

Cohomology fractals, Cannon–Thurston maps, and the geodesic flow

David Bachman, Matthias Goerner, Saul Schleimer & Henry Segerman

To cite this article: David Bachman, Matthias Goerner, Saul Schleimer & Henry Segerman (2022) Cohomology fractals, Cannon–Thurston maps, and the geodesic flow, *Experimental Mathematics*, 31:4, 1047–1085, DOI: 10.1080/10586458.2021.1994059

To link to this article: <https://doi.org/10.1080/10586458.2021.1994059>



© 2022 The Author(s). Published with
license by Taylor and Francis Group, LLC



Published online: 15 Sep 2022.



Submit your article to this journal



Article views: 793

[View related articles](#) View Crossmark data 

Cohomology fractals, Cannon–Thurston maps, and the geodesic flow

David Bachman^a, Matthias Goerner^b, Saul Schleimer^c, and Henry Segerman^d

^aPitzer College, 1050 N. Mills Ave. Claremont, CA 91711, USA; ^bPixar Animation Studios, 1200 Park Avenue, Emeryville, CA 94608, USA; ^cMathematics Institute, University of Warwick, Coventry CV4 7AL, United Kingdom; ^dDepartment of Mathematics, Oklahoma State University, Stillwater, OK 74078, USA

ABSTRACT

Cohomology fractals are images naturally associated to cohomology classes in hyperbolic three-manifolds. We generate these images for cusped, incomplete, and closed hyperbolic three-manifolds in real-time by ray-tracing to a fixed visual radius. We discovered cohomology fractals while attempting to illustrate Cannon–Thurston maps without using vector graphics; we prove a correspondence between these two, when the cohomology class is dual to a fibration. This allows us to verify our implementations by comparing our images of cohomology fractals to existing pictures of Cannon–Thurston maps.

In a sequence of experiments, we explore the limiting behaviour of cohomology fractals as the visual radius increases. Motivated by these experiments, we prove that the values of the cohomology fractals are normally distributed, but with diverging standard deviations. In fact, the cohomology fractals do not converge to a function in the limit. Instead, we show that the limit is a distribution on the sphere at infinity, only depending on the manifold and cohomology class.

1. Introduction

Cannon and Thurston discovered that Peano curves arise naturally in hyperbolic geometry [13]. They proved that for every closed hyperbolic three-manifold, equipped with a fibration over the circle, there is a map from the circle to the sphere that is continuous, finite to one, and surjective. Furthermore this *Cannon–Thurston map* is equivariant with respect to the action of the fundamental group. We review their construction in Section 3; Figure 1(a) shows an approximation.

In a previous expository paper [4], we introduced *cohomology fractals*; these are images arising from a hyperbolic three-manifold M equipped with a cohomology class $[\omega] \in H^1(M; \mathbb{R})$. See Figure 1(b). In that paper we gave an overview of the construction; we also discussed some of the features of the three-manifold and cohomology class that can be seen in its cohomology fractal. We have also written an open-source [5] real-time web application for exploring these fractals. This is available at https://henryseg.github.io/cohomology_fractals/.

In the present work we give rigorous definitions of cohomology fractals, we relate them to Cannon–Thurston maps (see Figure 1(b)), we give technical details of our implementation, and we discuss their limiting behaviour.

We now outline the contents of each section of the paper. Note that we include a glossary of notation in Appendix A. We begin by reviewing the definitions of ideal and material triangulations, and their hyperbolic geometry in Section 2. In Section 3 we define Cannon–Thurston maps. In Section 4 we discuss the differences between vector and raster graphics. We also recall a vector graphics algorithm (Algorithm 4.2) used in previous work to illustrate Cannon–Thurston maps.

In Section 5 we give several equivalent definitions of the cohomology fractal. It depends on choices beyond the manifold M and the cohomology class $[\omega]$: there is a choice of viewpoint $p \in M$ and a choice of a visual radius R . The cohomology fractal is a function $\Phi_R^{\omega, p}: \text{UT}_p M \rightarrow \mathbb{R}$. Roughly, for each vector $v \in \text{UT}_p M$ we build the geodesic arc γ of length R from p in the direction of v and compute $\Phi_R^{\omega, p}(v) = \omega(\gamma)$. (Note that we repeatedly generalise the definition of the cohomology fractal throughout the paper; the decorations alter to remind the reader of the desired context.)

In Figure 1, we see a cohomology fractal closely matching an approximation of a Cannon–Thurston map, as produced by Algorithm 4.2. In Section 6 we prove the following.

Proposition 6.2. *Cohomology fractals are dual to approximations of the Cannon–Thurston map.*

Thus we have a new representation of Cannon–Thurston maps. We also compare cohomology fractals with the *lightning curves* of Dicks and various coauthors. (The name is due to Wright [32, page 324].) We experimentally observe that the lightning curve corresponds to some of the brightest points of the cohomology fractal.




Figure 1. Matching up Cannon–Thurston map images with the cohomology fractal for $m004$, the figure-eight knot complement. Compare our Figure 1(c) with Figure 10.11 of *Indra’s Pearls* [32, page 335], which was produced by paint-filling a vector graphics image [48].

In Section 7 we describe the algorithms we use to produce images of cohomology fractals. Adding the ability to move through the manifold leads us to separate the viewpoint p from a *basepoint*, denoted b , of the cohomology fractal. We still trace rays starting at p , but then evaluate ω on any path in $\tilde{M}$ from b to the endpoint of γ .

We also generalise the above *material* view (with vectors v in $UT_p \tilde{M}$) to the *ideal* and *hyperideal* views (with vectors v being perpendicular to a horosphere or geodesic plane, respectively). Each view is a subset $D \subset UT \tilde{M}$; our notation for the cohomology fractal becomes $\Phi_R^{\omega, b, D}: D \rightarrow \mathbb{R}$.

In Section 8 we discuss cohomology fractals for incomplete and closed manifolds. We draw cohomology fractals in the closed case in two ways. First, we deform the cohomology fractal for a surgery parent through Thurston’s *Dehn surgery space*. Second, we reimplement our algorithms using material triangulations. We also discuss possible sources of numerical error in our implementations.

In Section 9 we give a sequence of experiments exploring the dependence of cohomology fractals on the visual radius R . For any fixed R , the cohomology fractal is constant on regions with sizes roughly proportional to $\exp(-R)$. As R increases, these regions subdivide, and intricate patterns come into focus. This suggests that there is a limiting object. The following shows that such a limit cannot be a function.

Theorem 9.2. *Suppose that M is a finite volume, oriented hyperbolic three-manifold. Suppose that F is a transversely oriented surface. Then the limit*

$$\lim_{R \rightarrow \infty} \Phi_R(v)$$

does not exist for almost all $v \in UT \tilde{M}$.

Indeed, experimentally, increasing R leads to noisy pictures. However, this is due to undersampling. A heuristic argument (see Remark 9.7) shows that we can avoid noise if we increase the screen resolution as we increase R . We simulate this by computing *supersampled* images. These, and further experiments, indicate that in contrast with Theorem 9.2, the *mean* of the cohomology fractal, taken over a pixel, converges. Its values appear to be normally distributed with standard deviation growing like $\sqrt{R}$.

Motivated by this, in Section 10, we show that the cohomology fractal obeys a central limit theorem.

Theorem 10.7. *Fix a connected, orientable, finite volume, complete hyperbolic three-manifold M and a closed, non-exact, compactly supported one-form $\omega \in \Omega_c^1(M)$. There is $\sigma > 0$ such that for all basepoints b , all views D with area measure μ_D , for all probability measures $\nu_D \ll \mu_D$, and for all $\alpha \in \mathbb{R}$, we have*

$$\lim_{T \rightarrow \infty} \nu_D \left[v \in D : \frac{\Phi_T(v)}{\sqrt{T}} \leq \alpha \right] = \int_{-\infty}^{\alpha} \frac{1}{\sigma \sqrt{2\pi}} e^{-(s/\sigma)^2/2} ds$$

where $\Phi_T = \Phi_T^{\omega, b, D}$ is the associated cohomology fractal.

That is, if we regard the cohomology fractal across a pixel as a random variable, divide it by $\sqrt{T}$, and take the limit, the result is a normal distribution of mean zero. The standard deviation of the normal distribution only depends on the manifold and cohomology class. The proof uses Sinai's central limit theorem for geodesic flows.

In Section 11, we prove that treating the cohomology fractals as *distributions* gives a well-defined limit. In this introduction, for simplicity, we focus on the case where D is a material view. The *pixel theorem* (Theorem 11.4) states that the limit

$$\Phi^{\omega, b, D}(\eta) = \lim_{T \rightarrow \infty} \int_D \Phi_T^{\omega, b, D} \cdot \eta$$

is well-defined for any two-form $\eta \in \Omega^2(D)$. Theorem 11.4 also states various transformation laws relating, for example, the distributions corresponding to different views. Thus there is a view-independent distribution related to the view-dependent distributions via the conformal isomorphism i_D from D to $\partial \tilde{M}$.

Corollary 11.5. *Suppose that M is a connected, orientable, finite volume, complete hyperbolic three-manifold. Fix a closed, compactly supported one-form $\omega \in \Omega_c^1(M)$ and a basepoint $b \in \tilde{M}$. Then there is a distribution $\Phi^{\omega, b}$ on $\partial_\infty \tilde{M}$ so that, for any material view D and for any $\eta \in \Omega^2(D)$, we have*

$$\Phi^{\omega, b, D}(\eta) = \Phi^{\omega, b}((i_D^{-1})^* \eta)$$

The above discussion addresses smooth test functions. We can also prove convergence for a wider class of test functions; these include the indicator functions of regions with piecewise smooth boundary. However, we do not know whether or not the cohomology fractal converges to a measure.

We conclude with a few questions and directions for future work in Section 12.

Acknowledgements

This material is based in part upon work supported by the National Science Foundation under Grant No. DMS-1439786 and the Alfred P. Sloan Foundation award G-2019-11406 while the authors were in residence at the Institute for Computational and Experimental Research in Mathematics in Providence, RI, during the Illustrating Mathematics program. The fourth author was supported in part by National Science Foundation grant DMS-1708239.

We thank François Guéritaud for suggesting we use ray-tracing to generate cohomology fractals. We thank Curt McMullen for suggesting that Theorem 11.4 should be true and also for giving us permission to reproduce Figure 15. We thank Mark Pollicott and Alex Kontorovich for guiding us through the literature on exponential mixing of the geodesic flow. We thank Ian Melbourne for enlightening conversations on central limit theorems. We thank the anonymous referee for many helpful comments and corrections.

2. Triangulations

We briefly review the notions of material and ideal triangulations of three-manifolds.

2.1. Combinatorics

Suppose that M is a compact, connected, oriented three-manifold. We will consider two cases. Either

- the boundary ∂M is empty; here we call M *closed*, or
- the boundary is non-empty, consisting entirely of tori; here we call M *cusped*.




Figure 2. The figure-eight knot complement. This manifold is known as `m004` in the SnapPy census. The black lines in Figure 2(a) cut the tube into triangles, corresponding to the eight green triangles in Figure 2(b).

Suppose that $\mathcal{T}$ is a *triangulation*: that is, a collection of oriented model tetrahedra together with a collection of orientation-reversing face pairings. We allow distinct faces of a tetrahedron to be glued, but we do not allow a face to be glued to itself. The quotient space, denoted $|\mathcal{T}|$, is thus a CW-complex which is an oriented three-manifold away from its zero-skeleton. We say that $\mathcal{T}$ is a *material* triangulation of M if there is an orientation-preserving homeomorphism from $|\mathcal{T}|$ to M . We say that $\mathcal{T}$ is an *ideal* triangulation of M if there is an orientation-preserving homeomorphism from $|\mathcal{T}|$, minus a small open neighbourhood of its vertices, to M . Equivalently, $|\mathcal{T}|$ minus its vertices is homeomorphic to M° , the interior of M .

Example 2.2. Suppose that M is obtained from S^3 by removing a small open neighbourhood of the figure-eight knot. See Figure 2(a). As discussed in [42, Chapter 1], the knot exterior M has an ideal triangulation with two tetrahedra. See Figure 2(b). Here we have not truncated the model tetrahedra. Instead we draw the vertex link; in this case it is a torus in M .

2.3. Geometry

We deal with the geometry of the two types of triangulations separately.

2.3.1. Ideal triangulations

We give each model ideal tetrahedron t a hyperbolic structure. That is, we realise t as an ideal tetrahedron in $\mathbb{H}^3$ with geodesic faces. This can be constructed as the convex hull of four points on $\partial_\infty \mathbb{H}^3$. We require that the face pairings be orientation reversing isometries. For these hyperbolic tetrahedra to combine to give a complete hyperbolic structure on the manifold M° requires certain conditions to be satisfied. Very briefly: consider a loop in the dual one-skeleton of the triangulation. This visits the tetrahedra in some order. The product of the corresponding sequence of isometries must give the identity if the loop is trivial in the fundamental group. If the loop is peripheral then the product must be a parabolic element.

These conditions reduce to a finite set of algebraic constraints. These are Thurston's *gluing equations*, see [42, Section 4.2] and [35, Section 4.2]. Using the upper half-space model of $\mathbb{H}^3$ we define the *shape* of each ideal hyperbolic tetrahedron to be the cross-ratio of its four ideal points. The gluing equations impose a finite number of polynomial conditions on these shapes.

For our implementation, we also require that the shapes have positive imaginary part. This ensures that the ideal hyperbolic tetrahedra glue together to give a complete, finite volume hyperbolic structure on M° . Furthermore, the model orientations of all of the tetrahedra agree with the orientation on M . In particular, when a geodesic ray crosses a face, it has a sensible continuation.

2.3.2. Material triangulations

To find a hyperbolic structure for material triangulations, we replace Thurston's gluing equations with a construction due to Andrew Casson [8] and Damian Heard [20, 21]. To specify a hyperbolic structure on a material triangulation, it suffices to assign lengths to its edges. This is because the isometry class of an oriented material hyperbolic tetrahedron is determined by its six edge lengths.

There are two conditions that must be satisfied. First, for each model tetrahedron, there is a collection of inequalities that must be satisfied for its edges. Second, for each edge of the triangulation, the dihedral angles about it must sum to 2π .

If these inequalities and equalities hold, then we obtain a hyperbolic structure on the three-manifold. See [19, Section 2] for further details.

3. Cannon–Thurston maps

$$\begin{array}{ccc}
 \partial_\infty \tilde{F} & \xrightarrow{\Psi} & \partial_\infty \tilde{M} \\
 \downarrow & & \downarrow \\
 \bar{F} & \xrightarrow{\bar{\alpha}} & \bar{M} \\
 \uparrow & & \uparrow \\
 \tilde{F} & \xrightarrow{\tilde{\alpha}} & \tilde{M} \\
 \downarrow & & \downarrow \\
 F & \xrightarrow{\alpha} & M
 \end{array}$$

Figure 3. The various spaces and maps involved in constructing the Cannon–Thurston map Ψ .

Here we sketch Cannon and Thurston’s construction; see Figure 3 for an overview. We refer to [13] for the details. See also [31].

Suppose that $F = F^2$ and $M = M^3$ are connected, compact, oriented two- and three-manifolds. Suppose that F° and M° admit complete hyperbolic metrics of finite area and volume respectively. In an abuse of notation, we will conflate F with F° , and similarly M with M° . We call a proper embedding $\alpha: F \rightarrow M$ a *fibre* if there is a map $\rho: M \rightarrow S^1$ so that for all $t \in S^1$ the preimage $\rho^{-1}(t)$ is a surface properly isotopic to $\alpha(F)$.

Let $\tilde{F}$ and $\tilde{M}$ be the universal covers of F and M respectively. Since F and M are hyperbolic, their covers are identified with hyperbolic two- and three-space respectively. Let $\partial_\infty \tilde{F} \cong \partial_\infty \mathbb{H}^2 \cong S^1$ and $\partial_\infty \tilde{M} \cong \partial_\infty \mathbb{H}^3 \cong S^2$ be their ideal boundaries. We set

$$\bar{F} = \tilde{F} \cup \partial_\infty \tilde{F} \quad \text{and} \quad \bar{M} = \tilde{M} \cup \partial_\infty \tilde{M}$$

Each union is equipped with the unique topology that makes the group action continuous. Note that $\bar{F}$ and $\bar{M}$ are homeomorphic to a closed two- and three-ball, respectively.

We compose the covering map from $\tilde{F} \rightarrow F$ with the embedding α and then lift to obtain an equivariant map $\tilde{\alpha}: \tilde{F} \rightarrow \tilde{M}$. We call $\tilde{\alpha}$ an *elevation* of F . Figure 4 shows an elevation of the fibre of the figure-eight knot complement.

Cannon and Thurston gave the first proof of the following theorem in the closed case [13, page 1319]. The cusped case follows from work of Bowditch [3, Theorem 0.1].

Theorem 3.1. *Suppose M is a connected, oriented, finite volume hyperbolic three-manifold. Suppose that $\alpha: F \rightarrow M$ is a fibre of a surface bundle structure on M . Then there is an extension of $\tilde{\alpha}$ to a continuous and equivariant (with respect to the fundamental group of M) map $\bar{\alpha}: \bar{F} \rightarrow \bar{M}$. The restriction of $\bar{\alpha}$ to $\partial_\infty \bar{F}$ gives a sphere-filling curve.* $\square$

We will use the notation $\Psi: \partial_\infty \tilde{F} \rightarrow \partial_\infty \tilde{M}$ for the restriction of $\bar{\alpha}$ to S^1_∞ . We call this a *Cannon–Thurston map*. We now turn to the task of visualising Ψ .

4. Illustrating Cannon–Thurston maps

The standard joke (see [43, page 373] and [32, page 335]) is that it is straightforward to draw an accurate picture of a Cannon–Thurston map; it is solid black.


The first (more instructive) illustration of a Cannon–Thurston map is due to Thurston. He gives a sequence of approximations to the sphere-filling curve in [43, Figure 8]. We reproduce the last of these in Figure 4(b). A striking version of this image by Wright also appears in *Indra’s Pearls* [32, Figure 10.11]. In this example both M and F are non-compact; M is the complement of the figure-eight knot and F is a Seifert surface.

4.1. Vector and raster graphics

In this section, we outline the technique used by Thurston and Wright to generate images of Cannon–Thurston maps, in order to contrast it with our algorithm.

Our algorithm generates an image by producing a colour for each pixel on a screen. In other words, its output is a map from a grid of pixels in the *image plane* into a space of possible colours. We call such a map a *raster graphics image*.

In contrast, **Algorithm 4.2** (below) produces *vector graphics* – that is a description of an image as a collection of various primitive objects in the (euclidean) image plane. An example of a primitive is a line segment, specified by the coordinates of its end points. Other primitives include arcs, circles, and so on. Note that we generally need to convert vector graphics to raster graphics to make a physical representation of an image. To *rasterise* a vector graphics image, we need to decide which pixels are coloured by which



(a) An elevation of the fibre. The fibre is a pleated surface made from two ideal triangles. The three pleating angles are $\pi/3$, π , and $5\pi/3$.

