Geom. Funct. Anal. Vol. 32 (2022) 595–661
https://doi.org/10.1007/s00039-022-00602-x
Published online May 17, 2022
© 2022 The Author(s)
GAFA Geometric And Functional Analysis

A RANDOM COVER OF A COMPACT HYPERBOLIC SURFACE HAS RELATIVE SPECTRAL GAP $\frac{3}{16} - \varepsilon$

Michael Magee, Frédéric Naud and Doron Puder

Abstract. Let X be a compact connected hyperbolic surface, that is, a closed connected orientable smooth surface with a Riemannian metric of constant curvature -1. For each $n \in \mathbf{N}$, let X_n be a random degree-n cover of X sampled uniformly from all degree-n Riemannian covering spaces of X. An eigenvalue of X or X_n is an eigenvalue of the associated Laplacian operator Δ_X or Δ_{X_n}. We say that an eigenvalue of X_n is *new* if it occurs with greater multiplicity than in X. We prove that for any $\varepsilon > 0$, with probability tending to 1 as $n \to \infty$, there are no new eigenvalues of X_n below $\frac{3}{16} - \varepsilon$. We conjecture that the same result holds with $\frac{3}{16}$ replaced by $\frac{1}{4}$.

Contents

1 Introduction . . . 596
1.1 Related works. . . . 601
1.2 Subsequent results. . . . 603
1.3 Structure of the proofs and the issues that arise. . . . 603
1.4 Notation. . . . 607
2 The Proof of Theorem 1.5 Given Theorem 1.11 . . . 607
2.1 Selberg's trace formula and counting closed geodesics. . . . 607
2.2 Choice of function for use in Selberg's trace formula. . . . 609
2.3 Proof of Theorem 1.5. . . . 610
2.4 Proof of Theorem 1.8. . . . 613
3 Tiled Surfaces . . . 614
3.1 Tiled surfaces. . . . 614
3.2 Blocks and chains. . . . 618
3.3 Boundary reduced and strongly boundary reduced tiled surfaces. . . . 619
3.4 Pieces and ε-adapted tiled surfaces. . . . 620
3.5 The octagons-vs-boundary algorithm. . . . 622
3.6 Resolutions from the octagons-vs-boundary algorithm. . . . 625
4 Representation Theory of Symmetric Groups . . . 628
4.1 Background. . . . 628
4.2 Effective bounds for dimensions. . . . 630
4.3 Effective bounds for the zeta function of the symmetric group. . . . 630
5 Estimates for the Probabilities of Tiled Surfaces . . . 633
5.1 Prior results. . . . 634

5.2 Partitioning Ξ_n and preliminary estimates. 636
5.3 The zero regime of b_ν. 640
5.4 The intermediate regime of b_ν. 641
5.5 The large regime of $b_\nu, \check{b}_\nu$. 642
5.6 Assembly of analytic estimates for Ξ_n. 643
5.7 A new expression for $\Xi_n^{\nu=(n-\mathfrak{v})}$. 643
5.8 Understanding $|\mathbb{X}_n^*(Y,\mathcal{J})|$. 647
5.9 Bounds on $\mathbb{E}_n^{\mathrm{emb}}(Y)$ for ε-adapted Y. 650
6 Proof of Theorem 1.11 652
6.1 Setup. 653
6.2 Part I: the contribution from non-ε-adapted surfaces. 653
6.3 Part II: the contribution from ε-adapted surfaces. 655
References 658

1 Introduction

Spectral gap is a fundamental concept in mathematics and related sciences as it governs the rate at which a process converges towards its stationary state. The question that motivates this paper is whether random objects have large, or even optimal, spectral gaps. This will be made precise below.

One of the simplest examples of spectral gap is the spectral gap of a *graph.* The spectrum of a graph $\mathcal{G}$ on n vertices is the collection of eigenvalues of its adjacency matrix $A_{\mathcal{G}}$. Assuming that $\mathcal{G}$ is d-regular, the largest eigenvalue occurs at d and is simple if and only if $\mathcal{G}$ is connected. This means, writing

$$\lambda_0 = d \geq \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_{n-1}$$

for the eigenvalues of $A_{\mathcal{G}}$, then there is a *spectral gap* between λ_0 and λ_1 (i.e. $\lambda_0 > \lambda_1$) if and only if $\mathcal{G}$ is connected. In fact, the Cheeger inequalities for graphs due to Alon and Milman [AM85] show that the size of the spectral gap (i.e. $\lambda_0 - \lambda_1$) quantifies how difficult it is, roughly speaking, to separate the vertices of $\mathcal{G}$ into two sets, each not too small, with few edges between them. This is in tension with the fact that a d-regular graph is *sparse.* Sparse yet highly-connected graphs are called *expander graphs* and are relevant to many real-world examples.[1]

However, a result of Alon and Boppana [Nil91] puts a sharp bound on what one can achieve: for a sequence of d-regular graphs $\mathcal{G}_n$ on n vertices, as $n \to \infty$, $\lambda_1(\mathcal{G}_n) \geq 2\sqrt{d-1} - o(1)$. The trivial eigenvalues of a graph occur at d, and if $\mathcal{G}$ has a bipartite component, at $-d$. A connected d-regular graph with all its non-trivial eigenvalues in the interval $[-2\sqrt{d-1}, 2\sqrt{d-1}]$ is called a *Ramanujan graph* after Lubotzky, Phillips, and Sarnak [LPS88].

In the rest of the paper, if an event depending on a parameter n holds with probability tending to 1 as $n \to \infty$, then we say it holds asymptotically almost surely (a.a.s.). A famous conjecture of Alon [Alo86], now a theorem due to Friedman

[1] Following Barzdin and Kolmogorov [BK93], consider the network of neurons in a human brain.

[Fri08], states that for any $\varepsilon > 0$, a.a.s. a random d-regular graph on n vertices, chosen uniformly amongst such graphs, has all its non-trivial eigenvalues bounded in absolute value by $2\sqrt{d-1}+\varepsilon$. In other words, almost all d-regular graphs have almost optimal spectral gaps. In [Bor], Bordenave has given a shorter proof of Friedman's theorem. A first result about uniform spectral gap for random regular graphs is due to Broder and Shamir [BS87b] who proved a.a.s. $\lambda_1 \leq 3d^{3/4}$. The approach in the current paper is similar to the direct trace method introduced by Broder–Shamir, which was subsequently improved by Puder and Friedman-Puder [Pud15, FP22] to show a.a.s. that $\lambda_1 \leq 2\sqrt{d-1} + \frac{2}{\sqrt{d-1}}$.

Friedman conjectured in [Fri03] that the following extension of Alon's conjecture holds. Given any finite graph $\mathcal{G}$ there is a notion of a degree-n cover[2] $\mathcal{G}_n$ of the graph. Elements of the spectrum[3] of $\mathcal{G}_n$ that are not elements of the spectrum of $\mathcal{G}$ are called *new eigenvalues* of $\mathcal{G}_n$. Friedman conjectured that for a fixed finite graph $\mathcal{G}$, for any $\varepsilon > 0$ a random degree-n cover of $\mathcal{G}$ a.a.s. has no new eigenvalues of absolute value larger than $\rho(\mathcal{G}) + \varepsilon$, where $\rho(\mathcal{G})$ is the spectral radius of the adjacency operator of the universal cover of $\mathcal{G}$, acting on ℓ^2 functions. For d even, the special case where $\mathcal{G}$ is a bouquet of $\frac{d}{2}$ loops recovers Alon's conjecture. Friedman's conjecture has recently been proved in a breakthrough by Bordenave and Collins [BC19].

The focus of this paper is the extension of Alon's and Friedman's conjectures to compact hyperbolic surfaces.

A *hyperbolic surface* is a Riemannian surface of constant curvature -1 without boundary. In this paper, all surfaces will be orientable. By uniformization [Bea84, Section 9.2], a connected compact hyperbolic surface can be realized as $\Gamma\backslash\mathbb{H}$ where Γ is a discrete subgroup of $\mathrm{PSL}_2(\mathbf{R})$ and

$$\mathbb{H} = \{\, x + iy \,:\, x, y \in \mathbf{R},\, y > 0 \,\}$$

is the hyperbolic upper half plane, upon which $\mathrm{PSL}_2(\mathbf{R})$ acts via Möbius transformations preserving the hyperbolic metric

$$\frac{dx^2 + dy^2}{y^2}.$$

Let $X = \Gamma\backslash\mathbb{H}$ be a connected compact hyperbolic surface. Topologically, X is a connected closed surface of some genus $g \geq 2$.

Since the Laplacian $\Delta_{\mathbb{H}}$ on $\mathbb{H}$ is invariant under $\mathrm{PSL}_2(\mathbf{R})$, it descends to a differential operator on $C^\infty(X)$ and extends to a non-negative, unbounded, self-adjoint operator Δ_X on $L^2(X)$. The spectrum of Δ_X consists of real eigenvalues

$$0 = \lambda_0(X) \leq \lambda_1(X) \leq \cdots \leq \lambda_n(X) \leq \cdots$$

[2] The precise definition of a cover of a graph is not important here; only that it is analogous to a covering space of a surface.

[3] We also take multiplicities into account in this statement.

with $\lambda_i \to \infty$ as $i \to \infty$. The same discussion also applies if we drop the condition that X is connected.[4] We have $\lambda_0(X) < \lambda_1(X)$ if and only if X is connected, as for graphs. With Friedman's conjecture in mind, we also note that the spectrum of $\Delta_{\mathbb{H}}$ is absolutely continuous and supported on the interval $[\frac{1}{4}, \infty)$ (e.g. [Bor16, Thm. 4.3]). There is also an analog of the Alon–Boppana Theorem in this setting: a result of Huber [Hub74] states that for any sequence of compact hyperbolic surfaces X_i with genera $g(X_i)$ tending to infinity,

$$\limsup_{i\to\infty} \lambda_1(X_i) \leq \frac{1}{4}.$$

To state an analog of the Alon/Friedman conjecture for surfaces, we need a notion of a *random cover.* Suppose X is a compact connected hyperbolic surface, and suppose $\tilde{X}$ is a degree-n Riemannian cover of X. Fix a point $x_0 \in X$ and label the fiber above it by $[n] \stackrel{\text{def}}{=} \{1, \ldots, n\}$. There is a monodromy map

$$\pi_1(X, x_0) \to S_n$$

that describes how the fiber of x_0 is permuted when following lifts of a closed loop from X to $\tilde{X}$. Here S_n is the symmetric group of permutations of the set $[n]$. The cover $\tilde{X}$ is uniquely determined by the monodromy homomorphism. Let g denote the genus of X. We fix an isomorphism

$$\pi_1(X, x_0) \cong \Gamma_g \stackrel{\text{def}}{=} \langle a_1, b_1, a_2, b_2, \ldots, a_g, b_g \mid [a_1, b_1] \cdots [a_g, b_g] = 1 \rangle . \tag{1.1}$$

Now, given any

$$\phi \in \mathbb{X}_{g,n} \stackrel{\text{def}}{=} \mathrm{Hom}(\Gamma_g, S_n)$$

we can construct a cover of X whose monodromy map is ϕ as follows. Using the fixed isomorphism of (1.1), we have a free properly discontinuous action of Γ_g on $\mathbb{H}$ by isometries. Define a new action of Γ_g on $\mathbb{H} \times [n]$ by

$$\gamma(z, i) = (\gamma z, \phi[\gamma](i)).$$

The quotient of $\mathbb{H} \times [n]$ by this action is named X_ϕ and is a hyperbolic cover of X with monodromy ϕ. This construction establishes a one-to-one correspondence between $\phi \in \mathbb{X}_{g,n}$ and degree-n covers with a labeled fiber X_ϕ of X. See also Example 3.4.

As for graphs, any eigenvalue of Δ_X will also be an eigenvalue of Δ_{X_ϕ}: every eigenfunction of Δ_X can be pulled back to an eigenfunction of Δ_{X_ϕ} with the same eigenvalue. We say that an eigenvalue of Δ_{X_ϕ} is *new* if it is not one of Δ_X, or more generally, appears with greater multiplicity in X_ϕ. To pick a random cover of X, we simply use the uniform probability measure on the finite set $\mathbb{X}_{g,n}$. Recall we say an event that pertains to any n holds a.a.s. if it holds with probability tending to one as $n \to \infty$. The analog of Friedman's conjecture for surfaces is the following.

