Comment. Math. Helv. 97 (2022), 729–768
DOI 10.4171/CMH/539

Commentarii Mathematici Helvetici

A model for random three-manifolds

Bram Petri and Jean Raimbault

Abstract. We study compact three-manifolds with boundary obtained by randomly gluing together truncated tetrahedra along their faces. We prove that, asymptotically almost surely as the number of tetrahedra tends to infinity, these manifolds are connected and have a single boundary component. We prove a law of large numbers for the genus of this boundary component, we show that the Heegaard genus of these manifolds is linear in the number of tetrahedra and we bound their first Betti number.

We also show that, asymptotically almost surely as the number of tetrahedra tends to infinity, our manifolds admit a unique hyperbolic metric with totally geodesic boundary. We prove a law of large numbers for the volume of this metric, prove that the associated Laplacian has a uniform spectral gap and show that the diameter of our manifolds is logarithmic as a function of their volume. Finally, we determine the Benjamini–Schramm limit of our sequence of random manifolds.

Mathematics Subject Classification (2020). 57K32; 60C05.

Keywords. Hyperbolic 3-manifold, random simplicial complex.

1. Introduction

1.1. Context. Random constructions of compact manifolds can be seen as an analogue of the well-established theory of random graphs and serve similar purposes. First of all, they make the notion of a "typical" manifold rigorous. Secondly, they can be used as a testing ground for conjectures of which the proof is still out of reach. Finally, there is what is often called the *probabilistic method* – using probability theory to prove the existence of objects with extremal properties. In this paper we are mostly interested in the first aspect.

Let us be more specific about what kind of objects we are interested in. As is the case for graphs, there are countably many homeomorphism types of compact manifolds. Thus, a random manifold consists not in one random variable but rather a family of random variables – say, $M_n, n \geq 1$ – where n is some measure of "complexity", usually in relation with a particular construction that is used to define the random objects. For graphs this will often be the number of vertices. Random models for three-manifolds that have been well-studied are random Heegaard

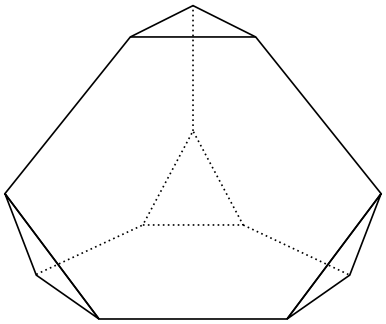

Figure 1. A truncated tetrahedron. M_n is built by randomly gluing n copies of this polyhedron together along their hexagonal faces.

splittings and random fibered manifolds; here the complexity depends on two integers: the genus g of the handlebody or the fiber, and the number of steps k used to generate a random mapping class [20, Section 2.10]. A basic property should be that the union of the support of the M_n is the whole set of the manifolds one is interested in studying. Random Heegaard splittings satisfy this requirement if one takes both $k, g \to +\infty$, but the asymptotic results pertaining to them (in particular, hyperbolicity) are mostly studied in terms of the mapping class (that is, when $k \to +\infty$). If one is interested in studying typical three-manifolds this does not seem satisfying.

A more direct measure of the complexity of a three-manifold is the minimal number of tetrahedra needed to triangulate it. A natural way to construct random three-manifolds is thus to start with a model for a random triangulation on n tetrahedra and condition it to be a manifold. However, studying such a model of a random manifold is hard because if one randomly glues the faces of n tetrahedra together in pairs, the probability that the result is a manifold tends to 0 as $n \to \infty$ (see, for instance, [20, Proposition 2.8]). So we cannot rely on the study of a generic triangulation to establish "asymptotically almost surely" properties of the manifolds and we have to instead study probabilities conditioned on a set of conditions that is hard to manage. We note that even counting the number of triangulations is a hard problem (the best known bounds we are aware of are due to Chapuy–Perarnau [15]).

We will not address this issue in this paper, but instead consider the compact three-manifolds with boundary that are naturally associated with random triangulations. The only points in a 3-dimensional triangulation that might not admit neighbourhoods homeomorphic to open sets in $\mathbb{R}^3$ are the vertices. As such, we obtain a random three-manifold with boundary by randomly gluing together n tetrahedra that are truncated near the vertices (see Figure 1). Moreover, all compact three-manifolds with non-empty boundary can be obtained in this way (see, for example, [17, Corollary 1.3]).

Our goal will be to understand the asymptotic behavior as $n \to +\infty$ of geometric and topological properties of M_n. We are particularly interested in finding properties whose probability of occurrence is asymptotically 1 (for regular graphs this can for instance be connectivity or expansion, depending on the model). For three-manifolds the most obvious candidate for such a property is hyperbolicity. We will prove that, asymptotically almost surely, our manifolds are hyperbolic (with totally geodesic boundary) with volume proportional to the number of tetrahedra, their Heegaard genus goes to infinity, and get estimates on their Betti numbers. We also prove some finer results about their geometry: they are expanders and we show that they converge to an explicit limit in a probabilistic version of the Gromov–Hausdorff topology.

1.2. Results. We will impose one further condition: we condition on two tetrahedra sharing at most one face and every face being incident to two distinct tetrahedra. This is strictly weaker than asking that the complex be simplicial so it is not a very unnatural constraint, though our main reason for working with this assumption is a technical one: we need it in the proof of Lemma 3.6. The resulting random manifold will be called M_n. A detailed description of the model can be found in Section 2.1.

The first question now is what the topology of the resulting manifold is. We prove the following theorem.

Theorem 1.1 (Topology). (a) *We have*

$$\lim_{n\to\infty} \mathbb{P}[M_n \text{ is connected and has a single boundary component}] = 1.$$

(b) *The genus $g(\partial M_n)$ of the boundary of M_n satisfies*

$$\lim_{n\to+\infty} \mathbb{P}[n - \theta(n) \le g(\partial M_n) \le n + 1] = 1,$$

for any function $\theta\colon \mathbb{N} \to \mathbb{R}$ that grows super-logarithmically, i.e.

$$\lim_{n\to\infty} \frac{\theta(n)}{\log(n)} = +\infty.$$

(c) *Let $\mathcal{D}M_n$ denotes the double of M_n along its boundary and $g(\mathcal{D}M_n)$ its Heegard genus. Then*

$$\lim_{n\to+\infty} \mathbb{P}[n - \theta(n) \le g(\mathcal{D}M_n) \le n + \theta(n)] = 1,$$

for any function $\theta\colon \mathbb{N} \to \mathbb{R}$ that grows super-logarithmically.

(d) *There exists C such that the Betti numbers $b_1(M_n)$ and $b_1(M_n, \partial M_n)$ satisfy*

$$\lim_{n\to+\infty} \mathbb{P}[b_1(M_n, \partial M_n) \le \theta(n)] = 1, \quad \lim_{n\to+\infty} \mathbb{P}[\,|b_1(M_n) - n| \le \theta(n)] = 1$$

for any function $\theta\colon \mathbb{N} \to \mathbb{R}$ that grows super-logarithmically.

This is a combination of Corollary 2.2, and Sections 2.5 and 2.6. Moreover, in Theorem 2.4 below we prove various combinatorial properties of the interior edges in our random complex. In item (c) above we look at the Heegaard genus of the double rather than the usual notion of Heegaard genus of the manifold itself (defined in terms of decompositions with compression bodies, cf. [41, Section 2.2]) because the latter is bounded below by the genus of the boundary, so (c) says something that is not covered by (b).

In low dimensions it turns out that typical objects are often hyperbolic and we prove that this is satisfied by our model. Note that it follows from Mostow rigidity that if M_n carries a hyperbolic metric with totally geodesic boundary, then this metric is unique up to isometry. As such, one can also ask for the geometric properties of this metric. We prove the following properties (recall that the *systole* $\mathrm{sys}(X)$ of a compact Riemannian manifold X is the smallest length of a closed geodesic; which does not include lengths of arcs with endpoints on the boundary).

Theorem 1.2 (Geometry). *We have*

$$\lim_{n\to+\infty} \mathbb{P}[M_n \text{ carries a hyperbolic metric with totally geodesic boundary}] = 1.$$

This metric has the following properties:

(a) *The hyperbolic volume* $\mathrm{vol}(M_n)$ *of* M_n *satisfies*[1]

$$\mathrm{vol}(M_n) \sim n \cdot v_O \quad \text{as } n \to \infty$$

in probability. Here v_O *denotes the volume of the regular right angled ideal hyperbolic octahedron.*

(b) *There exists a constant* $c_\lambda > 0$ *so that the first discrete Laplacian eigenvalue* $\lambda_1(M_n)$ *of* M_n *satisfies*

$$\lim_{n\to+\infty} \mathbb{P}[\lambda_1(M_n) > c_\lambda] = 1.$$

(c) *There exists a constant* $c_d > 0$ *such that the diameter* $\mathrm{diam}(M_n)$ *of* M_n *satisfies*

$$\lim_{n\to+\infty} \mathbb{P}[\mathrm{diam}(M_n) < c_d \log(\mathrm{vol}(M_n))] = 1.$$

(d) *There exists a constant* $c_s > 0$ *such that the systole* $\mathrm{sys}(M_n)$ *of* M_n *satisfies*

$$\lim_{n\to+\infty} \mathbb{P}[\mathrm{sys}(M_n) > c_s] = 1.$$

[1] If $X_n : \Omega_n \to \mathbb{R}$ is a sequence of random variables and $f : \mathbb{R} \to \mathbb{R}$ a function, we say that $X_n \sim f(n)$ in probability if $\mathbb{P}(X_n/f(n) \in [1-\varepsilon, 1+\varepsilon]) \to 1$ for all $\varepsilon > 0$.

(e) *For every $\varepsilon > 0$, we have*

$$\lim_{n\to+\infty} \mathbb{P}\left[\frac{1-\varepsilon}{4n} < \operatorname{sys}(\mathcal{D}M_n) < \frac{1}{n^{1-\varepsilon}}\right] = 1.$$

The same holds for the minimal length among arcs in M_n that are homotopically non-trivial relative to ∂M_n.

Some remarks about these results:

• Our proof of hyperbolisation for M_n does not rely on Perelman's proof of the Geometrisation conjecture. Instead, we use Andreev's theorem [40] and recent work by Futer–Purcell–Schleimer [23] on Dehn fillings. Note that there is also a "Ricci-flow-free" proof of hyperbolisation of random Heegaard splittings [21].

• Item (b) admits a more geometric reformulation, as it follows from it together with classical work by Buser [14] that the Cheeger constant of M_n is also (asymptotically almost surely) uniformly bounded from below. It also implies (with (a) and a theorem of Lackenby [31]) a weaker version of (c) in our topological theorem.

• Item (c) also implies that $\mathcal{D}M_n$ has logarithmic diameter (that is, $\operatorname{diam}(\mathcal{D}M_n) = O(\log \operatorname{vol}(M_n)) = O(\log n)$ in probability), and so M_n does as well. On the other hand it follows from an easy volume argument that the diameter of any closed hyperbolic three-manifold M satisfies

$$\operatorname{diam}(M) \geq \frac{1}{2}\log(\operatorname{vol}(M)) - C$$

for some uniform constant $C > 0$.

• The stark difference in behavior between the systoles of M_n and $\mathcal{D}M_n$ (items (d) and (e)) comes from the fact that M_n contains very short geodesic arcs with endpoints in ∂M_n that glue together into short loops in $\mathcal{D}M_n$.

Besides expansion, another way of looking efficiently at the global geometry of a (possibly random) compact Riemannian manifold that has recently seen much interest is the so-called Benjamini–Schramm topology (see [25] for a survey). We determine the Benjamini–Schramm limit of the sequence $(M_n)_n$ as a consequence of our proof of hyperbolisation. This is more technical than our other results, so we will not give precise statements here but just a sketch of what this means. Very roughly, a sequence of finite volume random hyperbolic manifolds $(M_n)_n$ converges in the Benjamini–Schramm sense to a limit M_∞ (a random pointed manifold) if for every fixed $R > 0$, the distribution of the R-neighbourhood of a uniformly random point in M_n converges (with respect to pointed Gromov–Hausdorff topology) to that of the R-neighbourhood of a random point in M_∞. It turns out that the Benjamini–Schramm limit of M_n can be identified with a tree of right angled octahedra pointed at a uniform random point (which makes sense since this manifold has a cofinite group of isometries). A rigorous exposition of these notions, and a precise statement for the result discussed above, are given in Section 3.6.2.

1.3. Notes and references. Various models for random surfaces have been studied (e.g. [11, 13, 26, 34, 38]) and all three types of questions mentioned above have been explored for these: the models in [11] and [38] are plausible as models of typical (hyperbolic) surfaces and [26] and [13] are applications of the probabilistic method to produce hyperbolic surfaces without short pants decompositions and hyperbolic surfaces with near-minimal diameter respectively. Finally, it is still an open question whether the analogue of Selberg's $1/4$ conjecture holds for various random models (see [34, 45]).

The original motivation for the introduction of random Heegaard splittings by Dunfield and Thurston in [20] was to study the (at that point still unsolved) virtual Haken conjecture. They discuss various models before focusing on this one, and it remains the most studied model for random three-manifolds together with random mapping tori. Both of these are hyperbolic with probability 1 [35,36]. Moreover, like our manifolds (Theorem 1.2 (b)) they satisfy a law of large numbers for volume [43], with a constant depending on the underlying random walk. Their spectral gap behaves differently: it is inversely quadratic in volume [28]. Their injectivity radius has been studied in [42] and torsion in their homology in [3]. Moreover, even if random Heegaard splittings turn out to be hyperbolic with probability 1 [35], unlike for instance random regular graphs [7] and random hyperbolic surfaces [11, 38, 39], they do not Benjamini–Schramm converge to their universal cover – i.e. already at a bounded scale, the geometry of these manifolds ceases to be that of $\mathbb{H}^3$. Finally, [33] gives a probabilistic proof of the existence of infinitely many closed hyperbolic homology three-spheres with a fixed Casson invariant and Heegaard genus.

Whether the model studied in this paper is plausible as a model for a "typical" three-manifold with boundary, we leave to the reader. We note that the question of studying this model has been evoked before (see, for instance, [19, Question 6.2]) but we are not aware of any prior other results on it.