(b) An approximation of the Cannon-Thurston map, reproduced from Figure 8 of [43].

Figure 4. Views in the universal cover of the figure-eight knot complement.

primitives. For example, a disk colours all of the pixels whose coordinates are close enough to the centre of the disk. Rasterisation is necessary for most output devices, such as screens or printers. The exceptions include plotters, laser cutters, and cathode-ray oscilloscopes. There are advantages to deferring rasterisation and saving the vector graphics to a file (in the PDF, PostScript, or SVG format, for example). For example, deferred rasterisation can take the resolution of the output device into account. Design programs such as Inkscape and Adobe Illustrator allow editing the geometric primitives in a vector graphics image. Rasterisation is usually carried out by a black-box general purpose algorithm, the details of which are hidden from the user.

Algorithm 4.2. (Approximate a Cannon-Thurston map) We are given a fibre F of the three-manifold M . We choose an elevation $\tilde{F} \subset \tilde{M}$ of F . As described in Section 3, the map

$$\Psi: S^1 \cong \partial_\infty \tilde{F} \rightarrow \partial_\infty \tilde{M} \cong S^2$$

is sphere-filling.

To approximate Ψ , we first choose a large disk $D \subset \tilde{F}$. Typically, M is described with an ideal triangulation $\mathcal{T}$, with F realised as a surface carried by the two-skeleton $\mathcal{T}^{(2)}$. Therefore $\tilde{F}$ is given as a surface carried by $\tilde{\mathcal{T}}^{(2)}$. The disk D then consists of some finite collection of ideal hyperbolic triangles in $\tilde{\mathcal{T}}^{(2)}$. The boundary of D consists of a loop of geodesics in $\mathbb{H}^3 \cup \partial_\infty \mathbb{H}^3$. We now define Ψ_D to be the loop in $\partial_\infty \mathbb{H}^3$ obtained by projecting each arc of ∂D to an arc in $\partial_\infty \mathbb{H}^3$.

Note that the algorithm produces a circularly ordered collection of points in $\partial_\infty \mathbb{H}^3$ spanning geodesics in $\mathbb{H}^3$. However, conventional vector graphics require primitives to be in the euclidean plane. Thus, we must make two choices of projections. The first projection from $\mathbb{H}^3$ to $\partial_\infty \mathbb{H}^3$ takes the geodesics to arcs in $\partial_\infty \mathbb{H}^3$ and the second projection takes these arcs in $\partial_\infty \mathbb{H}^3$ to arcs in the euclidean image plane.

We draw our pictures in the “ideal view”. That is, we use the upper half space model of $\mathbb{H}^3$ and project ∂D down to $\mathbb{C}$ (viewed as the boundary of $\mathbb{H}^3$). The arcs between vertices now simply become straight lines in $\mathbb{C}$. This is also the choice made by Thurston in Figure 4, as well as Wada in his program *OPTi* [46]. Other depictions by McMullen [28] and Calegari [7, Figure 1.14] project outward from the origin to the boundary of the Poincaré ball model of $\mathbb{H}^3$ (and then to the image plane using a perspective or orthogonal projection).

Remark 4.3. The images in *Indra’s Pearls* [32] have been rasterised using a customised rasteriser that illustrates further features of the Cannon-Thurston map. For example, one side of the polygonal path of line segments is filled, or the line segments are coloured using some combinatorial condition. See Figures 10.11 and 10.13 of [32].

4.4. Motivating raster graphics

Our work here began when we asked if we could avoid vector graphics when illustrating Cannon-Thurston maps. This is less natural, but would allow us to take advantage of extremely fast graphics processing unit (GPU) calculation. We were inspired in part by work of Vladimir Bulatov [6] and also of Roice Nelson and the fourth author [33], using reflection orbihedra. (See also the work of Peter Stampfli [39].) They all use raster graphics strategies to draw tilings of $\mathbb{H}^2$ and $\mathbb{H}^3$.

A number of others have also used raster graphics to explore kleinian groups, outside of the setting of reflection orbihedra. They include Peter Liepa [26], Jos Leys [25], and Abdelaziz Nait Merzouk [30] (also see [12]).

5. Cohomology fractals

We give a sequence of more-or-less equivalent definitions of cohomology fractals, beginning with the conceptually simplest (for us), and moving towards versions that are most convenient for our implementation or our proofs. Fixing notation, we take M to be a riemannian manifold and $\tilde{M}$ its universal cover. We also take $X = X_M = \text{UT } M$ and $\tilde{X} = \text{UT } \tilde{M}$ to be their unit tangent bundles. Take $\pi: \tilde{X} \rightarrow \tilde{M}$ to be the bundle map.

Definition 5.1. Suppose that we are given the following data.

- A connected, complete, oriented riemannian manifold M^n ,
- a cocycle $\omega \in Z^1(M; \mathbb{Z})$,
- a point $p \in M$, and
- a radius $R \in \mathbb{R}_{>0}$.

From these, we define the *cohomology fractal* on the unit sphere $\text{UT}_p M \cong S^{n-1}$. This is a function $\Phi_R = \Phi_R^{\omega, p}: \text{UT}_p M \rightarrow \mathbb{Z}$ defined as follows.

Suppose that $v \in \text{UT}_p M$ is a unit tangent vector. Let γ be the unique geodesic segment starting at p with initial direction v and of length R . Let q be the endpoint of γ . Choose any shortest path γ' from q to p (on a set of full measure γ' is unique). Thus $\gamma \cup \gamma'$ is a one-cycle. We define $\Phi_R(v) = \omega(\gamma \cup \gamma')$.

Our next definition moves in the direction of concrete examples:

Definition 5.2. Here we further assume that M is a three-manifold. Let $F \subset M$ be a properly embedded, transversely oriented surface. We choose F so that $p \notin F$. We define γ and q as above. Now take γ' to be the shortest path from q to p in the complement of F . We now define $\Phi_R(v) = \Phi_R^{F, p}(v)$ to be the algebraic intersection number between F and $\gamma \cup \gamma'$.

We modify once again to obtain a definition very close to our implementation.

Definition 5.3. We equip M with a material (or ideal) triangulation $\mathcal{T}$. We properly homotope the surface F to lie in the two-skeleton $\mathcal{T}^{(2)}$. For each face f this gives us a weight $\omega(f)$. This is the signed number of sheets of F running across f . We dispense with γ' ; we take $\Phi_R(v)$ to be the sum of the weighted intersections between γ and the faces of the triangulation.

To aid in comparing cohomology fractals to Cannon–Thurston maps (in Section 6), we lift to the universal cover, $\tilde{M}$.

Definition 5.4. Since cochains pull back, let $\tilde{\omega}$ be the lift of ω . Let $\tilde{p}$ be a fixed lift of the point p . Since $\tilde{\omega}$ is a coboundary, it has a primitive, say W ; we choose W so that $W(\tilde{p}) = 0$. We form $\tilde{\gamma}$ as before and let $\tilde{q}$ be its endpoint. We define $\Phi_R(v) = W(\tilde{q})$.

To analyse the behaviour of the cohomology fractal as R tends to infinity, we rephrase our definition in a dynamical setting. Here, the radius R is replaced by a time T .

Definition 5.5. Suppose that $\omega \in \Omega^1(M, \mathbb{R})$ is a closed one-form. Let $\varphi_t: \text{UT } M \rightarrow \text{UT } M$ be the geodesic flow for time t . We define

$$\Phi_T(v) = \Phi_T^{\omega, p}(v) = \int_0^T \omega(\varphi_t(v)) dt$$

In a slight abuse of notation, we also use φ_t to denote the geodesic flow on $\text{UT } \tilde{M}$.

In Section 7 we will discuss how to calculate Φ_R in practice. Before giving those details, we show the reader what Φ_R looks like for a few values of R . See Figure 5. The map Φ_R maps into $\mathbb{R}$; we indicate the value of $\Phi_R(v)$ by brightness. For each value of R , we draw the value of $\Phi_R(v)$ for a small square subset of the unit tangent vectors, $\text{UT}_p M$.

Here we are using Definition 5.3, our manifold M is the complement of the figure-eight knot, and the surface F is a fibre of M . Note that when R is small, as in Figure 5(a), Φ_R is constant on large regions of the sphere. As R increases, the value of Φ_R on nearby rays becomes less correlated, and we see a fractal structure come into focus.

Remark 5.6. This complicated behaviour is a consequence of the hyperbolic geometry of our manifold. Consider instead the example where M is the three-torus $S^1 \times S^1 \times S^1$ and the surface F is an essential torus embedded in M . Again, Φ_R counts the number of elevations of F the ray γ passes through. Since the elevations are parallel planes in $\tilde{M} \cong \mathbb{R}^3$, the value of Φ_R is constant on circles in $\text{UT}_p M$ parallel to these planes. Here Φ_R is much simpler.




Figure 5. Cohomology fractals for m004, with various values of R .

Remark 5.7. Some of the geometry and topology of the manifold M can be seen from the cohomology fractal Φ_R . Our recent expository paper on cohomology fractals [4] gives many such examples, including the appearances of cusps, totally geodesic subsurfaces, and loxodromic elements of the fundamental group.

6. Matching figures

In this section we compare our cohomology fractals to Cannon–Thurston maps.

Example 6.1. Suppose that M is the figure-eight knot complement. Suppose that F is a fibre of the fibration of M . Suppose that Ψ is the associated Cannon–Thurston map. Figure 1(a) shows an approximation $\Psi_D: S^1 \rightarrow S^2 \cong \partial_\infty \mathbb{H}^3$ of Ψ ; we produced this image using Algorithm 4.2. (Note that the vector graphics image in Figure 1(a) has been converted to a raster graphics image to save on file size and rendering time.)

Figure 1(b) shows a cohomology fractal Φ_R corresponding to F and looking towards the same part of $\partial_\infty \mathbb{H}^3$. Figure 1(c) shows Φ_R again, but with the contrast increased and colour scheme simplified. Here the colour associated to a vector v is either white or grey, according to whether $\Phi(v)$ is negative or not.

Figure 1(d) shows 1(a) overlaid on 1(c). We see that the red curve of Ψ_D is almost the common boundary of the white and grey regions of Φ_R . There are several small areas where Ψ_D does not track the boundary. These only appear close to fairly large cusps; they exist because implementations of Algorithm 4.2 generally have trouble approaching cusps from the side. In the cohomology fractal Φ_R we see that there are chains of “octopus heads” that reach almost all the way towards each cusp.

This behaviour is generally true for fibrations, as follows. Suppose $\tilde{p} \in \tilde{M}$ and define $\rho: \partial_\infty \tilde{M} \rightarrow \text{UT}_{\tilde{p}} \tilde{M}$ to be the inward central projection. Here we follow Definition 5.4.

Proposition 6.2. Suppose that M is a connected, oriented, finite volume hyperbolic three-manifold. Suppose that $\alpha: F \rightarrow M$ is a fibre of a surface bundle structure on M . Fix p in F , and a lift $\tilde{p} \in \tilde{M}$. Fix any $R > 0$ and let $\Phi_R = \Phi_R^{F,p}$ be the resulting cohomology fractal for F . Let $Z = \Phi_R^{-1}(\mathbb{R}_{\geq 0})$.

Then there is a disk $D \subset \tilde{F}$, containing $\tilde{p}$, so that (the image of) the Cannon–Thurston map approximation $\rho \circ \Psi_D$ is a component of ∂Z (with error at most $\exp(-R)$).

Proof sketch. We assume we are in the setting of Definition 5.3. Let ω be the one-cocycle dual to F . Let W be a primitive for $\tilde{\omega}$. Consider the two regions of $\mathbb{H}^3 \cong \tilde{M}$ where W is negative or, respectively, non-negative. The common boundary of these is exactly an elevation $\tilde{F}$ of the fibre F .

Let B_R^3 be the ball in $\mathbb{H}^3$ of radius R with centre $\tilde{p}$. Let S_R^2 be the boundary of B_R^3 . Thus the intersection $\tilde{F} \cap S_R^2$ is a collection of curves; these separate the points $\tilde{q} \in S_R^2$ where W is negative from those where it is non-negative. Finally, note that $\pi \circ \varphi_R: \text{UT}_{\tilde{p}} \mathbb{H}^3 \rightarrow S_R^2$ is the exponential map. Thus $\Phi_R = W \circ \pi \circ \varphi_R$.

Recall that $\tilde{F}$ is a union of triangles. Assume that one of these contains $\tilde{p}$. Let E be the collection of triangles of $\tilde{F}$ having at least one edge meeting the ball B_R^3 . Let D be the connected component of (the union over) E that contains $\tilde{p}$.

In a slight abuse of notation, let $\rho: \mathbb{H}^3 - \{\tilde{p}\} \rightarrow \text{UT}_{\tilde{p}} \mathbb{H}^3$ be the inward central projection. For any set $K \subset \mathbb{H}^3 - \{\tilde{p}\}$ we call the diameter of $\rho(K)$ the *visual diameter* of K . This is measured with respect to the fixed metric on the unit two-sphere $\text{UT}_{\tilde{p}} \mathbb{H}^3$.

Suppose that e is a bi-infinite geodesic in $\mathbb{H}^3$. If e lies outside of B_R^3 then the visual diameter of e is small; in fact, for large R the visual diameter of e is less than $2 \exp(-R)$. Likewise, if e meets S_R^2 , then for either component e' of $e - B_R^3$ the visual diameter of e' is less than $\exp(-R)$.

Now suppose that T is a triangle of E . Using the above, we deduce that the visual diameter of each component of $T - B_R^3$ is small. Thus the inward central projections of ∂D and $D \cap S_R^2$ have Hausdorff distance bounded by a small multiple of $\exp(-R)$.

So let Ψ_D be the curve in $\partial_\infty \mathbb{H}^3$ obtained by projecting ∂D outwards. By the above, each arc of the resulting polygonal curve has small visual diameter. Also by the above, the Hausdorff distance between the curves $\rho \circ \Psi_D$ and $\rho \circ (\tilde{F} \cap S_R^2)$ is small. $\square$

Remark 6.3. Suppose that F is totally geodesic or, more generally, quasi-fuchsian. In this case, the Cannon–Thurston map Ψ is a circle or quasi-circle, respectively. Note however that different elevations now give distinct Cannon–Thurston maps. It is natural to take their union and obtain a circle (or quasi-circle) packing.

Now, if F is also Thurston-norm minimising then we still obtain matches. For example in [4, Figure 6], we see how, for a totally geodesic surface in the Whitehead link complement, the cohomology fractal matches the associated circle packing. On the other hand, if $[F]$ is trivial in $H_2(M, \partial M)$ then the cohomology fractal Φ_R is bounded and oscillates as R tends to infinity.

6.4. Lightning curves

Suppose that M is a cusped, fibred three-manifold, with fibre F . Dicks with various co-authors defines and studies the *lightning curves* [1, 9, 10, 14, 15, 17]; these are certain fractal arcs in the plane. In more detail; suppose that c and d are distinct cusps of an elevation $\tilde{F}$ of F . Let $[c, d]$ be the arc of $\partial_\infty \tilde{F}$ that is between c and d and anti-clockwise of c . The Cannon–Thurston map Ψ sends the arc $[c, d]$ to a union of disks in $\partial_\infty \tilde{M}$ meeting only along points. The boundary of any one of these disks is a *lightning curve*.

Since the lightning curve is defined in terms of the Cannon–Thurston map, it is not too surprising that we can also see something of the lightning curve in the cohomology fractal. In Figure 6(a) we show a segment of the lightning curve for the figure-eight knot complement generated by Cannon and Dicks [10, Figure 7] overlaid on Figure 1(b). The lightning curve seems to follow some of the brightest pixels in the cohomology fractal. Figure 6(b) is a black and white version of Figure 1(b), with a relatively high threshold set for a pixel to be white – the lightning curve seems to be there, but this is nowhere near as clear as it was for the approximations to the Cannon–Thurston map. We do not fully understand the correspondence here.

We note that for clarity, Figure 6(b) shows only one segment of the lightning curve. There is another segment, symmetrical with the shown segment under a 180 degree rotation about the centre of the image. This second segment seems to follow the darkest pixels of the cohomology fractal.

7. Implementation

In this section we give an overview of an implementation of cohomology fractals. Our implementation is written in Javascript and GLSL; the code is available at [5]. We have also made cohomology fractals available in SnapPy [11]. Note that SnapPy cannot find hyperbolic structures on finite triangulations.

We now follow Definition 5.3. Suppose that M is a connected, oriented, finite volume hyperbolic three-manifold. Let $\mathcal{T}$ be a material or ideal triangulation of M . We are given a weighting $\omega: \mathcal{T}^{(2)} \rightarrow \mathbb{R}$ for the faces of the two-skeleton.

We represent the triangulation $\mathcal{T}$ as a collection $\{t_i\}$ of model hyperbolic tetrahedra. Each tetrahedron t_i has four faces f_m^i lying in four geodesic planes P_m^i in the hyperboloid model of $\mathbb{H}^3$. Suppose that $t_j \in \mathcal{T}$ is another model tetrahedron, with faces f_n^j . If the face f_m^i is glued to f_n^j , then we have isometries g_m^i and g_n^j realising the gluings. Note that g_m^i and g_n^j are inverses.




Figure 6. Matching up the lightning curve, reproduced from [10, Figure 7], with the cohomology fractal for $m004$.

We are given a camera location p in M ; this is realised as a point (again called p) in some tetrahedron t_i .

Remark 7.1. The reader familiar with computer graphics will note that we also require a frame at the camera location. To simplify the exposition, we will mostly suppress this detail.

We are also given a radius R as well as a maximum allowed step count S .

7.2. Ray-tracing

For each pixel of the screen, we generate a corresponding unit tangent vector u in the tangent space to the current tetrahedron t_i . We then *ray-trace* through $\mathcal{T}$. That is, we travel along the geodesic starting at p , in the direction u , for distance R , taking at most S steps. Figure 7(a) shows a toy example, where we replace the three-dimensional hyperbolic triangulation $\mathcal{T}$ of M with a two-dimensional euclidean triangulation of the two-torus.

It is perhaps most natural to think of ray-tracing as occurring in $\tilde{M}$, the universal cover of the manifold, as shown in Figure 7(a). However, the naïve floating-point implementation in the hyperboloid model quickly loses precision. We instead ray-trace in the manifold, as illustrated in Figure 7(b). Thus, all points we calculate lie within our fixed collection of ideal hyperbolic tetrahedra $\{t_i\}$.

For each pixel, we do the following.