[4] In which case X is a finite union of connected compact hyperbolic surfaces, each of which can be realized as a quotient of $\mathbb{H}$.

Conjecture 1.1. *Let X be a compact connected hyperbolic surface. Then for any $\varepsilon > 0$, a.a.s.*

$$\operatorname{spec}\left(\Delta_{X_\phi}\right) \cap \left[0, \frac{1}{4} - \varepsilon\right] = \operatorname{spec}\left(\Delta_X\right) \cap \left[0, \frac{1}{4} - \varepsilon\right]$$

and the multiplicities on both sides are the same.

Remark 1.2. The analog of Conjecture 1.1 for finite area non-compact surfaces appeared previously in the work of Golubev and Kamber [GK19, Conj. 1.6(1)].

Remark 1.3. We have explained the number $\frac{1}{4}$ in terms of the spectrum of the Laplacian on the hyperbolic plane and as an asymptotically optimal spectral gap in light of Huber's result [Hub74]. The number $\frac{1}{4}$ also features prominently in Selberg's eigenvalue conjecture [Sel65], that states for $X = \mathrm{SL}_2(\mathbf{Z})\backslash\mathbb{H}$, the (deterministic) family of congruence covers of X never have new eigenvalues below $\frac{1}{4}$. Although Selberg's conjecture is for a finite-area, non-compact hyperbolic orbifold, the Jacquet-Langlands correspondence [JL70] means that it also applies to certain arithmetic compact hyperbolic surfaces.

Remark 1.4. In [Wri20, Problem 10.4], Wright asks, for random compact hyperbolic surfaces sampled according to the Weil–Petersson volume form on the moduli space of genus g closed hyperbolic surfaces, whether $\liminf_{g\to\infty}(\mathbb{P}(\lambda_1 > \frac{1}{4})) > 0$. See Section 1.1 for what is known in this setting. It is not even known [Wri20, Problem 10.3] whether there is a sequence of Riemann surfaces X_n with genus tending to ∞ such that $\lambda_1(X_n) \to \frac{1}{4}$. Conjecture 1.1 offers a new route to resolving this problem via the probabilistic method, since it is known by work of Jenni [Jen84] that there exists a genus 2 hyperbolic surface X with $\lambda_1(X) > \frac{1}{4}$ and this X can be taken as the base surface in Conjecture 1.1. (See Section 1.2 for important developments in this area after the current paper was written.)

The main theorem of the paper, described in the title, is the following.

Theorem 1.5. *Let X be a compact connected hyperbolic surface. Then for any $\varepsilon > 0$, a.a.s.*

$$\operatorname{spec}\left(\Delta_{X_\phi}\right) \cap \left[0, \frac{3}{16} - \varepsilon\right] = \operatorname{spec}\left(\Delta_X\right) \cap \left[0, \frac{3}{16} - \varepsilon\right]$$

and the multiplicities on both sides are the same.

Remark 1.6. The appearance of the number $\frac{3}{16}$ in Theorem 1.5 is essentially for the same reason that $\frac{3}{4}$ appears in [MN20] (note that $\frac{3}{16} = \frac{3}{4}(1 - \frac{3}{4})$, and eigenvalues of the Laplacian are naturally parameterized as $s(1-s)$). Ultimately, the appearance of $\frac{3}{4}$ can be traced back to the method of Broder and Shamir [BS87b] who proved that a.a.s. a random $2d$-regular graph on n vertices has $\lambda_1 \leq O\left(d^{3/4}\right)$, using an estimate analogous to Theorem 1.11 below.

Remark 1.7. More mysteriously, $\frac{3}{16}$ is also the lower bound that Selberg obtained for the smallest new eigenvalue of a congruence cover of the modular curve $\mathrm{SL}_2(\mathbf{Z})\backslash\mathbb{H}$, in the same paper [Sel65] as his eigenvalue conjecture. In this context, the number arises ultimately from bounds on Kloosterman sums due to Weil [Wei48] that follow from Weil's resolution of the Riemann hypothesis for curves over finite fields. The state of the art on Selberg's eigenvalue conjecture, after decades of intermediate results [GJ78, Iwa89, LRS95, Iwa96, KS02], is due to Kim and Sarnak [Kim03] who produced a spectral gap of size $\frac{975}{4096}$ for congruence covers of $\mathrm{SL}_2(\mathbf{Z})\backslash\mathbb{H}$.

It was pointed out to us by A. Kamber that our methods also yield the following estimate on the density of new eigenvalues of a random cover.

Theorem 1.8. *Let*

$$0 \le \lambda_{i_1}(X_\phi) \le \lambda_{i_2}(X_\phi) \le \cdots \le \lambda_{i_{k(\phi)}}(X_\phi) \le \frac{1}{4}$$

denote the collection of new eigenvalues of Δ_{X_ϕ} of size at most $\frac{1}{4}$, included with multiplicity. For each of these, we write $\lambda_{i_j} = s_{i_j}(1-s_{i_j})$ with $s_{i_j} = s_{i_j}(X_\phi) \in \left[\frac{1}{2}, 1\right]$. For any $\varepsilon > 0$ and $\sigma \in \left(\frac{1}{2}, 1\right)$, a.a.s.

$$\#\left\{1 \le j \le k(\phi) \,:\, \lambda_{i_j} < \sigma\,(1-\sigma)\right\} = \#\left\{1 \le j \le k(\phi) \,:\, s_{i_j} > \sigma\right\} \le n^{3-4\sigma+\varepsilon}. \tag{1.2}$$

Remark 1.9. The estimate (1.2) was established by Iwaniec [Iwa02, Thm 11.7] for congruence covers of $\mathrm{SL}_2(\mathbf{Z})\backslash\mathbb{H}$. Although Iwaniec's theorem has been generalized in various directions [Hux86, Sar87, Hum18], as far as we know, Iwaniec's result has not been directly improved, so speaking about density of eigenvalues, Theorem 1.8 establishes for random covers the best result known in the arithmetic setting for eigenvalues above the Kim–Sarnak bound $\frac{975}{4096}$ [Kim03]. Density estimates such as Theorem 1.8 have applications to the cutoff phenomenon on hyperbolic surfaces by work of Golubev and Kamber [GK19].

We prove Theorems 1.5 and 1.8 using Selberg's trace formula in Section 2. We use as a 'black-box' in this method a statistical result (Theorem 1.11) about the expected number of fixed points of a fixed $\gamma \in \Gamma_g$ under a random ϕ.

If $\pi \in S_n$ then we write $\mathsf{fix}(\pi)$ for the number of fixed points of the permutation π. Given an element $\gamma \in \Gamma_g$, we let fix_γ be the function

$$\mathsf{fix}_\gamma : \mathbb{X}_{g,n} \to \mathbf{Z}, \quad \mathsf{fix}_\gamma(\phi) \stackrel{\text{def}}{=} \mathsf{fix}(\phi(\gamma)).$$

We write $\mathbb{E}_{g,n}[\mathsf{fix}_\gamma]$ for the expected value of fix_γ with respect to the uniform probability measure on $\mathbb{X}_{g,n}$. In [MP20], the first and third named authors proved the following theorem.

Theorem 1.10. *Let $g \ge 2$ and $1 \ne \gamma \in \Gamma_g$. If $q \in \mathbf{N}$ is maximal such that $\gamma = {\gamma_0}^q$ for some $\gamma_0 \in \Gamma_g$, then, as $n \to \infty$,*

$$\mathbb{E}_{g,n}[\mathsf{fix}_\gamma] = d(q) + O_\gamma\left(n^{-1}\right),$$

where $d(q)$ is the number of divisors of q.

In the current paper, we need an effective version of Theorem 1.10 that controls the dependence of the error term on γ. We need this estimate only for γ that are not a proper power. For $\gamma \in \Gamma_g$, we write $\ell_w(\gamma)$ for the *cyclic-word-length* of γ, namely, for the length of a shortest word in the generators $a_1, b_1, \ldots, a_g, b_g$ of Γ_g that represents an element in the conjugacy class of γ in Γ_g. The effective version of Theorem 1.10 that we prove here is the following.

Theorem 1.11. *For each genus $g \geq 2$, there is a constant $A = A(g)$ such that for any $c > 0$, if $1 \neq \gamma \in \Gamma_g$ is not a proper power of another element in Γ_g and $\ell_w(\gamma) \leq c \log n$ then*

$$\mathbb{E}_{g,n}[\mathsf{fix}_\gamma] = 1 + O_{c,g}\left(\frac{(\log n)^A}{n}\right).$$

The implied constant in the big-O depends only on c and g.

Remark 1.12. In the rest of the paper, just to avoid complications in notation and formulas that would obfuscate our arguments, we give the proof of Theorem 1.11 when $g = 2$. The extension to arbitrary genus is for the most part obvious: if it is not at some point, we will point out the necessary changes.

The proof of Theorem 1.11 takes up the bulk of the paper, spanning Section 4–Section 6. The proof of Theorem 1.11 involves delving into the proof of Theorem 1.10 and refining the estimates, as well as introducing some completely new ideas.

1.1 Related works.

The Brooks–Makover model. The first study of spectral gap for random surfaces in the literature is due to Brooks and Makover [BM04] who form a model of a random compact surface as follows. Firstly, for a parameter n, they glue together n copies of an ideal hyperbolic triangle where the gluing scheme is given by a random trivalent ribbon graph. Their model for this random ribbon graph is a modification of the Bollobás bin model from [Bol88]. This yields a random finite-area, non compact hyperbolic surface. Then they perform a compactification procedure to obtain a random compact hyperbolic surface $X_{\mathrm{BM}}(n)$. The genus of this surface is not deterministic, however. Brooks and Makover prove that for this random model, there is a non-explicit constant $C > 0$ such that a.a.s. (as $n \to \infty$)

$$\lambda_1(X_{\mathrm{BM}}(n)) \geq C.$$

Theorem 1.5 concerns a different random model, but improves on the Brooks–Makover result in two important ways: the bound on new eigenvalues is explicit, and this bound is independent of the compact hyperbolic surface X with which we begin.

It is also worth mentioning a recent result of Budzinski, Curien, and Petri [BCP21, Thm. 1] who prove that the ratios

$$\frac{\mathrm{diameter}(X_{\mathrm{BM}}(n))}{\log n}$$

converge to 2 in probability as $n \to \infty$; they also observe that this is not the optimal value by a factor of 2.

The Weil–Petersson model. Another reasonable model of random surfaces comes from the Weil–Petersson volume form on the moduli space $\mathcal{M}_g$ of compact hyperbolic surfaces of genus g. Let $X_{\mathrm{WP}}(g)$ denote a random surface in $\mathcal{M}_g$ sampled according to the (normalized) Weil–Petersson volume form. Mirzakhani proved in [Mir13, Section 1.2.I] that with probability tending to 1 as $g \to \infty$,

$$\lambda_1(X_{\mathrm{WP}}(g)) \geq \frac{1}{4}\left(\frac{\log 2}{2\pi + \log 2}\right)^2 \approx 0.00247.$$

We also note recent work of Monk [Mon22] who gives estimates on the density of eigenvalues below $\frac{1}{4}$ of the Laplacian on $X_{\mathrm{WP}}(g)$.

Prior work of the authors. In some sense, the closest result to Theorem 1.5 in the literature is due to the first and second named authors of the paper [MN20], but it does not apply to compact surfaces, rather to infinite area convex co-compact hyperbolic surfaces. Because these surfaces have infinite area, their spectral theory is more involved. We will focus on one result of [MN20] to illustrate the comparison with this paper.