1.4. Proof ideas. Let us briefly describe the main ideas behind the proofs of our results. The two big steps consist of understanding the combinatorial properties of the complex we build and then using those to understand the geometry and topology.

The first observation is that all results above are of the form $\mathbb{P}_n[P_n] \to 1$ as $n \to \infty$ for some sequence of properties P_n of M_n. It follows from classical results in graph theory [7, 44] that for such statements it is sufficient to prove the analogous statement for the random manifolds N_n obtained by randomly gluing the building blocks together (without setting the condition on faces we set for M_n).

In what follows, we will try to avoid repeating the phrase "asymptotically almost surely" and will often just say that M_n has this or that property when we mean it has the given property asymptotically almost surely.

The proofs now start with the combinatorics (in Section 2). Using the observation above, the idea is to study the properties of N_n and then turn these into properties of M_n. First of all, we prove, using elementary but tedious combinatorics, that the

number of boundary components of N_n (and hence M_n) is 1. The next step, which is responsible for the largest part of the combinatorial arguments, is to study the combinatorics of interior edges in N_n. We ask two questions: how many edges are there? And, given some number $k \in \mathbb{N}$, how many edges are there that are incident to k truncated tetrahedra? To answer these questions, we will use *peeling* techniques. These are techniques coming from the world of random planar maps (see, for instance, [18]). The basic idea is to explore the random cell complex N_n using a specific algorithm – adapted to the problem at hand – to determine in which order cells are explored. These lead to bounds on the expected number of interior edges (the total and the number that is incident to a fixed number of 3-cells) that we think might be interesting in their own right (see Theorem 2.4). This in turn yields the Euler characteristic and hence the genus of the boundary of N_n (and hence M_n). Note that all these combinatorial results can also be interpreted in terms of the cell complex obtained from gluing tetrahedra according to the same pattern – a pseudo-manifold.

After this, we deal with the geometric questions in Section 3. Our first goal is hyperbolisation. The main idea behind our proof of this is to see our manifold as a Dehn filling of a non-compact hyperbolic three-manifold (similar ideas were used, with very different objectives, in [17]). This non-compact manifold is obtained by gluing hyperbolic right-angled octahedra, using four alternating faces out of eight per octahedron, along the same pattern as M_n. We then first apply Andreev's theorem to fill cusps with "few" octahedra around them. The number such cusps is controlled by our combinatorial bounds. This creates another non-compact hyperbolic manifold, but without "small" cusps. We then fill the remaining cusps and rely on results by Futer–Purcell–Schleimer [23] to guarantee the result is hyperbolic. These same results also give us information about the way the geometry changes between the non-compact and the compact manifold.

Once we have proved that M_n is hyperbolic, we use results on random regular graphs together with a version of the Brooks–Burger transfer principle to show that $\lambda_1(M_n)$ can be uniformly bounded from below. The law of large numbers for the volumes of our manifolds is essentially a direct consequence of our proof of hyperbolisation. We also prove the logarithmic bound on the diameter of M_n by combining the fact that random 4-regular graphs have logarithmic diameter with the geometric control we have over the change of geometry during Dehn filling. Finally, we prove Benjamini–Schramm convergence, again using the geometric comparison between the non-compact hyperbolic manifold and M_n and results on random regular graphs.

1.5. Further questions. We finish this introduction with some questions about our model.

Question 1 (Poisson–Dirichlet distribution for edges). Let us write

$$L = (L_1, L_2, \ldots)$$

for the random vector that contains the *lengths* of all the interior edges in M_n. Here the length of an edge is the number of 3-cells incident to it and is counted with multiplicity – i.e. if an interior edges is incident to a 3-cell in multiple places then the 3-cell is counted multiple times.

If we order this vector so that $L_1 \geq L_2 \geq \cdots$ and normalise it by dividing by the total length $(6n)$, does the resulting partition of the interval $[0, 1]$ converge in distribution to a Poisson–Dirichlet distributed variable?

The analogous result is known to hold for surfaces obtained by randomly gluing polygons together [12, 16, 24].

Question 2 (Explicit measures of expansion). Determine the optimal spectral gaps and Cheeger constants that hold a.a.s. for M_n. For instance, do we have

$$\forall \varepsilon > 0, \quad \lim_{n \to +\infty} \mathbb{P}(\lambda_1(M_n) > 1 - \varepsilon) = 1?$$

An analogue of this is conjectured for random hyperbolic surfaces [45, Problem 10.3], [34, Conjecture 1.1] and holds for random regular graphs [22].

Finally, we can also ask for sharp estimates for the systoles of M_n and $\mathcal{D}M_n$.

Question 3. Give an explicit sequence (s_n) such that $\mathrm{sys}(\mathcal{D}M_n) \sim s_n$ in probability. Compute (if it exists) $\lim \mathbb{E}(\mathrm{sys}(M_n))$.

1.5.1. Finer behavior of homology and L^2-invariants of the limit. In part (d) of Theorem 1.1, we get good bounds for the typical Betti numbers of the M_n. However, particularly in view of the fact that the random Heegaard splittings of [20] typically have vanishing first Betti numbers, we ask the following question.

Question 4. Does $b_1(M_n, \partial M_n) = 0$ hold a.a.s.?

A positive answer is suggested by computer experiments conducted by Nathan Dunfield using Regina. His results also suggest the following conjecture about the behavior of the full integral homology, denoting by $H_1(M_n)_{\mathrm{tors}}$ the torsion subgroup of $H_1(M_n)$:

Question 5. Is $H_1(M_n)_{\mathrm{tors}}$ trivial a.a.s.? Or a weaker variant: do we have

$$\lim_{n \to +\infty} \log \frac{|H_1(M_n)_{\mathrm{tors}}|}{n} = 0$$

in probability?

Note that for random Heegaard splittings the opposite behavior occurs: $H_1(M_n)_{\mathrm{tors}}$ is as large as possible, i.e. of exponential size in the number of tetrahedra (see [30, Section 2.2]).

In view of the convergence discussed in Section 3.6.2, the last question could be related to the L^2-invariants of the infinite cover $O^\infty \to O$ (see [29] for an introduction to this topic). Our result on Betti numbers implies, via generalisations of the Lück Approximation theorem (see [29, Section 5.4.3]) that the L^2-Betti numbers of $O^\infty \to O$ relative to the boundary vanish. We can ask about other L^2-invariants as follows:

Question 6. What are the Novikov–Shubin invariants of $O^\infty \to O$? Is its L^2-torsion equal to 0?

In view of the approximation conjecture for torsion (which is wide open at present, see [29, 6.5] for a survey, but much simpler to deal with in 3 dimensions when the torsion vanishes, see [32]), the vanishing of L^2-torsion would likely imply an affirmative answer to the weaker form of Question 5. We note that our expansion results implies (via the proof of Lück approximation) that the zeroth Novikov–Shubin invariant is ∞^+.

1.5.2. Variations on the model: Killing symmetry. Finally, it is possible to derive models for random closed manifolds from our random manifolds with boundary. The simplest (which we study to some extent here) would be obtained by just doubling it; however this is supported only on manifolds admitting an involution with codimension-1 fixed locus. There are various ways to break this symmetry: the fact that the boundary is triangulated allows us to identify it to a fixed (depending only on the genus) model surface up to a bounded ambiguity. So we can talk about random mapping classes of the boundary almost as usual, and these allow us to perform various more complicated constructions such as gluing back a copy using a random mapping class. We can also glue the appropriate handlebody (or indeed any other manifold with connected boundary of the correct genus). It follows from our results and geometrisation that all these models are hyperbolic; however we do not know (say for a choice of the mapping class with few random steps) how their volume behaves, whether they are expanders, or whether they admit a Benjamini–Schramm limit. The investigation of such questions would certainly require a different set of tools than what we use here.

Acknowledgements. We worked on and off on this project for several years and as such it benefited from various grants. B. P. thanks the Max Planck Institute for Mathematics in Bonn and the ERC grant "Moduli". J. R. thanks the ANR for support through the project ANR-16-CE40-0022-01 – AGIRA and the Hausdorff institute through the junior trimester "Topologie" which was held there in 2016.

We also thank François Costantino, Juan Souto and Gabriele Viaggi for useful discussions. B. P. thanks Thomas Budzinski and Nicolas Curien for teaching him how to peel a surface. We are indebted to Nathan Dunfield for various comments on a preliminary version, suggesting a more elementary and efficient approach to Heegaard genus and sharing the results of his simulation of the model with us. We thank Ara Basmajian for useful conversations that lead to an improved bound in Theorem 1.2 (e). We thank the anonymous referee for useful comments that improved the exposition.

2. Combinatorics

In this section, we formally describe the combinatorial model we use. Moreover, we determine the combinatorial structure of the random cell complex underlying our manifolds and derive some basic topological properties of our manifolds from it.

2.1. The topological model. In what follows, $\mathcal{T}_n$ will denote a Δ-complex obtained by randomly gluing the faces of n tetrahedra together in pairs. The gluing that is used is picked at random among the three orientation reversing euclidean isometries between the faces.

More formally, this goes as follows:

(1) We start with n *labelled* tetrahedra. Here "labelled" means that every the vertex of these tetrahedra carries a unique label in $\{1, \dots, 4n\}$. If a face of a tetrahedron has vertices $v_1, v_2, v_3 \in \{1, \dots, 4n\}$, that moreover in this particular cyclic order induce an outward orientation on that particular face, then we will denote the face by the cycle $(v_1\ v_2\ v_3)$.

(2) The faces are partitioned into $2n$ pairs, uniformly at random. We will denote the resulting partition by $\omega_n = (\omega_n^{(i)})_{i=1}^{2n}$.

(3) Per pair of faces $\omega_n^{(i)} = \{(v_1\ v_2\ v_3),\ (w_1\ w_2\ w_3)\}$ in this partition, one of three cyclic-order-reversing pairings between the vertices is chosen uniformly at random. The resulting $2n$-tuple of pairings will be denoted $\sigma_n = (\sigma_n^{(i)})_{i=1}^{2n}$.

(4) We identify each pair of faces $\omega_n^{(i)} = \{(v_1\ v_2\ v_3),\ (w_1\ w_2\ w_3)\}$ in the partition ω_n using the unique orientation reversing simplicial map that sends v_j to $\sigma_n^{(i)}(v_j)$ for $j = 1, 2, 3$. The resulting simplicial complex is called $\mathcal{T}_n$.

Let us write $(\Omega_n, \mathbb{P}_n)$ for the corresponding probability space. So Ω_n is the finite set of all possibilities for ω_n and σ_n and $\mathbb{P}_n$ is the uniform probability measure on it. Note that

$$|\Omega_n| = (4n)!!\ 3^{2n},$$

where for an even number $k \in \mathbb{N}$, $k!! = (k-1) \cdot (k-3) \cdots 3 \cdot 1$.

The dual graph $\mathcal{G}_n$ to $\mathcal{T}_n$ - the 4-valent graph whose vertices correspond to the tetrahedra of $\mathcal{T}_n$ who share an edge per face that they have in common - is a random 4-regular graph. The model this induces is exactly the *configuration model*, one of the most studied models of random regular graphs (see e.g. [7, 44]).

Note that besides the number of tetrahedra (n), the number of 2-faces is also deterministic in this model ($2n$). The numbers of vertices and edges are random variables.

2.2. The results. Let N_n denote the manifold with boundary obtained by truncating $\mathcal{T}_n$ at the vertices. Figure 1 in the introduction shows what the basic building block of N_n looks like.

In most of this section, we will think in terms of $\mathcal{T}_n$. However, since we are eventually interested in N_n, we will describe some of the results in terms of N_n.

Theorem 2.1 (Topology). (a) *We have*

$$\lim_{n\to\infty} \mathbb{P}[N_n \text{ has a single boundary component}] = 1.$$

(b) *We have*

$$\mathbb{E}[\chi(\partial N_n)] = \log(n) - 2n + O(1)$$

as $n \to \infty$. In particular, if we write $g(\partial N_n)$ for the genus of the single boundary component of $\mathcal{M}_n$ we have that

$$\lim_{n\to+\infty} \mathbb{P}[n - \theta(n) \leq g(\partial M_n) \leq n + 1] = 1,$$

for any function $\theta\colon \mathbb{N} \to \mathbb{R}$ that grows super-logarithmically.

Part (a) follows from Proposition 2.3, and part (b) from Theorem 2.4. In the latter, we also prove bounds on the expected number of edges incident to a given number of tetrahedra in $\mathcal{T}_n$. We will need these in the geometric part of the paper.

We will write M_n for the random manifold we obtain if we condition on $\mathcal{G}_n$ being simple, i.e. not having loops or multiple edges. Bollobás [6] proved that

$$\lim_{n\to\infty} \mathbb{P}[\mathcal{G}_n \text{ is simple}] > 0.$$

In particular, this implies that if $(\mathcal{P}_n)_n$ is a sequence of properties of M_n and N_n then as $n \to \infty$, we have that

$$\text{if } \mathbb{P}[N_n \text{ has } \mathcal{P}_n] \to 1, \text{ then } \mathbb{P}[M_n \text{ has } \mathcal{P}_n] \to 1. \tag{2.1}$$

Combining this with the theorem above, we get the following corollary.

Corollary 2.2 (Topology of the boundary). (a) *We have*

$$\lim_{n\to\infty} \mathbb{P}[M_n \text{ has a single boundary component}] = 1.$$

(b) *We have*

$$\lim_{n\to+\infty} \mathbb{P}[n - \theta(n) \leq g(\partial M_n) \leq n + 1] = 1,$$

for any function $\theta\colon \mathbb{N} \to \mathbb{R}$ that grows super-logarithmically.

2.3. The number of vertices. We start by studying the number of vertices. Let us write V for the number of vertices of $\mathcal{T}_n$. Note that V is also the number of boundary components of N_n. As such the following proposition implies Theorem 2.1 (a).

Proposition 2.3. *We have*

$$\mathbb{P}_n[V = 1] \to 1,$$

as $n \to \infty$.