1. The following initial data are given: an index i of a tetrahedron, a point p in t_i , and a tangent vector u at p . Initialise the following variables.
 - The total distance travelled: $r \leftarrow 0$.
 - The number of steps taken: $s \leftarrow 0$.
 - The current tetrahedron index: $j \leftarrow i$.
 - The current position: $q \leftarrow p$.
 - The current tangent vector: $v \leftarrow u$.
2. Let γ be the geodesic ray starting at p in the direction of u . Find the index n so that γ exits t_j through the face f_n^j . Let t_k be the other tetrahedron glued to face f_n^j .
3. Calculate the position q' and tangent vector v' where γ intersects f_n^j . Let r' be the distance from q to q' . Set $r \leftarrow r + r'$ and set $s \leftarrow s + 1$.
4. If $r > R$ or $s > S$ then stop.
5. Set $j \leftarrow k$, set $q \leftarrow g_n^j(q')$, and set $v \leftarrow Dg_n^j(v')$.
6. Go to step (2).

This implements the ray-tracing part of the algorithm. In our toy example, this is shown in Figure 7(b).




Figure 7. A toy example of developing a ray through a tiling of a euclidean torus. Note that the geodesic segments passing through a tile are parallel; this is only because the geometry is euclidean. In a hyperbolic tiling the segments are much less ordered.

7.3. Integrating

To determine the colour of the pixel, we also track the total signed weight we accumulate along the ray. For this, we add the following steps to the loop above.

(1b) An initial weight w_0 is given. Initialise the following.

- The current weight: $w_c \leftarrow w_0$.

(5b) Let f be the face between t and t' , co-oriented towards t' . Set $w_c \leftarrow w_c + \omega(f)$.

At the end of the loop, the value of w_c gives the brightness of the current pixel. (In fact, we apply a function very similar to the arctangent function to remap the possible values of w_c to a bounded interval. We then apply a gradient that passes through a number of different colours. This helps the eye see finer differences between values than a direct map to brightness.)

7.4. Moving the camera

In our applications, we enable the user to fly through the manifold M . Depending on the keys pressed by the user at each time step, we apply an isometry g to p . We also track an orthonormal frame for the user; this determines how tangent vectors correspond to pixels of the screen. We also apply the isometry g to this frame. When the user flies out of a face f of the tetrahedron they are in, we apply the corresponding isometry g_k^i to the position p and the user's frame. We also add $w(f)$ to the initial weight w_0 . Without this last step, the overall brightness of the image would change abruptly as the user flies through a face with non-zero weight.

Remark 7.5. With this last modification, the cohomology fractal depends on a choice of basepoint $b \in \tilde{M}$. The point $p \in M$ must now also be replaced by $p \in \tilde{M}$ (abusing notation, we use the same symbol for both points). We add b to the notation, and now write the cohomology fractal as

$$\Phi_R^{\omega, b, p}: \text{UT}_p \tilde{M} \rightarrow \mathbb{R}$$

Remark 7.6. The dependence of the cohomology fractal on b is minor: If we change b to b' , then the value of $\Phi_R^{\omega, b, p}(v)$ changes by the weight we pick up along any path from b' to b .




Figure 8. Comparison between different views of the cohomology fractal for m004.

7.7. Material, ideal, and hyperideal views

The above discussion describes the *material view*; the geodesic rays emanate radially from p . To render an image, we place a rectangle in the tangent space at p . For each pixel of the screen, we take the tangent vector u to be the corresponding point of the rectangle. See Figure 8(a).

Definition 7.8. The *field of view* of a material image is the angle between the tangent vectors pointing at the midpoints of the left and right sides of the image.

Remark 7.9. The material view suffers from perspective distortion. This is most noticeable towards the edges of the image, and is worse when the field of view is large.

To generalise the material view to the ideal and hyperideal, we introduce the following terminology. We say that a subset $D \subset \text{UT } \tilde{M}$ is a *view* if it is one of the following.

1. In the *material view*, D is a fibre $\text{UT}_p \tilde{M} \cong S^2$.
2. In the *ideal view*, we take D to be the collection of outward normals to a horosphere H . That is, the vectors point away from $\partial_\infty H$. To render an image we place a rectangle in D . For each pixel of the screen we set the initial vector v to be the corresponding point of the rectangle. The starting point is then $\pi(v) \in \tilde{M}$, the basepoint of v . Finally, we set w_0 to be the total weight accumulated, along the arc from b to $\pi(v)$, as we pass through faces of the triangulation. See Figure 8(b).
3. In the *hyperideal view*, we take D to be the collection of normals to a transversely oriented geodesic plane P . We draw P on the euclidean rectangle of the screen using the Klein model. The algorithm is otherwise identical to the ideal view case. See Figure 8(c).

Remark 7.10. The ideal view in hyperbolic geometry is the analogue of an orthogonal view in euclidean geometry. In both cases this is the limit of backing the camera away from the subject while simultaneously zooming in.

Remark 7.11. The hyperideal view suffers from an “inverse” form of perspective distortion. Towards the edges of the image, round circles look like ellipses, with the minor axis along the radial direction.

Definition 7.12. Let $D \subset \text{UT } \tilde{M}$ be a view, as discussed above. In the notation for the cohomology fractal, we replace p by D :

$$\Phi_R^{\omega, b, D}: D \rightarrow \mathbb{R}$$

7.13. Edges

We give the user the option to see the edges of the triangulation. The user selects an edge thickness $\varepsilon > 0$. The web application implements this in a lightweight fashion: In step (3), if the distance from the point q' to one of the three edges of the face we have intersected is less than ε , then we exit the loop early. Depending on user choice, the pixel is either coloured by the weight w_c or by the distance d . See Figure 9. In SnapPy, we compute the intersection of the ray with a cylinder about the edge in addition to the intersection with the faces.

7.14. Elevations

We also give the user the option to see several elevations of the surface F . The user selects a weight $w_{\max} > 0$. In step (5b), if $w_0 < 0$ but $w_c > 0$, then we have crossed the elevation at weight zero. In this case we exit the loop, and colour the pixel by the distance d . Similarly, if $w_0 > w_{\max}$ but $w_c < w_{\max}$, then we have crossed the elevation at weight $w_{\max}$, and again we stop and colour by distance. Finally, if $0 < w_0 < w_{\max}$, then we stop if w_c has changed from w_0 . Figure 4(a) shows a single elevation.




Figure 9. Edges of the ideal triangulation of $m004$, as seen in the material view.

7.15. Triangulations, geometry, and cocycles

We obtain our triangulations and their hyperbolic shapes from the SnapPy census. We put some effort into choosing good representative cocycles; the choice here makes very little difference to the appearance of the cohomology fractal, but it makes a large difference to the appearance of the elevations. That is, a poor choice of cocycle gives a “noisy” elevation. For example, adding the boundary of a tetrahedron to the Poincaré dual surface may perform a one-three move to its triangulation. This adds unnecessary “spikes” to the elevations.

When our manifold has Betti number one, there is only one cohomology class of interest. Here we searched for taut ideal structures dual to this class [24]. When the SnapPy triangulation did not admit such a taut structure, we randomly searched for one that did. A taut structure gives a Poincaré dual surface with the minimum possible Euler characteristic.

When the Betti number is larger than one, we used `tnorm` [47] to find initial simplicial representatives of vertices of the Thurston norm ball [44] in $H_2(M, \partial M)$. We then greedily performed Pachner moves to reduce the complexity of the cocycles. We often, but not always, realised the minimum possible Euler characteristic.

7.16. Discussion

Any visualisation of a hyperbolic tiling suffers from the mismatch between the hyperbolic metric of the tiling and the euclidean metric of the image. The tools for generating more of the tiling involve applying hyperbolic isometries. The tiles thus shrink exponentially in size while growing exponentially in number. This makes it difficult for the tiles to cleanly approach $\partial_\infty \mathbb{H}^2$ or $\partial_\infty \mathbb{H}^3$. Approaching a “parabolic” point at infinity is even more difficult.

In the vector graphics approach, one must be careful to avoid wasting time generating huge numbers of invisible objects: tiles may be too small or their aspect ratios too large.

The ray-tracing approach (and any similar raster graphics approach) deals with this mismatch directly. Here we start with the pixel that is to be coloured and then generate only the hyperbolic geometry needed to determine its colour.

A disadvantage of the ray-tracing approach is that we generate the hyperbolic geometry necessary for each pixel independently, meaning that much work is duplicated. However, the massive parallelism in modern graphics processing units mitigates, and is in fact designed to deal with, this kind of issue. It often turns out to be faster to duplicate work in many parallel processes rather than compute once then transmit the result to all processes requiring it.

8. Incomplete structures and closed manifolds

Suppose that M is a cusped hyperbolic manifold. Recall that we generate cohomology fractals for M by using an ideal triangulation $\mathcal{T}$. Associated to $\mathcal{T}$ there is the *shape variety*; that is we impose the gluing equations outlined in Section 2.3, omitting the peripheral ones. This gives us a space of deformations of the complete hyperbolic structure to incomplete hyperbolic structures; see [42, Section 4.4] and [35, Section 6.2]. If we deform correctly, we reach an incomplete structure whose completion has the structure of a hyperbolic manifold. The result is a *hyperbolic Dehn filling* of the original cusped manifold.

8.1. Incomplete structures

Suppose that $(M, \mathcal{T})$ is an ideally triangulated manifold. Let Z^s be a path in the shape variety, where Z^∞ is the complete structure and the completion of Z^1 is a closed hyperbolic three-manifold obtained by Dehn filling M . Between the two endpoints, we have *incomplete structures* M_s on the manifold M .




Figure 10. Cohomology fractals for $m1.2.2(4s, -s)$ as s varies.

In an incomplete geometry, there are geodesic segments that cannot be extended indefinitely. Suppose that, as in our algorithm, we only consider geodesic segments emanating from p of length at most R . The endpoints of the rays that do not extend to distance R form the *incompleteness locus* Σ_s in the ball $B_R^3 \subset \mathbb{H}^3$. It follows from work of Thurston that Σ_s is a discrete collection of geodesic segments, for generic values of s [43].

Suppose that $\omega: \mathcal{T}^{(2)} \rightarrow \mathbb{R}$ is the given weight function dual to a properly embedded surface F in M . We assume that the boundary of F (if any) gives loops in the filled manifold that, there, bound disks. Thus F also gives a cohomology fractal in the filled manifold.

Remark 8.2. Note that there is no canonical way of transferring a base point b or view D between two different geometric structures M_s and $M_{s'}$. However, we can choose b and D for each M_s in a way that gives us continuously varying pictures. We do not dwell on the details here.

Figure 10 shows cohomology fractals for various M_s . We see a kind of branch cut in the background to either side of the incompleteness locus Σ_s . As we vary s , the background appears to bend along the geodesic. Other paths in the shape variety will give shearing as well as (or instead of) bending.

When we reach a Dehn filling, the two sides again match, and we see the structure of the closed filled manifold. See Figure 11. (The two sides can also match before we reach the Dehn filling due to symmetries of the cusped manifold lining up with the cone structure.)

8.3. Numerical instability near the incompleteness locus

Our algorithm, given in Section 7, does not require completeness. However, a ray from p to Σ_s necessarily meets infinitely many tetrahedra. This is because near Σ_s we are far from the thick part of any tetrahedron, and the thin parts of the tetrahedra are almost “parallel” to Σ_s . Thus the innermost loop of the algorithm will always halt by reaching the maximum step count; it follows that we




Figure 11. Cohomology fractal for the Dehn filling $m122(4, -1)$. This gives a final image for Figure 10, with $s = 1$.




Figure 12. Cohomology fractal for the Dehn filling $m122(4, -1)$ drawn with an incomplete structure on an ideal triangulation. Here the maximum number of steps S is 55. Compare with Figure 11.

cannot “see through” a neighbourhood of Σ_s . Figure 12 shows the cohomology fractal drawn with a small maximum step count, making such a neighbourhood visible.

Increasing the maximum step count shrinks the opaque neighbourhood of Σ_s . However, as a ray approaches Σ_s , its segments within the model tetrahedra tend to their ideal vertices. Thus the coordinates blow up; this appears to lead to numerically unstable behaviour. See Figure 13(a). In the next section we describe a method to eliminate these numerical defects; we use this to produce Figure 13(b).

Note that numerical instability caused by rays approaching the ideal vertices also occurs for the complete structure on a cusped manifold. It is less noticeable in this case however, because these errors occur in a small part of the visual sphere for typical positions.




Figure 13. A view of the cohomology fractal for the manifold $m122(4, -1)$ near the incompleteness locus. On the left we have taken the maximum number of steps S sufficiently large to ensure that all rays reach distance R .

8.4. Material triangulations

In order to remove instability around the incompleteness locus, we remove it. That is, we abandon (spun) ideal triangulations in favour of material triangulations. There is no change to the algorithm in [Section 7](#); we only alter the input data (the planes P_k^i and face-pairing matrices g_k^i):

Given the edge lengths (see [Section 2.3.2](#)) for a material triangulation, [\[19, Lemma 3.4\]](#) assigns hyperbolic isometries to the edges of a doubly truncated simplex (also known as permutahedron). These can be used to switch a tetrahedron between different standard positions (as defined in [\[19, Definition 3.2\]](#)) where one of its faces is in the $\mathbb{H}^2 \subset \mathbb{H}^3$ plane. We assume that every tetrahedron is in $(0, 1, 2, 3)$ -standard position. Given a face-pairing, we apply the respective isometries to each of the two tetrahedra such that the faces in question line up in the $\mathbb{H}^2 \subset \mathbb{H}^3$ plane. The face-pairing matrix g_k^i is now given by composing the inverse of the first isometry with the second isometry. For example, let face 3 of one tetrahedron be paired with face 2 of another tetrahedron via the permutation $(0, 1, 2, 3) \mapsto (0, 1, 3, 2)$. To line up the faces, we need to bring the second tetrahedron from the default $(0, 1, 2, 3)$ -standard position into $(0, 1, 3, 2)$ -standard position by applying γ_{012} from [\[19, Lemma 3.4\]](#) which will thus be the face-pairing matrix, see [\[19, Figure 4\]](#). It is left to compute the planes P_k^i . Note that P_3^i (for each i) is the canonical copy of $\mathbb{H}^2 \subset \mathbb{H}^3$. All other P_k^i can be obtained by applying the isometries from [\[19, Lemma 3.4\]](#) again.


8.5. Cannon–Thurston maps in the closed case

Cannon and Thurston’s original proof was in the closed case. Thurston’s original images and all subsequent renderings, with one notable exception, are in the cusped case. With some minor modifications, [Proposition 6.2](#) applies in the closed case; thus the cohomology fractals again approximate Cannon–Thurston maps.

We are aware of only one previous example in the closed case, due to McMullen [\[28\]](#). In [Figure 14](#) we give a rasterisation of his original vector graphics image [\[29\]](#), and our version of the same view. The filling $m004(0, 2)$ of the figure-eight knot complement has an incomplete hyperbolic metric. The completion is a hyperbolic orbifold $\mathcal{O}$ with angle π about the orbifold locus; the universal cover is $\mathbb{H}^3$.

Since the filling is a multiple of the longitude, the orbifold $\mathcal{O}$ is again fibred. An elevation of this fibre to $\mathbb{H}^3$ gives a Cannon–Thurston map. Our image, [Figure 14\(a\)](#) is the cohomology fractal for the fibre in $\mathcal{O}$, in the hyperideal view. This is implemented using a material triangulation of an eight-fold cover M . Since M with its fibre, is commensurable with $\mathcal{O}$ with its fibre, we obtain the same image.

McMullen’s image, reproduced in [Figure 14\(a\)](#) was generated using his program `lim` [\[27\]](#). Briefly, let $\mathcal{O}^\infty$ be the infinite cyclic cover of $\mathcal{O}$. McMullen produces a sequence $\mathcal{O}^n$ of quasi-fuchsian orbifolds that converge in the geometric topology to $\mathcal{O}^\infty$. In each of these the convex core boundary is a pleated surface. The supporting planes of this pleated surfaces give round circles in $\partial\mathbb{H}^3$. His image then is obtained by taking n fairly large, passing to the universal cover of $\mathcal{O}^n$, and drawing the boundaries of many supporting planes [\[29\]](#).



(a) McMullen's illustration [29].
See also [28].

(b) Cohomology fractal - hyperideal
view.

Figure 14. Views of $m004(0, 2)$.

8.6. Accumulation of floating point errors

Our implementation uses single-precision floating point numbers. As we saw in Section 8.3, this can cause problems when rays approach the vertices of ideal tetrahedra. However, floating point errors can accumulate for large values of R whether or not rays approach the vertices. This can therefore also affect material triangulations.

With these problems in mind, we cannot claim that our images are rigorously correct. However, for small values of R we can be confident that our images are accurate. For very small values the endpoints of our rays all sit within the same tetrahedron, and so all pixels are the same colour. As we increase R (as in Figure 5), we see regions of constant colour, separated by arcs of circles. This is provably correct: (horo-)spheres meet the totally geodesic faces of tetrahedra in circles.

If we zoom in whilst increasing R , eventually floating point errors become visible. Figure 15 shows the results of an experiment to determine when this happens, for a material triangulation. At around $R = 11$, the circular arcs separating regions of the same colour become stippled. At around $R = 13$, the regions are no longer distinct.

Remark 8.7. Perhaps surprisingly, this accumulation of error does not mean that our pictures are inaccurate. Suppose that the side lengths of our pixels are on a somewhat larger scale than the precision of our floating point numbers. For each pixel, our implementation produces a piecewise geodesic, starting in the direction through the centre of the pixel, but with small angle defect at each vertex. Due to the nature of hyperbolic geometry, this piecewise geodesic cannot curve away from the true geodesic fast enough to leave the visual cone on the desired pixel. Thus, as long as the pixel size is not too small, each pixel is coloured according to some sample within that pixel.

9. Experiments

The sequence of images in Figure 5 suggests that some form of fractal object is coming into focus. When R is small, the function $\Phi_R = \Phi_R^{F,b,D}$ is constant on large regions of D . As R increases, these regions subdivide, producing intricate structures.

As we have defined it so far, the cohomology fractal depends on R . A natural question is whether or not there is a limiting object that does not depend on R . In this section we describe a sequence of experiments we undertook to explore this question. Inspired by these, in Sections 10 and 11 we provide mathematical explanations of our observations.

9.1. These pictures do not exist

A naïve guess might be that the cohomology fractal converges to a function as R tends to infinity. However, consider a ray following a closed geodesic γ in M that has positive algebraic intersection with the surface F . Choosing D so that it contains a tangent direction ν along γ , we see that $\Phi_R^F(\nu)$ diverges to infinity as R tends to infinity. The issue is not restricted to the measure zero set of rays along closed geodesics. Suppose that ν is a generic vector in a material view D . Recall that the geodesic flow is ergodic [22, Hauptsatz 7.1]. Thus the ray starting from ν hits F infinitely many times. So $\Phi_R^F(\nu)$ again diverges. Thus we have the following theorem.