Suppose X is a connected non-elementary, non-compact, convex co-compact hyperbolic surface. The spectral theory of X is driven by a critical parameter $\delta = \delta(X) \in (0,1)$. This parameter is both the critical exponent of a Poincaré series and the Hausdorff dimension of the limit set of X. If $\delta > \frac{1}{2}$ then results of Patterson [Pat76] and Lax-Phillips [LP81] say that the bottom of the spectrum of X is a simple eigenvalue at $\delta(1-\delta)$ and there are finitely many eigenvalues in the range $[\delta(1-\delta), \frac{1}{4})$. In [MN20], a model of a random degree-n cover of X was introduced that is completely analogous to the one used here; the only difference in the construction is that the fundamental group of X is a free group $\mathbf{F}_r$ and hence one uses random $\phi \in \mathrm{Hom}(\mathbf{F}_r, S_n)$ to construct the random surface X_ϕ. The following theorem was obtained in [MN20, Thm. 1.3.].

Theorem 1.13. *Assume that $\delta = \delta(X) > \frac{1}{2}$. Then for any $\sigma_0 \in \left(\frac{3}{4}\delta, \delta\right)$, a.a.s.*

$$\mathrm{spec}\left(\Delta_{X_\phi}\right) \cap \left[\delta\left(1-\delta\right), \sigma_0(1-\sigma_0)\right] = \mathrm{spec}\left(\Delta_X\right) \cap \left[\delta\left(1-\delta\right), \sigma_0(1-\sigma_0)\right] \quad (1.3)$$

and the multiplicities on both sides are the same.

Although Theorem 1.13 is analogous to Theorem 1.5 (for compact X, $\delta(X) = 1$), the methods used in [MN20] have almost no overlap with the methods used here. For infinite area X, the fundamental group is free, so the replacement of Theorem 1.11 was already known by results of Broder–Shamir [BS87b] and the third named author [Pud15]. The challenge in [MN20] was to develop bespoke analytic machinery to access these estimates.

Conversely, in the current paper, the needed analytic machinery already exists (Selberg's trace formula) and rather, it is the establishment of Theorem 1.11 that is the main challenge here, stemming from the non-free fundamental group Γ_g.

1.2 Subsequent results. Since the preprint version of the current paper appeared in March 2020, several important results have been obtained in the area of spectral gap of random surfaces. Independently of each other, Wu and Xue [WX21] and Lipnowski and Wright [LW21] proved that for any $\varepsilon > 0$, a Weil–Petersson random compact hyperbolic surface of genus g has spectral gap of size at least $\frac{3}{16} - \varepsilon$ with probability tending to one as $g \to \infty$. This result has been extended to the case of Weil–Petersson random surfaces with not too many cusps by Hide in [Hid21].

In [MN21], the results of [MN20] have been strengthened by the first and second named author to an (essentially optimal) analog of Friedman's theorem for bounded frequency resonances on infinite area Schottky surfaces.

Hide and the first named author have recently proved in [HM21] that the analog of Conjecture 1.1 for finite area non-compact hyperbolic surfaces holds true, and by combining this result with a cusp removal argument of Buser, Burger, and Dodziuk [BBD88], in [HM21] it is also proved that there exist compact hyperbolic surfaces with genera tending to infinity and $\lambda_1 \to \frac{1}{4}$. (We have chosen to preserve Remark 1.4 as originally written here for posterity.)

1.3 Structure of the proofs and the issues that arise.

Proof of Theorem 1.5 given Theorem 1.11. First, we explain the outline of the proof of Theorem 1.5 from Theorem 1.11. Theorem 1.8 also follows from Theorem 1.11 using the same ideas. Both proofs are presented in full in Section 2.

Our method of proving Theorem 1.5 is analogous to the method of Broder and Shamir [BS87b] for proving that a random $2d$-regular graph has a large spectral gap. For us, the Selberg trace formula replaces a more elementary formula for the trace of a power of the adjacency operator of a graph in terms of closed paths in the graph.

Let Γ denote the fundamental group of X. By taking the difference of the Selberg trace formula for X_ϕ and that for X we obtain a formula of the form

$$\sum_{\text{new eigenvalues } \lambda \text{ of } X_\phi} F(\lambda) = \sum_{[\gamma]\in C(\Gamma)} G(\gamma)\left(\mathsf{fix}_\gamma(\phi) - 1\right), \tag{1.4}$$

where $C(\Gamma)$ is the collection of conjugacy classes in Γ, and F and G are interdependent functions that depend on n. The way we choose F and G together is to ensure

- $F(\lambda)$ is non-negative for any possible λ, and large if λ is an eigenvalue we want to forbid, and
- $G(\gamma)$ localizes to γ with $\ell_w(\gamma) \le c \log n$ for some $c = c(X)$.

By taking expectations of (1.4) we obtain

$$\mathbb{E}\left[\sum_{\text{new eigenvalues } \lambda \text{ of } X_\phi} F(\lambda)\right] = \sum_{[\gamma]\in C(\Gamma)} G(\gamma)\mathbb{E}\left[\mathsf{fix}_\gamma(\phi) - 1\right]. \tag{1.5}$$

The proof will conclude by bounding the right hand side and applying Markov's inequality to conclude that there are no new eigenvalues in the desired forbidden region. Since G is well-controlled in our proof, it remains to estimate each term $\mathbb{E}\left[\mathsf{fix}_\gamma(\phi) - 1\right]$. To do this, we echo Broder–Shamir [BS87b] and partition the summation on the right-hand side of (1.5) according to three groups.

- If γ is the identity, then $G(1)$ is easily analyzed, and $\mathbb{E}\left[\mathsf{fix}_\gamma(\phi) - 1\right] = n - 1$.
- If γ is a proper power of a non-trivial element of Γ, then we use a trivial bound $\mathbb{E}\left[\mathsf{fix}_\gamma(\phi) - 1\right] \leq n - 1$, so we get no gain from the expectation. On the other hand, the contribution to
$$\sum_{[\gamma]\in C(\Gamma)} G(\gamma)$$
from these elements is negligible. Intuitively, this is because the number of elements of Γ with $\ell_w(\gamma) \leq L$ and that are proper powers is (exponentially) negligible compared to the total number of elements.
- If γ is not a proper power and not the identity, then we use Theorem 1.11 to obtain $\mathbb{E}\left[\mathsf{fix}_\gamma(\phi) - 1\right] = O_X\left(\frac{(\log n)^A}{n}\right)$. Thus for 'most' summands in the right-hand side of (1.5) we obtain a significant gain from the expectation.

Assembling all these estimates together gives a sufficiently upper strong bound on (1.5) to obtain Theorem 1.5 via Markov's inequality.

Proof of Theorem 1.11. To understand the proof of Theorem 1.11, we suggest that the reader first read the overview below, then Section 6 where all the components of the proof are brought together, and then Section 3–Section 5 where the technical ingredients are proved. As throughout the paper, we assume $g = 2$ in this overview and we will forgo precision to give a bird's-eye view of the proof.

Fixing an octagonal fundamental domain for X, any X_ϕ is tiled by octagons; this tiling comes with some extra labelings of edges corresponding to the generators of Γ. Any labeled 2-dimensional CW-complex that can occur as a subcomplex of some X_ϕ is called a tiled surface. For any tiled surface Y, we write $\mathbb{E}_n^{\mathrm{emb}}(Y)$ for the expected number, when ϕ is chosen uniformly at random in $\mathrm{Hom}(\Gamma, S_n)$, of embedded copies of Y in X_ϕ.

In the previous paper [MP20], we axiomatized certain collections $\mathcal{R}$ of tiled surfaces, depending on γ, that have the property that
$$\mathbb{E}_{2,n}[\mathsf{fix}_\gamma] = \sum_{Y\in\mathcal{R}} \mathbb{E}_n^{\mathrm{emb}}(Y). \tag{1.6}$$
These collections are called *resolutions.* Here we have oversimplified the definitions to give an overview of the main ideas.

In [MP20], we chose a resolution, depending on γ, that consisted of two special types of tiled surfaces: those that are boundary reduced or strongly boundary reduced. The motivation for these definitions is that they make our methods for

estimating $\mathbb{E}_n^{\mathrm{emb}}(Y)$ more accurate. To give an example, if Y is strongly boundary reduced then we prove that for Y fixed and $n \to \infty$, we obtain[5]

$$\mathbb{E}_n^{\mathrm{emb}}(Y) = n^{\chi(Y)} \left(1 + O_Y\left(n^{-1}\right)\right). \tag{1.7}$$

However, the implied constant depends on Y, and in the current paper we have to control uniformly all γ with $\ell_w(\gamma) \leq c \log n$. The methods of [MP20] are not good enough for this goal. To deal with this, we introduce in Definition 3.12 a new type of tiled surface called 'ε-adapted' (for some $\varepsilon \geq 0$) that directly generalizes, and quantifies, the concept of being strongly boundary reduced. We will explain the benefits of this definition momentarily. We also introduce a new algorithm called the *octagons-vs-boundary algorithm* that given γ, produces a finite resolution $\mathcal{R}$ as in (1.6) such that every $Y \in \mathcal{R}$ is either

- ε-adapted for some $\varepsilon > 0$, or
- boundary reduced, with the additional condition that $\mathfrak{d}(Y) < \mathfrak{f}(Y) < -\chi(Y)$, where $\mathfrak{d}(Y)$ is the length of the boundary of Y and $\mathfrak{f}(Y)$ is the number of octagons in Y.

Any $Y \in \mathcal{R}$ has $\mathfrak{d}(Y) \leq c'(\log n)$ and $\mathfrak{f}(Y) \leq c'(\log n)^2$ given that $\ell_w(\gamma) \leq c \log n$ (Corollary 3.25). The fact that we maintain control on these quantities during the algorithm is essential. However, a defect of this algorithm is that we lose control of how many ε-adapted $Y \in \mathcal{R}$ there are of a given Euler characteristic. In contrast, in the algorithm of [MP20] we control, at least, the number of elements in the resolution of Euler characteristic zero. We later have to work to get around this.

We run the octagons-vs-boundary algorithm for a fixed $\varepsilon = \frac{1}{32}$ to obtain a resolution $\mathcal{R}$. Let us explain the benefits of this resolution we have constructed. The ε-adapted $Y \in \mathcal{R}$ contribute the main contributions to (1.6), and the merely boundary reduced Y contribute something negligible.

Indeed, we prove for any boundary reduced $Y \in \mathcal{R}$ in the regime of parameters we care about, that

$$\mathbb{E}_n^{\mathrm{emb}}(Y) \ll (A_0 \mathfrak{f}(Y))^{A_0 \mathfrak{f}(Y)} n^{\chi(Y)}, \tag{1.8}$$

where $A_0 > 0$. This bound (1.8) appears in (6.5) as the result of combining Corollary 4.5, Theorem 5.1, Proposition 5.11 and Lemma 3.6; the proof is by carefully effectivizing the arguments of [MP20].

While the bound (1.8) is quite bad (for example, using it on all terms in (1.6) would not even recover the results of [MP20]), the control of the dependence on $\mathfrak{d}(Y)$ is enough so that when combined with $\mathfrak{d}(Y) < \mathfrak{f}(Y) < -\chi(Y)$ we obtain

$$\mathbb{E}_n^{\mathrm{emb}}(Y) \ll (A_0 \mathfrak{f}(Y))^{A_0 \mathfrak{f}(Y)} n^{-\mathfrak{f}(Y)} \ll \left(\frac{\left(c'(\log n)^2\right)^{A_0}}{n} \right)^{\mathfrak{f}(Y)}.$$

[5] Some of the notation we use here is detailed in Section 1.4.

This is good enough that it can simply be combined with counting *all possible* Y with $\mathfrak{d}(Y) \leq c'(\log n)$ and $\mathfrak{f}(Y) \leq c'(\log n)^2$ to obtain that the non-ε-adapted surfaces in $\mathcal{R}$ contribute $\ll \frac{(\log n)^A}{n}$ to (1.6) for $A > 0$. This is Proposition 6.1.