Proof. Let us write $V_{\text{small}}\colon \Omega_n \to \mathbb{N}$ for the random variable that counts the number of vertices of $\mathcal{T}_n$ incident to most $2n$ tetrahedra (with multiplicity – i.e. if a tetrahedron is incident to a vertex in multiple corners, it is counted multiple times). We will prove that

$$\mathbb{E}_n[V_{\text{small}}] \to 0,$$

as $n \to \infty$. Note that this is sufficient to prove the proposition.

We will write

$$\mathbb{E}_n[V_{\text{small}}] = \sum_a \mathbb{E}_n[\mathbb{1}_a].$$

Here the sum runs over "labelled" vertices a. A labelled vertex a is the data of the gluings of all the labelled faces incident to a given vertex. Here $\mathbb{1}_a\colon \Omega_n \to \{0,1\}$ is the indicator for the event that a appears in ω. As such

$$\mathbb{E}[\mathbb{1}_a] = \frac{1}{3^f (4n-1)(4n-3)\cdots(4n-2f+1)},$$

where f is the number of faces incident to the given vertex.

Given a, write n_1, n_2, n_3, n_4 for the number of tetrahedra that are incident in 1, 2, 3 and 4 of their vertices to a, respectively. Note that the number of tetrahedron faces involved in such a gluing is given by

$$2f = 3n_1 + 4(n_2 + n_3 + n_4).$$

This implies that n_1 must be even and that the number of gluings with the same selection of vertices is

$$3^{\frac{n_3}{2}+2n_4}(2f)!!\,.$$

The power of 3 comes from the fact that only the faces with three vertices adjacent to them can be rotated. This number of faces is equal to $n_3 + 4n_4$. As such we obtain

$$\mathbb{E}_n[V_{\text{small}}] = \sum_{\substack{0<n_1+n_2+n_3+n_4\le n,\\ n_1+2n_2+3n_3+4n_4\le 2n,\\ n_1 \text{ even}}} \binom{n}{n_1, n_2, n_3, n_4} \frac{3^{\frac{n_3}{2}+2n_4}(2f)!!}{3^f (4n-1)(4n-3)\cdots(4n-2f+1)},$$

where we write

$$\binom{n}{n_1, n_2, n_3, n_4} = \frac{n!}{n_1! \cdot n_2! \cdot n_3! \cdot n_4! \cdot (n - n_1 - n_2 - n_3 - n_4)!}$$

counts the number of subsets of the n tetrahedra with the appropriate number of vertices in them.

We claim that this implies that $\mathbb{E}_n[V_{\text{small}}] = O(n^{-1})$. This follows by analyzing the terms in the sum. The largest terms in this sum are when

$$2f = 3n_1 + 4(n_2 + n_3 + n_4)$$

is smallest. Indeed, using that $n_1 + 2n_2 + 3n_3 + 4n_4 \leq 2n$, an elementary but tedious computation shows that a term decreases when one of the n_i's is increased. Given that the number of terms is quartic in n and the number of terms that are larger than $O(n^{-5})$ is bounded, we obtain the estimate. □

2.4. The number of edges. The random variable that counts the number of edges that are incident to k tetrahedra will be denoted by

$$E_k \colon \Omega_n \to \mathbb{N}.$$

In this variable, tetrahedra are counted *with* multiplicity. That is, if an edge appears multiple times in the boundary of a given tetrahedron, this tetrahedron adds to its "length" each time. Note that

$$\sum_{k \geq 0} k \cdot E_k = 6n.$$

We will also write

$$E = \sum_{k \geq 0} E_k \colon \Omega_n \to \mathbb{N}$$

for the total number of edges. The goal of this section is to study the distribution of $(E_k)_k$ as $n \to \infty$.

We will also count edges that we will call *simple*. These are edges that neighbour each tetrahedron at most once. Figure 2 shows an example. We will denote the number of simple edges by E° and the number of simple edges adjacent to k tetrahedra by E_k°. Note that $E_k^\circ = 0$ for all $k > n$, which is not at all necessary for non-simple edges.

Concretely, we will prove the following estimates, where $k = o(f(n))$ for some function $f \colon \mathbb{N} \to \mathbb{N}$, means that the statement holds for any function $k \colon \mathbb{N} \to \mathbb{N}$ such that $k(n) = o(f(n))$.

Theorem 2.4 (Combinatorics of edges). (a) *We have*

$$\mathbb{E}_n[E] = \frac{1}{2} \log(n) + O(1)$$

as $n \to \infty$.

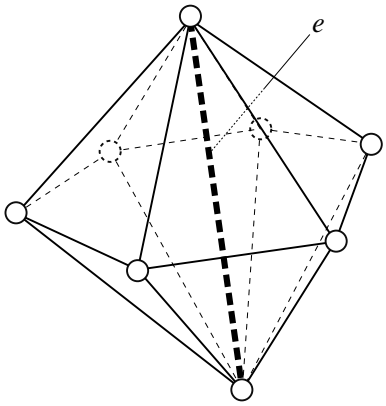

Figure 2. A simple edge e.

(b) *For all $k = o(\sqrt{n})$, we have*

$$\mathbb{E}[E_k^\circ] = \frac{1}{2k} \cdot (1 + o(1))$$

as $n \to \infty$. Moreover, the error is uniform in k.

(c) *For all $k = o(n)$ and $k_1 \leq k_2 \leq k$, we have*

$$\mathbb{E}[E_k] \leq \frac{1}{2k} + O(n^{-1}) \quad \textit{and} \quad \mathbb{E}\left[\sum_{l=k_1}^{k_2} E_l\right] \leq \frac{1}{2}\log(k_1/k_2) + O(1)$$

as $n \to \infty$. In particular, for all $k = o(\sqrt{n})$ and $k_1 \leq k_2 \leq k$, we have

$$\mathbb{E}[E_k] = \frac{1}{2k} \cdot (1 + o(1)), \quad \mathbb{E}\left[\sum_{l=k_1}^{k_2} E_l\right] = \frac{1}{2}\log(k_1/k_2)(1 + o(1))$$

and

$$\mathbb{E}[E_k - E_k^\circ] = o(1)$$

as $n \to \infty$.

(d) *For all $K, L = o(n^{1/3})$, we have*

$$\mathbb{E}[E_{KL}] = o(1)$$

as $n \to \infty$. Here E_{KL} counts the number of pairs of edges of size $\leq K$ and $\leq L$, respectively, that are incident to a common tetrahedron.

Before we get to the proof of this theorem, let us briefly note how to derive the Euler characteristic of ∂N_n from it.

Proof of Theorem 2.1 (b). Writing v, e and f for the number of vertices, edges and faces of the triangulation on ∂N_n, we have

$$v = 2E, \quad e = 6n, \quad \text{and} \quad f = 4n.$$

As such, Theorem 2.4 (a) implies our claim. □

2.4.1. Peeling. In order to prove Theorem 2.4 (a), (c) and (d), we will use peeling (see, for instance, [18]). Before we get to the proof, we need to do some preparation.

The main idea behind peeling is to build our random cell complex $\mathcal{T}_n$ in a specific order. In particular, we will describe a peeling algorithm that determines a sequence of cell complexes

$$\mathcal{T}_n^{(0)}, \mathcal{T}_n^{(1)}, \ldots, \mathcal{T}_n^{(2n)},$$

where $\mathcal{T}_n^{(0)}$ consists of n disjoint tetrahedra, $\mathcal{T}_n^{(2n)} = \mathcal{T}_n$ (in the sense that it has the same distribution) and in general $\mathcal{T}_n^{(i+1)}$ can be obtained from $\mathcal{T}_n^{(i)}$ by identifying exactly one pair of faces of $\mathcal{T}_n^{(i)}$.

We will use two different peeling algorithms. The first to prove part (a) and the second to prove parts (c) and (d).

2.4.2. Algorithm 1. The first algorithm is very simple and closely resembles that of [11, Section 8].

Peeling algorithm 1:
Initialisation:
Objects: faces $f^{(0)}, f^{(1)}, \ldots, f^{(2n)}$, $f'^{(0)}, f'^{(1)}, \ldots, f'^{(2n)}$ and $t \in \mathbb{Z}$.

- Set $t = 0$.
- Set $\mathcal{T}_n^{(0)}$ equal to a disjoint union of n tetrahedra.
- Set $f^{(0)}$ equal to a face in $\partial\, \mathcal{T}_n^{(0)}$, picked uniformly at random.

Iteration: while $t < 2n$, repeat the following steps:

(1) Glue the face $f^{(t)}$ to a uniformly random face $f'^{(t)}$ in $\partial\, \mathcal{T}_n^{(t)} \setminus f^{(t)}$, with a uniformly random gluing. Call the result $\mathcal{T}_n^{(t+1)}$.
(2) If $t < 2n$: Pick a uniformly random face $f^{(t+1)} \subset \partial\, \mathcal{T}_n^{(t)}$
(3) Add 1 to t.

Note that the distribution of $\mathcal{T}_n^{(2n)}$ is the same as that of $\mathcal{T}_n$.

2.4.3. Closing off edges. The reason for setting up the peeling algorithm is that we can now control the number of edges in $\mathcal{T}_n$ by bounding the number of edges that are *closed off* – i.e. that disappear from the boundary – during each step of the process.

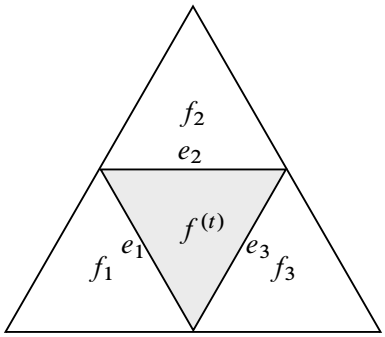

Figure 3. The face $f^{(t)}$ and its neighbours.

As such, let us define random variables $E^{(t)}$ that count the number of edges that are closed off when $\mathcal{T}_n^{(t)}$ is created. This is the number of edges that lie in $\partial\,\mathcal{T}_n^{(t-1)}$ but not in $\partial\,\mathcal{T}_n^{(t)}$. Note that

$$E = \sum_{t=1}^{2n} E^{(t)}.$$

Moreover, since any edge that gets closed off at the tth step necessarily lies in $f^{(t)}$, we have $E^{(t)} \leq 3$. One of the things we will argue below is that most of the time, we actually have $E^{(t)} \leq 1$. To this end, consider Figure 3. It shows a schematic overview of the situation around the face $f^{(t)}$ at time t. Note that some of the faces f_1, f_2, f_3 and $f^{(t)}$ may coincide. However, if they do not, only one edge can be created at step t: e_i is then closed off if and only if $f^{(t)}$ is glued to f_i with exactly one out of the three possible face identifications.

So, in what follows, we will make a distinction between *singular faces* – faces that are their own neighbour or of which some of the neighbours coincide – and *regular faces* – faces that are not singular. We will write $F_{\text{sing}}^{(t)}$ for the random variable that counts the number of singular faces in $\partial\,\mathcal{T}_n^{(t)}$. Likewise, we define two sequences of random variables $E_{\text{sing}}^{(t)}$, $E_{\text{reg}}^{(t)}$, where

$$E_{\text{sing}}^{(t)} = \begin{cases} E^{(t)} & \text{if } f^{(t)} \text{ is singular,} \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad E_{\text{reg}}^{(t)} = E^{(t)} - E_{\text{sing}}^{(t)}.$$

Since there are $4n - 2t$ faces left when the tth step starts, we have

$$\begin{aligned} \mathbb{E}\big[E_{\text{reg}}^{(t)} \mid F_{\text{sing}}^{(t)}\big] &= \frac{4n - 2t - F_{\text{sing}}^{(t)}}{4n - 2t} \frac{3}{3(4n - 2t - 1)} \\ &= \frac{4n - 2t - F_{\text{sing}}^{(t)}}{(4n - 2t)(4n - 2t - 1)}, \end{aligned} \tag{2.2}$$

and since at most three edges (those on $f^{(t)}$) are closed off during tth step, and moreover there are at most 3 choices for $f'^{(t)}$ and 3 gluings per choice that result in an edge closure, we have

$$\mathbb{E}\big[E^{(t)}_{\text{sing}} \mid F^{(t)}_{\text{sing}}\big] \leq \frac{F^{(t)}_{\text{sing}}}{4n-2t}\,\frac{3\cdot 3\cdot 3}{3(4n-2t-1)} = \frac{9\cdot F^{(t)}_{\text{sing}}}{(4n-2t)(4n-2t-1)}. \tag{2.3}$$

So, we need to control $F^{(t)}_{\text{sing}}$.

Lemma 2.5. *For all $t \in \mathbb{N}$ so that $t < 2n$, we have*

$$\mathbb{E}_n\big[F^{(t)}_{\text{sing}}\big] \leq 12\cdot\log\Big(\frac{2n}{2n-t}\Big) + o(1)$$

as $n \to \infty$. The implied error is independent of t

Proof. Let us write $\Delta F^{(t)}_{\text{sing}}$ for the random variable that counts the number of singular faces that is created at step t – since we are only interested in an upper bound, we will ignore singular faces that disappear. Because $F^{(0)}_{\text{sing}} = 0$, we have

$$F^{(t)}_{\text{sing}} \leq \sum_{s=1}^{t} \Delta F^{(s)}_{\text{sing}}.$$

So, let us try to control $\mathbb{E}[\Delta F^{(s)}]$. The only faces whose neighbourhood changes during step s are the neighbours of $f^{(s)}$ and $f'^{(s)}$. If such a neighbour of say $f^{(s)}$ is regular, it becomes singular only if one of its neighbours is also a neighbour of $f'^{(s)}$. Likewise, a regular neighbour of $f'^{(s)}$ becomes singular only if it shares a neighbour with $f^{(s)}$. In other words, $\Delta F^{(s)}$ can only be positive if the combinatorial distance – the number of edges that needs to be traversed in $\partial\,\mathcal{T}^{(s)}_n$ in order to move from one face to the other – between $f^{(s)}$ and $f'^{(s)}$ is at most 3. Note that this is independent of whether or not $f^{(s)}$ and $f'^{(s)}$ are regular.