Figure 15. We zoom into the cohomology fractal for $m122(4, -1)$ while increasing R . The field of view of the image is proportional to e^{-R} . Noise due to rounding errors becomes visible at $R \sim 10.5$ and completely dominates the picture when $R \sim 13.5$.

Theorem 9.2. *Suppose that M is a finite volume, oriented hyperbolic three-manifold. Suppose that p is any point of M . Suppose that F is a compact, transversely oriented surface. Then the limit*

$$\lim_{R \rightarrow \infty} \Phi_R(v)$$

does not exist for almost all $v \in \text{UT}_p M$.

□

Remark 9.3. To generalise [Theorem 9.2](#) from finite volume to infinite volume manifolds, we must replace Hopf’s ergodicity theorem by some other dynamical property. For example, Rees [36, Theorem 4.7] proves the ergodicity of the geodesic flow on the infinite cyclic cover of a hyperbolic surface bundle. This is generalised to the bounded geometry case by Bishop and Jones [2, Corollary 1.4]. Both of these works rely in a crucial fashion on Sullivan’s equivalent criteria for ergodicity [41, page 172].

One might hope that as R tends to infinity, nearby points diverge in similar ways. If so, we might be able to rescale and have, say, Φ_R/R or $\Phi_R/\sqrt{R}$ converge. However, increasing R in our implementation produces the sequence of images shown in [Figure 16](#). We see that, as we increase R , the images become noisy as neighbouring pixels appear to decorrelate. Eventually the fractal structure is washed away. Dividing the cohomology fractal by, say, some power of R only changes the contrast. Depending on this power, the limit is either almost always zero or does not exist.

[Figure 16](#) also demonstrates that [Remark 8.7](#), while valid, is misleading; it is true that for large R , every ray ends up somewhere within its pixel, but the colour one obtains is random noise. This noise is due to undersampling. In our images each pixel U is coloured using a single ray passing (almost, as we saw in [Section 8.6](#)) through its centre. When R is small relative to the side length of U the function $\Phi_R|_U$ is generally constant; thus any sample is a good representative. As R becomes larger the function $\Phi_R|_U$ varies more and more wildly; thus a single sample does not suffice.




Figure 16. Cohomology fractals for $m004$, with larger values of R . Each image here and in [Figure 5](#) has 1000×1000 pixels.




Figure 17. Side view of “screens” (in red) for two ideal views, drawn in the upper half space model. The outward pointing normals to each horosphere point down in the figure.

9.4. Take a step back and look from afar

Let D be an ideal view in the sense of [Section 7.7](#). We identify $\pi(D)$, isometrically, with the euclidean plane $\mathbb{C}$. Using this identification, we may refer to the vectors of D as z_D for $z \in \mathbb{C}$. Let E be the ideal view obtained from D by flowing outwards by a distance d . Thus, $\varphi_d(D) = E$. We similarly identify $\pi(E)$ with the euclidean plane, in such a way that for each $z \in \mathbb{C}$, we have $\varphi_d(z_D) = (e^d z)_E$. We may now state the following.

Lemma 9.5. Suppose that D is an ideal view and $E = \varphi_d(D)$. Then the cohomology fractal based at b satisfies

$$\Phi_{R+d}^{\omega,b,D}(z_D) = \Phi_R^{\omega,b,E}((e^d z)_E)$$

Proof. Consider [Figure 17](#). □

Said another way, if we fly backwards a distance d and replace R with $R + d$, we see the exact same image, scaled down by a factor of e^d . As a consequence, in the ideal view we have the following.

Remark 9.6. Each small part of a cohomology fractal with large R is the same as the cohomology fractal for a smaller R with a different view.

Remark 9.7. Since we know that we can make non-noisy images for small enough values of R , we can therefore make a non-noisy image of a cohomology fractal for any value of R , as long as we are willing to use a screen with high enough resolution.

The natural question then is how the *perceived* image changes as we simultaneously increase the resolution and increase R . This convergence question is different from the convergence of the cohomology fractal to a function as in [Theorem 9.2](#): when we look at a very large screen from far away, our eyes average the colours of nearby pixels. Thus, we move away from thinking of the limit as a function evaluated at points, towards thinking of it as a measure evaluated by integrating over a region. As we will see later, in fact the correct limiting object is a distribution.

9.8. Supersampling

To investigate this without requiring ever larger screens to view the results, we sample the cohomology fractal at many vectors in a grid within each pixel and average the results to give a colour for the pixel. That is, we employ supersampling. See [Figure 18](#). Here we draw cohomology fractals with R ranging from 4 to 12, and with either 1, 2^2 , or 128^2 subsamples for each pixel. Each image has resolution 128×128 .




Figure 18. $m_{122}(4, -1)$. Field of view: 12.8° . 128×128 pixels. For each image, the visual radius R is given at the start of its row, while the number of samples per pixel is given at the top of its column.

Remark. Note that some pdf readers do not show individual pixels with sharp boundaries: they automatically blur the image when zooming in. To combat this blurring and see the pixels clearly, we have scaled each image by a factor of three, so each pixel of our result is represented by nine pixels in these images.

With one sample per pixel, as we increase R the fractal structure comes into focus but then is lost to noise. This matches our observations in Figures 5 and 16. Taking subsamples and averaging makes little difference for small R : the only advantage is an anti-aliasing effect on the boundaries between regions of constant value. However, subsamples help greatly with reducing noise for larger R . With 2×2 subsamples, we see much less noise at $R = 10$, becoming more noticeable at $R = 12$. Taking 128×128 samples seems




Figure 19. The graph of the average value of the cohomology fractal for $m122(4, -1)$ for various square regions U with field of view 0.1° . Thus, these are the same size as the pixels of Figure 18. These are each computed by taking 1000×1000 samples. We also show the envelopes of 0.5, 1.0 and 1.5 standard deviations.

to be very stable: there is almost no difference between the images with $R = 10$ and $R = 12$. This suggests that the perceived images converge.

9.9. Mean and variance within a pixel

To better understand how subsampling interacts with increasing R , in Figure 19 we graph the average value within a selection of pixel-sized regions as R increases.

When R is small, the graphs are more-or-less step functions, as much of the time the pixel U is inside of a constant value region of the cohomology fractal. The graphs are also very similar for small R . This is because the pixels are close to each other, so all of their rays initially cross the same sequence of faces of the triangulation. Around $R = 6$, we reach the “last step” of the step function, then the regions of constant value become smaller than U . For $R \geq 10$, the mean seems to settle down, while the standard deviation appears to grow like $\sqrt{R}$.

Again this suggests that the perceived images converge. However, if the standard deviation continues to increase with R , then eventually any number of subsamples within each pixel will succumb to noise.

9.10. Histograms

We have looked at the standard deviation of a sample of values within a pixel. Next, we analyse the distribution of these values in more detail. See Figure 20.

We fix $R = e^2$. We sample Φ_R at each point of a 1000×1000 grid within a square of a material view with field of view 20° . We chose a relatively large field of view here so that we get an “in focus” image of the cohomology fractal with a relatively small value of R . Here we are being cautious to get good data, avoiding potential problems that our implementation has with large values of R as discussed in Section 8.6.

We histogram the resulting data with appropriate choices of bucket widths. In Figure 20(a) we show the histogram and the normal distribution with the same mean and standard deviation for our closed example, $m122(4, -1)$. In Figure 20(b) we show the sample data as a 1000 by 1000 pixel image. We also draw the normal distribution with the same mean and standard deviation; the data seems to fit this well.




Figure 20. Statistics for a cohomology fractal of $m122(4, -1)$ for a square region with field of view 20° and $R = e^2$.




Figure 21. Statistics for a cohomology fractal of $s789$ for a class vanishing on cusp.




Figure 22. Statistics for a cohomology fractal of $s789$ for a class not vanishing on cusp.

We repeat this experiment back in the cusped case with $s789$. See Figures 21 and 22. Here we show the cohomology fractal for two different cohomology classes $[\omega] \in H^1(M)$. The cohomology class shown in Figure 21 vanishes when restricted to ∂M , while in Figure 22 it does not. The distribution appears to be normal when the cohomology class vanishes on ∂M . When $[\omega]$ does not vanish on ∂M , something more complicated appears to be happening. One feature here is that the tails are much too long for a normal

distribution. A heuristic explanation for this is that in a neighbourhood of the cusp, a geodesic ray crosses the surface repeatedly in the same direction. This allows it to gain a linear weight in logarithmic distance.

10. The central limit theorem

In this section, we prove a central limit theorem for the values of the cohomology fractal Φ_T across a pixel. That is, the distribution of the values of the scaled cohomology fractal $R_T = \Phi_T/\sqrt{T}$ converges to a normal distribution with mean zero.

10.1. Setup

We recall the framework introduced in Section 7.7 that unifies the material, ideal, and hyperideal views. Let $\tilde{M}$ be a connected, orientable, finite volume, complete hyperbolic three-manifold. As above we set $X = \text{UT } \tilde{M}$ and $\tilde{X} = \text{UT } \tilde{M}$. We call a two-dimensional subset $D \subset \tilde{X}$ a *view* if it is of one of the following.

- For a material view, fix a basepoint $p \in \tilde{M}$ and let $D = \text{UT}_p \tilde{M}$. Note that D can be identified isometrically with S^2 .
- For the ideal view, fix a horosphere $H \subset \tilde{M}$ and let D be the set of outward normals to H .
- For the hyperideal view, fix a hyperbolic plane $H \subset \tilde{M}$ and let D be the set of normals to H facing one of the two possible directions.

Note that $D \subset \tilde{X}$ has a riemannian metric induced from the riemannian metric on $\tilde{X}$. This metric also endows D with an area two-form and associated measure denoted by $\zeta = \zeta_D$ and $\mu = \mu_D$. Recall that $\pi: \tilde{X} \rightarrow \tilde{M}$ is the projection to the base space.

Remark 10.2. Note that there is another riemannian metric on D for the ideal and hyperideal view coming from isometrically identifying the horosphere or hyperbolic plane $H = \pi(D)$ with $\mathbb{E}^2$ or with, respectively, $\mathbb{H}^2$. Up to a constant factor, this metric is the same as the above metric. The factor is trivial for the hyperideal view; it is $\sqrt{2}$ for the ideal view. This arises as $\sqrt{1 + K_H^2}$ from the extrinsic curvature K_H of H . We have $K_H = 1$ so that adding K_H to the ambient curvature -1 of $\mathbb{H}^3$ gives zero, the horosphere's intrinsic curvature.

In this notation, the definition of the cohomology fractal, for a given closed one-form $\omega \in \Omega^1(M)$ and basepoint $b \in \tilde{M}$, becomes the following. For $v \in D$, we have

$$\Phi_T^{\omega, b, D}(v) = \int_0^T \omega(\varphi_t(v)) dt + \int_b^{\pi(v)} \omega \quad (10.3)$$

For the second integral, any path from b to $\pi(v)$ in $\tilde{M}$ can be chosen as ω is closed. This integral is constant in v for the material view since $\pi(D) = p$. Choosing $W \in \Omega^0(\tilde{M})$ so that $dW = \tilde{\omega}$ and $W(b) = 0$, we can simply write

$$\Phi_T(v) = \Phi_T^{\omega, b, D}(v) = W \circ \pi \circ \varphi_T(v)$$

The central limit theorem will apply to probability measures ν_D that are absolutely continuous with respect to the area measure μ_D on D . We use the usual notation $\nu_D \ll \mu_D$ for absolute continuity. By the Radon-Nikodym theorem, this is equivalent to saying that the measure ν_D is given by

$$\nu_D(U) = \int_U h \cdot d\mu_D \quad (10.4)$$

where $h \geq 0$ is measurable with $\int_D h \cdot d\mu_D = 1$.

Remark 10.5. In Section 11, we will switch from measures ν_D to forms η_D , for the following reason. Here, in Section 10 we follow the well-established notation of [37, 49]. However, the transformation laws in Section 11 are better stated in the language of forms. In both cases we consider probability measures or two-forms that are “products”: namely of a suitable function $h: D \rightarrow \mathbb{R}$ with the area measure μ_D or two-form ζ_D respectively. The function h should be thought of as an indicator (or a kernel) function for a pixel.

10.6. The statement of the central limit theorem

The goal of this section is to prove the following.

Theorem 10.7. Fix a connected, orientable, finite volume, complete hyperbolic three-manifold M and a closed, non-exact, compactly supported one-form $\omega \in \Omega_c^1(M)$. There is $\sigma > 0$ such that for all basepoints b , for all views D with area measure μ_D , for all probability measures $\nu_D \ll \mu_D$, and for all $\alpha \in \mathbb{R}$, we have

$$\lim_{T \rightarrow \infty} \nu_D \left[v \in D : \frac{\Phi_T(v)}{\sqrt{T}} \leq \alpha \right] = \int_{-\infty}^{\alpha} \frac{1}{\sigma \sqrt{2\pi}} e^{-(s/\sigma)^2/2} ds$$

where $\Phi_T = \Phi_T^{\omega, b, D}$ is the associated cohomology fractal.

Let us recall some notions from probability to clarify what this means. Let (P, ν) be a probability space. For each $T \in \mathbb{R}_{>0}$, let $R_T: P \rightarrow \mathbb{R}$ be a measurable function. For each T , the probability measure ν on P induces a probability measure $\nu \circ R_T^{-1}$ on $\mathbb{R}$ telling us how the values of the random variable R_T are distributed when sampling P with respect to ν . Let ψ be a probability measure on $\mathbb{R}$.

Definition 10.8. We say that the random variables R_T converges in distribution to ψ if the measures $\nu \circ R_T^{-1}$ converge in measure to ψ . This is denoted by $\nu \circ R_T^{-1} \Rightarrow \psi$.

Here, by the Portmanteau theorem, we can use any of several equivalent definitions of weak convergence of measures. We are only interested in the case where ψ is absolutely continuous with respect to Lebesgue measure on $\mathbb{R}$; that is, we can write $\psi(V) = \int_V p(x) dx$ for any measurable $V \subset \mathbb{R}$. Note that here $p: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ is the *probability density function* for ψ . Convergence in distribution $\nu \circ R_T^{-1} \Rightarrow \psi$ is then equivalent to saying that for all α we have

$$\lim_{T \rightarrow \infty} \nu[x \in P : R_T(x) \leq \alpha] = \int_{-\infty}^{\alpha} p(s) \cdot ds$$

We define

$$n_{\sigma}(s) = \frac{1}{\sigma \sqrt{2\pi}} e^{-(s/\sigma)^2/2} \quad \text{and} \quad \psi_{\sigma}(V) = \int_V n_{\sigma}(x) dx$$

The latter is the *normal distribution* with mean zero and standard deviation σ .

Example 10.9. The process of flipping coins can be modelled as follows. Set $P = \{\text{head}, \text{tail}\}^{\mathbb{N}}$. We define a measure ν_P on P as follows. Given any prefix ν of length n , the set of all infinite words in P starting with ν has measure 2^{-n} . Let $S_i: P \rightarrow \mathbb{R}$ be the random variable $S_i(w) = \pm 1$ as w_i is heads or tails respectively. Define $\Sigma_N = S_0 + S_1 + \cdots + S_{N-1}$. The classical central limit theorem states that

$$\nu_P \circ R_N^{-1} \Rightarrow \psi_1 \quad \text{where} \quad R_N = \frac{\Sigma_N}{\sqrt{N}}: P \rightarrow \mathbb{R}$$

We can now restate [Theorem 10.7](#) as

$$\nu_D \circ R_T^{-1} \Rightarrow \psi_{\sigma} \quad \text{where} \quad R_T = \frac{\Phi_T}{\sqrt{T}}: D \rightarrow \mathbb{R}$$

10.10. Sinai's theorem

Our proof of [Theorem 10.7](#) starts with Sinai's central limit theorem for geodesic flows [37]. We use the following version of Sinai's theorem which is adopted from [18, Theorem VIII.7.1 and subsequent Nota Bene]. This applies to functions that are not derivatives in the following sense. Recall that $X = \text{UT } M$.

Definition 10.11. Let $f: X \rightarrow \mathbb{R}$ be a smooth function. We say that f is a *derivative* if there is a smooth function $F: X \rightarrow \mathbb{R}$ such that

$$f(v) = \left. \frac{dF(\varphi_t(v))}{dt} \right|_{t=0}$$

Let $\mu_X = \mu_{\text{Haar}} / \mu_{\text{Haar}}(X)$ be the normalised Haar measure.

Theorem 10.12 (Sinai-Le Jan's Central Limit Theorem). Fix a connected, orientable, finite volume, complete hyperbolic three-manifold M . Let $f: X = \text{UT } M \rightarrow \mathbb{R}$ be a compactly supported, smooth function with $\int_X f \cdot d\mu_X = 0$. Assume f is not a derivative. Let

$$R_T(v) = \frac{\int_0^T f(\varphi_t(v)) dt}{\sqrt{T}}$$

Then there is a $\sigma > 0$ such that $\mu_X \circ R_T^{-1} \Rightarrow \psi_{\sigma}$. □

In fact, the constant σ appearing in [Theorem 10.12](#) is the square root of the *variance* of f which Franchi–Le Jan denote by $\mathcal{V}(f)$. They give a formula for $\mathcal{V}(f)$ in [18, Theorem VIII.7.1] and state that $\mathcal{V}(f)$ vanishes if and only if f is a derivative.

Remark 10.13. To relate [Theorem 10.12](#) to [18, Theorem VIII.7.1], note that Franchi–Le Jan think of f as a function on the frame bundle of $\tilde{M}$ that is both Γ and $\text{SO}(2)$ -invariant. Since f is smooth and compactly supported, it satisfies the hypotheses of their theorem. Note that they also require f to not be a derivative (denoted by $\mathcal{L}_0 h$, see [18, (VIII.1)]) of a function h but allow h to be a function on the frame bundle. However, if an $\text{SO}(2)$ -invariant f is the derivative of a function h on the frame bundle, it is also the derivative of an $\text{SO}(2)$ -invariant function on the frame bundle.

We deduce [Theorem 10.7](#) from Sinai's theorem in three steps.

1. [Theorem 10.18](#) generalises Sinai's theorem to arbitrary probability measures $\nu_X \ll \mu_X$ on the five-dimensional $X = \text{UT } M$.
2. [Theorem 10.27](#) restricts from X to the two-dimensional view D using a measure $\nu_D \ll \mu_D$.
3. Finally, we show that the term $\int_b^{\pi(v)} \omega$ from (10.3) can be added to obtain Φ_T .

10.14. Generalising Sinai's Theorem

We begin with a definition.