So from now on assume $Y \in \mathcal{R}$ is ε-adapted and we explain how to control the contributions to (1.6) from these remaining Y. We first prove that there is a rational function Q_Y such that

$$\mathbb{E}_n^{\mathrm{emb}}(Y) = n^{\chi(Y)} \left(Q_Y(n) + O\left(\frac{1}{n}\right)\right) \left(1 + O\left(\frac{(\log n)^2}{n}\right)\right), \tag{1.9}$$

where the implied constants hold for any ε-adapted $Y \in \mathcal{R}$ as long as $\ell_w(\gamma) \leq c \log n$ (Theorem 5.1, Proposition 5.12 and Corollary 5.21). In fact, this expression remains approximately valid for the same Y if n is replaced throughout by m with $m \approx (\log n)^B$ for some $B > 0$; this will become relevant momentarily.

The rational function Q_Y is new to this paper; it appears through Corollary 5.15 and Lemma 5.20 and results from refining the representation-theoretic arguments in [MP20]. The description of Q_Y is in terms of Stallings core graphs [Sta83], and related to the theory of expected number of fixed points of words in the free group. In the notation of the rest of the paper,

$$Q_Y(n) = \frac{(n)_{\mathfrak{v}(Y)}}{n^{\chi(Y)}} \sum_{H \in \mathcal{Q}(Y)} \frac{(n)_{\mathfrak{v}(H)}}{\prod_{f \in \{a,b,c,d\}} (n)_{\mathfrak{e}_f(H)}}, \tag{1.10}$$

where $\mathcal{Q}(Y)$ is a collection of core graphs obtained by adding handles to the one-skeleton of Y, performing 'folding' operations, and taking quotients in a particular way (see Section 5.8 for details).

The argument leading to (1.9) involves isolating some of the terms that contribute to $\mathbb{E}_n^{\mathrm{emb}}(Y)$, and reinterpreting these as related to the size of a set $\mathbb{X}_n^*(Y, \mathcal{J})$ of maps $\mathbf{F}_4 \to S_n$ that contain, in an appropriate sense, an embedded copy of Y but only satisfy the relation of Γ modulo $S_{n-\mathfrak{v}(Y)}$ rather than absolutely (Proposition 5.13). Then by topological arguments the set $\mathbb{X}_n^*(Y, \mathcal{J})$ is counted in terms of core graphs leading to Lemma 5.20 that gives (1.10) here.

One unusual thing is that our combinatorial description of Q_Y does not immediately tell us the order of growth of $Q_Y(n)$, because we do not know much about $\mathcal{Q}(Y)$. On the other hand, we know enough about Q_Y (for example, for what range of parameters it is positive) so that we can 'black-box' results from [MP20] to learn that if Y is fixed and $n \to \infty$, $Q_Y(n) \to 1$. (We also learn from this argument the interesting topological fact that there is exactly one element of $\mathcal{Q}(Y)$ of maximal Euler characteristic.)

This algebraic properties of Q_Y, together with a priori facts about Q_Y, allow us to use (1.9) to establish the two following important inequalities:

$$\mathbb{E}_n^{\mathrm{emb}}(Y) = n^{\chi(Y)}\left(1 + O_c\left(\frac{(\log n)^4}{n}\right) + O\left(\frac{m}{n}\frac{\mathbb{E}_m^{\mathrm{emb}}(Y)}{m^{\chi(Y)}}\right)\right) \tag{1.11}$$

$$n^{\chi(Y)} \ll \frac{m}{n}\mathbb{E}_m^{\mathrm{emb}}(Y), \quad \text{if } \chi(Y) < 0 \tag{1.12}$$

where $m \approx (\log n)^B$ is <u>much smaller</u> than n. These inequalities are provided by Proposition 5.27 and Corollary 5.25 (see also Remark 5.26). While (1.12) may look surprising, its purpose is for running our argument in reverse with decreased parameters as explained below.

Let us now explain precisely the purpose of (1.12) and (1.11). By black-boxing the results of [MP20] one more time, we learn that there is exactly one ε-adapted $Y \in \mathcal{R}$ with $\chi(Y) = 0$, and none with $\chi(Y) > 0$. This single Y with $\chi(Y) = 0$ contributes the main term of Theorems 1.10 and 1.11 through (1.11). Any other term coming from ε-adapted Y can be controlled in terms of $\mathbb{E}_m^{\mathrm{emb}}(Y)$ using (1.11) and (1.12). These errors could accumulate, but we can control them <u>all at once</u> by using (1.6) in reverse with n replaced by m to obtain

$$\sum_{Y\in\mathcal{R}} \mathbb{E}_m^{\mathrm{emb}}(Y) = \mathbb{E}_{2,m}[\mathsf{fix}_\gamma] \le m \approx (\log n)^B.$$

Putting the previous arguments together proves Theorem 1.11.

1.4 Notation. The commutator of two group elements is $[a,b] \overset{\mathrm{def}}{=} aba^{-1}b^{-1}$. For $m, n \in \mathbf{N}$, $m \le n$, we use the notation $[m,n]$ for the set $\{m, m+1, \dots, n\}$ and $[n]$ for the set $\{1, \dots, n\}$. For $q, n \in \mathbf{N}$ with $q \le n$ we use the Pochammer symbol

$$(n)_q \overset{\mathrm{def}}{=} n(n-1)\cdots(n-q+1).$$

For real-valued functions f, g that depend on a parameter n we write $f = O(g)$ to mean there exist constants $C, N > 0$ such that for $n > N$, $|f(n)| \le Cg(n)$. We write $f \ll g$ if there are $C, N > 0$ such that $f(n) \le Cg(n)$ for $n > N$. We add constants as a subscript to the big O or the $\ll$ sign to mean that the constants C and N depend on these other constants, for example, $f = O_\varepsilon(g)$ means that both $C = C(\varepsilon)$ and $N = N(\varepsilon)$ may depend on ε. If there are no subscripts, it means the implied constants depend only on the genus g, which is fixed throughout most of the paper. We use the notation $f \asymp g$ to mean $f \ll g$ and $g \ll f$; the use of subscripts is the same as before.

2 The Proof of Theorem 1.5 Given Theorem 1.11

2.1 Selberg's trace formula and counting closed geodesics. Here we describe the main tool of this Section 2: Selberg's trace formula for compact hyperbolic surfaces. Let $C_c^\infty(\mathbf{R})$ denote the infinitely differentiable real functions on $\mathbf{R}$

with compact support. Given an *even* function $\varphi \in C_c^\infty(\mathbf{R})$, its Fourier transform is defined by

$$\widehat{\varphi}(\xi) \stackrel{\text{def}}{=} \int_{-\infty}^{\infty} \varphi(x) e^{-ix\xi} dx$$

for any $\xi \in \mathbf{C}$. As $\varphi \in C_c^\infty(\mathbf{R})$, the integral above converges for all $\xi \in \mathbf{C}$ to an entire function.

Given a compact hyperbolic surface X, we write $\mathcal{L}(X)$ for the set of closed oriented geodesics in X. A geodesic is called *primitive* if it is not the result of repeating another geodesic q times for $q \geq 2$. Let $\mathcal{P}(X)$ denote the set of closed oriented primitive geodesics on X. Every closed geodesic γ has a length $\ell(\gamma)$ according to the hyperbolic metric on X. Every closed oriented geodesic $\gamma \in \mathcal{L}(X)$ determines a conjugacy class $[\tilde{\gamma}]$ in $\pi_1(X, x_0)$ for any basepoint x_0. Clearly, a closed oriented geodesic in X is primitive if and only if the elements of the corresponding conjugacy class are not proper powers in $\pi_1(X, x_0)$. For $\gamma \in \mathcal{L}(X)$ we write $\Lambda(\gamma) = \ell(\gamma_0)$ where γ_0 is the unique primitive closed oriented geodesic such that $\gamma = \gamma_0^q$ for some $q \geq 1$.

We now give Selberg's trace formula for a compact hyperbolic surface in the form of [Bus10, Thm. 9.5.3] (see Selberg [Sel56] for the original appearance of this formula and Hejhal [Hej76, Hej83] for an encyclopedic treatment).

Theorem 2.1. (Selberg's trace formula). *Let X be a compact hyperbolic surface and let*

$$0 = \lambda_0(X) \leq \lambda_1(X) \leq \cdots \leq \lambda_n(X) \leq \cdots$$

denote the spectrum of the Laplacian on X. For $i \in \mathbf{N} \cup \{0\}$ let

$$r_i(X) \stackrel{\text{def}}{=} \begin{cases} \sqrt{\lambda_i(X) - \frac{1}{4}} & if\ \lambda_i(X) > 1/4 \\ i\sqrt{\frac{1}{4} - \lambda_i(X)} & if\ \lambda_i(X) \leq 1/4 \end{cases}.$$

Then for any even $\varphi \in C_c^\infty(\mathbf{R})$

$$\sum_{i=0}^{\infty} \widehat{\varphi}(r_i(X)) = \frac{\text{area}(X)}{4\pi} \int_{-\infty}^{\infty} r\widehat{\varphi}(r) \tanh(\pi r) dr + \sum_{\gamma \in \mathcal{L}(X)} \frac{\Lambda(\gamma)}{2\sinh\left(\frac{\ell(\gamma)}{2}\right)} \varphi(\ell(\gamma)).$$

(Both sides of the formula are absolutely convergent.)

We will also need a bound on the number of closed oriented geodesics with length $\ell(\gamma) \leq T$. In fact we only need the following very soft bound from e.g. [Bus10, Lem. 9.2.7].

Lemma 2.2. *For a compact hyperbolic surface X, there is a constant $C = C(X)$ such that*

$$|\{\gamma \in \mathcal{L}(X) : \ell(\gamma) \leq T\}| \leq Ce^T.$$

Much sharper versions of this estimate are known, but Lemma 2.2 suffices for our purposes.

Suppose that X is a connected compact hyperbolic surface. We fix a basepoint $x_0 \in X$ and an isomorphism $\pi_1(X, x_0) \cong \Gamma_g$ as in (1.1) where $g \geq 2$ is the genus of X. If γ is a closed oriented geodesic, by abuse of notation we let $\ell_w(\gamma)$ denote the minimal word-length of an element in the conjugacy class in Γ_g specified by γ (on page 5 we used the same notation for an element of Γ_g). We want to compare $\ell(\gamma)$ and $\ell_w(\gamma)$. We will use the following simple consequence of the Švarc–Milnor lemma [BH99, Prop. 8.19].

Lemma 2.3. *With notations as above, there exist constants $K_1, K_2 \geq 0$ depending on X such that*

$$\ell_w(\gamma) \leq K_1 \ell(\gamma) + K_2.$$

2.2 Choice of function for use in Selberg's trace formula. We now fix a function $\varphi_0 \in C_c^\infty(\mathbf{R})$ which has the following key properties:

1. φ_0 is non-negative and even.
2. $\mathrm{Supp}(\varphi_0) = (-1, 1)$.
3. The Fourier transform $\widehat{\varphi_0}$ satisfies $\widehat{\varphi_0}(\xi) \geq 0$ for all $\xi \in \mathbf{R} \cup i\mathbf{R}$.