Since there are at most $3\cdot 2\cdot 2 = 12$ faces at combinatorial distance at most 3 from $f^{(s)}$, we have

$$\mathbb{E}\big[\Delta F^{(s)}_{\text{sing}} \mid F^{(s-1)}_{\text{sing}}\big] \leq \frac{12}{4n-2s-1}.$$

This implies that

$$\mathbb{E}_n\big[F^{(t)}_{\text{sing}}\big] \leq \sum_{1\leq s\leq t} \frac{12}{4n-2s-1} \leq 12 \sum_{k=4n-2t-1}^{4n-1} \frac{1}{k} = 12\log\Big(\frac{4n}{4n-2t}\Big) + o(1)$$

as $n \to \infty$. □

2.4.4. The total edge count. We now have all the set-up we need for the first part of Theorem 2.4.

Proof of Theorem 2.4 (a). Fix any $\alpha \in (0,1)$. We have

$$\mathbb{E}_n[E^{(t)}] = \mathbb{E}_n\big[E^{(t)} \cdot \mathbb{1}_{F^{(t)}_{\mathrm{sing}} \geq t^\alpha}\big] + \mathbb{E}_n\big[E^{(t)} \cdot \mathbb{1}_{F^{(t)}_{\mathrm{sing}} < t^\alpha}\big]$$

We start with the first term. By Lemma 2.5, (2.2) and Markov's inequality, we have

$$0 \leq \mathbb{E}_n\big[E^{(t)}_{\mathrm{reg}} \cdot \mathbb{1}_{F^{(t)}_{\mathrm{sing}} \geq t^\alpha}\big] \leq \frac{12\log(t)}{n^\alpha}\frac{1}{4n-2t-1},$$

for all n large enough. Likewise, using Lemma 2.5, (2.3) and Markov's inequality, we obtain

$$0 \leq \mathbb{E}_n\big[E^{(t)}_{\mathrm{reg}} \cdot \mathbb{1}_{F^{(t)}_{\mathrm{sing}} \geq t^\alpha}\big] \leq \frac{12\log(t)}{n^\alpha}\frac{9}{4n-2t-1}. \tag{2.4}$$

So, we obtain

$$\sum_{t=1}^{2n} \mathbb{E}_n\big[E^{(t)} \cdot \mathbb{1}_{F^{(t)}_{\mathrm{sing}} \geq t^\alpha}\big] = O\left(\frac{\log(n)^2}{n^\alpha}\right)$$

as $n \to \infty$. In other words, most edges are created when few singular faces are present.

So, let us control this term. Again using (2.2), we have

$$\frac{4n-2t-t^\alpha}{4n-2t}\frac{1}{4n-2t-1} \leq \mathbb{E}_n\big[E^{(t)}_{\mathrm{reg}} \cdot \mathbb{1}_{F^{(t)}_{\mathrm{sing}} < t^\alpha}\big] \leq \frac{1}{4n-2t-1}.$$

Using (2.3),

$$0 \leq \mathbb{E}_n\big[E^{(t)}_{\mathrm{sing}} \cdot \mathbb{1}_{F^{(t)}_{\mathrm{sing}} < t^\alpha}\big] \leq \frac{9t^\alpha}{(4n-2t)(4n-2t-1)}.$$

These last two bounds give

$$\sum_{t=1}^{2n} \mathbb{E}_n\big[E^{(t)} \cdot \mathbb{1}_{F^{(t)}_{\mathrm{sing}} < t^\alpha}\big] = \sum_{t=1}^{2n} \frac{1}{4n-2t-1} + O(1) = \frac{1}{2}\log(n) + O(1).$$

Together with (2.4), this proves our claim. □

2.4.5. Simple edges. The proof of part (b) of our theorem – the count of the expected number of simple edges – will not use a peeling algorithm.

Proof of Theorem 2.4 (b). We write

$$\mathbb{E}_n[E^\circ_k] = \sum_{(c_1\ c_2\ \ldots\ c_k)} \mathbb{E}_n[\mathbb{1}_{(c_1\ c_2\ \ldots\ c_k)}],$$

where the sum runs over all cycles $(c_1\ c_2\ \ldots\ c_k)$ of length k so that

(i) c_i is an ordered pair of faces of a tetrahedron, and

(ii) at most one such pair of faces of any given tetrahedron appears in the sequence.

Finally, given $\omega \in \Omega_n$,

$$\mathbb{1}_{(c_1\ c_2\ \ldots\ c_k)}(\omega) = \begin{cases} 1 & (c_1\ c_2\ \ldots\ c_k) \text{ appears around an edge in } \omega, \\ 0 & \text{otherwise.} \end{cases}$$

It follows from (ii) that every cycle $(c_1\ c_2\ \ldots\ c_k)$ in the sum corresponds to an identification of k pairs of faces. As such

$$\mathbb{E}_n[\mathbb{1}_{(c_1\ c_2\ \ldots\ c_k)}] = \frac{1}{3^k(4n-1)(4n-3)\cdots(4n-2k+1)}.$$

So, all that remains is counting the number of possible cycles $(c_1\ c_2\ \ldots\ c_k)$ of pairs of faces. We have

$$|\{(c_1\ c_2\ \ldots\ c_k)\}| = \frac{1}{2k}\binom{n}{k} \cdot 6^k \cdot 2^k \cdot k!\,.$$

The reason for this expression is as follows. First we count sequences instead of cycles:

- This gives a total of $\binom{n}{k}$ choices for the tetrahedra.
- Per tetrahedron out of which pair of faces is used, we have a choice of 6 pairs (corresponding to the edges of the tetrahedron). This gives rise to a factor 6^k.
- Per pair of faces that is used, we have two choices for the order in which those two faces appear. This leads to a factor 2^k.
- Finally, there are $k!$ ways to order the sequence.

Since we do not want to make a difference between sequences that differ by a cyclic permutation or are each other's inverse, we divide by $2k$.

So, we get

$$\begin{aligned} \mathbb{E}_n[E_k^\circ] &= \frac{1}{2k}\prod_{i=1}^{k}\frac{4n-4i+4}{4n-2i+1} = \frac{1}{2k}\exp\Big(\sum_{i=1}^{k}\log\Big(1-\frac{2i-3}{4n-2i+1}\Big)\Big) \\ &= \frac{1}{2k}\exp\Big(-\sum_{i=1}^{k}\Big(\frac{2i-3}{4n-2i+1}+O(n^{-2})\Big)\Big) = \frac{1}{2k}\exp(o(1)). \end{aligned}$$

Where we used the fact that $k = o(\sqrt{n})$ in the last line. □

2.4.6. Algorithm 2. The second algorithm is actually a collection of algorithms, tailored towards counting the number of edges incident to two edges in a fixed tetrahedron. As such, it starts peeling our random triangulation around a fixed starting edge e and then continuous to peel around another fixed edge e' once e is closed.

Peeling algorithm 2:
Input:
A labelled tetrahedron τ and two fixed labelled oriented edges $e, e' \subset \tau$.
Initialisation:
Objects: oriented edges $e^{(0)}, e^{(1)}, \dots, e^{(2n)}$, faces $f^{(0)}, f^{(1)}, \dots, f^{(2n)}$, $f'^{(0)}, f'^{(1)}, \dots, f'^{(2n)}$ and $t \in \mathbb{Z}$.

- Set $t = 0$.
- Set $\mathcal{T}_n^{(0)}$ equal to a disjoint union of n tetrahedra, containing τ.
- Set $e^{(0)} = e$.

Iteration: while $t < 2n$, repeat the following steps:

(1) Glue the face $f^{(t)}$ to the right of $e^{(t)}$ to a uniformly random face $f'^{(t)}$ in $\partial \mathcal{T}_n^{(t)} \setminus f$, with a uniformly random gluing. Call the result $\mathcal{T}_n^{(t+1)}$.

(2) If $t + 1 < 2n$ and $e^{(t)} \not\subseteq \partial \mathcal{T}_n^{(t+1)}$, then

- If $e' \subset \partial \mathcal{T}^{(t+1)}$, set $e^{(t+1)} = e'$.
- Else, pick a uniformly random edge $e^{(t+1)} \subset \partial \mathcal{T}_n^{(0)}$ and orient it randomly.

Else, $e^{(t+1)} = e^{(t)}$.

(3) Add 1 to t.

Again note that the distribution of $\mathcal{T}_n^{(2n)}$ is the same as that of $\mathcal{T}_n$. Figure 4 shows what the initial set-up looks like.

2.4.7. All edges.

Proof of Theorem 2.4 (c). Since we already have Theorem 2.4 (b), we only need an upper bound on $\mathbb{E}[E_k]$. In order to prove such a bound, we will write

$$\mathbb{E}[E_k] = \frac{1}{2k} \sum_e \mathbb{E}\big[\mathbb{1}_e^k\big],$$

where the sum runs over labelled oriented edges e in our collection of n tetrahedra (so the sum has $12n$ terms in total) and $\mathbb{1}_e^k$ is the indicator for the event that e is incident to exactly k corners in the complex $\mathcal{T}_n$.

In order to bound $\mathbb{E}[\mathbb{1}_e^k]$ from above, we use our second algorithm with e as input. The second oriented edge e' that the algorithm uses does not play a role in this proof,

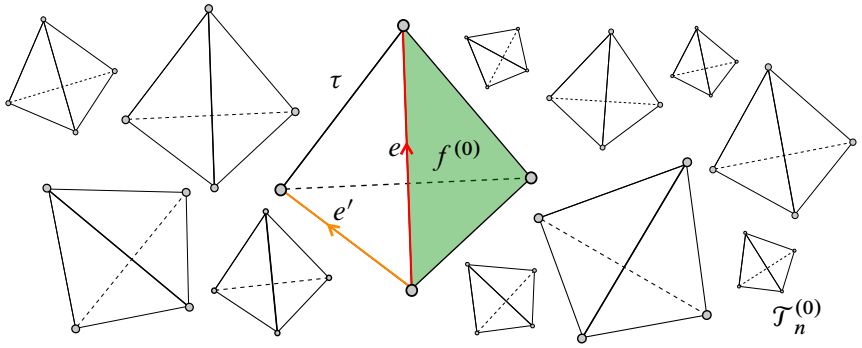

Figure 4. The initial set-up.

so we pick an arbitrary edge. We will also only care what happens in the first k steps of the process.

Just like before, if during every step of the process, $f^{(t)}$ has three distinct neighbours, none of which is $f^{(t)}$ itself – i.e. if $f^{(t)}$ is regular –, it is easy to control the probability that our edge closes up in exactly k steps. So, just like before, we need to bound the probability that in the first k steps, our face becomes singular.

With the same argument as in Lemma 2.5, we have

$$\mathbb{E}\big[\mathbb{1}_{f^{(t)}\text{ is singular}}\big] \leq \frac{12}{4n-2t+1} + \mathbb{E}\big[\mathbb{1}_{f^{(t-1)}\text{ is singular}}\big].$$

In particular,

$$\mathbb{E}\big[\mathbb{1}_{f^{(k)}\text{ is singular}}\big] \leq \frac{12k}{4n-2k+1}.$$

After $k-1$ gluings, there are $4n-2k+1$ faces left, and even if a singular face is involved in the kth gluing, there are at most 3 possible gluings that result in a closure. So, we get

$$\begin{aligned}\mathbb{E}\big[\mathbb{1}_e^k\big] &= \mathbb{E}\big[\mathbb{1}_e^k \mathbb{1}_{f^{(k)}\text{ is singular}}\big] + \mathbb{E}\big[\mathbb{1}_e^k \mathbb{1}_{f^{(k)}\text{ is regular}}\big] \\ &\leq \frac{1}{3(4n-2k+1)} + \frac{12k}{4n-2k+1}\frac{3}{4n-2k+1}.\end{aligned}$$

This means that

$$\mathbb{E}[E_k] \leq \frac{12n}{2k}\left(\frac{1}{3(4n-2k+1)} + \frac{12k}{4n-2k+1}\frac{3}{4n-2k}\right),$$

which proves our claim. □

Proof of Theorem 2.4 (d). Let $k \leq K$ and $l \leq L$. Moreover, let E'_{kl} denote the number of pairs of edges of sizes exactly k and l, respectively, that are incident to a

common tetrahedron. Just like in the proof above, we will write

$$\mathbb{E}\big[E'_{kl}\big] = \frac{1}{4kl}\sum_{e,e'}\mathbb{E}\big[\mathbb{1}^{kl}_{e,e'}\big],$$

where the sum runs over pairs of labelled oriented edges e, e' that are incident to a single tetrahedron (so the sum has $12n \cdot 10$ terms in total) and $\mathbb{1}^{kl}_{e,e'}$ is the indicator for the event that e is incident to exactly k and e' to l corners in the complex $\mathcal{T}_n$.