Definition 10.15. Let (P, μ) be a finite measure space. For $n \in \mathbb{N}$, let $Q_n: P \rightarrow \mathbb{R}$ be a measurable function. We say that Q_n converges to zero in probability and write $\mu \circ Q_n^{-1} \rightarrow 0$ if for all $\varepsilon > 0$ we have

$$\lim_{n \rightarrow \infty} (\mu \circ Q_n^{-1})((-\infty, -\varepsilon) \cup (\varepsilon, \infty)) \rightarrow 0$$

The following result is called *strong distributional convergence*, see [49, Proposition 3.4].

Theorem 10.16. Let (P, μ) be a finite measure space and $T: P \rightarrow P$ be an ergodic, measure-preserving transformation. For all $n \in \mathbb{N}$, let $R_n: P \rightarrow \mathbb{R}$ be a measurable function. Let $Q_n = R_n \circ T - R_n$ and suppose that $\mu \circ Q_n^{-1} \rightarrow 0$. Let ψ be a probability measure on $\mathbb{R}$. If we have $\nu \circ R_n^{-1} \Rightarrow \psi$ for some probability measure $\nu \ll \mu$, then we have $\nu \circ R_n^{-1} \Rightarrow \psi$ for all probability measures $\nu \ll \mu$. $\square$

Remark 10.17. We have specialised Zweimüller's Proposition 3.4 in [49] to finite measure spaces. To obtain Zweimüller's result for σ -finite measure spaces (P, μ) , we need to replace the requirement $\mu \circ Q_n^{-1} \rightarrow 0$ by the following weaker requirement denoted by $Q_n \xrightarrow{\mu} 0$ in [49, Footnote 3]: for all probability measures $\nu \ll \mu$ we have $\nu \circ Q_n^{-1} \rightarrow 0$. To see that $Q_n \xrightarrow{\mu} 0$ is weaker, we can use the following standard result: for any $\nu \ll \mu$ we have

$$\sup\{\nu(A) : A \text{ measurable with } \mu(A) \leq \varepsilon\} \rightarrow 0 \quad \text{as } \varepsilon \rightarrow 0$$

The requirements $\mu \circ Q_n^{-1} \rightarrow 0$ and $Q_n \xrightarrow{\mu} 0$ are equivalent if μ is finite.

Using this, we now give our first variant of Sinai's theorem.

Theorem 10.18. With the same hypotheses as in [Theorem 10.12](#), we have the following. There is a $\sigma > 0$ such that for any probability measure $\nu_X \ll \mu_X$ we have $\nu_X \circ R_T^{-1} \Rightarrow \psi_\sigma$.

Proof. By [Theorem 10.12](#), there is a σ such that $\mu_X \circ R_T^{-1} \Rightarrow \psi_\sigma$ as $T \rightarrow \infty$. Note that the random variables in [Theorem 10.16](#) are indexed by $n \in \mathbb{N}$ instead of $T \in \mathbb{R}$ but it is easy to see that sequential convergence and convergence in distribution are equivalent. In other words, it suffices to show that for any sequence (T_n) with $T_n \rightarrow \infty$ the random variables $S_n = R_{T_n}$ satisfy $\nu_X \circ S_n^{-1} \Rightarrow \psi_\sigma$. In order to apply [Theorem 10.16](#), we need the following two claims.

Claim 10.19. The time-one map φ_1 for the geodesic flow is ergodic.

Proof. By [18, Theorem V.3.1] the geodesic flow is mixing. It follows that the time-one map φ_1 is also mixing, and thus ergodic. $\square$

Claim 10.20. Let $Q_n = S_n \circ \varphi_1 - S_n$. Then, $\nu_X \circ Q_n^{-1} \rightarrow 0$.

Proof. We will prove a stronger statement: $\|Q_n\|_\infty \rightarrow 0$.

$$\begin{aligned} |Q_n(v)| &= \left| \frac{\int_1^{T_n+1} f(\varphi_t(v)) dt}{\sqrt{T_n}} - \frac{\int_0^{T_n} f(\varphi_t(v)) dt}{\sqrt{T_n}} \right| \\ &= \left| \frac{\int_{T_n}^{T_n+1} f(\varphi_t(v)) dt}{\sqrt{T_n}} - \frac{\int_0^1 f(\varphi_t(v)) dt}{\sqrt{T_n}} \right| \\ &\leq \frac{2\|f\|_\infty}{\sqrt{T_n}} \end{aligned}$$

$\square$

We can now finish the proof of [Theorem 10.18](#). We apply [Theorem 10.16](#) with R_n replaced by S_n and T replaced by φ_1 . $\square$




Figure 23. Coordinates for $\tilde{X} = \text{UT } \tilde{M} = \text{UT } \mathbb{H}^3$ and flow of a box $(x_u, (-\varepsilon, \varepsilon), B_\varepsilon^s(\pi(x_u)))$.

10.21. Coordinates

Given a view D , we introduce coordinates for a neighbourhood of D in $\tilde{X} = \text{UT } \tilde{M} \cong \text{UT } \mathbb{H}^3$ as follows; it may be helpful to consult Figure 23(a). Fix $v \in \tilde{X}$. If v is close enough to D in $\tilde{X}$, then there is an $x_u \in D$ such that the rays emanating from x_u and v converge to the same ideal point in $\partial_\infty \tilde{M}$. Consider the set $H = H^s(x_u) \subset \tilde{X}$ such that

- $x_u \in H$,
- $\pi(H)$ is a horosphere, and
- H are the “inward pointing” normals to $\pi(H)$.

Let x_s be the intersection of H with the line through v . Let x_f be the signed distance from x_s to v along this line. The triple

$$(x_u, x_f, x_s) \quad \text{with } x_u \in D, x_f \in \mathbb{R}, \text{ and } x_s \in H^s(x_u)$$

determines the vector $v \in \tilde{X}$ uniquely. In an abuse of notation, we will simply write $v = (x_u, x_f, x_s)$.

Suppose N is a submanifold. Let $d_N(p, q)$ denote the length of the shortest curve in N connecting p and q . Given $x_u \in D$, let

$$B_\varepsilon^s(x_u) = \{x_s \in H : d_H(x_u, x_s) \leq \varepsilon\} \quad \text{where } H = H^s(x_u)$$

Let

$$D_\varepsilon = \{(x_u, x_f, x_s) : x_u \in D, x_f \in (-\varepsilon, \varepsilon), x_s \in B_\varepsilon^s(x_u)\} \subset \tilde{X}$$

Remark 10.22. Note that the subscripts appearing in the coordinates $v = (x_u, x_f, x_s)$ refer to the unstable, flow, and stable foliations.

- The points $(x_u, 0, x_u)$ give a copy of D , which is unstable.
- If we fix x_u and x_s , and vary x_f , then we obtain a geodesic flow line.
- Also, if we fix x_u and x_f , and vary x_s , then we obtain a stable horosphere.

Note that each $H = H^s(x_u)$ is isometric to a (pointed) copy of $\mathbb{C}$. However, the coordinates above do not live in a geometric product $D \times \mathbb{R} \times \mathbb{C}$. They instead form a smooth fibre bundle over D . (In the material case, the view D is a copy of S^2 . If we factor away the flow direction from our coordinates, what remains is isomorphic to the non-trivial bundle TS^2 .) Thus we will only locally appeal to a “product structure” on these coordinates.

The following lemma is deduced from the exponential convergence inside of stable leaves. See Figure 23(b).

Lemma 10.23. For all $t \geq 0$, $1 > \varepsilon > 0$ and all $v = (x_u, x_f, x_s) \in D_\varepsilon$, we have

$$d_{\tilde{X}}(\varphi_t(x_u), \varphi_t(v)) \leq 2\varepsilon$$

Proof. In our coordinates we have $\varphi_t(x_u) = (x_u, t, x_u)$ and $\varphi_t(v) = (x_u, x_f + t, x_s)$. We take $H = H^s(x_u)$. Then we have

$$\begin{aligned} d_{\tilde{X}}((x_u, t, x_u), (x_u, t, x_s)) &\leq d_{\varphi_t(H)}((x_u, t, x_u), (x_u, t, x_s)) \\ &= e^{-t} d_H((x_u, 0, x_u), (x_u, 0, x_s)) \\ &\leq e^{-t} \varepsilon \leq \varepsilon \end{aligned}$$

and

$$d_{\tilde{X}}((x_u, t, x_s), (x_u, t + x_f, x_s)) = |x_f| \leq \varepsilon$$

Applying the triangle inequality gives the result. $\square$

10.24. Proof of the central limit theorem

We now use these coordinates to continue with the proof of [Theorem 10.7](#).

Lemma 10.25. *Let (P, ν) be a probability space. Let $R_T, S_T: P \rightarrow \mathbb{R}$ be a pair of one-parameter families of measurable functions. Assume that $\nu \circ R_T^{-1} \Rightarrow \psi$ where ψ is a probability measure on $\mathbb{R}$ with bounded probability density function $p: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$. Assume that there is a monotonically growing family $U_T \subset P$ of measurable sets such that $\|(R_T - S_T)|_{U_T}\|_{\infty} \rightarrow 0$ and $\nu(P - U_T) \rightarrow 0$. Then $\nu \circ S_T^{-1} \Rightarrow \psi$.*

Proof. Fix α and let

$$P_T = \nu[x \in P : S_T(x) \leq \alpha] \quad \text{and} \quad Q_T = \nu[x \in P : S_T(x) > \alpha]$$

We need to show that for every $\varepsilon > 0$, there is a T_0 such that for all $T \geq T_0$ we have

$$\int_{-\infty}^{\alpha} p(s) \cdot ds - \varepsilon \leq P_T \leq \int_{-\infty}^{\alpha} p(s) \cdot ds + \varepsilon$$

We only deal with the second inequality since the first inequality can be derived in an analogous way using $P_T + Q_T = 1$. We have the following estimate.

$$P_T \leq \nu[x \in P : R_T(x) \leq \alpha + \|(R_T - S_T)|_{U_T}\|_{\infty}] + \nu(P - U_T)$$

Fix $\varepsilon > 0$. Let $\delta = \varepsilon/(3\|p\|_{\infty})$. By hypothesis, we have for all large enough T

$$P_T \leq \nu[x \in P : R_T(x) \leq \alpha + \delta] + \frac{\varepsilon}{3}$$

Because $\nu \circ R_T^{-1} \Rightarrow \psi$, we furthermore have for all large enough T

$$\begin{aligned} P_T &\leq \frac{\varepsilon}{3} + \int_{-\infty}^{\alpha+\delta} p(s) \cdot ds + \frac{\varepsilon}{3} \\ &\leq \frac{\varepsilon}{3} + \int_{-\infty}^{\alpha} p(s) \cdot ds + \|p\|_{\infty}\delta + \frac{\varepsilon}{3} = \int_{-\infty}^{\alpha} p(s) \cdot ds + \varepsilon \end{aligned} \quad \square$$

We return to the case of interest where f is given by a one-form ω .

Lemma 10.26. *Let $\omega \in \Omega^1(M)$ be closed but not exact. Then $\omega: X \rightarrow \mathbb{R}$ is not a derivative in the sense of [Definition 10.11](#).*

Proof. We prove the contrapositive: that is, if ω is a derivative in the sense of [Definition 10.11](#) then $\omega = dW$ for a function $W: M \rightarrow \mathbb{R}$. Fix a basepoint $p \in M$. We define $W(q) = \int_{\gamma} \omega$. Here γ is a path from p to q . All that is left is to show that W is well-defined.

So, suppose that γ' is another path from p to q . Thus $z = \gamma - \gamma'$ is a cycle. Let z^* be the geodesic representative of z . Since ω is closed we have $\int_z \omega = \int_{z^*} \omega$. Since ω is a derivative we have $\int_{z^*} \omega = 0$ and we are done. $\square$

Theorem 10.27. *With the same hypotheses as in [Theorem 10.7](#), we have the following. There is $\sigma > 0$ such that for all views D with area measure μ_D , for all probability measures $\nu_D \ll \mu_D$, and for all $\alpha \in \mathbb{R}$, we have*

$$\lim_{T \rightarrow \infty} \nu_D[\nu \in D : R_T(\nu) \leq \alpha] = \int_{-\infty}^{\alpha} \frac{1}{\sigma\sqrt{2\pi}} e^{-(s/\sigma)^2/2} ds$$

where $R_T(\nu) = \int_0^T \omega(\varphi_t(\nu)) dt / \sqrt{T}$.

Proof. The one-form ω is not a derivative by [Lemma 10.26](#). Taking $f = \omega$, let σ be as in [Theorem 10.18](#).

Fix a probability measure $\nu_D \ll \mu_D$. We define a measure $\nu_{\tilde{X}}$ on $\tilde{X}$ using the coordinates $\nu = (x_u, x_f, x_s)$ by taking the product of, in order,

- ν_D
- the Lebesgue measure on $\mathbb{R}$ restricted to $[-1, 1]$
- the Lebesgue measure on $\mathbb{C} \cong H^5(x_u)$ restricted to the unit disk $B_1^S(x_u)$

We scale $\nu_{\tilde{X}}$ to be a probability measure. Note that the Lebesgue measure on $H^5(x_u)$ does not depend on the isometric identification of $\mathbb{C}$ with $H^5(x_u)$. Thus $\nu_{\tilde{X}}$ is well-defined.

By summing over fundamental domains, the probability measure $\nu_{\tilde{X}}$ descends to a probability measure $\nu_X \ll \mu_X$ on X . Given that $R_T: \tilde{X} \rightarrow \mathbb{R}$ is $\pi_1(M)$ -invariant, [Theorem 10.18](#) yields $\nu_{\tilde{X}} \circ R_T^{-1} \Rightarrow \psi_\sigma$.

Note that $\nu_{\tilde{X}}$ is supported in the closure of D_1 (as defined before [Remark 10.22](#)). We have a projection $p: D_1 \rightarrow D$ where $p(x_u, x_f, x_s) = (x_u, 0, x_u)$. By construction, we have $\nu_{\tilde{X}}(p^{-1}(U)) = \nu_D(U)$ for any measurable set $U \subset D$.

Claim 10.28. We have

$$\nu_{\tilde{X}} \circ S_T^{-1} \Rightarrow \psi_\sigma \quad \text{where} \quad S_T = R_T \circ p$$

Proof. We take $P = D$ and we take $U_T = D$ for all T . Applying [Lemma 10.25](#), it is left to show that $\|R_T - S_T\|_\infty \rightarrow 0$. Let $W: \tilde{M} \rightarrow \mathbb{R}$ be a primitive of $\tilde{\omega}$. That is $dW = \tilde{\omega}$. In an abuse of notation, we abbreviate $W \circ \pi: \tilde{X} \rightarrow \mathbb{R}$ as W . Recall that $v = (x_u, x_f, x_s)$. We can now write

$$R_T(v) = \frac{W(\varphi_T(v)) - W(v)}{\sqrt{T}} \quad \text{and} \quad S_T(v) = \frac{W(\varphi_T(x_u)) - W(x_u)}{\sqrt{T}}$$

Since $1/\sqrt{T} \rightarrow 0$, it is sufficient to show that both of

$$W(\varphi_T(v)) - W(\varphi_T(x_u)) \quad \text{and} \quad W(v) - W(x_u)$$

are bounded by twice the Lipschitz constant of W . This follows from [Lemma 10.23](#) when replacing ε by 1 and setting t to either T or 0. $\square$

Fix α . [Theorem 10.27](#) follows from

$$\begin{aligned} \nu_D[v \in D : R_T(v) \leq \alpha] &= \nu_D[v \in D : S_T(v) \leq \alpha] \\ &= \nu_{\tilde{X}}(p^{-1}(\{v \in D : S_T(v) \leq \alpha\})) \\ &= \nu_{\tilde{X}}[v \in D_1 : S_T(v) \leq \alpha] \end{aligned}$$

converging to $\int_{-\infty}^{\alpha} n_\sigma(s) ds$ by [Claim 10.28](#). $\square$

Proof of Theorem 10.7. Note that [Theorem 10.27](#) shows convergence in distribution for

$$R_T(v) = \frac{\int_0^T \omega(\varphi_t(v)) dt}{\sqrt{T}}$$

However we need to show convergence in distribution for $Q_T(v) = \Phi_T(v)/\sqrt{T}$, the difference being

$$\Delta(v) = R_T(v) - Q_T(v) = \frac{\int_b^{\pi(v)} \omega}{\sqrt{T}} = \frac{\int_b^{\pi(u)} \omega + \int_{\pi(u)}^{\pi(v)} \omega}{\sqrt{T}}$$

where $u \in D$ is a fixed basepoint. Thus, we need to show that [Lemma 10.25](#) applies when taking $P = D$. Denote the constant $|\int_b^{\pi(u)} \omega|$ by C . Let C' be a bound on the absolute value of $\omega: X \rightarrow \mathbb{R}$. It is convenient to let

$$U_T = \{v \in D : d_D(u, v) \leq \sqrt[4]{T}\}$$

Then, $\|\Delta|_{U_T}\|_\infty \leq (C + C' \sqrt[4]{T})/\sqrt{T} \rightarrow 0$ for $v \in U_T$. Since U_T exhausts D and ν is a finite measure, $\nu(D - U_T) \rightarrow 0$. $\square$

11. The pixel theorem

In this section, we prove that the cohomology fractal gives rise to a distribution at infinity. That is, integrating against the cohomology fractal then taking the limit as R tends to infinity, gives a continuous linear functional on smooth, compactly supported functions.

11.1. Motivation

Throughout the paper, we have drawn many images of cohomology fractals, always depending on a visual radius R . The obvious question is whether there is a limiting image as R tends to infinity.

It turns out that the answer critically hinges on the question of what a pixel is. As we showed in [Theorem 9.2](#), thinking of a pixel as a sampled point does not work. After realising this, our next thought was that the cohomology fractal might converge to a signed measure μ . We managed to prove this for squares (as well as for regions with piecewise smooth boundary). However our proof does not generalise to arbitrary measurable sets. See [Section 11.35](#) for a discussion.

We finally arrived at the notion of thinking of a pixel as a smooth test function; see [\[38\]](#). The cohomology fractal now assigns to a pixel its weighted “average value”; in other words, we obtain a well-defined *distribution*. This distribution satisfies various transformation laws; these describe how it changes as we alter the chosen cocycle, basepoint, or view. To prove these we rely heavily on the *exponential mixing* of the geodesic flow.

11.2. Background and statement

Before stating the theorem we establish our notation. We define ω , b , D , and T as in [Section 10.1](#). However, as mentioned in [Remark 10.5](#), we switch from using the area measure μ_D to the area two-form ζ_D and from a probability measure ν_D to a compactly supported two-form η_D . To obtain $\eta = \eta_D \in \Omega_c^2(D)$, we set $\eta_D = h \cdot \zeta_D$; here $h \in \Omega_c^0(D)$ is compactly supported and smooth. That is, h is Hodge dual to η .

The function h should be thought of as the kernel function for a pixel. The discussion below could be phrased completely in terms of h . However, using η allows us to neatly express the transformation laws between different views.