Proof that such a function exists. Let ψ_0 be a C^∞, even, real-valued non-negative function whose support is exactly $(-\frac{1}{2}, \frac{1}{2})$. Let $\varphi_0 \stackrel{\text{def}}{=} \psi_0 \star \psi_0$ where

$$\psi_0 \star \psi_0(x) \stackrel{\text{def}}{=} \int_{\mathbf{R}} \psi_0(x-t)\psi_0(t)dt.$$

Then φ_0 has the desired properties. □

We now fix a function φ_0 as above and for any $T > 0$ define

$$\varphi_T(x) \stackrel{\text{def}}{=} \varphi_0\left(\frac{x}{T}\right).$$

Lemma 2.4. *For all $\varepsilon > 0$, there exists $C_\varepsilon > 0$ such that for all $t \in \mathbf{R}_{\geq 0}$ and for all $T > 0$*

$$\widehat{\varphi_T}(it) \geq C_\varepsilon T e^{T(1-\varepsilon)t}.$$

Proof. First observe that

$$\widehat{\varphi_T}(it) = T\widehat{\varphi_0}(Tit) = T\int_{\mathbf{R}} \varphi_0(x)e^{Txt}dx.$$

Using $t \geq 0$ and $\mathrm{Supp}(\varphi_0) = (-1, 1)$ with φ_0 non-negative, we have for some $C_\varepsilon > 0$

$$\widehat{\varphi_T}(it) \geq T\int_{1-\varepsilon}^{1} \varphi_0(x)e^{Txt}dx \geq TC_\varepsilon e^{T(1-\varepsilon)t}.$$ □

2.3 Proof of Theorem 1.5. Let X be a genus g compact hyperbolic surface and let X_ϕ be the cover of X corresponding to $\phi \in \mathrm{Hom}(\Gamma_g, S_n)$ constructed in the introduction. In what follows we let

$$T = 4 \log n.$$

For every $\gamma \in \mathcal{L}(X)$, we pick $\tilde{\gamma} \in \Gamma_g$ in the conjugacy class in Γ_g corresponding to γ (so in particular $\ell_w(\tilde{\gamma}) = \ell_w(\gamma)$). Every closed oriented geodesic δ in X_ϕ covers, via the Riemannian covering map $X_\phi \to X$, a unique closed oriented geodesic in X that we will call $\pi(\delta)$. This gives a map

$$\pi : \mathcal{L}(X_\phi) \to \mathcal{L}(X).$$

Note that $\ell(\delta) = \ell(\pi(\delta))$. We claim that $|\pi^{-1}(\gamma)| = \mathsf{fix}_{\tilde{\gamma}}(\phi)$, recalling that $\mathsf{fix}_{\tilde{\gamma}}(\phi)$ is the number of fixed points of $\phi(\tilde{\gamma})$. Indeed, by its very definition, X_ϕ is a fiber bundle over X with fiber $[n]$. If $\gamma \in \mathcal{P}(X)$, and we fix some regular point $o \in \gamma$ (not a self-intersection point), then in X_ϕ, the fiber of o can be identified with $[n]$. The oriented geodesic path $\gamma \backslash \{o\}$ lifts to n oriented geodesic paths with start and end points equal to $[n]$. The permutation of $[n]$ obtained by following these from start to end is (up to conjugation) $\phi(\tilde{\gamma})$ and hence, the δ's with $\pi(\delta) = \gamma$ are precisely the paths that close up, or in other words, the δ's with $\pi(\delta) = \gamma$ correspond to fixed points of $\phi(\tilde{\gamma})$. For general $\gamma \in \mathcal{L}(X)$, assume that $\gamma = {\gamma_0}^q$ with $q \geq 1$ and $\gamma_0 \in \mathcal{P}(X)$. A similar argument shows that the elements in $\pi^{-1}(\gamma)$ are in bijection with fixed points of $\tilde{\gamma_0}^q$ which we may take as our $\tilde{\gamma}$.

We also have $\mathrm{area}(X_\phi) = n \cdot \mathrm{area}(X)$. Now applying Theorem 2.1 to X_ϕ with the function φ_T gives

$$\begin{aligned}
\sum_{i=0}^{\infty} \widehat{\varphi_T}(r_i(X_\phi)) &= \frac{\mathrm{area}(X_\phi)}{4\pi} \int_{-\infty}^{\infty} r\widehat{\varphi_T}(r) \tanh(\pi r) dr + \sum_{\delta \in \mathcal{L}(X_\phi)} \frac{\Lambda(\delta)}{2\sinh\left(\frac{\ell(\delta)}{2}\right)} \varphi_T(\ell(\delta)) \\
&= \frac{n \cdot \mathrm{area}(X)}{4\pi} \int_{-\infty}^{\infty} r\widehat{\varphi_T}(r) \tanh(\pi r) dr \\
&\quad + \sum_{\gamma \in \mathcal{L}(X)} \sum_{\delta \in \pi^{-1}(\gamma)} \frac{\Lambda(\delta)}{2\sinh\left(\frac{\ell(\gamma)}{2}\right)} \varphi_T(\ell(\gamma)) \\
&= \frac{n \cdot \mathrm{area}(X)}{4\pi} \int_{-\infty}^{\infty} r\widehat{\varphi_T}(r) \tanh(\pi r) dr \\
&\quad + \sum_{\gamma \in \mathcal{P}(X)} \frac{\mathsf{fix}_{\tilde{\gamma}}(\phi)\ell(\gamma)}{2\sinh\left(\frac{\ell(\gamma)}{2}\right)} \varphi_T(\ell(\gamma)) \\
&\quad + \sum_{\gamma \in \mathcal{L}(X) - \mathcal{P}(X)} \sum_{\delta \in \pi^{-1}(\gamma)} \frac{\Lambda(\delta)}{2\sinh\left(\frac{\ell(\gamma)}{2}\right)} \varphi_T(\ell(\gamma)),
\end{aligned}$$

where in the second equality we used the fact that for $\delta \in \mathcal{L}(X_\phi)$, $\ell(\delta) = \ell(\pi(\delta))$, and in the third equality we used that if $\gamma \in \mathcal{P}(X)$, then $\delta \in \mathcal{P}(X_\phi)$ for all

$\delta \in \pi^{-1}(\gamma)$, so $\Lambda(\delta) = \Lambda(\gamma) = \ell(\gamma)$. Let $i_1, i_2, i_3, \dots$ be a subsequence of $1, 2, 3, \dots$ such that

$$0 \le \lambda_{i_1}(X_\phi) \le \lambda_{i_2}(X_\phi) \le \cdots$$

are the *new* eigenvalues of X_ϕ. Thus $\lambda_{i_1}(X_\phi)$ is the smallest new eigenvalue of X_ϕ. Taking the difference of the above formula with the trace formula for X (with the same function φ_T) gives

$$\begin{aligned}\sum_{j=1}^{\infty} \widehat{\varphi_T}(r_{i_j}(X_\phi)) = \frac{(n-1)\cdot \mathrm{area}(X)}{4\pi} \int_{-\infty}^{\infty} r\widehat{\varphi_T}(r)\tanh(\pi r)dr \\ + \sum_{\gamma\in\mathcal{P}(X)} \frac{(\mathsf{fix}_{\tilde{\gamma}}(\phi)-1)\ell(\gamma)}{2\sinh\left(\frac{\ell(\gamma)}{2}\right)}\varphi_T(\ell(\gamma)) \\ + \sum_{\gamma\in\mathcal{L}(X)-\mathcal{P}(X)} \frac{\varphi_T(\ell(\gamma))}{2\sinh\left(\frac{\ell(\gamma)}{2}\right)} \left(\left(\sum_{\delta\in\pi^{-1}(\gamma)} \Lambda(\delta)\right) - \Lambda(\gamma)\right). \quad (2.1)\end{aligned}$$

Since φ_T is non-negative and for any $\gamma \in \mathcal{L}(X)$, $|\pi^{-1}(\gamma)| \le n$, and $\Lambda(\delta) \le \ell(\delta) = \ell(\gamma)$ for all $\delta \in \pi^{-1}(\gamma)$, the sum on the bottom line of (2.1) is bounded from above by

$$n \sum_{\gamma\in\mathcal{L}(X)-\mathcal{P}(X)} \frac{\varphi_T(\ell(\gamma))}{2\sinh\left(\frac{\ell(\gamma)}{2}\right)} \cdot \ell(\gamma) = n \sum_{\gamma\in\mathcal{P}(X)} \sum_{k=2}^{\infty} \frac{\varphi_T(k\ell(\gamma))}{2\sinh\left(\frac{k\ell(\gamma)}{2}\right)} k\ell(\gamma). \quad (2.2)$$

We have

$$\sum_{k=2}^{\infty} \frac{\varphi_T(k\ell(\gamma))}{2\sinh\left(\frac{k\ell(\gamma)}{2}\right)} k\ell(\gamma) \overset{(*)}{\ll_X} \ell(\gamma) \sum_{k=2}^{\infty} k e^{-\frac{k\ell(\gamma)}{2}} \overset{(**)}{\ll_X} \ell(\gamma) e^{-\ell(\gamma)}, \quad (2.3)$$

where in $(*)$ we relied on that φ_T is bounded, and in both $(*)$ and $(**)$ on that there is a positive lower bound on the lengths of closed geodesics in X. As φ_T is supported on $(-T, T)$, the left hand side of (2.3) vanishes whenever $\ell(\gamma) \ge T/2$. Using Lemma 2.2 we thus get

$$\begin{aligned} n \sum_{\gamma\in\mathcal{P}(X)} \sum_{k=2}^{\infty} \frac{\varphi_T(k\ell(\gamma))}{2\sinh\left(\frac{k\ell(\gamma)}{2}\right)} k\ell(\gamma) &\ll_X n \sum_{\gamma\in\mathcal{P}(X):\ell(\gamma)\le T} \ell(\gamma)e^{-\ell(\gamma)} \\ &\le n \sum_{m=0}^{T} \sum_{\gamma\in\mathcal{L}(X):\, m\le\ell(\gamma)<m+1} (m+1)\, e^{-m} \\ &\ll_X n \sum_{m=0}^{T} (m+1)e^{m+1}e^{-m} \ll nT^2. \quad (2.4)\end{aligned}$$

We also have

$$\begin{aligned}\int_{-\infty}^{\infty} r\widehat{\varphi_T}(r)\tanh(\pi r)dr &= T\int_{-\infty}^{\infty} r\widehat{\varphi_0}(Tr)\tanh(\pi r)dr\\ &= \frac{1}{T}\int_{-\infty}^{\infty} r'\widehat{\varphi_0}(r')\tanh\left(\pi\frac{r'}{T}\right)dr'\\ &\le \frac{2}{T}\int_{0}^{\infty} |r'||\widehat{\varphi_0}(r')|dr' \ll \frac{1}{T}. \end{aligned} \tag{2.5}$$

The final estimate uses that, since φ_0 is compactly supported, $\widehat{\varphi_0}$ is a Schwartz function and decays faster than any inverse of a polynomial. Combining (2.1), (2.2), (2.4) and (2.5) gives

$$\begin{aligned}\sum_{j=1}^{\infty}\widehat{\varphi_T}(r_{i_j}(X_\phi)) &= O\left(\frac{(n-1)\cdot \mathrm{area}(X)}{4\pi}\cdot\frac{1}{T}\right)\\ &\quad+\sum_{\gamma\in\mathcal{P}(X)}\frac{(\mathsf{fix}_{\tilde\gamma}(\phi)-1)\,\ell(\gamma)}{2\sinh\left(\frac{\ell(\gamma)}{2}\right)}\varphi_T\left(\ell(\gamma)\right)+O_X\left(T^2n\right)\\ &= \sum_{\gamma\in\mathcal{P}(X)}\frac{(\mathsf{fix}_{\tilde\gamma}(\phi)-1)\,\ell(\gamma)}{2\sinh\left(\frac{\ell(\gamma)}{2}\right)}\varphi_T\left(\ell(\gamma)\right)+O_X\left(T^2n\right), \end{aligned} \tag{2.6}$$

where in the last equality we used $T>1$.