We again write

$$\begin{aligned}\mathbb{E}\big[\mathbb{1}^{kl}_{e,e'}\big] = \mathbb{E}\big[\mathbb{1}^{kl}_{e,e'}\mathbb{1}_{f^{(k)}\text{ is singular}}\big] &+ \mathbb{E}\big[\mathbb{1}^{kl}_{e,e'}\mathbb{1}_{f^{(k)}\text{ is regular and }f^{(k+l)}\text{ is singular}}\big]\\ &+ \mathbb{E}\big[\mathbb{1}^{kl}_{e,e'}\mathbb{1}_{f^{(k)}\text{ and }f^{(k+l)}\text{ are regular}}\big].\end{aligned}$$

So, with exactly the same arguments as in the previous proof, we obtain

$$\begin{aligned}\mathbb{E}\big[\mathbb{1}^{kl}_{e,e'}\big] \leq\ & \frac{12k}{4n-2k+1}\frac{3}{4n-2k+1}\\ &+ \frac{1}{3(4n-2k+1)}\frac{12l}{4n-2k-2l+1}\frac{3}{4n-2k-2l+1}\\ &+ \frac{1}{3(4n-2k+1)}\frac{1}{3(4n-2k-2l+1)}\end{aligned}$$

So, multiplying this with $120n$ and using that $k, l = o(n^{1/3})$ uniformly gives

$$\mathbb{E}[E'_{kl}] = o(n^{-2/3})$$

as $n \to \infty$, where the implied constant is uniform over k, l. Now summing over k and l gives

$$\mathbb{E}[E_{KL}] = o(1)$$

as $n \to \infty$. □

2.5. Betti numbers. We give here the proof of our estimates on Betti numbers of M_n (or N_n) in Theorem 1.1 (d). First, a generating family for $H_1(M_n, \partial M_n)$ is given by images of the edges of $\mathcal{T}_n$. Applying Markov's inequality to Theorem 2.4 (a), we get that

$$b_1(M_n, \partial M_n) = o(\theta(n)) \quad \text{as } n \to \infty$$

for any function $\theta\colon \mathbb{N} \to \mathbb{R}$ such that $\lim_{n\to\infty} \theta(n)/\log(n) = +\infty$, establishing the first estimate. From the connectedness of the boundary Theorem 2.1 (a), it follows that with asymptotic probability 1, we get an exact sequence

$$0 \to H_2(M_n) \to H_2(M_n, \partial M_n) \to H_1(\partial M_n) \to H_1(M_n) \to H_1(M_n, \partial M_n) \to 0$$

with asymptotic probability 1. Now this exact sequence and Poincaré duality (the "half lives, half dies" argument) imply that

$$b_1(M_n) = \frac{1}{2} b_1(\partial M_n) + b_1(M_n, \partial M_n),$$

and together with the estimate for $b_1(M_n, \partial M_n)$ Theorem 2.1 (b), we can conclude that with asymptotic probability 1, we have

$$|b_1(M_n) - n| = o(\theta(n))$$

for any function $\theta\colon \mathbb{N} \to \mathbb{R}$ that grows super-logarithmically, which is the second estimate.

2.6. Heegaard genus. Here we prove the estimates on Heegaard genus of the double $\mathcal{D}M_n$ (or $\mathcal{D}N_n$) of Theorem 1.1 (c), following an argument of Nathan Dunfield. Recall that E is the number of edges in the triangulation $\mathcal{T}_n$ (equivalently the number of interior edges in the cellulation of M_n). We will first prove that

$$g(\mathcal{D}M_n) \leq n + 1 + E, \tag{2.5}$$

which in view of Theorem 2.4 implies the upper bound we are after.

To prove (2.5) we observe that M_n minus regular neighbourhoods of its interior edges is a handlebody of genus $n + 1$, as it is a regular neighbourhood of the dual graph to the cellulation of M_n in truncated tetrahedra, which is a 4-valent graph on n vertices. For each edge e, we write its regular neighbourhood U_e as $D_e \times e$, where D_e is a disc. We split it as $D_e = D_e^1 \cup D_e^2$, where D_e^i are half-discs, and we put $U_e^i = D_e^i \times e$; note that $M_n \cup \bigcup_e U_e^i$ is still a handlebody of genus $n + 1$. We consider the two copies $M_n, \bar{M}_n$ of M_n in $\mathcal{D}M_n$ (so $\mathcal{D}M_n = M_n \cup \bar{M}_n$), for a subset $W \subset M_n$ we denote by $\bar{W}$ its image in $\bar{M}_n$, and we put

$$H_n^1 = \left(M_n \setminus \left(\bigcup_e U_e^1 \right) \right) \cup \left(\bigcup_e \bar{U}_e^2 \right),$$

which is just $M_n \cup \bigcup_e U_e^1$ with 1-handles attached (one for each edge), so it is a handlebody of genus $n + E + 1$. Similarly

$$H_n^2 = \left(\bar{M}_n \setminus \left(\bigcup_e U_e^2 \right) \right) \cup \left(\bigcup_e \bar{U}_e^1 \right)$$

is a handlebody of the same genus. Now $\mathcal{D}M_n = H_n^1 \cup H_n^2$ and this proves (2.5).

For the lower bound, we observe that the long exact sequence associated with $\partial M_n \to M_n \cup M_n \to \mathcal{D}M_n$ reduces to

$$\begin{aligned} 0 \to H_2(M_n) \oplus H_2(M_n) \to H_2(\mathcal{D}M_n) \to H_1(\partial M_n) \\ \xrightarrow{i} H_1(M_n) \oplus H_1(M_n) \to H_1(\mathcal{D}M_n) \to 0. \end{aligned}$$

Since $b_2(\mathcal{D}M_n) = b_1(\mathcal{D}M_n)$ (Poincaré duality) and by our results on Betti numbers and genus of the boundary $2b_1(M_n) = b_1(\partial M_n)$ up to super-logarithmic error, it follows that

$$b_1(\mathcal{D}M_n) = \mathrm{rank}(i) + \theta_0(n)$$

with $\theta_0(n)$ super-logarithmic. As i is a diagonal embedding, we have

$$\mathrm{rank}(i) \leq b_1(M_n) = n$$

up to super-logarithmic error, so we can conclude that $b_1(\mathcal{D}M_n) \geq n - \theta_1(n)$ with θ_1 super-logarithmic. On the other hand $g(\mathcal{D}M_n) \geq b_1(\mathcal{D}M_n)$, and this proves the lower bound. □

3. Geometry

In this section we combine the combinatorial results from the previous section with hyperbolic geometry. Recall that M_n denotes the compact manifold with boundary associated to $\mathcal{T}_n$. Our main goal is to prove the following theorem.

Theorem 3.1 (Geometry). *We have*

$$\lim_{n\to+\infty} \mathbb{P}[M_n \textit{ carries a hyperbolic metric with totally geodesic boundary}] = 1.$$

This metric has the following properties

(a) *The hyperbolic volume* $\mathrm{vol}(M_n)$ *of* M_n *satisfies*

$$\mathrm{vol}(M_n) \sim n \cdot v_O \quad \textit{as } n \to \infty$$

in probability.

(b) *There exists a constant* $c_\lambda > 0$ *so that the first discrete Laplacian eigenvalue* $\lambda_1(M_n)$ *of* M_n *satisfies*

$$\lim_{n\to+\infty} \mathbb{P}[\lambda_1(M_n) > c_\lambda] = 1.$$

(c) *There exists a constant* $c_d > 0$ *such that the diameter* $\mathrm{diam}(M_n)$ *of* M_n *satisfies*

$$\lim_{n\to+\infty} \mathbb{P}[\mathrm{diam}(M_n) < c_d \log(\mathrm{vol}(M_n))] = 1.$$

(d) *There exists a constant* $c_s > 0$ *such that the systole* $\mathrm{sys}(M_n)$ *of* M_n *satisfies*

$$\lim_{n\to+\infty} \mathbb{P}[\mathrm{sys}(M_n) > c_s] = 1.$$

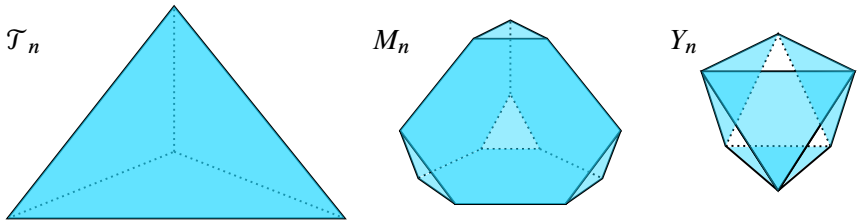

Figure 5. The three building blocks for $\mathcal{T}_n$, M_n and Y_n, respectively. The faces that the gluing is performed along are shaded.

(e) *For every $\varepsilon > 0$, we have*

$$\lim_{n\to+\infty} \mathbb{P}\left[\operatorname{sys}(\mathcal{D}M_n) < \frac{1}{n^{1/4-\varepsilon}}\right] = 1.$$

The same holds for the minimal length among arcs in M_n that are homotopically non-trivial relative to ∂M_n.

We will prove this theorem in multiple steps. The first is hyperbolisation, which follows from Lemmas 3.6 and 3.7). The asymptotic behavior of the volume is then determined in Proposition 3.10 and the spectral gap is proven in Proposition 3.8. We prove the bounds on the diameter and systole in Proposition 3.11.

Finally, we will also prove that the Benjamini–Schramm limit of the sequence $(M_n)_n$ is the "octatree" that we introduce here.

3.1. Random models for hyperbolic manifolds. We will first describe two manifolds associated to an element $\omega \in \Omega_n$ in our probability space Ω_n. The first is a cusped hyperbolic manifold Y_n and the second is the manifold M_n that we saw in the previous section, but now viewed as a Dehn filling of Y_n.

Figure 5 gives a topological picture of what is going on. We already associated a manifold M_n to $\mathcal{T}_n$ by truncating all the tetrahedra involved at their vertices. If we now contract the edges in the interior of this manifold and remove the resulting vertices, we obtain a new manifold Y_n that is built out of a gluing of octahedra. The link of the octahedra's ideal vertices in this manifold are annuli, and we can fill them with cylinders to go back to the compact manifold.

In what follows we describe this in some more detail.

3.1.1. Manifolds with cusps and boundary. Let O be the ideal regular octahedron in $\mathbb{H}^3$ (it can be realised as the convex hull of the vertices of a regular octahedron on the boundary at infinity $\mathbb{S}^2$). Its dihedral angles are right angles and its faces are ideal triangles. We orient each face with its outward normal.

We take n copies of O which we label as follows: for each copy we attribute a label in $\{1, \dots, 4n\}$ to four of its faces so that no two of them are adjacent (and we ask

that the labelling map be injective). Each of the unlabelled faces is then determined by the labels of the three faces adjacent to it and since it is oriented we can identify it with a 3-cycle on their labels.

This setting is similar to that of Section 2.1 and we can perform the same random construction from it: we partition the non-labelled faces uniformly randomly into pairs, and we glue the two faces in a pair in a uniformly randomly chooses orientation-reversing way.

The resulting octahedral complex is a non-compact manifold with boundary, and by endowing each O with its hyperbolic structure we obtain that the result is a complete orientable hyperbolic manifold with totally geodesic boundary. We denote by X_n the random hyperbolic manifold with boundary we constructed. We will also consider Y_n where we condition on there not being any loops or bigons in the graph dual to the tessellation by octahedra. We record the hyperbolic structure in a lemma in order to be able to refer to it later on.

Lemma 3.2. *The manifolds X_n and Y_n carry complete hyperbolic metrics of finite volume with totally geodesic boundary.*

We denote by Θ_n the (finite) set of all hyperbolic manifolds obtained by gluing n octagons in this fashion. Let Y be in some Θ_n. Each cusp c of Y is tessellated by squares. For any given c let $\ell(c)$ be their number (we will also call this the *length* of c) and for $k \in \mathbb{N}$ let $B_k(Y)$ be the number of cusps of Y with $\ell = k$.

Lemma 3.3. *Let $C_k := B_k(Y_n)$, then as random variables $C_k = E_k$ (where E_k is the random variable introduced at the beginning of Section 2.1 above).*

3.1.2. Compact manifolds. Recall from Section 2.1 that N_n is a manifold with boundary obtained from randomly gluing truncated tetrahedra along their faces. Moreover, M_n is a random manifold that has the distribution of N_n, conditioned on the dual graph $\mathcal{G}_n$ not having any loops and multiple edges.

Let us now describe how M_n (and N_n respectively) can be obtained from Y_n (and X_n, respectively) via Dehn filling.

Let $Y \in \Theta_n$. Its boundary S is a hyperbolic surface with cusps. Moreover, there is a pairing on the cusps where we associate two cusps of S if they are asymptotic in Y. If we remove a horospherical neighbourhood of each cusp we obtain a compact manifold $\overline{Y}$ whose boundary is made up of S together with closed annuli linking paired cusps, and by the thick-thin decomposition Y is homeomorphic to $\overline{Y}$ minus the annuli. We can then perform surgery on $\overline{Y}$ as follows: to each annulus we glue a cylinder $D \times [0,1]$ along the boundary $[0,1] \times \partial D$. Another way to describe it is that it is the restriction to Y of the unique Dehn surgery on the double $\mathcal{D}Y$ which is equivariant with respect to the reflection in ∂Y. We obtain a compact manifold M with boundary $\overline{S}$. We denote by Ξ_n the set of such manifolds M obtained from $Y \in \Theta_n$.

Note that M_n (and N_n, respectively) has the same distribution as the Dehn filling of Y_n (and X_n, respectively) described above. We again record this in a proposition to be able to refer to it later.

Proposition 3.4. *The variables M_n (respectively, N_n) and the Dehn filling of Y_n (respectively, X_n) are homeomorphic*[2].

3.2. Bounds for Dehn surgeries and hyperbolicity. In this section we prove that M_n is hyperbolic with asymptotic probability 1 (this is part of the statement of Lemma 3.7) and we give precise bounds for the variation in geometry between Y and M.

Our proof goes in two steps. First we use Andreev's theorem to control what happens when "small" cusps of Y are filled and after this we use recent results by Futer–Purcell–Schleimer to control the change in geometry when the "large" cusps are filled.

3.2.1. Andreev's theorem. To construct explicit hyperbolisations we will need Andreev's theorem describing acute-angled polyhedra in $\mathbb{H}^3$. We refer to [40] for a proof of this result. Before giving the statement we recall that given a combinatorial 2-polyhedron P, with dual graph (recall that this is the graph on 2-dimensional faces of P with an edge between two faces for each edge they share) H, a circuit in H is said to be *prismatic* if for any edge in the circuit, the endpoints of the corresponding edge of P are distinct. The following is a combination of Theorem 1.4 and Proposition 1.5 in [40].

Theorem 3.5 (Andreev's theorem). *Let P be an abstract polyhedron with at least six faces, and α a function from edges of P to $(0, \pi/2]$. Then there exists a realisation of P in $\mathbb{H}^3$ which is of finite volume and whose dihedral angles are given by α if and only if the following conditions are satisfied:*

- *For any three edges e_1, e_2, e_3 meeting in one vertex, we have*
$$\sum \alpha(e_i) \geq \pi$$
(equality occurring if and only if the vertex is ideal).
- *If $(e_1, \dots, e_k)$ is a prismatic k-circuit with $k = 3, 4$ then*
$$\sum \alpha(e_i) < (k-2)\pi.$$

3.2.2. Filling small cusps. We start by filling the small cusps of Y_n. These will be cusps of length $\leq n^{1/4}$. Note that the resulting manifold, that we call Z_n is rigid – i.e. if we can find a complete hyperbolic metric of finite volume on it, it is unique up to isometry – by Mostow–Prasad rigidity of its double. We prove the following properties for Z_n; note that this is the only part of our proof of hyperbolisation where we use the assumption that the dual graph $\mathcal{G}_n$ is simple.

