S_W$ which correspond to the same point $x \in X$ is bounded above by $N_W(K)$ (the size of a closed K -ball in W). It follows that

$$|\partial S_W| \leq N_W(K)|X|. \tag{A.5}$$

For each $x \in X$, $f(x)$ is within distance K of a point in ∂S_W . Each point in ∂S_W is within distance 1 of a point S_W , and each point in S_W is within distance K of some point in $f(S_T)$. Therefore, for each $x \in X$, there exists some $x' \in S_T$ such that

$$d_W(f(x), f(x')) \leq 2K + 1.$$

Since f is a (C, K) -quasi-isometry, it follows that

$$d_T(x, x') \leq C(3K + 1).$$

Therefore, X is contained in the closed $C(3K + 1)$ -neighborhood around S_T ; in other words, we have

$$X \subseteq \bigcup_{x \in S_T} B_T(x, C(3K + 1)).$$

Denote this union by $B_T(S_T, C(3K + 1))$.

On the other hand, we claim that $X \cap S_T$ is empty. Suppose, for a contradiction, that there exists a point $x \in X \cap S_T$. Since x is in X , there exists some $y \in \partial S_W$ such that $d_W(f(x), y) \leq K$. However, since x is also in S_T , and S_W contains all the points in the closed K -neighborhood around $f(S_T)$, it follows that y is in S_W . But this contradicts the assumption that $y \in \partial S_W$.

Since $X \cap S_T$ is empty, we have

$$X \subseteq B_T(S_T, C(3K + 1)) \setminus S_T.$$

For each j , let $\partial^j S_T = \partial B_T(S_T, j - 1)$. Then $\partial^j S_T$ contains all the points whose distance to the nearest point in S_T is exactly j , hence we have

$$B_T(S_T, C(3K + 1)) \setminus S_T = \bigcup_{j=1}^{C(3K+1)} \partial^j S_T.$$

Moreover, every element of $\partial^{j+1} S_T$ is adjacent to $\partial^j S_T$, which implies that $|\partial^{j+1} S_T| \leq \deg(T) |\partial^j S_T|$. Therefore,

$$\begin{aligned} \left| \bigcup_{j=1}^{C(3K+1)} \partial^j S_T \right| &\leq \sum_{j=1}^{C(3K+1)} |\partial^j S_T| \\ &\leq \sum_{j=1}^{C(3K+1)} \deg(T)^{j-1} |\partial S_T| \\ &\leq \deg(T)^{C(3K+1)} F_T(n). \end{aligned}$$

Since X is a subset of $\bigcup_{j=1}^{C(3K+1)} \partial^j S_T$, it follows that

$$|X| \leq \deg(T)^{C(3K+1)} F_T(n) \quad (\text{A.6})$$

as well. Furthermore, since $f(S_T) \subseteq S_W$, it follows from Lemma A.2 that

$$n = |S_T| \leq N_T(CK) |f(S_T)| \leq N_T(CK) |S_W|.$$

Using the fact that F_W is sublinear, it follows that

$$F_W(n) \leq F_W(N_T(CK) |S_W|) \leq N_T(CK) F_W(|S_W|) \leq N_T(CK) |\partial S_W|.$$

Using (A.5) and (A.6), we conclude that

$$F_W(n) \leq N_T(CK) N_W(K) \deg(T)^{C(3K+1)} F_T(n).$$

In particular, $F_W \lesssim F_T$.

Since T and W are quasi-isometric, there also exists a quasi-isometry $g: W \rightarrow T$. The proof above then shows that $F_T \lesssim F_W$, hence $F_T \sim F_W$. \square

References

- Auffinger, A., Gorski, C.: Asymptotic shapes for stationary first passage percolation on virtually nilpotent groups. *Probab. Theory Relat. Fields* **186**(1–2), 285–326 (2023)
- Auffinger, A., Damron, M., Hanson, J.: 50 Years of First-Passage Percolation, University Lecture Series, vol. 68. American Mathematical Society, Providence, RI (2017)

- Beem, J.K.: Pseudo-Riemannian manifolds with totally geodesic bisectors. *Proc. Am. Math. Soc.* **49**, 212–215 (1975)
- Benjamini, I., Schramm, O.: Percolation beyond \mathbf{Z}^d , many questions and a few answers. *Electron. Commun. Probab.* **1**(8), 71–82 (1996)
- Benjamini, I., Tessler, R.: First passage percolation on nilpotent Cayley graphs. *Electron. J. Probab.* **20**(99), 20 (2015)
- Benjamini, I., Tessler, R.: First passage percolation on a hyperbolic graph admits bi-infinite geodesics. *Electron. Commun. Probab.* **22**, 14, 8 (2017)
- Clay, M., Margalit, D. (eds.): *Office Hours with a Geometric Group Theorist*. Princeton University Press, Princeton (2017)
- Cox, J.Th., Durrett, R.: Some limit theorems for percolation processes with necessary and sufficient conditions. *Ann. Probab.* **9**(4), 583–603 (1981)
- Damron, M., Gold, J., Lam, W.-K., Shen, X.: On the number and size of holes in the growing ball of first-passage percolation, [arXiv:2205.09733](https://arxiv.org/abs/2205.09733) [math.PR] (2022)
- Damron, M., Hanson, J., Lam, W.-K.: The size of the boundary in first-passage percolation. *Ann. Appl. Probab.* **28**(5), 3184–3214 (2018)
- Diestel, R., Leader, I.: A conjecture concerning a limit of non-Cayley graphs. *J. Algebraic Combin.* **14**(1), 17–25 (2001)
- Druţu, C., Kapovich, M.: *Geometric group theory*, American Mathematical Society Colloquium Publications, vol. 63, American Mathematical Society, Providence, RI, With an appendix by Bogdan Nica (2018)
- Engelking, R.: *Dimension theory*, North-Holland Mathematical Library, vol. 19, North-Holland Publishing Co., Amsterdam-Oxford-New York; PWN—Polish Scientific Publishers, Warsaw, (1978), Translated from the Polish and revised by the author
- Erschler, A.: On isoperimetric profiles of finitely generated groups. *Geom. Dedicata* **100**, 157–171 (2003)
- Eskin, A., Fisher, D., Whyte, K.: Coarse differentiation of quasi-isometries I: Spaces not quasi-isometric to Cayley graphs. *Ann. Math. (2)* **176**(1), 221–260 (2012)
- Følner, E.: On groups with full Banach mean value. *Math. Scand.* **3**, 243–254 (1955)
- Gorski, C.: Strict monotonicity for first passage percolation on graphs of polynomial growth and quasi-trees, [arXiv:2208.13922](https://arxiv.org/abs/2208.13922) [math.PR] (2022)
- Hatcher, A.: *Algebraic Topology*. Cambridge University Press, Cambridge (2002)
- Kosinski, A.A.: *Differential Manifolds*, Pure and Applied Mathematics, vol. 138. Academic Press Inc, Boston (1993)
- Löh, C.: *Geometric Group Theory*. Universitext, Springer, Cham (2017). (**An introduction**)
- Manin, F., Roldán, E., Schweinhart, B.: Topology and local geometry of the Eden model. *Discrete Comput. Geom.* **69**(3), 771–799 (2023)
- Munkres, J.R.: *Topology*, 2nd edn. Prentice Hall Inc, Upper Saddle River (2000)
- Richardson, D.: Random growth in a tessellation. *Proc. Camb. Philos. Soc.* **74**, 515–528 (1973)
- Thom, R.: Quelques propriétés globales des variétés différentiables. *Comment. Math. Helv.* **28**, 17–86 (1954)