Definition 11.3. For a compactly supported two-form $\eta \in \Omega_c^2(D)$, we define

$$\Phi^{\omega,b,D}(\eta) = \lim_{T \rightarrow \infty} \int_D \Phi_T^{\omega,b,D} \cdot \eta$$

As we shall see, $\Phi^{\omega,b,D}$ is a *distribution*: a continuous linear functional on $\Omega_c^2(D)$. We recall the topology on $\Omega_c^2(D)$ in the proof of [Theorem 11.4](#). We will use $\int_D$ to denote the *canonical distribution* $\eta \mapsto \int_D \eta$.

To give a transformation law between views D and E , we will need a way to relate one to the other. Recall that $\tilde{M}$ is isometric to $\mathbb{H}^3$; thus we have $\partial_\infty \tilde{M} \cong \mathbb{CP}^1$. As t tends to infinity, the flow φ_t takes a unit tangent vector $v \in D$ to some point $\varphi_\infty(v) \in \partial_\infty \tilde{M}$. This induces a conformal embedding i_D of D into $\partial_\infty \tilde{M}$. We define i_E similarly. We take $i_{E,D} = i_D^{-1} \circ i_E$ where it is defined. This is a *conformal isomorphism* from (a subset of) E to (a subset of) D . We can now state the main result of this section.

Theorem 11.4 (Pixel theorem). Suppose that M is a connected, orientable, finite volume, complete hyperbolic three-manifold. Fix a closed, compactly supported one-form $\omega \in \Omega_c^1(M)$, a basepoint $b \in \tilde{M}$, and a view D .

1. Then $\Phi^{\omega,b,D}$ is well-defined and is a distribution.
2. Given $\omega' \in \Omega_c^1(M)$ with $[\omega] = [\omega']$, there is a constant C so that we have

$$\Phi^{\omega',b,D} - \Phi^{\omega,b,D} = C \cdot \int_D$$

3. Given another basepoint $b' \in \tilde{M}$, we have

$$\Phi^{\omega,b',D} - \Phi^{\omega,b,D} = \left[\int_b^{b'} \omega \right] \cdot \int_D$$

4. Given another view E and a two-form $\eta \in \Omega_c^2(\text{image}(i_{E,D}))$, we have

$$\Phi^{\omega,b,D}(\eta) = \Phi^{\omega,b,E}(i_{E,D}^* \eta)$$

The last property gives us a distribution at infinity as follows.

Corollary 11.5. Suppose that M is a connected, orientable, finite volume, complete hyperbolic three-manifold. Fix a closed, compactly supported one-form $\omega \in \Omega_c^1(M)$ and a basepoint $b \in \tilde{M}$. Then there is a distribution $\Phi^{\omega,b}$ on $\partial_\infty \tilde{M}$ so that, for any view D and any $\eta \in \Omega_c^2(D)$, we have

$$\Phi^{\omega,b,D}(\eta) = \Phi^{\omega,b}((i_D^{-1})^* \eta)$$

□

11.6. Proof of the pixel theorem

We now describe some background necessary to prove the theorem. Throughout this section, we fix a connected, orientable, finite volume, complete hyperbolic three-manifold M . Again we take $X = \text{UT } M$ and $\tilde{X} = \text{UT } \tilde{M}$. Let

$$Y_t^{\omega, D}(\eta) = \int_D (\omega \circ \varphi_t) \cdot \eta \quad (11.7)$$

Fubini's theorem implies that

$$\int_D \Phi_T^{\omega, b, D} \cdot \eta = \int_D \Phi_0^{\omega, b, D} \cdot \eta + \int_0^T Y_t^{\omega, D}(\eta) dt \quad (11.8)$$

so much of the proof boils down to obtaining exponential decay of $Y_t = Y_t^{\omega, D}(\eta)$.

In this section we will use the Haar measure μ_{Haar} for integrals over $X = \text{UT } M$. We will use the shorthand $dv = d\mu_{\text{Haar}}(v)$ throughout.

We also need to introduce the Sobolev norm $\mathcal{S}_m = \mathcal{S}_{m, \infty}$ for smooth functions f on homogeneous spaces. First consider functions $f: G \rightarrow \mathbb{R}$ where G is a Lie group. Fix a basis for the Lie algebra of G ; we think of the elements in this basis as left-invariant vector fields on G . The Sobolev norm $\mathcal{S}_m(f)$ is the maximum of all L_∞ -norms of functions obtained by differentiating f up to m times using these vector fields in any order. Suppose that Γ and H are (respectively) discrete and compact subgroups of G . The Sobolev norm of $f: \Gamma \backslash G/H \rightarrow \mathbb{R}$ is the Sobolev norm of the lift of f to G .

As usual, we have $M = \Gamma \backslash \mathbb{H}^3 \cong \Gamma \backslash \text{PSL}(2, \mathbb{C}) / \text{PSU}(2)$. Likewise, we have $X = \text{UT } M \cong \Gamma \backslash \text{PSL}(2, \mathbb{C}) / \text{PSO}(2)$. For any of the three views, material, ideal, or hyperideal, we can also express D in this fashion. For example, in the material view we have $D \cong S^2 \cong \text{PSU}(2) / \text{PSO}(2)$.

For $\eta \in \Omega_c^2(D)$, define $\mathcal{S}_m(\eta) = \mathcal{S}_m(h)$ where h is the Hodge dual of η . Note that the Sobolev norm depends on our choice of basis; however, changing the basis changes the resulting norm by a bounded factor and thus only changes the constant C in the following lemma.

Lemma 11.9. *Let M be a connected, orientable, finite volume, complete hyperbolic three-manifold. There is a constant $m \in \mathbb{N}$ such that the following is true. Fix a view D in M and a smooth, compactly supported function $f: X = \text{UT } M \rightarrow \mathbb{R}$ with $\int_X f(v) dv = 0$. Fix a compact set $K \subset D$. There are constants $C > 0$ and $c > 0$ such that for all two-forms $\eta \in \Omega_K^2(D)$ supported in K and for all $t \geq 0$, we have*

$$\left| \int_D (f \circ \varphi_t) \cdot \eta \right| \leq C e^{-ct} \mathcal{S}_m(\eta)$$

To prove this, we use the exponential decay of correlation coefficients for geodesic flows. This is a much studied area. We will rely on [23] because they explicitly give the dependence of the decay on the Sobolev norms of the functions involved. For hyperbolic, finite volume three-manifolds, their theorem can be simplified to the following.

Theorem 11.10. *Let M be a connected, orientable, finite volume, complete hyperbolic three-manifold. Then there exists $m \in \mathbb{N}$, $C > 0$ and $c > 0$ with the following property. For any smooth functions $f, g: X = \text{UT } M \rightarrow \mathbb{R}$ with $\int_X f(v) dv = 0$ and for all $t \geq 0$, we have*

$$\left| \int_X f(\varphi_t(v)) g(v) dv \right| \leq C e^{-ct} \mathcal{S}_m(f) \mathcal{S}_m(g) \quad (11.11)$$

Proof. The more general [23, Theorem 3.1] relates to Theorem 11.10 as follows. They integrate over the frame bundle $\Gamma \backslash \text{PSL}(2, \mathbb{C})$ using the Bowen-Margulis-Sullivan-measure. However, we can think of a function $X \rightarrow \mathbb{R}$ as an $\text{PSO}(2)$ -invariant function on the frame bundle and the BMS-measure is simply the Haar measure in the case of a hyperbolic, finite volume three-manifold $\Gamma \backslash \mathbb{H}^3$. Note that [23, Theorem 3.1] requires the functions f and g to be supported on a unit neighbourhood of the preimage of the convex core of M . However, for finite volume M , the convex core is just M . Furthermore, conventions for the Sobolev norm $\mathcal{S}_m$ differ in whether to take the sum or maximum of the L_∞ -norms of derivatives; however the resulting norms are equivalent because they differ by a constant factor. $\square$

Let $f: X \rightarrow \mathbb{R}$ be a compactly supported, smooth function and $\tilde{f} = f \circ \pi: \tilde{X} \rightarrow \mathbb{R}$ its lift. To prove Lemma 11.9 using Theorem 11.10, we construct test functions $h_\varepsilon: \tilde{X} \rightarrow \mathbb{R}$ that tend to the given two-form $\eta \in \Omega_c^2(D)$ in the sense that

$$Y_t^{f, D}(\eta) = \int_{\tilde{X}} (\tilde{f} \circ \varphi_t) \cdot \eta \quad (11.12)$$

can be approximated by

$$Y_{t, \varepsilon}^{f, D}(\eta) = \int_{\tilde{X}} \tilde{f}(\varphi_t(v)) h_\varepsilon(v) dv \quad (11.13)$$

Note that there are several incompatibilities between η and h_ε :

1. $\eta \in \Omega_c^2(D)$ is a two-form but h_ε has to be a function.
2. $D \subset \tilde{X}$ but the integral [Theorem 11.10](#) is over X .
3. D is two-, not five-dimensional.

The first issue is solved by using the Hodge dual $h \in \Omega_c^0(D)$. That is, $\eta = h \cdot \zeta$ where $\zeta = \zeta_D$ is the area form on D .

For the second issue, we reformulate [Theorem 11.10](#) as follows:

Theorem 11.14. *Let M be a connected, orientable, finite volume, complete hyperbolic three-manifold. There is a constant $m \in \mathbb{N}$ such that the following is true. Fix a smooth, compactly supported function $f: X = \text{UT } M \rightarrow \mathbb{R}$ with $\int_X f(v) dv = 0$ and a compact set $K \subset \tilde{X} = \text{UT } \tilde{M}$. There exists $C > 0$ and $c > 0$ such that for all smooth functions $g: \tilde{X} \rightarrow \mathbb{R}$ supported in K and all $t \geq 0$, we have*

$$\left| \int_{\tilde{X}} \tilde{f}(\varphi_t(v)) g(v) dv \right| \leq C e^{-ct} \mathcal{S}_m(g) \quad (11.15)$$

Proof. Note that

$$\int_{\tilde{X}} \tilde{f}(\varphi_t(v)) g(v) dv = \int_X f(\varphi_t(v)) g_{\Sigma}(v) dv$$

where $g_{\Sigma}(v)$ is the sum of all $g(\tilde{v})$ where $\tilde{v} \in \tilde{X}$ is a preimage of $v \in X$. Since K is covered by a finite number of copies of a fundamental domain of M , the sum $g_{\Sigma}(v)$ has a finite and bounded number of terms. Since f has compact support, $\mathcal{S}_m(f)$ is finite and the result follows from [Theorem 11.10](#). $\square$

To address the third issue, we make the following definition.

Definition 11.16. Define an ε -bump function by

$$b_\varepsilon(x) = \begin{cases} \exp\left(\frac{1}{(x/\varepsilon)^2 - 1}\right) & \text{if } |x| < \varepsilon, \\ 0 & \text{otherwise} \end{cases}$$

and set $B = [\int_{-\infty}^{\infty} b_1(x) dx][\int_0^{\infty} b_1(r) 2\pi r dr]$.

We again use the coordinates on $\tilde{X}$ already introduced in [Section 10.21](#). Recall that $H = H^s(x_u)$. We define $v_s = d_H(x_u, x_s)$. Again, see [Figure 23](#). We now define

$$h_\varepsilon(x_u, x_f, x_s) = h(x_u) \cdot \frac{b_\varepsilon(x_f) b_\varepsilon(v_s)}{B\varepsilon^3} \quad (11.17)$$

Using Fubini's theorem, we can now write

$$\begin{aligned} Y_{t,\varepsilon}^{f,D}(\eta) &= Y_{t,\varepsilon}^{f,D}(h \cdot \zeta) \\ &= \int_D \int_{\mathbb{R}} \int_{H^s(x_u)} \tilde{f}(\varphi_t(v)) h_\varepsilon(v) J_D(v) dx_s dx_f dx_u \end{aligned}$$

where $v = (x_u, x_f, x_s)$, where dx_u and dx_s are using the area measures on D and $H^s(x_u)$, respectively, and where J_D is the smooth function such that

$$d\mu_X = J_D(v) dx_s dx_f dx_u \quad (11.18)$$

Note that, by construction, J_D is invariant under isometries fixing D . In particular, $J_D(x_u) = J_D(x_u, 0, x_u)$ is a positive constant. We set $J_0 = J_D(x_u)$.

Defining

$$Z_{t,\varepsilon}^{f,D}(x_u) = \int_{\mathbb{R}} \int_{H^s(x_u)} \tilde{f}(\varphi_t(v)) J_D(v) \frac{b_\varepsilon(x_f) b_\varepsilon(v_s)}{B\varepsilon^3} dx_s dx_f, \quad (11.19)$$

where again $v = (x_u, x_f, x_s)$, we can write

$$Y_{t,\varepsilon}^{f,D}(\eta) = Y_{t,\varepsilon}^{f,D}(h \cdot \zeta) = \int_D Z_{t,\varepsilon}^{f,D}(x_u) h(x_u) dx_u$$

Lemma 11.20. *For any smooth, compactly supported function $f: X \rightarrow \mathbb{R}$ and any view D , there is a $C > 0$ such that for all $t \geq 0$, for all $1 > \varepsilon > 0$, and for all $x_u \in D$, we have*

$$\left| \tilde{f}(\varphi_t(x_u)) \cdot J_0 - Z_{t,\varepsilon}^{f,D}(x_u) \right| \leq C\varepsilon$$

Proof. Note that the support of h_ε is contained in the closure of D_ε and that the second factor in (11.17) is normalised such that

$$\int_{\mathbb{R}} \int_{H^s(x_u)} \frac{b_\varepsilon(x_f)b_\varepsilon(v_s)}{B\varepsilon^3} dx_s dx_f = 1$$

for all $x_u \in D$. Thus, it suffices to show that there is a $C > 0$ such that for all $t \geq 0$, for all $1 > \varepsilon > 0$, and for all $v = (x_u, x_f, x_s) \in D_\varepsilon$, we have

$$\left| \tilde{f}(\varphi_t(x_u)) \cdot J_0 - \tilde{f}(\varphi_t(v)) \cdot J_D(v) \right| \leq C\varepsilon \quad (11.21)$$

Claim 11.22. Let $f: X \rightarrow \mathbb{R}$ be a smooth, compactly supported function. Let $L = L(f)$ be a Lipschitz constant for f and let $\tilde{f}: \tilde{X} \rightarrow \mathbb{R}$ be its lift. For all $t \geq 0$, for all $1 > \varepsilon > 0$, and for all $v = (x_u, x_f, x_s) \in D_\varepsilon$, we have

$$\left| \tilde{f}(\varphi_t(x_u)) - \tilde{f}(\varphi_t(v)) \right| \leq 2L\varepsilon \quad \text{and} \quad \left| \tilde{f}(\varphi_t(x_u)) \right| \leq \|f\|_\infty$$

Proof. The first inequality follows from Lemma 10.23. The second is by definition. $\square$

Claim 11.23. The function J_D has a finite Lipschitz constant L' when restricted to D_1 . Thus, for all $1 > \varepsilon > 0$ and for all $v = (x_u, x_f, x_s) \in D_\varepsilon$, we have

$$|J_0 - J_D(v)| \leq 2L'\varepsilon$$

Proof. Since J_D is invariant under isometries preserving D , we can assume that x_u is fixed. Then, the domain $(x_u, (-1, 1), B_1^s(x_u))$ has compact closure and thus J_D has a finite Lipschitz constant L' when restricted to it. The claim now follows from Lemma 10.23. $\square$

Using the above two claims, the left hand side of (11.21) is bounded by $2L\varepsilon \cdot J_0 + \|f\|_\infty \cdot 2L'\varepsilon + 2L\varepsilon \cdot 2L'\varepsilon$. Thus setting $C = 2LJ_0 + 2L'\|f\|_\infty + 4LL'$ suffices to prove Lemma 11.20. $\square$

Lemma 11.24. Let M be a connected, orientable, finite volume, complete hyperbolic three-manifold and $f: X \rightarrow \mathbb{R}$ a smooth, compactly supported function. Fix a view D and a compact set $K \subset D$. Let $\zeta = \zeta_D$ be the area form on D . There is a $C > 0$ such that for all smooth $h: D \rightarrow \mathbb{R}$ supported in K and for all $t \geq 0$, $1 > \varepsilon > 0$, we have

$$\left| Y_t^{f,D}(h \cdot \zeta) J_0 - Y_{t,\varepsilon}^{f,D}(h \cdot \zeta) \right| \leq C\varepsilon \|h\|_\infty$$

Proof. We have

$$Y_t^{f,D}(h \cdot \zeta) J_0 - Y_{t,\varepsilon}^{f,D}(h \cdot \zeta) = \int_D \left(\tilde{f}(\varphi_t(x_u)) J_0 - Z_{t,\varepsilon}^{f,D}(x_u) \right) h(x_u) dx_u$$

which is bounded by

$$C\varepsilon \int_D |h(x_u)| dx_u$$

by Lemma 11.20. The result now follows since K has finite area. $\square$

Lemma 11.25. Let M be a connected, orientable, finite volume, complete hyperbolic three-manifold. Fix $m \in \mathbb{N}$. Fix a view D and a compact set $K \subset D$. There is a $C > 0$ such that for all smooth $h: D \rightarrow \mathbb{R}$ supported in K and all $1 > \varepsilon > 0$, we have

$$\mathcal{S}_m(h_\varepsilon) \leq C\varepsilon^{-(m+3)} \mathcal{S}_m(h)$$

where h_ε is as defined in (11.17).

Proof. We estimate $\mathcal{S}_m(h_\varepsilon)$ by using that h_ε is separable as defined in (11.17). In suitable coordinates, the second factor can be written as

$$g_\varepsilon: \mathbb{R}^3 \rightarrow \mathbb{R}, \quad (x, y, z) \mapsto \frac{b_\varepsilon(x)b_\varepsilon(\sqrt{y^2 + z^2})}{B\varepsilon^3}$$

We have $g_\varepsilon(u) = g_1(u/\varepsilon)/\varepsilon^3$ so $\mathcal{S}_m(g_\varepsilon) = \varepsilon^{-(m+3)} \mathcal{S}_m(g_1)$. Using that all h_ε are supported in a common compact set, the lemma follows from the following fact about Sobolev norms.