We are now in a position to use Theorem 1.11. The contributions to the sum above come from γ with $\ell(\gamma)\le T$. By Lemma 2.3, this entails $\ell_w(\tilde\gamma)=\ell_w(\gamma)\le K_1T+K_2\le c\log n$ for some $c=c(X)>0$ and n sufficiently large. Moreover, if $\gamma\in\mathcal{P}(X)$, then $\tilde\gamma$ is not a proper power in Γ_g. Thus for each γ appearing in (2.6), Theorem 1.11 applies to give

$$\mathbb{E}_{g,n}\left[\mathsf{fix}_{\tilde\gamma}(\phi)-1\right]\ll_X\frac{(\log n)^A}{n}$$

where $A=A(g)>0$ and the implied constant depends on X. Now using that $\widehat{\varphi_T}$ is non-negative on $\mathbf{R}\cup i\mathbf{R}$, we take expectations of (2.6) with respect to the uniform measure on $\mathbb{X}_{g,n}$ to obtain

$$\begin{aligned}&\mathbb{E}_{g,n}\left[\widehat{\varphi_T}\left(r_{i_1}(X_\phi)\right)\right]\\ &\quad\le\sum_{\gamma\in\mathcal{P}(X)}\frac{\mathbb{E}_{g,n}\left[\mathsf{fix}_{\tilde\gamma}(\phi)-1\right]\ell(\gamma)}{2\sinh\left(\frac{\ell(\gamma)}{2}\right)}\varphi_T(\ell(\gamma))+O_X\left(T^2n\right)\\ &\quad\overset{\text{Theorem 1.11}}{\ll_X}\frac{(\log n)^A}{n}\sum_{\gamma\in\mathcal{P}(X)}\frac{\ell(\gamma)}{2\sinh\left(\frac{\ell(\gamma)}{2}\right)}\varphi_T(\ell(\gamma))+T^2n\\ &\quad\ll_X\frac{(\log n)^A}{n}\sum_{\gamma\in\mathcal{P}(X):\,\ell(\gamma)\le T}\ell(\gamma)e^{-\ell(\gamma)/2}+T^2n\end{aligned}$$

$$
\begin{aligned}
&\leq \frac{(\log n)^A}{n} \sum_{m=0}^{\lceil T-1 \rceil} \sum_{\gamma \in \mathcal{L}(X)\,:\, m \leq \ell(\gamma) < m+1} (m+1)e^{-m/2} + T^2 n \\
&\overset{\text{Lemma 2.2}}{\ll_X} \frac{(\log n)^A}{n} \sum_{m=0}^{\lceil T-1 \rceil} (m+1)e^{-m/2}e^{m+1} + T^2 n \\
&\ll \frac{(\log n)^A}{n} T e^{T/2} + T^2 n \\
&\overset{T=4\log n}{\ll_\varepsilon} n^{1+\varepsilon/3},
\end{aligned} \tag{2.7}
$$

where ε is the parameter given in Theorem 1.5. The third inequality above used that on a compact hyperbolic surface, the lengths of closed geodesics are bounded below away from zero (by the Collar Lemma [Bus10, Thm. 4.1.1]), together with the fact that φ_T is supported in $[-T, T]$. So $\mathbb{E}_{g,n}\left[\widehat{\varphi_T}\left(r_{i_1}(X_\phi)\right)\right] \leq n^{1+\varepsilon/2}$ for large enough n, and for these values of n, by Markov's inequality

$$
\mathbb{P}\left[\widehat{\varphi_T}(r_{i_1}(X_\phi)) > n^{1+\varepsilon}\right] \leq n^{-\varepsilon/2}. \tag{2.8}
$$

Lemma 2.4 implies that if $\lambda_{i_1}(X_\phi) \leq \frac{3}{16} - \varepsilon$, in which case $r_{i_1}(X_\phi) = it_\phi$ with $t_\phi \in \mathbf{R}$ and $t_\phi \geq \sqrt{\frac{1}{16} + \varepsilon} \geq \frac{1}{4} + \varepsilon$ for ε sufficiently small, then

$$
\widehat{\varphi_T}(r_{i_1}(X_\phi)) \geq C_\varepsilon T e^{T(1-\varepsilon)t_\phi} \geq C_\varepsilon n^{4(1-\varepsilon)(1/4+\varepsilon)} \geq C_\varepsilon n^{1+2\varepsilon} > n^{1+\varepsilon}, \tag{2.9}
$$

by decreasing ε if necessary, and then assuming n is sufficiently large. Combining (2.8) and (2.9) gives

$$
\mathbb{P}\left[X_\phi \text{ has a new eigenvalue } \leq \frac{3}{16} - \varepsilon\right] \leq \mathbb{P}\left[\widehat{\varphi_T}(r_{i_1}(X_\phi)) > n^{1+\varepsilon}\right] \leq n^{-\varepsilon/2}
$$

completing the proof of Theorem 1.5, under the assumption of Theorem 1.11. □

2.4 Proof of Theorem 1.8. We continue using the same notation as in the previous section, including the choice of $T = 4\log n$. We let

$$
0 \leq \lambda_{i_1}(X_\phi) \leq \lambda_{i_2}(X_\phi) \leq \cdots \leq \lambda_{i_{k(\phi)}} \leq \frac{1}{4}
$$

denote the collection of new eigenvalues of X_ϕ of size at most $\frac{1}{4}$, with multiplicities included. For each such eigenvalue we write $\lambda_{i_j} = s_{i_j}(1 - s_{i_j})$ with $s_{i_j} \in \left[\frac{1}{2}, 1\right]$; this has the result that $r_{i_j} = i(s_{i_j} - \frac{1}{2})$.

Again taking expectations of (2.6) with respect to the uniform measure on $\mathbb{X}_{g,n}$, but this time, keeping more terms, gives

$$\mathbb{E}_{g,n}\left[\sum_{j=1}^{k(\phi)}\widehat{\varphi_T}\left(r_{i_j}(X_\phi)\right)\right] \leq \sum_{\gamma\in\mathcal{P}(X)}\frac{\mathbb{E}_{g,n}\left[\mathsf{fix}_{\tilde{\gamma}}(\phi)-1\right]\ell(\gamma)}{2\sinh\left(\frac{\ell(\gamma)}{2}\right)}\varphi_T(\ell(\gamma))+O_X\left(T^2n\right)$$
$$\ll_{X,\varepsilon} n^{1+\varepsilon/3}$$

by (2.7). On the other hand, Lemma 2.4 implies that for $\varepsilon\in(0,1)$

$$\sum_{j=1}^{k(\phi)}\widehat{\varphi_T}\left(r_{i_j}\left(X_\phi\right)\right)\gg_\varepsilon\sum_{j=1}^{k(\phi)}Te^{T(1-\varepsilon)\left(s_{i_j}(X_\phi)-\frac{1}{2}\right)}\gg\sum_{j=1}^{k(\phi)}n^{4(1-\varepsilon)\left(s_{i_j}(X_\phi)-\frac{1}{2}\right)}.$$

Therefore

$$\mathbb{E}_{g,n}\left[\sum_{j=1}^{k(\phi)}n^{4(1-\varepsilon)\left(s_{i_j}(X_\phi)-\frac{1}{2}\right)}\right]\leq n^{1+\varepsilon/2}$$

for n sufficiently large. Markov's inequality therefore gives

$$\mathbb{P}\left[\sum_{j=1}^{k(\phi)}n^{4(1-\varepsilon)\left(s_{i_j}(X_\phi)-\frac{1}{2}\right)}\geq n^{1+\varepsilon}\right]\leq n^{-\varepsilon/2}$$

so a.a.s. $\sum_{j=1}^{k(\phi)}n^{4(1-\varepsilon)\left(s_{i_j}(X_\phi)-\frac{1}{2}\right)}<n^{1+\varepsilon}$. This gives that for any $\sigma\in\left(\frac{1}{2},1\right)$, a.a.s.

$$\#\left\{1\leq j\leq k(\phi)\,:\,s_{i_j}>\sigma\right\}\leq n^{1+\varepsilon-4(1-\varepsilon)(\sigma-\frac{1}{2})}\leq n^{3-4\sigma+3\varepsilon}.$$

This finishes the proof of Theorem 1.8 assuming Theorem 1.11. □

3 Tiled Surfaces

3.1 Tiled surfaces. Here we assume $g=2$, and let $\Gamma\overset{\text{def}}{=}\Gamma_2$. We write $\mathbb{X}_n\overset{\text{def}}{=}\mathbb{X}_{2,n}$ throughout the rest of the paper. Consider the construction of the surface Σ_2 from an octagon by identifying its edges in pairs according to the pattern $aba^{-1}b^{-1}cdc^{-1}d^{-1}$. This gives rise to a CW-structure on Σ_2 consisting of one vertex (denoted o), four oriented 1−cells (labeled by a,b,c,d) and one 2-cell which is the octagon glued along eight 1-cells.[6] See Figure 1. We identify Γ_2 with $\pi_1\left(\Sigma_2,o\right)$, so that in the presentation (1.1), words in the generators a,b,c,d correspond to the homotopy classes of the corresponding closed paths based at o along the 1-skeleton of Σ_2.

Note that every covering space $p\colon\Upsilon\to\Sigma_2$ inherits a CW-structure from Σ_2: the vertices are the pre-images of o, and the open 1-cells (2-cells) are the connected

[6] We use the terms vertices, edges and octagons interchangeably with 0-cells, 1-cells and 2-cells, respectively.

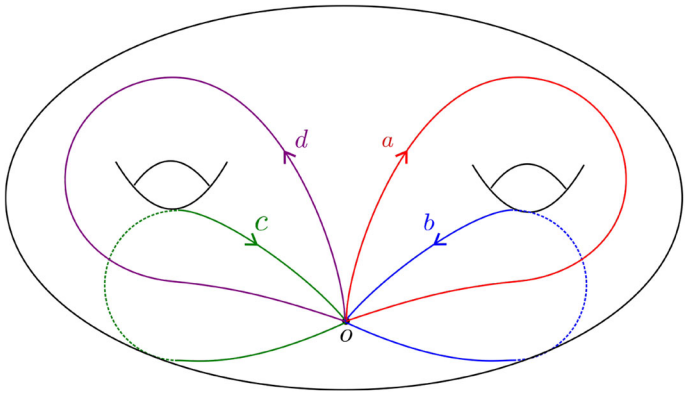

Figure 1: The CW-structure we give to the surface Σ_2 with fundamental group $\Gamma = \Gamma_2 = \langle a, b, c, d \mid [a, b]\,[c, d]\rangle$: it consists of a single vertex (0-cell), four edges (1-cells) and one octagon (a 2-cell).

components of the pre-images of the open 1-cells (2-cells, respectively) in Σ_2. In particular, this is true for the universal covering space $\widetilde{\Sigma_2}$ of Σ_2, which we can now think of as a CW-complex. A sub-complex of a CW-complex is a subspace consisting of cells such that if some cell belongs to the subcomplex, then so are the cells of smaller dimension at its boundary.

Definition 3.1. *(Tiled surface)* [MP21, Def. 3.1]. *A* tiled surface *Y is a sub-complex of a (not-necessarily-connected) covering space of Σ_2. In particular, a tiled surface is equipped with the restricted covering map $p\colon Y \to \Sigma_2$ which is an immersion. We denote by $Y^{(0)}$ the set of vertices and by $Y^{(1)}$ the 1-skeleton of Y. If Y is compact, we write $\mathfrak{v}(Y)$ for the number of vertices of Y, $\mathfrak{e}(Y)$ for the number of edges and $\mathfrak{f}(Y)$ for the number of octagons.*

Alternatively, instead of considering a tiled surface Y to be a complex equipped with a restricted covering map, one may consider Y to be a complex as above with directed and labeled edges: the directions and labels (a, b, c, d) are pulled back from Σ_2 via p. These labels uniquely determine p as a combinatorial map between complexes.

Note that a tiled surface is not always a surface: for example, it may also contain vertices or edges with no 2-cells incident to them. However, as Y is a sub-complex of a covering space of Σ_2, namely, of a surface, any neighborhood of Y inside the cover is a surface, and it is sometimes beneficial to think of Y as such.

Definition 3.2. *(Thick version of a tiled surface)* [MP21, Def. 3.2]. *Given a tiled surface Y which is a subcomplex of the covering space Υ of Σ_2, consider a small, closed, regular neighborhood of Y in Υ. This neighborhood is a closed surface, possibly with boundary, which is referred to as the* thick version of Y.

We let ∂Y denote the boundary of the thick version of Y and $\mathfrak{d}(Y)$ denote the number of edges along ∂Y (so if an edge of Y does not border any octagon, it is counted twice).

We stress that we do not think of Y as a sub-complex, but rather as a complex for its own sake, which happens to have the capacity to be realized as a subcomplex of a covering space of Σ_2. In particular, if Y is compact, it is a combinatorial object given by a finite amount of data. See [MP21, Section 3] for a more detailed discussion.

Definition 3.3. *(Morphisms of tiled surfaces). Let $p_i\colon Y_i \to \Sigma_2$ be tiled surfaces for $i = 1, 2$. A map $f\colon Y_1 \to Y_2$ is a* morphism *of tiled surfaces if it is a combinatorial map of CW-complexes that commutes with the restricted covering maps.*

$$\begin{tikzcd} Y_1 \arrow[rr, "f"] \arrow[dr, "p_1"'] & & Y_2 \arrow[dl, "p_2"] \\ & \Sigma_2 & \end{tikzcd}$$

In other words, a morphism of tiled surfaces is a combinatorial map of CW-complexes sending i-cells to i-cells and which respects the directions and labels of edges.