[2] That is, for every $\omega \in \Omega$ there exists an homeomorphism between $M_n(\omega)$ and the filling of $Y_n(\omega)$ (resp., between $N_n(\omega)$ and the filling of $X_n(\omega)$).

Lemma 3.6. *There exists $J_0 > 0$ such that the following holds for any Margulis constant $\delta > 0$ and any $\varepsilon > 0$. For any $Y \in \Theta_n$,*

- *let $c_1, \dots, c_m$ be the cusps of Y of length at most $n^{1/4}$,*
- *let Z be the manifold obtained by filling $c_1, \dots, c_m$,*
- *let Z_1 be the union of all octahedra of Y containing one of the c_i and Z_2 its complement.*

Then with probability at least $1 - \varepsilon$ in the model Y_n for n large enough, we have that

- *Viewing Y as the complement of the core arcs in Z, the δ-thick part of the image of Z_1 in Z is J_0-bi-Lipschitz to that of Z_1;*
- *The image of Z_2 in Z_σ is isometric to Z_2.*

Proof. Let $O_1, \dots, O_n$ be the octahedra tessellating Y and O_i, $1 \leq i \leq k$ those containing a cusp c_i with $\ell(c_i) \leq n^{1/4}$ and $O_{k+1}, \dots, O_n$ the remaining ones. By Theorem 2.4 (d), we may assume that each O_i contains exactly one such cusp. We have $Z_1 = O_1 \cup \cdots \cup O_k$ and $Z_2 = O_{k+1} \cup \cdots \cup O_n$. Then the part of boundary of Z_1 and Z_2 along which they are glued is a disjoint union of ideal regular squares.

To construct the hyperbolic structure on the filled manifold we replace $O_1, \dots, O_k$ by polyhedra constructed as follows. First we assume that $\ell(c_i) \geq 4$. Consider the following polyhedron, which is an octahedron on which 1 vertex has been replaced by an edge (marked red in the picture):

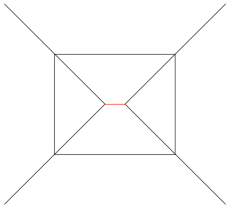

There are no prismatic 3- or 4-circuits in the dual graph so it follows from Andreev's theorem (Theorem 3.5) that for $l \geq 4$, this has the structure of an hyperbolic polyhedron P_l with right angles at all edges except the red one which has angle $2\pi/l$ (we need $l \geq 4$ for this not to be obtuse).

If $\ell(c_i) = 3$ then we can still construct P_3 as follows: the combinatorial polyhedron has a symmetry along the red edge, which decomposes it as the double of the following polyhedron along the blue-coloured face:

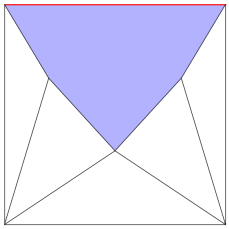

and the latter has no prismatic circuits, so it admits a hyperbolic structure with right angles on the black edges and $\pi/3$ on the red edge by Andreev's theorem.

For $l \geq 4$, let Q_l be the hyperbolic manifold obtained by gluing l copies of O in a circular pattern, along disjoint faces sharing an ideal vertex. The faces opposite to the glued faces form a union of disjoint ideal triangles in Q_l. Dehn surgery on Q_l amounts to replacing each copy of O by a copy of the polyhedron P_l (the edge with angle $2\pi/l$ replacing the ideal vertex on which surgery is done), to obtain a polyhedron Q_l' whose boundary is two l-gons, l regular ideal squares and $2l$ ideal triangles meeting at right angles. This is illustrated in the following figure, where these are coloured blue and the central edge red:

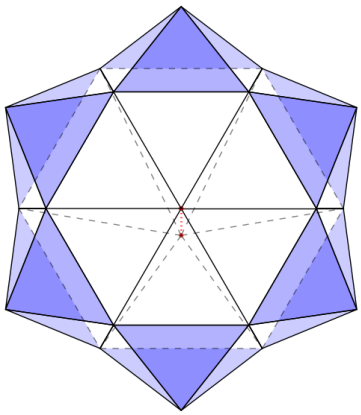

Now if M is generic in the sense of Theorem 2.4 (d), and furthermore all edges of length at most $K(n)$ (where $K(n)$ is any $o(n^{1/4})$) are simple (which is generic by Theorem 2.4 (b), (c)) then Z_1 is a disjoint union of Q_ls and filling the small cusps amounts to replacing each of these with a Q_l'. In particular, we can glue the rest of the octahedra in the pattern given by $\mathcal{G}$ to obtain Z. This proves that the image of Z_2 in Z is isometric to Z_2. It follows from Thurston's hyperbolic Dehn surgery theorem (for example) that P_l converges in Gromov–Hausdorff topology (pointed anywhere in the thick part) to O so that the δ-thick parts of P_l are uniformly (independently of l) bi-Lipschitz to the δ-thick part of O, and it follows immediately that Z_1 is uniformly (independently of generic $\mathcal{G}$) bi-Lipschitz to its image in Z. □

3.2.3. Filling large cusps. In order to fill the cusps remaining in Z, we rely on results of Futer–Purcell–Schleimer. Again using Z to denote the manifold obtained from Y by filling its small cusps, we will think of $Y \subset Z \subset M$. We have the following lemma.

Lemma 3.7. *Let Y, Z be as in Lemma 3.6. Then, for any $\delta > 0$ (smaller than the Margulis constant for $\mathbb{H}^3$) and any $\eta > 0$, with probability for the Y_n model converging to 1 as $n \to +\infty$, the following hold:*

- *There exists Riemannian metrics $g_0, g_{2\pi}$ on M such that (Y, g_0) is the complete hyperbolic structure and the completion of $(M, g_{2\pi})$ is a compact hyperbolic manifold with totally geodesic boundary which is diffeomorphic to the Dehn filling of Y.*
- *Moreover, $Z_{\geq 6\delta/5} \subset M_{\geq \delta} \subset Z_{\geq 5\delta/6}$ and these inclusions are $(1+\eta)$-Lipschitz.*

Proof. Let $c_{m+1}, \dots, c_h$ be all remaining cusps in Z (recall that $c_1, \dots, c_m$ were the cusps of length at most $n^{1/4}$). Realising them as arbitrary horosphere quotients in Z these cusps have an area a_j and a length of the vertical curve l_j. Let $L_j = l_j/\sqrt{a_j}$; this does not depend on the arbitrary choice of horospheres (as long as their quotients are homeomorphic to 2-tori). Following [23, Definition 1.3], we define $L > 0$ by

$$\frac{1}{L^2} = \sum_{j=m+1}^{h} \frac{1}{L_j^2}.$$

For $m+1 \leq j \leq h$, let k_j be the number of cusps c_i, $1 \leq i \leq m$, which share an octahedron with c_j – in other words, the number of small cusps that share an octahedron with c_j. We claim that with probability tending to 1, we have

$$\max(k_j) \ll \frac{\log(n)}{n^{1/24}} \ell(c_j).$$

To prove this, we separate two cases: First, when $n^{1/4} \leq \ell(c_j) \leq n^{7/24}$ (note $7/24 = 1/3 - 1/24 = 1/4 + 1/24$), by Theorem 2.4 (d) with probability tending to 1, we have $k_j = 0$ for all these j. In the remaining cases, we have that $k_j \ll h \cdot n^{1/4}$ and since $h = E$ (the number of edges in the original triangulation) and $E \ll \log(n)$ by Theorem 2.4 (a), it follows that if $\ell(c_j) \geq n^{7/24}$, then

$$k_j \ll \log(n) n^{1/4} \leq \frac{\log(n)}{n^{1/24}} \ell(c_j),$$

which finishes the proof of the claim.

Now it follows from Lemma 3.6, if $m+1 \leq j \leq h$ we have at least $\ell(c_j) - k_j$ regular squares in the tessellation of c_j in Z, and so

$$\max_{m+1 \leq j \leq h} \left(\frac{1}{L_j^2} \right) \leq \max_{m+1 \leq j \leq h} \left(\frac{1}{\ell(c_j) - k_j} \right) \ll \frac{1}{n^{1/4}}.$$

Using Theorem 2.4 (a) again, with asymptotic probability 1, we have

$$\sum_{j=m+1}^{h} \frac{1}{L_j^2} \leq \frac{\alpha \log(n)^2}{n^{1/4}},$$

and in particular, $L^2 \geq n^{1/8}$ with asymptotic probability 1. With this, our lemma is an immediate consequence of [23, Theorem 9.30] as this implies that, for any fixed δ with asymptotic probability 1, we have that Z satisfies the hypothesis (9.32) in this statement (with $\epsilon = \delta$). □

3.3. Expansion. For a non-necessarily compact Riemannian manifold V we denote by $\lambda_1(V)$ the bottom of the discrete spectrum of the Laplace–Beltrami operator of V (if V has non-empty boundary we take it to be the minimum between the spectra with Neumann or Dirichlet conditions). Using the results from the preceding section and comparison results in spectral geometry due to Mantuano and Hamenstädt, we prove the following.

Proposition 3.8. *There exists $c_\lambda > 0$ such that*

$$\lim_{n\to+\infty} \mathbb{P}[\lambda_1(M_n) \geq c_\lambda] = 1.$$

Proof. One way to prove Proposition 3.8 would be a minor modification of the argument in [10, Section 4], based on the Cheeger constant of M (see also the appendix to [9]). We will instead work with the double $N = \mathcal{D}M$ to be able to apply directly the results by Hamenstädt and Mantuano.

In this proof we work with a $M \in \Xi_n$, which carries a hyperbolic metric with totally geodesic boundary (which we proved happens with asymptotic probability 1).

Let $N = \mathcal{D}M$ be the double of M along its (totally geodesic) boundary, which is a closed hyperbolic manifold. The space $L^2(\mathcal{D}M)$ decomposes into the direct sum of ± 1 eigenspaces for the symmetry in ∂M and these spaces correspond to spaces of functions on M satisfying Neumann or Dirichlet conditions on ∂M. So we have that $\lambda_1(\mathcal{D}M) = \lambda_1(M)$, and in the rest of the proof we will be concerned with establishing that $\lambda_1(\mathcal{D}X)$ is bounded away from 0.

We fix a Margulis constant δ. By [27, Theorem 1] we have that $\lambda_1(N) > \lambda_1(N_{\geq\delta})/3$ (or $\lambda_1(N)$ is uniformly bounded away from zero, in which case we are finished). So we must bound $\lambda_1(N_{\geq\delta})$ from below.

To do so we will use the following result which is an immediate application of [37, Theorem 3.7]:

- let V be a compact hyperbolic three-manifold with $\mathrm{inj}(V) \geq \delta$ (for example the δ-thick part of a manifold of finite volume if its boundary is smooth),
- let $\mathcal{X}$ be a maximal $\delta/2$-separated subset of V, on which we put the graph structure where there is an edge between $x, y \in \mathcal{X}$ if they are at distance at most 2δ from each other in V ($\mathcal{X}$ is called a discretisation of V, and δ its mesh).

Then there is $c > 0$ depending only on δ such that $\lambda_1(V) \geq c\lambda_1(\mathcal{X})$.

We record the following well-known facts which we will use to compare between discretisations of our manifolds Y, Z and M.

Lemma 3.9. *Let E_1, E_2 be two metric geodesic spaces and $\mathcal{X}_i$ a discretisation of E_i.*

(1) *The inclusion $\mathcal{X}_i \subset E_i$ is a quasi-isometry with constants depending only on the mesh.*

(2) *If φ is a quasi-isometry from E_1 to E_2 and q is a nearest-neighbour projection from E_2 to $\mathcal{X}_2$ then $q \circ \varphi$ induces a quasi-isometry from $\mathcal{X}_1$ to $\mathcal{X}_2$, whose quasi-isometry constants depend only on those of φ and on the meshes of $\mathcal{X}_1$, $\mathcal{X}_2$.*

Proof. To prove the "quasi-isometric embedding" part of the first statement take $x, x' \in \mathcal{X}_i$, then the nerve of a cover of a geodesic in E_i between x, x' by δ-balls (where δ is the mesh) centred in $\mathcal{X}_i$ gives a path in X_i with length at most $2\delta^{-1} d_{E_i}(x, x')$; the reverse inequality is immediate. It is also immediate to check that a quasi-inverse is given by any nearest-point projection, which is a quasi-isometry whose constants also depend only on the mesh. The second point immediately follows. □

Let $\mathcal{G}_1$ be a discretisation of $N_{\geq\delta}$ with mesh $\delta/2$. Let $\varphi\colon N_{\geq\delta} \to \mathcal{D}Y_{\geq\delta}$ be defined as follows: by Lemma 3.7, we have $M_{\geq\delta} \supset Z_{\geq 6\delta/5}$, so we can define a retraction $\pi_1\colon M_{\geq\delta} \to Z_{\geq 6\delta/5}$ by following the geodesic flow in the direction orthogonal to the boundary $\partial Z_{\geq 5\delta/6}$ – if there are multiple possible directions to follow, i.e. if we are on a core geodesic, we choose one arbitrarily. We extend this to $N_{\geq\delta} = \mathcal{D}M_{\geq\delta}$ by symmetry. By the rest of the statement of the lemma, this is $(1+\eta)$-bi-Lipschitz on $N_{\geq\delta}$, and since $M_{\geq\delta} \subset Z_{\geq 5\delta/6}$ for $x \in N_{\geq\delta} \setminus \mathcal{D}Z_{\geq 6\delta/5}$, we have $d_M(x, \mathcal{D}Z_{\geq 6\delta/5}) \leq a\delta$ for an absolute a. It follows that π is a $(1+\eta, b\eta\delta)$-quasi-isometry, for some absolute b. Now we extend by symmetry the J_0-bi-Lipschitz map ψ from $Z_{\geq 6\delta/5}$ to $Y_{\geq 6\delta/5}$ given by Lemma 3.6 and put $\varphi = \psi \circ \pi$, which from what we said is a quasi-isometry from $N_{\geq\delta}$ to $\mathcal{D}Y_{\geq 6\delta/5}$ with constants depending only on δ, η.