Recall that a Sobolev norm requires a choice of vector fields that pointwise span the tangent space of the manifold (or a bundle over the manifold when using the Sobolev norm defined earlier by lifting a function $f: \Gamma \backslash G/H \rightarrow \mathbb{R}$ to G). However, any two such choices yield Sobolev norms that differ by a bounded factor when restricting to a small enough neighbourhood or compact set. In particular, up to a bounded factor, we can estimate the Sobolev norm $\mathcal{S}_m(h_\varepsilon)$ by Sobolev norms using local coordinates in which h_ε is separable. $\square$

Proof of Lemma 11.9. Let m be as in Theorem 11.14. Fix a smooth, compactly supported function f , a view D , and a compact $K \subset D$. Theorem 11.14 states that there is a C_0 and c_0 such that for all smooth $h: D \rightarrow \mathbb{R}$ supported in K and all $1 > \varepsilon > 0$, we may set $g = h_\varepsilon$ and have

$$\left| Y_{t,\varepsilon}^{f,D}(h \cdot \zeta) \right| = \left| \int_{\tilde{X}} \tilde{f}(\varphi_t(v)) h_\varepsilon(v) dv \right| \leq C_0 e^{-c_0 t} \mathcal{S}_m(h_\varepsilon)$$

Applying Lemma 11.24 to the left-hand side and Lemma 11.25 to the right hand-side, there are C_1 and C_2 such that

$$\left| Y_t^{f,D}(h \cdot \zeta) J_0 \right| \leq C_1 \varepsilon \|h\|_\infty + C_0 e^{-c_0 t} C_2 \varepsilon^{-(m+3)} \mathcal{S}_m(h)$$

Since this holds for all $0 < \varepsilon < 1$, we can set $\varepsilon = e^{-c_0 t/2(m+3)}$ and obtain

$$\left| Y_t^{f,D}(h \cdot \zeta) \right| \leq \frac{C_1 + C_0 C_2}{J_0} e^{-c_0 t/2(m+3)} \mathcal{S}_m(h) \quad \square$$

Proof of Theorem 11.4. Recall that

$$\Phi^{\omega,b,D}(\eta) = \int_0^\infty Y_t^{\omega,D}(\eta) dt + \int_D \left[\int_b^{\pi(v)} \omega \right] \cdot \eta \quad (11.26)$$

We now discuss the topology on $\Omega_c^2(D)$. Many textbooks on distributions define the topology on $C_c^\infty(\mathbb{R}^n)$. This suffices in the ideal and hyperideal cases, where the view D is diffeomorphic to $\mathbb{R}^2$. When D is a two-sphere, we instead rely on the general theory of distributions on smooth manifolds. We point the reader to [45, Section 2.3] where such distributions are called “generalised sections”. In particular, the definition of $\mathcal{D}(M; E)$ in [45] gives the correct topology on $\Omega_c^0(D) \cong C_c^\infty(D)$ and thus on $\Omega_c^2(D) \cong \Omega_c^0(D)$.

Using [45, Corollary 2.2.1] and the definition of the topology on $\Omega_c^2(D)$, it is not hard to see that statement (1) of the theorem is equivalent to the following claim.

Claim 11.27. For any compact $K \subset D$, there is a $C > 0$ such that for all $\eta \in \Omega_K^2(D)$, the integrals in $\Phi^{\omega,b,D}(\eta)$ exist and we have

$$|\Phi^{\omega,b,D}(\eta)| \leq C \mathcal{S}_m(\eta)$$

Proof. We apply Lemma 11.9 to $f = \omega: X \rightarrow \mathbb{R}$. The lemma implies that there is a $C_0 > 0$ such that for all $\eta \in \Omega_K^2(D)$, the integral $\int_0^\infty Y_t^{\omega,D}(\eta) dt$ exists and is bounded by $C_0 \mathcal{S}_m(\eta)$. Let $C_1 = \int_K \zeta_D$ be the area of K and C_2 be the maximum of $\int_b^{\pi(v)} \omega$ over $v \in K$. Recall that $h \in \Omega_K^0(D)$ denotes the Hodge dual to $\eta = h \cdot \zeta_D$. Then, by Equation 11.26, we have the following.

$$\left| \Phi^{\omega,b,D}(\eta) \right| \leq C_0 \mathcal{S}_m(\eta) + C_1 C_2 \|h\|_\infty \leq (C_0 + C_1 C_2) \mathcal{S}_m(\eta) \quad \square$$

Let $W: \tilde{M} \rightarrow \mathbb{R}$ be a primitive such that $dW = \omega$. We again abbreviate $W \circ \pi: \tilde{X} \rightarrow \mathbb{R}$ as W . We have

$$\begin{aligned} \Phi^{\omega,b,D}(\eta) &= \lim_{T \rightarrow \infty} \int_{v \in D} (W(\varphi_T(v)) - W(b)) \cdot \eta(v) \\ &= \left[\lim_{T \rightarrow \infty} \int_D (W \circ \varphi_T) \cdot \eta \right] - W(b) \cdot \int_D \eta \end{aligned} \quad (11.28)$$

Statement (2) is equivalent to the following claim.

Claim 11.29. If $[\omega] = 0$, there is a C such that $\Phi^{\omega,b,D} = C \cdot \int_D$

Proof. Since ω is a coboundary, its primitive descends to a map $W: M \rightarrow \mathbb{R}$. Add a constant to W to arrange $\int_M W dm = 0$; here we integrate with respect to the volume measure. Now taking $f = W$, Lemma 11.9 now applies to (11.28) and we have $\Phi^{\omega,b,D} = -W(b) \cdot \int_D$. $\square$

Statement (3) follows from Remark 7.6.

It is left to show Statement (4). Recall that we defined maps $i_D: D \rightarrow \partial_\infty \tilde{M}$ and $i_E: E \rightarrow \partial_\infty \tilde{M}$. Let $U = \text{image}(i_D) \cap \text{image}(i_E) \subset \partial_\infty \tilde{M}$. Furthermore, let us define two distributions on U by

$$\Phi_{\leftarrow D}^{\omega,b}(\delta) = \Phi^{\omega,b,D}(i_D^* \delta) \quad \text{and} \quad \Phi_{\leftarrow E}^{\omega,b}(\delta) = \Phi^{\omega,b,E}(i_E^* \delta)$$

For every $\eta \in \Omega_c^2(\text{image}(i_{E,D}))$, there is a compactly supported two-form δ on U such that $\eta = i_D^* \delta$ and $i_{E,D}^* \eta = i_E^* \delta$. Thus, it is enough to show that

$$\Phi_{\leftarrow D}^{\omega,b} = \Phi_{\leftarrow E}^{\omega,b} \quad (11.30)$$




Figure 24. The regions $\varphi_T(i_D^{-1}(N))$ and $\varphi_{T+\Delta T_p}(i_E^{-1}(N))$ for two material views. The distance between corresponding points in the two regions is bounded by a constant for all large enough T . The constant can be made arbitrarily small by making the neighbourhood N about p small enough.

Claim 11.31. For every $p \in U$ and $\varepsilon > 0$, there is a neighbourhood $N \subset U$ of p , a number $\Delta T \in \mathbb{R}$ and $T_0 \geq 0$ such that for all $s \in N$ and $T \geq T_0$, we have

$$|W(\varphi_T(i_D^{-1}(s))) - W(\varphi_{T+\Delta T}(i_E^{-1}(s)))| \leq \varepsilon$$

Remark 11.32. Note that the left hand side can be made arbitrary small by shrinking N . However, for fixed N increasing T might not suffice. This is why we will need to use a partition of unity argument to verify (11.30).

Proof of Claim 11.31. It might be helpful to consult Figure 24. Let L be a Lipschitz constant of $\tilde{W}: X \rightarrow \mathbb{R}$.

Suppose that s lies in $U \subset S^2$. Let γ_D and γ_E be the oriented, pointed, geodesics through $i_D^{-1}(s)$ and $i_E^{-1}(s)$ respectively. Both forward converge to s . Let H be the horosphere centered at s containing $\pi(i_D^{-1}(s))$. Note that γ_E intersects H orthogonally, say at v . Let ΔT_s be the signed distance, along γ_E , from $\pi(i_E^{-1}(s))$ to v . Let $v_s = d_H(v, \pi(i_D^{-1}(s)))$.

Recall that p lies in $U \subset S^2$. Recall that φ_T is the geodesic flow. Exponential convergence and the triangle inequality give the following estimate:

$$d_{\tilde{X}}(\varphi_T(i_D^{-1}(s)), \varphi_{T+\Delta T_p}(i_E^{-1}(s))) \leq \sqrt{2}v_s e^{-T} + |\Delta T_p - \Delta T_s|$$

We refer to Remark 10.2 to explain the factor of $\sqrt{2}$.

There is a neighbourhood N of p such that for all $s \in N$, we have $|\Delta T_p - \Delta T_s| \leq \varepsilon/2L$. Note that v_s is bounded on N so there is a T_0 such that $\sqrt{2}v_s e^{-T_0} \leq \varepsilon/2L$. We deduce that

$$d_{\tilde{X}}(\varphi_T(i_D^{-1}(s)), \varphi_{T+\Delta T_p}(i_E^{-1}(s))) \leq \varepsilon/L$$

Since L is a Lipschitz constant for $\tilde{W}$, the claim follows by setting $\Delta T = \Delta T_p$. □

We now define a norm $\|\cdot\|_1$ on $\Omega_c^2(U)$. Identify $\partial_\infty \tilde{M}$ with S^2 . Such an identification induces an area form ζ_∞ on $\partial_\infty \tilde{M}$. Given $\delta \in \Omega_c^2(U)$, let d be the Hodge dual, that is $\delta = d \cdot \zeta_\infty$. Let

$$\|\delta\|_1 = \int_U |d| \cdot \zeta_\infty$$

Note that $\|\delta\|_1$ does not depend on the identification of $\partial_\infty \tilde{M}$ with S^2 .

Claim 11.33. For every $p \in U$ and $\varepsilon > 0$, there is a neighbourhood $N \subset U$ of p such that for all $\delta \in \Omega_c^2(N)$, we have

$$|\Phi_{\leftarrow D}^{\omega, b}(\delta) - \Phi_{\leftarrow E}^{\omega, b}(\delta)| \leq \varepsilon \|\delta\|_1$$

Proof. Let N , ΔT , and T_0 as in [Claim 11.31](#). Add a constant to the primitive W of ω such that $W(b) = 0$. Then, the difference $\Phi_{\leftarrow D}^{\omega,b}(\delta) - \Phi_{\leftarrow E}^{\omega,b}(\delta)$ can be expressed using W as follows.

$$\begin{aligned} & \lim_{T \rightarrow \infty} \int_U (W \circ \varphi_T \circ i_D^{-1}) \cdot \delta - \lim_{T \rightarrow \infty} \int_U (W \circ \varphi_T \circ i_E^{-1}) \cdot \delta = \\ & \lim_{T \rightarrow \infty} \int_U (W \circ \varphi_T \circ i_D^{-1}) \cdot \delta - \lim_{T \rightarrow \infty} \int_U (W \circ \varphi_{T+\Delta T} \circ i_E^{-1}) \cdot \delta = \\ & \lim_{T \rightarrow \infty} \int_U (W \circ \varphi_T \circ i_D^{-1} - W \circ \varphi_{T+\Delta T} \circ i_E^{-1}) \cdot \delta \end{aligned}$$

[Claim 11.31](#) implies that the integral is bounded by $\varepsilon \|\delta\|_1$, for all $T_0 \geq T$. Thus the limit is bounded by $\varepsilon \|\delta\|_1$. $\square$

We now verify [\(11.30\)](#). Fix a smooth, compactly supported $\delta \in \Omega_c^2(U)$. Also fix $\varepsilon > 0$. For each $p \in U$, pick a neighbourhood N as in [Claim 11.33](#). The support $\text{supp}(\delta)$ can be covered by finitely many of these neighbourhoods. Consider a partition of unity with respect to this finite cover. By multiplying δ by the partition functions, we obtain smooth two-forms $\delta_1, \dots, \delta_n$. Let d_i be the Hodge dual of δ_i . Since $\sum |d_i| = |d|$ (pointwise) we have $\sum \|\delta_i\|_1 = \|\delta\|_1$. Thus [Claim 11.33](#) implies

$$\left| \Phi_{\leftarrow D}^{\omega,b}(\delta) - \Phi_{\leftarrow E}^{\omega,b}(\delta) \right| \leq \varepsilon \|\delta\|_1$$

Since ε was arbitrary, we are done. $\square$

Remark 11.34. McMullen [29] suggests another approach to [Theorem 11.4](#). Arrange matters so that W , the primitive of ω , is harmonic with respect to the hyperbolic metric. Prove that W has suitably bounded growth as it approaches $\partial_\infty \mathbb{H}^3$, in terms of the euclidean metric in the Poincaré ball model. Now prove and then apply an appropriate version (of part (iii) implies part (i)) of [\[40, Theorem 1.1\]](#).

11.35. The cohomology fractal as a measure

We do not know whether the cohomology fractal converges to a signed measure; that is we do not know whether or not

$$\lim_{T \rightarrow \infty} \int_U \Phi_T^{\omega,b,D} \cdot d\mu_D$$

converges for every measurable $U \subset D$. Instead, we have the following partial result.

Theorem 11.36 (Square pixel theorem). *Suppose that M is a connected, orientable, finite volume, complete hyperbolic three-manifold. Fix a closed, compactly supported one-form $\omega \in \Omega_c^1(M)$, a basepoint $b \in \tilde{M}$, and a view D . Suppose that $U \subset D$ is bounded. Suppose further that ∂U has finite length. Then the following limit exists*

$$\lim_{T \rightarrow \infty} \int_U \Phi_T^{\omega,b,D} \cdot d\mu_D$$

Remark 11.37. Recall that the pictures in [Figure 18](#) were generated by uniformly sampling across a pixel square. [Theorem 11.36](#) finally proves that this technique (with enough samples) will give accurate images in some non-empty range of visual radii. Note that the earlier [Theorem 11.4](#) is not sufficient; it required a smooth filter function.

Remark 11.38. It seems difficult to generalise [Theorem 11.36](#) to measurable subsets. Our proof does not apply, for example, to an open set $U \subset D$ bounded by several Osgood arcs [\[34\]](#).

Before sketching a proof of [Theorem 11.36](#), we discuss *mollification*. Suppose that $h: D \rightarrow \mathbb{R}$ is a bounded measurable function with compact support. We define $h_\varepsilon: D \rightarrow \mathbb{R}$ by setting

$$h_\varepsilon(u) = \frac{1}{B_\varepsilon} \int_D h(v) \cdot b_\varepsilon(d_D(u, v)) \cdot d\mu_D(v) \quad (11.39)$$

Here B_ε is a normalisation factor such that $\int_D h \cdot d\mu_D = \int_D h_\varepsilon \cdot d\mu_D$. We call h_ε the ε -mollification (in the unstable direction) of h .

Lemma 11.40. *Let $h: D \rightarrow \mathbb{R}$ be a bounded measurable function with compact support. Assume that there are constants $c > 0$ and $C > 0$ such that for all $1 > \varepsilon > 0$, we have*

$$\|h - h_\varepsilon\|_1 \leq C\varepsilon^c$$

Then the following limit exists

$$\lim_{T \rightarrow \infty} \int_D \Phi_T^{\omega,b,D} \cdot h \cdot d\mu_D$$

Proof sketch. We will show that $Y_t^{\omega,D}(h \cdot d\mu_D) = \int_D (\omega \circ \varphi_t) \cdot h \cdot d\mu_D$ decays exponentially. Let $\eta_\varepsilon = h_\varepsilon \cdot \zeta_D$ and $Y_{t,\varepsilon}^{\omega,D}(h \cdot d\mu_D) = \int_D (\omega \circ \varphi_t) \cdot \eta_\varepsilon$. Regarding ω as function $X \rightarrow \mathbb{R}$, we have

$$\left| Y_t^{\omega,D}(h \cdot d\mu_D) - Y_{t,\varepsilon}^{\omega,D}(h \cdot d\mu_D) \right| \leq \|h - h_\varepsilon\|_1 \cdot \|\omega\|_\infty \leq C\varepsilon^c \|\omega\|_\infty$$

Since h is bounded, the Sobolev norm of the mollification h_ε of h can be estimated from the Sobolev norm of the mollification kernel. This can be done similarly to [Lemma 11.25](#). Thus, there is a $C_0 > 0$ such that for all $1 > \varepsilon > 0$, we have $\mathcal{S}_m(\eta_\varepsilon) \leq C_0 \varepsilon^{-(m+2)}$. Taking $f = \omega$ and $\eta = \eta_\varepsilon$, [Lemma 11.9](#) states that there is a $C_1 > 0$ and $c_1 > 0$ such that for all $t \geq 0$ and $1 > \varepsilon > 0$, we have

$$\left| Y_{t,\varepsilon}^{\omega,D}(h \cdot d\mu_D) \right| \leq C_1 e^{-c_1 t} \varepsilon^{-(m+2)}$$

Thus, we have

$$\left| Y_t^{\omega,D}(h \cdot d\mu_D) \right| \leq C\varepsilon^c \|\omega\|_\infty + C_1 e^{-c_1 t} \varepsilon^{-(m+2)}$$

We obtain exponential decay when setting $\varepsilon = e^{-c_1 t/2(m+2)}$. □

Proof sketch of Theorem 11.36. The theorem follows from [Lemma 11.40](#) by taking h to be the indicator function χ_U . It is left to show that there is a C such that $\|h - h_\varepsilon\|_1 \leq C\varepsilon$.

Note that $h - h_\varepsilon$ is bounded by one, thus it is enough to show that the area where $h - h_\varepsilon$ is non-trivial is bounded by $C\varepsilon$. The area of this neighbourhood is bounded above, up to multiplication by a universal constant, by the length of ∂U multiplied by ε . □

12. Questions and projects

Question 12.1. Suppose that F is a surface in a hyperbolic three-manifold M . [Theorem 10.7](#) tells us that the standard deviation σ is a topological invariant of the pair (M, F) . What are the number theoretic (or other) properties of σ ? Fixing M , does σ “see” the shape of the Thurston norm ball?

Question 12.2. Suppose that F is a fibre of a closed, connected hyperbolic surface bundle M . In [Proposition 6.2](#) we showed that approximations of the Cannon–Thurston map are (components of) level sets of the cohomology fractal. Is there some more precise sense in which the Cannon–Thurston map Ψ is a “level set” of the distributional cohomology fractal $\Phi^{F,b}$?

Question 12.3. Can the cohomology class $[\omega]$ be recovered from the distributional cohomology fractal $\Phi^{\omega,b}$?

Question 12.4. [Figure 19](#) suggests that in the example of `m122 (4, -1)` the mean has settled down at around $R = 10$ for a pixel size of 0.1° . We also see this in [Figure 18](#) in that there is hardly any difference between the images at $R = 10$ and $R = 12$ with 128×128 samples. [Theorem 11.36](#) tells us that given enough samples we can produce an accurate picture of the distributional cohomology fractal. Can one calculate effective bounds that would allow us to produce a provably correct image?

Question 12.5. In [Figure 18](#), the image with 1×1 samples and $R = 8$ is very similar to the image with 128×128 samples and $R = 12$. However, we do not understand how an image generated with only one sample per pixel can so closely approximate the limiting object. The manifold `m122 (4, -1)` is small; as a result, perhaps the geodesic flow mixes rapidly enough? Does this fail in larger manifolds?