Example 3.4. The fibered product construction gives a one-to-one correspondence between $\mathrm{Hom}(\Gamma, S_n)$ and topological degree-n covers of Σ_2 with a labeled fiber over the basepoint o. Explicitly, for $\phi \in \mathrm{Hom}(\Gamma, S_n)$, we can consider the quotient

$$X_\phi \stackrel{\text{def}}{=} \Gamma \backslash \left(\widetilde{\Sigma_2} \times [n] \right)$$

where $\widetilde{\Sigma_2}$ is the universal cover of Σ_2 (an open disc) and Γ acts on $\widetilde{\Sigma_2} \times [n]$ diagonally, by the usual action of Γ on $\widetilde{\Sigma_2}$ on the first factor, and via ϕ on the second factor. The covering map $X_\phi \to \Sigma_2$ is induced by the projection $\widetilde{\Sigma_2} \times [n] \to \widetilde{\Sigma_2}$.

Being a covering space of Σ_2, each X_ϕ is automatically also a tiled surface. The fiber of $o \in \Sigma_2$ is the collection of vertices of X_ϕ. We fix throughout the rest of the paper a vertex $u \in \widetilde{\Sigma_2}$ lying over $o \in \Sigma_2$. This identifies the fiber over o in X_ϕ with $\{u\} \times [n]$ and hence gives a fixed bijection between the vertices of X_ϕ and the numbers in $[n]$. The map $\phi \mapsto X_\phi$ is the desired one-to-one correspondence between $\mathrm{Hom}(\Gamma, S_n)$ and topological degree-n covers of Σ_2 with the fiber over o labeled bijectively by $[n]$.

Example 3.5. For any $1 \neq \gamma \in \Gamma$, pick a word $\tilde{\gamma}$ of minimal length in the letters a, b, c, d and their inverses that represents an element in the conjugacy class of γ in Γ. In particular, $\tilde{\gamma}$ is cyclically reduced. Now take a circle and divide it into $\{a, b, c, d\}$-labeled and directed edges separated by vertices, such that following around the circle from some vertex and in some orientation, and reading off the labels and directions, spells out $\tilde{\gamma}$. Call the resulting complex $\mathcal{C}_\gamma$. That $\mathcal{C}_\gamma$ is a tiled surface

follows from [MP21] (in particular, it is embedded in the core surface $\mathrm{Core}(\langle\gamma\rangle)$ which is itself a tiled surface, by [MP21, Thm. 5.10]). Note that generally $\mathcal{C}_\gamma$ is not uniquely determined by γ (e.g., [MP21, Figure 1.2 and Section 4, Section 5]), and we choose one of the options arbitrarily. We have $\mathfrak{v}(\mathcal{C}_\gamma) = \mathfrak{e}(\mathcal{C}_\gamma) = \ell_w(\gamma)$ and $\mathfrak{f}(\mathcal{C}_\gamma) = 0$.

If Y is a compact tiled surface, there are some simple relations between the quantities $\mathfrak{v}(Y)$, $\mathfrak{e}(Y)$, $\mathfrak{f}(Y)$, $\mathfrak{d}(Y)$, and $\chi(Y)$, the topological Euler characteristic of Y. We note the following relations, which are straightforward or standard. For example, $\mathfrak{e}(Y) \le 4\mathfrak{v}(Y)$ as each vertex is incident to at most 8 half-edges.[7]

$$\mathfrak{d}(Y) = 2\mathfrak{e}(Y) - 8\mathfrak{f}(Y). \tag{3.1}$$

$$4\mathfrak{f}(Y) \le \mathfrak{e}(Y) \le 4\mathfrak{v}(Y). \tag{3.2}$$

The following lemma will be useful later.

Lemma 3.6. *Let Y be a compact tiled surface without isolated vertices. Then*

$$\mathfrak{v}(Y) \le \mathfrak{f}(Y) + \mathfrak{d}(Y).$$

Proof. Let $\mathfrak{i}$ denote the number of *internal* vertices of Y, namely, vertices adjacent to 8 octagons, and let $\mathfrak{p}$ denote the number of the remaining, *peripheral* vertices. As there are no isolated vertices, $\mathfrak{p} \le \mathfrak{d}(Y)$ (when going through the boundary cycles, one edge at a time, one passes at every step exactly one peripheral vertex, and each peripheral vertex is traversed at least once, although possibly more than once). We have

$$\begin{aligned} 8\mathfrak{f}(Y) &= \sum_{O \text{ an octagon of } Y} \#\{\text{corners of } O\} \\ &= \sum_{v \text{ a vertex of } Y} \#\{\text{corners of octagons at } v\} \\ &\ge 8\mathfrak{i} = 8\mathfrak{v}(Y) - 8\mathfrak{p} \ge 8\mathfrak{v}(Y) - 8\mathfrak{d}(Y). \end{aligned}$$

□

The Euler characteristic $\chi(Y)$ is also controlled by $\mathfrak{f}(Y)$ and $\mathfrak{d}(Y)$.

Lemma 3.7. *Let Y be a compact tiled surface without isolated vertices. Then*[8]

$$\chi(Y) \le \frac{\mathfrak{d}(Y)}{2} - 2\mathfrak{f}(Y).$$

Proof. We have

$$\chi(Y) = \mathfrak{v}(Y) - \mathfrak{e}(Y) + \mathfrak{f}(Y) \overset{(3.1)}{=} \mathfrak{v}(Y) - 3\mathfrak{f}(Y) - \frac{\mathfrak{d}(Y)}{2} \overset{\text{Lemma 3.6}}{\le} \frac{\mathfrak{d}(Y)}{2} - 2\mathfrak{f}(Y).$$

□

[7] For general $g \ge 2$, (3.1) is $\mathfrak{d}(Y) = 2\mathfrak{e}(Y) - 4g\mathfrak{f}(Y)$, and (3.2) is $2g\mathfrak{f}(Y) \le \mathfrak{e}(Y) \le 2g\mathfrak{f}(Y)$.

[8] For arbitrary $g \ge 2$, the bound is $\chi(Y) \le \frac{\mathfrak{d}(Y)}{2} - (2g-2)\mathfrak{f}(Y)$.

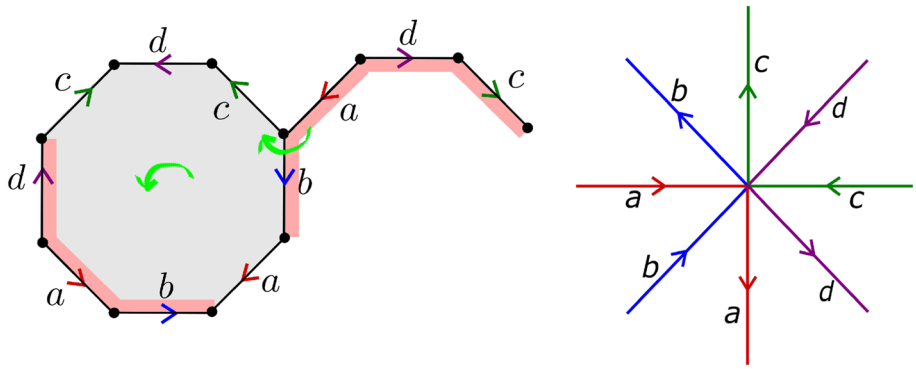

Figure 2: The right figure shows a vertex with 8 half-edges around it, ordered (clockwise) according to the fixed cyclic order induced from the CW-structure on Σ_2. On the left is a tiled surface with 11 vertices, 11 edges and one octagon. The orientation on the octagon is counter-clockwise, while around any vertex it is clockwise. The pink stripes describe blocks: a half-block spelling $c^{-1}d^{-1}ab$ and a block of length 3 spelling $d^{-1}ab$. The latter one can be extended in both ends.

3.2 Blocks and chains. Here we introduce language that was used in [MP21, MP20], based on terminology of Birman and Series from [BS87a]. Let Y denote a tiled surface throughout this Section 3.2. When we refer to directed edges of Y, they are not necessarily directed according to the definition of Y.

First of all, we augment Y by adding **half-edges**, which should be thought of as copies of $[0, \frac{1}{2})$. Of course, every edge of $Y^{(1)}$ is thought of as containing two half edges, each of which inherits a label in $\{a, b, c, d\}$ and a direction from their ambient edge. We add to Y $\{a, b, c, d\}$-labeled and directed half-edges to form Y_+ so that every vertex of Y_+ has exactly 8 emanating half-edges, with labels and directions given by 'a-outgoing, b-incoming, a-incoming, b-outgoing, c-outgoing, d-incoming, c-incoming, d-outgoing'. The cyclic order we have written here induces a fixed cyclic ordering on the half-edges at each vertex of Y_+. If a half-edge of Y_+ does not belong to an edge of Y (hence was added to Y_+), we call it a **hanging half-edge**. We may think of Y_+ as a surface too, by considering the thick version of Y and attaching a thin rectangle for every hanging half-edge. We call the resulting surface the **thick version of Y_+**, and mark its boundary by ∂Y_+. See Figure 2 for the cyclic ordering of half-edges around every vertex and Figure 4 for a piece of ∂Y_+.

For two directed edges $\vec{e_1}$ and $\vec{e_2}$ of Y with the terminal vertex v of $\vec{e}_1$ equal to the source of $\vec{e}_2$, the *half-edges between $\vec{e}_1$ and $\vec{e}_2$* are by definition the half edges of Y_+ at v that are strictly between $\vec{e}_1$ and $\vec{e}_2$ in the given cyclic ordering. There are m of these where $0 \leq m \leq 7$.

A **path** in Y is a sequence $\mathcal{P}=(\vec{e_1}, \ldots, \vec{e}_k)$ of directed edges in $Y^{(1)}$, such that for each $1 \leq i \leq k-1$ the terminal vertex of $\vec{e}_i$ is the initial vertex of $\vec{e}_{i+1}$. A **cycle** in Y is a cyclic sequence $\mathcal{C} =(\vec{e_1}, \ldots, \vec{e}_k)$ which is a path with the terminal vertex of $\vec{e}_k$ identical to the initial vertex of $\vec{e}_1$. A **boundary cycle** of Y is a cycle corresponding to a

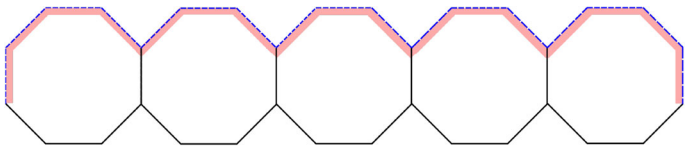

Figure 3: A long chain (the pink stripe) consisting of five consecutive blocks of lengths $4, 3, 3, 3, 4$.

boundary component of the thick version of Y. A boundary cycle is always oriented so that if Y is embedded in the full cover Z, the boundary reads successive segments of the boundaries of the neighboring octagons (in $Z - Y$) with the orientation of each octagon coming from $[a, b]\,[c, d]$ (and not from the inverse word). For example, the unique boundary cycle of the tiled surface in the left side of Figure 2, starting at the rightmost vertex, spells the cyclic word $c^{-1}d^{-1}abab^{-1}a^{-1}dcd^{-1}c^{-1}a^{-1}dc$.

If $\mathcal{P}$ is a path in $Y^{(1)}$, a **block** in $\mathcal{P}$ is a non-empty (possibly cyclic) subsequence of successive edges, each successive pair of edges having no half-edges between them (this means that a block reads necessarily a subword of the cyclic word $[a, b]\,[c, d]$). A **half-block** is a block of length 4 (in general, $2g$) and a **long block** is a block of length at least 5 (in general, $2g + 1$). See Figure 2.