Applying the lemma to φ we get that $\mathcal{G}_1$ and an arbitrary discretisation $\mathcal{G}_2$ of $Y_{\geq 6\delta/5}$ with mesh $\delta/2$ are quasi-isometric to each other, with constants depending only on δ. Let $\mathcal{D}\mathcal{G}$ be the graph dual to the tessellation of N by octahedra; it is obtained by taking two copies of $\mathcal{G}$ and adding four edges between every pair of corresponding vertices. On $\mathcal{G}_2$ we can define a map to $\mathcal{D}\mathcal{G}$ by mapping all vertices in a given octahedron of M to the centre of that octahedron (we choose arbitrarily for vertices on the boundary between two faces). This is a quasi-isometry with constants depending only on δ (via the diameter of $O_{\geq\delta}$). Composing φ with this we get a quasi-isometry from $\mathcal{G}_1$ to $\mathcal{G}$. By [37, Theorem 2.1] it follows that $\lambda_1(\mathcal{G}_1) \geq c'\lambda_1(\mathcal{D}\mathcal{G})$, where c depends only on δ. As $\mathcal{G}_1$ is a discretisation of $\mathcal{D}X_{\geq\delta}$, by loc. cit., Theorem 3.7 (see the statement at the beginning of the section), it finally follows that $\lambda_1(\mathcal{D}X_{\geq\delta}) \geq c''\lambda_1(\mathcal{D}\mathcal{G})$. It is well known that $\mathcal{G}_n$ is an expander asymptotically almost surely; the sharpest bounds on its spectral gap are due to Friedman [22]. The double $\mathcal{D}\mathcal{G}$ is quasi-isometric to $\mathcal{G}$ with uniform constants via the inclusion, so loc. cit., Theorem 2.1 gives us that it is also an expander a.a.s. We conclude that $\lambda_1(N_{\geq\delta})$, and hence also $\lambda_1(N)$, is bounded away from zero. □

3.4. Volumes. If the manifold $X \in \Xi_n$ is hyperbolic, then it has a hyperbolic volume $\mathrm{vol}(X)$. Otherwise, we take $\mathrm{vol}(X) = 0$. Recall that v_O denotes the volume of the right-angled hyperbolic octahedron.

Proposition 3.10. *We have*

$$\mathrm{vol}(M_n) \sim n \cdot v_O \quad \text{as } n \to \infty$$

in probability.

Proof. If $M \in \Xi_n$ is hyperbolic then it is a Dehn surgery on a manifold $Y \in \Theta_n$. The latter is a union of n copies of the octahedron O. As the hyperbolic volume decreases under hyperbolic Dehn surgery we get that $\mathrm{vol}(X) \leq n v_O$.

All statements in the following paragraph hold asymptotically almost surely. By Lemma 3.6, we have that

$$\mathrm{vol}(Z) \geq \mathrm{vol}(Y) - O(n^{1/4} \log n) = n \,\mathrm{vol}(O) - O(n^{1/4} \log n) \tag{3.1}$$

(at most $4n^{1/4} \log n$ octahedra are changed from Z to Y since this is an upper bound for the number of squares in the small cusps in a generic Y by Theorem 2.4 (a)). By Lemma 3.7, for any positive δ and η, we have, since if two Riemannian metrics on a manifold are η'-bi-Lipschitz to each other their volume forms are $O(\eta)$ pointwise close to each other, that

$$\mathrm{vol}(M) \geq (1 - c\eta)\,\mathrm{vol}(Z_{\geq 6\delta/5})$$

for some $c > 0$ independent of δ, η. The thin part of Z is made of $O(\log(n))$ tubes coming from the Dehn filling of small cusps, so constituting a volume $O(\log n)$, and the rest is cusps. For a cusp C, we have $\mathrm{vol}(C) = \mathrm{Area}(\partial C)$, and the boundary of the cusps of $Z_{\geq 6\delta/5}$ is made of $n - O(\log n)$ euclidean squares with edge length $O(\delta)$. It follows that

$$\mathrm{vol}(M) \geq (1 - c\eta)\,\mathrm{vol}(Z_{\geq 6\delta/5}) \geq (1 - c\eta)\,\mathrm{vol}(Z) - \delta O(n). \tag{3.2}$$

Taking δ and η to 0 we get the statement we want from (3.1) and (3.2). □

3.5. Diameter and systole. Lemmas 3.6 and 3.7 above, together with our combinatorial bounds (Theorem 2.4) and results by Futer–Purcell–Schleimer and Bollobás–Fernandez-de-la-Vega, imply the following bounds on the diameter and systole of M_n and $\mathcal{D}M_n$:

Proposition 3.11. (a) *There exists a constant $c_d > 0$ such that the diameter* $\mathrm{diam}(M_n)$ *of M_n satisfies*

$$\lim_{n \to +\infty} \mathbb{P}[\mathrm{diam}(M_n) < c_d \log(\mathrm{vol}(M_n))] = 1.$$

(b) *There exists a constant $c_s > 0$ such that the systole* $\mathrm{sys}(M_n)$ *of M_n satisfies*

$$\lim_{n \to +\infty} \mathbb{P}[\mathrm{sys}(M_n) > c_s] = 1.$$

(c) *For every $\varepsilon > 0$, we have*

$$\lim_{n\to+\infty} \mathbb{P}\left[\frac{1-\varepsilon}{4n} < \mathrm{sys}(\mathcal{D}M_n) < \frac{1}{n^{1-\varepsilon}}\right] = 1.$$

The same holds for the minimal length among arcs in M_n that are homotopically non-trivial relative to ∂M_n.

Proof. We start with item (a). Bollobás–Fernandez-de-la-Vega [8] proved that the diameter (in the graph distance) of a random 4-regular graph $\mathcal{G}_n$ on n vertices satisfies

$$\mathrm{diam}(\mathcal{G}_n) \leq \log_3(n) + o(\log(n))$$

in probability.

Again, using results from graph theory about the configuration model (see (2.1)), we may assume that $\mathcal{G}_n$ is conditioned to not have loops or multiple edges, so that $\mathcal{G}_n$ is uniformly quasi-isometric to the δ-thick part of Y_n (with constants that only depend on δ).

Now we pick $\delta > 0$ smaller than the Margulis constant for $\mathbb{H}^3$. Using Lemmas 3.6 and 3.7, plus the fact that the polytopes P_l described in the proof of Lemma 3.6 converge to O, this implies that the δ-thick part $(M_n)_{\geq\delta}$ of M_n is uniformly quasi-isometric to $\mathcal{G}_n$. Hence, there exists a constant $C_\delta > 0$ such that

$$\mathrm{diam}((M_n)_{\geq\delta}) \leq C_\delta \log(n)$$

asymptotically almost surely.

In order to control the diameter of the thin parts of M_n, it is easier to think in terms of Margulis tubes, so we will consider the double $\mathcal{D}M_n$ as a Dehn filling of $\mathcal{D}Y_n$. The Margulis lemma tells us that the thin part of $\mathcal{D}M_n$ consists of standard tubes (see for instance [5, Chapter D]) of the form

$$T_r = \{x \in \mathcal{D}M_n : \mathrm{d}(x,\gamma) < r\},$$

where γ is a simple closed geodesic. As such, the diameter of such a tube is at most

$$2r + \ell(\gamma) \leq 2r + \delta.$$

The length of a meridian on the boundary torus of a standard tube is $2\pi\sinh(r)$. In Y_n, the length of the meridian is a constant multiple of the combinatorial length of the corresponding cusp (and hence bounded by $6n$).

We want to estimate the lengths of meridians in M_n in terms of those in Y_n. To do so we first observe that these lengths are the same between Y_n and Z_n by Lemma 3.6. Then using Lemma 3.7 in the same way that we used to construct a retraction in the proof of Proposition 3.8, we see that there is a bi-Lipschitz map (with constants independent of n) between the boundaries of $M_{\geq\delta}$ and $Z_{\geq\delta}$, which sends meridian

to meridian. It follows that there exists a constant $D_\delta > 0$ such that length of the meridian of any tube in $\mathcal{D}M_n$ is at most $D_\delta \cdot n$. This in turn implies that the radius of each such tube can be bounded by $E_\delta \log(n)$ for some constant $E_\delta > 0$, depending on δ only. Combining this with our estimate on the diameter of the thick part and the estimates on volume from Proposition 3.10 this implies item (a).

We proceed to item (b). We observe that, for δ below the Margulis constant in $\mathbb{H}^3$, the δ-thin part of M_n is simply connected. In particular, any closed geodesic that passes through the $(\delta/2)$-thin part of M_n has length at least $\mathrm{d}((M_n)_{\geq\delta}, (M_n)_{<\delta/2})$, which is uniformly bounded from below (by applying Lemmas 3.6 and 3.7) to the $(\delta/2)$-thick part of M_n). The $(\delta/2)$-thick part of M_n is bi-Lipschitz to the $\widetilde{\delta}$-thick part of Y_n for some uniform $\widetilde{\delta} > 0$ (by Lemmas 3.6 and 3.7). The systole of the latter is bounded from below by the distance between two distinct faces of $O_{\geq\widetilde{\delta}}$, which gives us a lower bound on the systole of the $(\delta/2)$-thick part of M_n. Together with the uniform bound on the length of geodesics that pass through the thin part, this implies a lower bound on $\mathrm{sys}(M_n)$.

The lower bound in item (c) is deterministic and follows directly from [4, Theorem 1.1]. For the upper bound we use our combinatorial bounds again. Recall from Theorem 2.4 (a) that the expected number of interior edges in our complex is logarithmic in n. So, with probability tending to 1 as $n \to \infty$, it is less than n^ε for any ε. This implies there is at least one cusp in Y_n that is incident to more than $n^{1-\varepsilon}$ octahedra. Corollary 6.13 of [23] now immediately implies the result. □

3.6. Benjamini–Schramm convergence.

3.6.1. Coxeter groups. We will use yet another description of the manifolds Y_n using subgroups of a single Coxeter group which could be avoided but we feel simplifies the argument and might be of further interest.

Let T be ideal tetrahedron obtained by cutting O along all of its median hyperplanes. Let Γ_T be the associated reflection group. It is a Coxeter group with presentation

$$\Gamma_T = \langle \sigma, \tau_1, \tau_2, \tau_3 \mid \tau_i^2, \sigma^2, [\tau_i, \tau_j], (\sigma\tau_j)^4 \rangle$$

This group is useful for us because of the following lemma.

Lemma 3.12. *Any $X \in \Theta_n$ is an orbifold cover of T.*

Proof. Let $X \in \Theta_n$ and $\mathcal{G}$ the graph dual to its tessellation by octahedra. Then X is an orbifold cover of O if and only if $\mathcal{G}$ is 4-edge-colourable, which is not always the case.

However, if we replace each vertex of $\mathcal{G}$ by a cube with four outgoing edges placed at pairwise non-adjacent vertices we get a graph $\mathcal{G}'$. We colour its edges as follows: all edges between cubes are coloured with σ, and inside the cube we choose the unique colouring corresponding to the labels (in $\mathbb{Z}/3\mathbb{Z}$) – this makes sense since

the edge corresponds to a face of O and the adjacent edges of O are each specified by a τ_i. This specifies a unique map $X \to T$ which is an orbifold cover. □

3.6.2. Invariant random subgroups. Let G be a Lie group (we will only consider $G = \mathrm{PGL}_2(\mathbb{C})$). We recall from [1] that an *invariant random subgroup* of G is a Borel probability measure on the Chabauty space of closed subgroups of G (a compact Hausdorff topological space the definition of which can be found in loc. cit.) which is invariant under the action of G on this space by conjugation.

An important constructions of such is the following: if $\Lambda \leq G$ is a subgroup whose normaliser is a lattice Γ in G then the closure of the conjugacy class of Λ supports a unique invariant random subgroup (the image of Haar measure on G/Γ). We denote this by μ_Λ.

Using this we can associate an invariant random subgroup to the random variable M_n as follows: let Ξ_n^{hyp} be the subset of manifolds in Ξ_n which support a complete hyperbolic structure with totally geodesic boundary. For $M \in \Xi_n^{\mathrm{hyp}}$ we consider the hyperbolic orbifold on M whose singular locus is its boundary ∂M (a mirror) and the hyperbolic structure on the interior is that of M. This is a compact hyperbolic orbifold and we choose an arbitrary monodromy group $\Gamma_M \leq G$ for it and let

$$\mu_M = \mu_{\Gamma_M}.$$

If $M \notin \Xi_n^{\mathrm{hyp}}$, we take μ_M to be the Dirac mass at the trivial subgroup. We put

$$\mu_n = \sum_{M \in \Xi_n} \mathbb{P}(M_n = M)\mu_M. \tag{3.3}$$

We also need to define some other invariant random subgroups which will play a role in what follows. Consider the ideal octahedron O as a complete hyperbolic orbifold (all faces being mirrors). Let Γ_O be its orbifold fundamental group, which is generated by the reflections on the sides of O. Let Q be the group generated by the rotations of angle $2\pi/3$ in the faces. Then $P = Q \setminus O$ is an hyperbolic orbifold; let Γ_P be its orbifold fundamental group, which we view as a lattice in G (we need this larger group because not every manifold in Ξ_n is an orbifold cover of O).