Question 12.6 (Mark Pollicott). We consider lowering the dimension of F and M by one. Let F be a non-separating curve in a closed, connected hyperbolic surface M . Fix a point $p \in M$. Let P be a “pixel” – that is, a closed arc in $UT_p M \cong S^1$ with centre c_P and radius r_P . The distributional cohomology fractal Φ^F exists and in fact gives a “signed measure” to each such pixel P (see [Section 11.35](#)). We define a function $\mu^F: S^1 \times (0, \pi] \rightarrow \mathbb{R}$ by setting $\mu^F(c_P, r_P) = \Phi^F(P)$. What does the graph of μ^F look like? For example, what happens if we fix c_P and allow the radius to vary? How does the graph behave as r_P approaches zero?

Question 12.7. The histogram in [Figure 20\(a\)](#) is low near the mean. Increasing R (within the range that we trust our experiments, see [Section 8.6](#)) reduces, but does not remove, this gap. Why is it there? (This does not seem to happen in [Figure 20\(b\)](#), where the surface is closed but the manifold is not.)

Question 12.8. Consider the experiment shown in [Figure 22](#). Here the support of the cocycle ω is not compact. We see that the distribution of the cohomology fractal, over a pixel, appears to not be normal. Further experiments show that it depends sensitively on the choice of pixel. Can one verify rigorously that it is not normal?

We suspect that some version of “subtracting the largest excursion” (see the remarks immediately before [\[16, Theorem 1\]](#)) will yield a more reasonable distribution.

Question 12.9. Theorem 11.4 applies to cocycles ω with compact support. Consider a cusped manifold and a cocycle ω such that the pullback of $[\omega]$ to the cusp torus is non-trivial. In this case, the Sobolev norm of ω is infinite; thus Theorem 11.10 does not apply. Are there modifications, perhaps as indicated in Question 12.8, so that we again obtain a distribution at infinity for the cohomology fractal?

Question 12.10. In the fibred case, what is the relationship between the cohomology fractal and the lightning curve? See Section 6.4.

We end with some ideas for future software projects.

Project 12.11. One could use material triangulations in the closed case to draw an approximation Ψ_D to the Cannon–Thurston map, following Algorithm 4.2. By Proposition 6.2, these match the cohomology fractal. Motivated by Figures 11 and 14, we anticipate that Ψ_D will look significantly different from Cannon–Thurston maps in the cusped case; the “mating dendrites” that approximate Ψ have bounded branching at all points.

Project 12.12. In Section 8, we discussed cohomology fractals for incomplete structures along a line in Dehn surgery space. As discussed in Section 8.3, all of these suffer from numerical defects along the incompleteness locus Σ_s . These defects are visible (although small) in Figure 10. When the slope s is integral, we use material triangulations (Section 8.4) to remove these defects. For general s , material triangulations are not available. Instead, one could *accelerate* through tubes about Σ_s . That is, we modify the cellulation of the manifold by truncating each tetrahedron, replacing the lost volume with a solid torus “cell” around Σ_s .

Appendix A: Notation

For the convenience of the reader, we list some of the notation used in the paper.

$m004$	Snappy notation (for the figure-eight knot complement). We use $m004(p, q)$ to denote a (p, q) Dehn filling.
M	connected, oriented, finite volume hyperbolic three-manifold.
$\mathcal{T}$	triangulation of M , see Section 2.
F	connected, oriented, finite-area hyperbolic surface.
$\tilde{F}, \tilde{M}$	universal covers.
$\partial_\infty \tilde{F}$	ideal boundary of $\tilde{F}$.
$\partial_\infty \tilde{M}$	ideal boundary of $\tilde{M}$.
p, q, b	points of M or $\tilde{M}$.
Ψ	Cannon–Thurston map (Theorem 3.1).
Ψ_D	approximation to the Cannon–Thurston map (Algorithm 4.2).
UT	unit tangent bundle.
u, v	unit vectors in $UT M$ or $UT \tilde{M}$.
π	unit tangent bundle map (or ratio of circumference to diameter of circle in the euclidean plane).
ω	one-cocycle for M . We often replace ω by a (Poincaré dual) surface F in our notation.
R, T	radius and a time respectively. These serve the same purpose; we use one or the other as we are thinking geometrically (Definition 5.1) or dynamically (Definition 5.5).
$\Phi_R^{\omega, p}$	the cohomology fractal at radius R (Definition 5.1). We suppress superscripts which are not relevant to the discussion at hand. Several variants of this notation follow.
$\Phi_R^{F, p}$	(Definition 5.2).
$\Phi_R^{\omega, b, p}$	(Remark 7.5).
$\Phi_R^{\omega, b, D}$	(Definition 7.12 and Equation 10.3).
$\Phi_R^{F, b, D}$	(Section 9).
$\tilde{\omega}$	lift of ω to $\tilde{M}$ (Definition 5.4).
W	primitive for $\tilde{\omega}$ (Definition 5.4 and Section 10.1).
φ_t	geodesic flow for time t (Definition 5.5).
D, E	views in $UT \tilde{M}$ (Section 7.7).
U	subdomain of D : for example, a pixel.
$X, \tilde{X}$	abbreviations for $UT M$ and $UT \tilde{M}$ in Sections 10 and 11.
μ_D	induced area measure on D (Section 10.1).
ζ_D	area form on D (so $\zeta_D = d\mu_D$).
h	measurable (Equation 10.4) or smooth (Section 11.2) test function from D to $\mathbb{R}$.
h_ε	ε -mollification of h in $\tilde{X}$ (Equation 11.17) or in the unstable direction (Equation 11.39).
$\nu \ll \mu$	ν (always a probability measure) absolutely continuous with respect to μ (Section 10.1). We decorate with the measure space (as a subscript) when needed for context.
R_T, S_T	random variable, usually $R_T = \Phi_T / \sqrt{T}$ (Sections 10.6 and 10.24).
σ	standard deviation (square root of the variance).
ψ_σ	normal distribution with standard deviation σ (Section 10.6).

n_σ	probability density function for ψ_σ (Section 10.6).
$\Rightarrow$	convergence in measure (Definition 10.8).
$\rightarrow$	convergence in probability (Definition 10.15).
μ_X	Haar measure on X scaled to be a probability measure (Section 10.10).
x_u, x_f, x_s	coordinates for the unstable, flow, and stable manifolds in a neighbourhood of a view (Section 10.21).
$H^S(x_u)$	stable manifold through x_u (Section 10.21).
d_N	induced distance on a submanifold N (Section 10.21).
$B_\varepsilon^S(x_u)$	ε -ball in $H^S(x_u)$ (Section 10.21).
D_ε	product ε -neighbourhood of the view D (Section 10.21).
η	two-form in $\Omega_c^2(D)$ (Section 11.2).
$\Phi^{\omega,b,D}$	the distributional cohomology fractal (Definition 11.3).
$\Phi^{\omega,b}$	the cohomology fractal at infinity (Corollary 11.5).
$\int_D$	canonical distribution (Section 11.2).
i_D	conformal embedding of D into $\partial_\infty \tilde{M}$ (Section 11.2).
$i_{D,E}$	conformal isomorphism from (a subset of) E to (a subset of) D (Section 11.2).
$Y_t^{\omega,D}(\eta)$	slice of the integral defining $\Phi^{\omega,b,D}(\eta)$ (Equation 11.7).
$Y_t^{f,D}(\eta)$	$Y_t^{\omega,D}(\eta)$ generalised to $f: X = \text{UT } M \rightarrow \mathbb{R}$ (Equation 11.12).
$Y_{t,\varepsilon}^{f,D}(\eta)$	mollified slice (Equation 11.13).
$Z_{t,\varepsilon}^{f,D}(x_u)$	mollified sample (Equation 11.19).
S_m	Sobolev norm (Section 11.6).
f, g	functions in the mixing theorems (Lemma 11.9, Theorem 11.10, and Theorem 11.14).
b_ε, B	bump function and normalising factor (Definition 11.16).
J_D	ratio between $d\mu_X$ and $dx_s dx_f dx_u$ (Equation 11.18).
J_0	equals $J_D(x_u, 0, x_u)$ (Equation 11.18).

Acknowledgments

This work is in the public domain except for Figures 4 and 6.

References

- [1] Roger C. Alperin, Warren Dicks, and Joan Porti. The boundary of the Gieseking tree in hyperbolic three-space. *Topology Appl.*, 93(3):219–259, 1999. doi:10.1016/S0166-8641(97)00270-8.
- [2] Christopher J. Bishop and Peter W. Jones. The law of the iterated logarithm for Kleinian groups. In *Lipa's legacy (New York, 1995)*, volume 211 of *Contemp. Math.*, pages 17–50. Amer. Math. Soc., Providence, RI, 1997. doi:10.1090/conm/211/02813.
- [3] Brian H. Bowditch. The Cannon–Thurston map for punctured-surface groups. *Math. Z.*, 255(1):35–76, 2007. doi:10.1007/s00209-006-0012-4.
- [4] David Bachman, Saul Schleimer, and Henry Segerman. Cohomology fractals. In Carolyn Yackel, Robert Bosch, Eve Torrence, and Kristóf Fenyvesi, editors, *Proceedings of Bridges 2020: Mathematics, Art, Music, Architecture, Education, Culture*, pages 175–182, Phoenix, Arizona, 2020. Tessellations Publishing. <http://archive.bridgesmathart.org/2020/bridges2020-175.html>, arXiv:2002.00239.
- [5] David Bachman, Saul Schleimer, and Henry Segerman. Cohomology fractals, 2020. https://github.com/henryseg/cohomology_fractals.
- [6] Vladimir Bulatov. 3d hyperbolic tiling and horosphere cross section, 2018. Joint AMS/MAA meeting, San Diego, January 10, 2018. <http://bulatov.org/math/180110/index.html>.
- [7] Danny Calegari. *Foliations and the geometry of 3-manifolds*. Oxford Mathematical Monographs. Oxford University Press, 2007. <http://math.uchicago.edu/~dannyc/books/foliations/foliations.html>.
- [8] Andrew Casson. geo, software for geometrizing three-manifolds. <https://faculty.math.illinois.edu/~nmd/computop/software/casson.tar.gz>.
- [9] James W. Cannon and Warren Dicks. On hyperbolic once-punctured-torus bundles. *Geom. Dedicata*, 94:141–183, 2002. doi:10.1023/A:1020956906487.
- [10] James W. Cannon and Warren Dicks. On hyperbolic once-punctured-torus bundles. II. Fractal tessellations of the plane. *Geom. Dedicata*, 123:11–63, 2006. doi:10.1007/s10711-006-9070-3.
- [11] Marc Culler, Nathan M. Dunfield, Matthias Goerner, and Jeffrey R. Weeks. SnapPy, software for studying the geometry and topology of three-manifolds, (version 3.0). <http://snappy.computop.org>.
- [12] Mikael Hviðfeldt Christensen. Building 4d polytopes. <https://syntopia.github.io/Polytopia/polytopes.html>.
- [13] James W. Cannon and William P. Thurston. Group invariant Peano curves. *Geom. Topol.*, 11:1315–1355, 2007. doi:10.2140/gt.2007.11.1315.
- [14] Warren Dicks and Joan Porti. On the Hausdorff dimension of the Gieseking fractal. *Topology Appl.*, 126(1-2):169–186, 2002. doi:10.1016/S0166-8641(02)00077-9.
- [15] Warren Dicks and Makoto Sakuma. On hyperbolic once-punctured-torus bundles. III. Comparing two tessellations of the complex plane. *Topology Appl.*, 157(12):1873–1899, 2010. arXiv:0811.1678, doi:10.1016/j.topol.2010.02.010.
- [16] Harold G. Diamond and Jeffrey D. Vaaler. Estimates for partial sums of continued fraction partial quotients. *Pacific J. Math.*, 122(1):73–82, 1986. doi:10.2140/pjm.1986.122.73.
- [17] Warren Dicks and David J. Wright. On hyperbolic once-punctured-torus bundles IV: Automata for lightning curves. *Topology Appl.*, 159(1):98–132, 2012. doi:10.1016/j.topol.2011.08.018.
- [18] Jacques Franchi and Yves Le Jan. *Hyperbolic dynamics and Brownian motion*. Oxford University Press, 2012. doi:10.1093/acprof:oso/9780199654109.001.0001.

- [19] Matthias Goerner. Verified computations for closed hyperbolic 3-manifolds. *Bull. Lond. Math. Soc.*, 53(2):596–618, 2021. [arXiv:1904.12095](https://arxiv.org/abs/1904.12095), doi:10.1112/blms.12445.
- [20] Damian Heard. Orb, software for studying hyperbolic three-orbifolds and three-manifolds (2018-06-30). <https://github.com/DamianHeard/orb>.
- [21] Damian Heard. *Computation of hyperbolic structures on 3-dimensional orbifolds*. PhD thesis, University of Melbourne, 2005. <https://github.com/DamianHeard/orb-thesis/>.
- [22] Eberhard Hopf. Statistik der geodätischen Linien in Mannigfaltigkeiten negativer Krümmung. *Ber. Verh. Sächs. Akad. Wiss. Leipzig Math.-Phys. Kl.*, 91:261–304, 1939.
- [23] Dubi Kelmer and Hee Oh. Exponential mixing and shrinking targets for geodesic flow on geometrically finite hyperbolic manifolds. 2018. [arXiv:1812.05251v3](https://arxiv.org/abs/1812.05251v3).
- [24] Marc Lackenby. Taut ideal triangulations of 3-manifolds. *Geom. Topol.*, 4:369–395, 2000. [arXiv:math/0003132](https://arxiv.org/abs/math/0003132), doi:10.2140/gt.2000.4.369.
- [25] Jos Leys. A fast algorithm for limit sets of kleinian groups with the Maskit parametrisation, 2017. http://www.josleys.com/articles/Kleinian%20escape-time_3.pdf.
- [26] Peter Liepa. Visualizing hyperbolic symmetries using WebGL (aka a hyperbolic doodler in your browser). <https://brainjam.home.blog/2011/02/19/visualizing-hyperbolic-symmetries-using-webgl/>.
- [27] Curtis T. McMullen. lim. <http://people.math.harvard.edu/~ctm/programs/home/prog/lim/src/lim.tar>.
- [28] Curtis T. McMullen. Figure-eight orbifold fiber, 1995. <http://people.math.harvard.edu/~ctm/gallery/index.html>.
- [29] Curtis T. McMullen. 2019–2020. personal communication.
- [30] Abdelaziz Nait Merzouk. Knighty collection. <https://github.com/Syntopia/Fragmentarium/tree/master/Fragmentarium-Source/Examples/Knighty%20Collection>.
- [31] Mahan Mj. Cannon–Thurston maps. In *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. II. Invited lectures*, pages 885–917. World Sci. Publ., Hackensack, NJ, 2018. [arXiv:1712.00760](https://arxiv.org/abs/1712.00760), doi:10.1142/11060.
- [32] David Mumford, Caroline Series, and David Wright. *Indra’s pearls*. Cambridge University Press, New York, 2002. The vision of Felix Klein. doi:10.1017/CBO9781107050051.024.
- [33] Roice Nelson and Henry Segerman. Visualizing hyperbolic honeycombs. *Journal of Mathematics and the Arts*, 11(1):4–39, 2017. [arXiv:1511.02851](https://arxiv.org/abs/1511.02851), doi:10.1080/17513472.2016.1263789.
- [34] William F. Osgood. A Jordan curve of positive area. *Trans. Amer. Math. Soc.*, 4(1):107–112, 1903. doi:10.2307/1986455.
- [35] Jessica S. Purcell. *Hyperbolic knot theory (Graduate Studies in Mathematics)*. American Mathematical Society, Providence, RI, 2020. [arXiv:2002.12652](https://arxiv.org/abs/2002.12652).
- [36] Mary Rees. Checking ergodicity of some geodesic flows with infinite Gibbs measure. *Ergodic Theory Dynam. Systems*, 1(1):107–133, 1981. doi:10.1017/S0143385700001206.
- [37] Ja. G. Sinaĭ. The central limit theorem for geodesic flows on manifolds of constant negative curvature. *Soviet Math. Dokl.*, 1:983–987, 1960.
- [38] Alvy Smith. A pixel is not a little square, a pixel is not a little square, a pixel is not a little square, 1995. http://alvyray.com/Memos/CG/Microsoft/6_pixel.pdf.
- [39] Peter Stampfli. Fractal images from multiple inversion in circles. In Susan Goldstine, Douglas McKenna, and Kristóf Fenyvesi, editors, *Proceedings of Bridges 2019: Mathematics, Art, Music, Architecture, Education, Culture*, pages 263–270, Phoenix, Arizona, 2019. Tessellations Publishing. <http://archive.bridgesmathart.org/2019/bridges2019-263.pdf>.
- [40] Emil J. Straube. Harmonic and analytic functions admitting a distribution boundary value. *Ann. Scuola Norm. Sup. Pisa Cl. Sci. (4)*, 11(4):559–591, 1984. http://www.numdam.org/item?id=ASNSP_1984_4_11_4_559_0.
- [41] Dennis Sullivan. The density at infinity of a discrete group of hyperbolic motions. *Inst. Hautes Études Sci. Publ. Math.*, (50):171–202, 1979. doi:10.1007/BF02684773.
- [42] William P. Thurston. *Geometry and topology of 3-manifolds*. 1978. <http://msri.org/publications/books/gt3m/>.
- [43] William P. Thurston. Three-dimensional manifolds, kleinian groups and hyperbolic geometry. *Bull. Amer. Math. Soc. (N.S.)*, 6(3):357–381, 1982. doi:10.1090/S0273-0979-1982-15003-0.
- [44] William P. Thurston. A norm for the homology of 3-manifolds. *Mem. Amer. Math. Soc.*, 59(339):i–vi and 99–130, 1986.
- [45] Erik van den Ban and Marius Crainic. Analysis on manifolds. <https://webpace.science.uu.nl/~crain101/AS-2013/chapters-1-2-3.pdf>, <https://webpace.science.uu.nl/~ban00101/anman2009/lecture2.pdf>.
- [46] Masaaki Wada. OPTi, software for visualising quasi-conformal deformations of once-punctured torus groups (version 4.0.5). <https://apps.apple.com/us/app/opti/id1111993523>.
- [47] William Worden. tnrm, software for computing the thurston norm unit ball of finite volume orientable hyperbolic three-manifolds, 2018–2020. (version 0.1.14). <https://pypi.python.org/pypi/tnrm>.
- [48] David Wright. 2019. personal communication.
- [49] Roland Zweimüller. Infinite measure preserving transformations with compact first regeneration. *J. Anal. Math.*, 103:93–131, 2007. doi:10.1007/s11854-008-0003-y.