Since O is right-angled, mapping four reflections on non-adjacent faces to the identity gives a map

$$\pi\colon \Gamma_O \to *_{i=1}^4 \mathbb{Z}/2\mathbb{Z}$$

(each remaining face maps to the generator of one of the free factors), and the latter is isomorphic to $D_\infty * D_\infty$, where $D_\infty = \mathbb{Z}/2\mathbb{Z} * \mathbb{Z}/2\mathbb{Z}$ is the infinite dihedral group. Let O^∞ be the associated cover; it is the infinite hyperbolic polyhedron obtained by gluing copies of O in a 4-valent tree pattern, along non-adjacent faces ($D_\infty * D_\infty$ acts via its action on the 4-valent tree). Note that if we view the mirrors as a totally geodesic boundary O^∞ is the universal cover of any manifold in some Θ_n.

Since Q respects the colouring of the faces of O we have that $\ker(\pi)$ is a normal subgroup in Γ_P. We let μ_{O^∞} be the invariant random subgroup of G associated to the normal subgroup $\ker(\pi) \leq \Gamma_P$. Our main result in this section is then the following.

Proposition 3.13. *The invariant random subgroup μ_n converges to μ_{O^∞}*

Proof. Let μ'_n be the invariant random subgroup associated to the random variable Y_n. We will first prove that the sequence μ'_n converges to μ_{O^∞}. Let Γ_P be the group defined above and $\Lambda = \ker(\pi) \leq \Gamma_P$. Every $M \in \Xi_n$ admits a (possibly non-continuous) piecewise isometric map to O (by mapping its marked octahedra to O). The non-continuity comes from the rotations made when gluing faces so the composition $M \to O \to P$ is continuous and hence a covering map. Let $\Gamma_n \leq \Gamma_P$ be the invariant random index-$12n$ subgroup corresponding to M_n.

The Schreier graph of Γ_P/Γ_n with loops removed is obtained from the graph dual to the tessellation of M_n by replacing each vertex with a fixed graph $\mathcal{Q}$. The dual graphs follow the same distribution as the configuration model, and as this model of graphs BS-converges to the tree (this follows from [6]) we get that the random variable Γ_n converges in distribution to Λ (since the Schreier graph of Γ_P/Λ is obtained from the tree by replacing vertices with $\mathcal{Q}$). Now μ_n is the IRS obtained by induction of Γ_n from Γ_0 to $\mathrm{PGL}_2(\mathbb{C})$, and μ_{O^∞} by induction of Λ (see [2, Section 11.1] for the definition of induction). As induction is continuous we get that μ'_n converges to μ_{O^∞}.

Now if μ''_n is the IRS associated to Z it follows immediately from the convergence of μ'_n together with Lemma 3.6 that we also have $\lim(\mu''_n) = \mu_{O^\infty}$.

We pass to the larger space of random pointed metric spaces with Benjamini–Schramm topology (see [25, Section 5]). If μ is an IRS of $\mathrm{PGL}_2(\mathbb{C})$ we denote by $\mu^{\geq\delta}$ the random pointed compact manifold with boundary which comes from conditioning the point to be in the thick part (note that doing so we lose all invariance properties). It follows from the previous paragraph that $(\mu''_n)^{\geq\delta}$ converges to $\mu_{O^\infty}^{\geq\delta}$ and from Lemma 3.7 that $\mu_n^{\geq\delta}$ also does.

Now the map $(M, x) \mapsto (M_{\geq\delta}, x)$ is an homeomorphism onto its image: it is continuous (immediate) and injective (the boundary of the thin part determines the complex length of the core geodesic if it is a tube, and the isometry class of the cross-section if it is a torus), and the space of hyperbolic manifolds pointed in their thick part is compact. We can thus conclude that μ_n converges to μ_{O^∞}. □

References

[1] M. Abert, N. Bergeron, I. Biringer, T. Gelander, N. Nikolov, J. Raimbault, and I. Samet, On the growth of L^2-invariants for sequences of lattices in Lie groups. *Ann. of Math. (2)* **185** (2017), no. 3, 711–790 Zbl 1379.22006 MR 3664810

[2] M. Abert, N. Bergeron, I. Biringer, T. Gelander, N. Nikolov, J. Raimbault, and I. Samet, On the growth of L^2-invariants of locally symmetric spaces, II: Exotic invariant random subgroups in rank one. *Int. Math. Res. Not. IMRN* (2020), no. 9, 2588–2625 Zbl 1483.22007 MR 4095420

[3] H. Baik, D. Bauer, I. Gekhtman, U. Hamenstädt, S. Hensel, T. Kastenholz, B. Petri, and D. Valenzuela, Exponential torsion growth for random 3-manifolds. *Int. Math. Res. Not. IMRN* (2018), no. 21, 6497–6534 Zbl 1448.57026 MR 3873535

[4] A. Basmajian, Tubular neighborhoods of totally geodesic hypersurfaces in hyperbolic manifolds. *Invent. Math.* **117** (1994), no. 2, 207–225 Zbl 0809.53052 MR 1273264

[5] R. Benedetti and C. Petronio, *Lectures on hyperbolic geometry*. Universitext, Springer, Berlin, 1992 Zbl 0768.51018 MR 1219310

[6] B. Bollobás, A probabilistic proof of an asymptotic formula for the number of labelled regular graphs. *European J. Combin.* **1** (1980), no. 4, 311–316 Zbl 0457.05038 MR 595929

[7] B. Bollobás, *Random graphs*. Second edn., Cambridge Stud. Adv. Math. 73, Cambridge Univ. Press, Cambridge, 2001 Zbl 0979.05003 MR 1864966

[8] B. Bollobás and W. Fernandez de la Vega, The diameter of random regular graphs. *Combinatorica* **2** (1982), no. 2, 125–134 Zbl 0505.05053 MR 685038

[9] E. Breuillard, Expander graphs, property (τ) and approximate groups. In *Geometric group theory*, pp. 325–377, IAS/Park City Math. Ser. 21, Amer. Math. Soc., Providence, RI, 2014 Zbl 1440.20010 MR 3329732

[10] R. Brooks, The spectral geometry of the Apollonian packing. *Comm. Pure Appl. Math.* **38** (1985), no. 4, 359–366 Zbl 0575.52009 MR 792395

[11] R. Brooks and E. Makover, Random construction of Riemann surfaces. *J. Differential Geom.* **68** (2004), no. 1, 121–157 Zbl 1095.30037 MR 2152911

[12] T. Budzinski, N. Curien, and B. Petri, Universality for random surfaces in unconstrained genus. *Electron. J. Combin.* **26** (2019), no. 4, Paper No. 4.2 Zbl 1422.05093 MR 4025406

[13] T. Budzinski, N. Curien, and B. Petri, On the minimal diameter of closed hyperbolic surfaces. *Duke Math. J.* **170** (2021), no. 2, 365–377 Zbl 1459.05295 MR 4202496

[14] P. Buser, A note on the isoperimetric constant. *Ann. Sci. École Norm. Sup. (4)* **15** (1982), no. 2, 213–230 Zbl 0501.53030 MR 683635

[15] G. Chapuy and G. Perarnau, On the number of coloured triangulations of d-manifolds. *Discrete Comput. Geom.* **65** (2021), no. 3, 601–617 Zbl 1461.52015 MR 4226484

[16] S. Chmutov and B. Pittel, On a surface formed by randomly gluing together polygonal discs. *Adv. in Appl. Math.* **73** (2016), 23–42 Zbl 1328.05003 MR 3433499

[17] F. Costantino, R. Frigerio, B. Martelli, and C. Petronio, Triangulations of 3-manifolds, hyperbolic relative handlebodies, and Dehn filling. *Comment. Math. Helv.* **82** (2007), no. 4, 903–933 Zbl 1180.57022 MR 2341844

[18] N. Curien, Peeling random planar maps. Saint-Flour lecture notes, 2019, `https://www.imo.universite-paris-saclay.fr/~curien/enseignement.html`

[19] K. Delp, D. Hoffoss, and J. F. Manning, Problems in groups, geometry, and three-manifolds. 2015, arXiv:1512.04620

[20] N. M. Dunfield and W. P. Thurston, Finite covers of random 3-manifolds. *Invent. Math.* **166** (2006), no. 3, 457–521 Zbl 1111.57013 MR 2257389

[21] P. Feller, A. Sisto, and G. Viaggi, Uniform models and short curves for random 3-manifolds. 2020, arXiv:1910.09486

[22] J. Friedman, A proof of Alon's second eigenvalue conjecture and related problems. *Mem. Amer. Math. Soc.* **195** (2008), no. 910 Zbl 1177.05070 MR 2437174

[23] D. Futer, J. S. Purcell, and S. Schleimer, Effective bilipschitz bounds on drilling and filling. *Geom. Topol.* **26** (2022), no. 3, 1077–1188 Zbl 7584041 MR 4466646

[24] A. Gamburd, Poisson–Dirichlet distribution for random Belyi surfaces. *Ann. Probab.* **34** (2006), no. 5, 1827–1848 Zbl 1113.60095 MR 2271484

[25] T. Gelander, A lecture on invariant random subgroups. In *New directions in locally compact groups*, pp. 186–204, London Math. Soc. Lecture Note Ser. 447, Cambridge Univ. Press, Cambridge, 2018 Zbl 1400.43003 MR 3793288

[26] L. Guth, H. Parlier, and R. Young, Pants decompositions of random surfaces. *Geom. Funct. Anal.* **21** (2011), no. 5, 1069–1090 Zbl 1242.32007 MR 2846383

[27] U. Hamenstädt, Small eigenvalues and thick-thin decomposition in negative curvature. *Ann. Inst. Fourier (Grenoble)* **69** (2019), no. 7, 3065–3093 Zbl 1454.58027 MR 4286830

[28] U. Hamenstädt and G. Viaggi, Small eigenvalues of random 3-manifolds. *Trans. Amer. Math. Soc.* **375** (2022), no. 6, 3795–3840 Zbl 1492.58007 MR 4419048

[29] H. Kammeyer, *Introduction to ℓ^2-invariants*. Lecture Notes in Math. 2247, Springer, Cham, 2019 Zbl 1458.55001 MR 3971279

[30] E. Kowalski, Crible en expansion. Séminaire Bourbaki 2010/2011. Exposés 1027–1042. *Astérisque* **348** (2012), Exp. No. 1028, 17–64 Zbl 1323.11070 MR 3050711

[31] M. Lackenby, Heegaard splittings, the virtually Haken conjecture and property (τ). *Invent. Math.* **164** (2006), no. 2, 317–359 Zbl 1110.57015 MR 2218779

[32] T. T. Q. Lê, Growth of homology torsion in finite coverings and hyperbolic volume. *Ann. Inst. Fourier (Grenoble)* **68** (2018), no. 2, 611–645 Zbl 1409.57016 MR 3803114

[33] A. Lubotzky, J. Maher, and C. Wu, Random methods in 3-manifold theory. *Tr. Mat. Inst. Steklova* **292** (2016), Algebra, Geometriya i Teoriya Chisel, 124–148 Zbl 1354.57026 MR 3628457

[34] M. Magee, F. Naud, and D. Puder, A random cover of a compact hyperbolic surface has relative spectral gap $\frac{3}{16} - \varepsilon$. *Geom. Funct. Anal.* **32** (2022), no. 3, 595–661 Zbl 7536482 MR 4431124

[35] J. Maher, Random Heegaard splittings. *J. Topol.* **3** (2010), no. 4, 997–1025 Zbl 1207.37027 MR 2746344

[36] J. Maher, Random walks on the mapping class group. *Duke Math. J.* **156** (2011), no. 3, 429–468 Zbl 1213.37072 MR 2772067

[37] T. Mantuano, Discretization of compact Riemannian manifolds applied to the spectrum of Laplacian. *Ann. Global Anal. Geom.* **27** (2005), no. 1, 33–46 Zbl 1077.58018 MR 2130531

[38] M. Mirzakhani, Growth of Weil–Petersson volumes and random hyperbolic surfaces of large genus. *J. Differential Geom.* **94** (2013), no. 2, 267–300 Zbl 1270.30014 MR 3080483

[39] J. Raimbault, Géométrie et topologie des variétés hyperboliques de grand volume. In *Actes de Séminaire de Théorie Spectrale et Géométrie. Année 2012–2014*, pp. 163–195, Séminaire de Théorie Spectrale et Géométrie 31, Université de Grenoble I, Institut Fourier, St. Martin d'Hères, 2014 Zbl 1375.22008

[40] R. K. W. Roeder, J. H. Hubbard, and W. D. Dunbar, Andreev's theorem on hyperbolic polyhedra. *Ann. Inst. Fourier (Grenoble)* **57** (2007), no. 3, 825–882 Zbl 1127.51012 MR 2336832

[41] M. Scharlemann, Heegaard splittings of compact 3-manifolds. In *Handbook of geometric topology*, pp. 921–953, North-Holland, Amsterdam, 2002 Zbl 0985.57005 MR 1886684

[42] A. Sisto and S. J. Taylor, Largest projections for random walks and shortest curves in random mapping tori. *Math. Res. Lett.* **26** (2019), no. 1, 293–321 Zbl 1453.20062 MR 3963985

[43] G. Viaggi, Volumes of random 3-manifolds. *J. Topol.* **14** (2021), no. 2, 504–537 Zbl 1470.57038 MR 4286047

[44] N. C. Wormald, Models of random regular graphs. In *Surveys in combinatorics (Canterbury, 1999)*, pp. 239–298, London Math. Soc. Lecture Note Ser. 267, Cambridge Univ. Press, Cambridge, 1999 Zbl 0935.05080 MR 1725006

[45] A. Wright, A tour through Mirzakhani's work on moduli spaces of Riemann surfaces. *Bull. Amer. Math. Soc. (N.S.)* **57** (2020), no. 3, 359–408 Zbl 1452.32003 MR 4108090

Received 28 September 2021

B. Petri, Institut de Mathématiques de Jussieu–Paris Rive Gauche, UMR7586,
Sorbonne Université - Campus Pierre et Marie Curie, 4, place Jussieu,
75252 Paris Cedex 05, France

E-mail: bram.petri@imj-prg.fr

J. Raimbault, Institut de Mathématiques de Toulouse, UMR5219, CNRS, UPS IMT,
Université de Toulouse, 31062 Toulouse Cedex 9; and
Institut de Mathématiques de Marseille, UMR7373, CNRS, Aix-Marseille Université,
Marseille, France

E-mail: jean.raimbault@univ-amu.fr