Two blocks $(\vec{e}_i, \dots, \vec{e}_j)$ and $(\vec{e}_k, \dots, \vec{e}_\ell)$ in a path $\mathcal{P}$ are called *consecutive* if $(\vec{e}_i, \dots, \vec{e}_j, \vec{e}_k, \dots, \vec{e}_\ell)$ is a (possibly cyclic) subsequence of $\mathcal{P}$ and there is precisely one half-edge between $\vec{e}_j$ and $\vec{e}_k$. A **chain** is a (possibly cyclic) sequence of consecutive blocks. Note that in a chain, an f-edge with some $f \in \{a^{+1}, \dots, d^{\pm 1}\}$ is followed by an edge labeled by the letter f' that follows f in the cyclic word $[a, b]\,[c, d]$, or by the letter that follows the inverse of f'. For example, a b^{-1}-edge is always followed in a chain by either a c-edge or a d^{-1}-edge. A *cyclic chain* is a chain whose blocks pave an entire cycle (with exactly one half-edge between the last block and the first blocks). A **long chain** is a chain consisting of consecutive blocks of lengths

$$4, 3, 3, \dots, 3, 4$$

(in general, $2g, 2g - 1, 2g - 1, \dots, 2g - 1, 2g$). See Figure 3. A **half-chain** is a cyclic chain consisting of consecutive blocks of length 3 (in general, $2g - 1$) each.

3.3 Boundary reduced and strongly boundary reduced tiled surfaces.

We recall the following definitions from [MP21, Def. 4.1, 4.2].

Definition 3.8. *(Boundary reduced). A tiled surface Y is* boundary reduced *if no boundary cycle of Y contains a long block or a long chain.*

Definition 3.9. *(Strongly boundary reduced). A tiled surface Y is* strongly boundary reduced *if no boundary cycle of Y contains a half-block or is a half-chain.*

Given a tiled surface Y embedded in a boundary reduced tiled surface Z, the BR-closure of Y in Z, denoted $\mathsf{BR}\,(Y \hookrightarrow Z)$ and introduced in [MP21, Def. 4.4], is defined as the intersection of all boundary reduced sub-tiled surfaces of Z containing Y. We compile some properties of the BR-closure into the following proposition.

Proposition 3.10. *Let $Y \hookrightarrow Z$ be an embedding of a compact tiled surface Y into a boundary reduced tiled surface Z, and denote $Y' \stackrel{\text{def}}{=} \mathsf{BR}(Y \hookrightarrow Z)$.*

1. [MP21, Prop. 4.5] *Y' is boundary reduced.*
2. [MP21, proof of Prop. 4.6] *Y' is compact, and $\mathfrak{d}(Y') \leq \mathfrak{d}(Y)$, with equality if and only if $Y' = Y$.*
3. [MP21, proof of Prop. 4.6] *Y' can be obtained from Y by initializting $Y' = Y$ and then repreatedly either (i) annexing an octagon of $Z \backslash Y'$ which borders a long block along $\partial Y'$, or (ii) annexing the octagons of $Z \backslash Y'$ bordering some long chain along $\partial Y'$, until Y' is boundary reduced.*
4. *We have*[9]

$$\mathfrak{f}(Y') \leq \mathfrak{f}(Y) + \frac{\mathfrak{d}(Y)^2}{6}.$$

Proof of item 4. Assume that Y' is obtained from Y by the procedure described in item 3. In each such step, $\mathfrak{d}(Y')$ decreases by at least two, so there are at most $\frac{\mathfrak{d}(Y)}{2}$ steps where octagons are added. We will be done by showing that at each step at most $\frac{\mathfrak{d}(Y)}{3}$ octagons are added. And indeed, in option (i) exactly one octagon is added (and $1 \leq \frac{\mathfrak{d}(Y)}{3}$ or otherwise Y is boundary reduced). In option (ii), if the long chain consists of ℓ blocks, it is of length $3\ell + 2 \leq \mathfrak{d}(Y')$, and at most $\ell \leq \frac{\mathfrak{d}(Y')-2}{3} < \frac{\mathfrak{d}(Y)}{3}$ new octagons are added. □

3.4 Pieces and ε-adapted tiled surfaces. For the proof of Theorem 1.11 we will need to quantify (strongly) boundary reduced tiled surfaces. This is captured by the notion of ε-adapted tiled surface we introduce in this Section 3.4. The following concepts of a piece and its defect play a crucial role here.

Definition 3.11. *(Piece, defect). A* piece *P of ∂Y_+ is a (possibly cyclic) path along ∂Y_+, consisting of whole directed edges and/or whole hanging half-edges. We write $\mathfrak{e}(P)$ for the number of full directed edges in P, $\mathfrak{he}(P)$ for the number of hanging half-edges in P, and $|P| \stackrel{\text{def}}{=} \mathfrak{e}(P) + \mathfrak{he}(P)$. We let*

$$\mathrm{Defect}(P) \stackrel{\text{def}}{=} \mathfrak{e}(P) - 3\mathfrak{he}(P).$$

(In general, $\mathrm{Defect}(P) \stackrel{\text{def}}{=} \mathfrak{e}(P) - (2g-1)\,\mathfrak{he}(P)$.) See Figure 4 for an illustration of a piece.

Definition 3.12. *(ε-adapted). Let $\varepsilon \geq 0$ and let Y be a tiled surface. A piece P of ∂Y_+ is ε-*adapted *if it satisfies*[10]

$$\mathrm{Defect}(P) \leq 4\chi(P) - \varepsilon|P|. \tag{3.3}$$

[9] For larger values of the genus g, we could get a tighter bound, but the stated bound holds and is good enough.

[10] In general, if $\mathrm{Defect}(P) \leq 2g \cdot \chi(P) - \varepsilon|P|$.

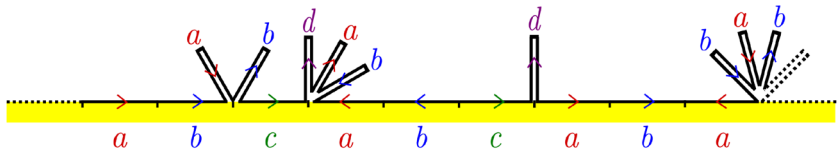

Figure 4: A piece P of ∂Y_+ is shown in black line. The broken black line marks parts of ∂Y_+ adjacent to but not part of P and the yellow stripe marks the side of the internal side of Y. This piece consists of 9 full directed edges and 9 hanging half-edges, so $\mathrm{Defect}\,(P) = -18$.

We have $\chi(P) = 0$ if P is a whole boundary component and $\chi(P) = 1$ otherwise. We say that a piece P is ε-bad if (3.3) does not hold, i.e., if $\mathrm{Defect}(P) > 4\chi(P) - \varepsilon|P|$. *We say that Y is ε-adapted if every piece of Y is ε-adapted.*

The following lemma shows that this notion indeed quantifies the notion of strongly boundary reduced tiled surfaces.

Lemma 3.13. *Let Y be a tiled surface.*

1. *Y is boundary reduced if and only if it is 0-adapted.*
2. *Y is strongly boundary reduced if and only if every piece of ∂Y is ε-adapted for some $\varepsilon > 0$. If Y is compact, this is equivalent to that Y is ε-adapted for some $\varepsilon > 0$.*

Proof. A block at ∂Y is a piece P with $\mathfrak{he}\,(P) = 0$. Assume that Y is 0-adapted. If P is a block at ∂Y, then $\mathfrak{e}\,(P) = \mathrm{Defect}\,(P) \leq 4\chi\,(P) \leq 4$, so P cannot be a long block. If P is a long chain at ∂Y consisting of k blocks ($k-2$ of length 3 and two of length 4) and the $k-1$ hanging edges between them, then $\mathrm{Defect}\,(P) = (3k+2) - 3\,(k-1) = 5 > 4 = 4\chi\,(P)$, which is a contradiction. Similarly, if P is a half-block or a half-chain, then $\mathrm{Defect}\,(P) = 4\chi\,(P)$, and so P is ε-bad for any $\varepsilon > 0$. The converse implications are not hard and can be found in [MP20, proof of Lem. 5.18]. □

We need the following lemma in the analysis of the next subsection.

Lemma 3.14. *If $0 \leq \varepsilon < 3$ and P is an ε-bad piece of a compact tiled surface Y, then*[11]

$$|P| < \frac{4\mathfrak{d}\,(Y)}{3-\varepsilon}. \tag{3.4}$$

Proof. If P is ε-bad, then by definition $\mathfrak{e}\,(P) - 3 \cdot \mathfrak{he}\,(P) > 4\chi\,(P) - \varepsilon\,|P|$. So

$$(3-\varepsilon)\,|P| < 3\,(\mathfrak{e}\,(P) + \mathfrak{he}\,(P)) + (\mathfrak{e}\,(P) - 3 \cdot \mathfrak{he}\,(P) - 4\chi\,(P)) \leq 4 \cdot \mathfrak{e}\,(P) \leq 4\mathfrak{d}\,(Y)\,.$$

□

[11] For arbitrary $g \geq 2$, the bound is $|P| < \frac{2g \cdot \mathfrak{d}(Y)}{2g-1-\varepsilon}$.

3.5 The octagons-vs-boundary algorithm. In this Section 3.5 we describe an algorithm whose purpose is to grow a given tiled surface in such a way that either

- the output Y' is ε-adapted for some fixed $\varepsilon > 0$, or alternatively,
- the number of octagons of Y' is larger than the length of the boundary of Y'.

If Y' is ε-adapted for a suitable ε, it is very well adapted to our methods, so that we can give an estimate for $\mathbb{E}_n^{\mathrm{emb}}(Y')$ with an effective error term (e.g., Proposition 5.27). If, on the other hand, $\mathfrak{f}(Y') > \mathfrak{d}(Y')$, then the Euler characteristic of Y' can be linearly comparable to the number of octagons in Y' by Lemma 3.7, and see Section 6.2 where it is used.

The algorithm depends on a positive constant $\varepsilon > 0$; we shall see below that fixing $\varepsilon = \frac{1}{32}$ works fine for our needs (for arbitrary $g \geq 2$ we shall fix $\varepsilon = \frac{1}{16g}$.) To force the algorithm to be deterministic, we a priori make some choices:

Notation 3.15. *For every compact tiled surface Y which is boundary reduced but not ε-adapted, we pick an ε-bad piece $P(Y)$ of ∂Y.*

With the ambient parameter ε fixed as well as the choices of ε-bad pieces, the octagons-vs-boundary (OvB) algorithm is as follows.

Input. An embedding of tiled surfaces $Y \hookrightarrow Z$ where Y is compact and Z has no boundary.

Output. A compact tiled surface Y' and a factorization of the input embedding $Y \hookrightarrow Z$ by $Y \hookrightarrow Y' \hookrightarrow Z$ where both maps are embeddings.

Algorithm. Let $Y' = Y$.
(a) Let $Y' = \mathsf{BR}(Y' \hookrightarrow Z)$. If

$$\theta(Y') \stackrel{\text{def}}{=} \mathfrak{f}(Y') - \mathfrak{d}(Y') > 0 \tag{3.5}$$

terminate the algorithm and return Y'.
(b) If Y' is not ε-adapted, add all the octagons of Z meeting[a] $P(Y')$ to Y', and go to **(a)**.
Return Y'.

[a] An octagon O in Z is said to meet $P(Y')$ if some directed edge or hanging-half-edge of $P(Y')$ lies at ∂O.

Note that the output Y' of the algorithm is always boundary reduced. Of course, we would like to know when/if this algorithm terminates.

In step **(a)**, if $\mathsf{BR}(Y' \hookrightarrow Z) \neq Y'$ then $\mathfrak{d}(Y')$ decreases by at least two, and $\mathfrak{f}(Y')$ increases by at least one. So $\theta(Y')$ increases by at least three.

In step **(b)**, if Y' changes, the following lemma shows that $\theta(Y')$ increases by at least one provided that $\varepsilon \leq \frac{1}{16}$.

Lemma 3.16. *With notation as above, if Y' is modified in step* **(b)**, *then*

1. $\mathfrak{d}(Y')$ *increases by less than* $2\varepsilon\left|P\left(Y'\right)\right|$.
2. $\theta(Y')$ *increases by more than* $\left(\frac{1}{8}-2\varepsilon\right)\left|P\left(Y'\right)\right|$, *so the increase is positive when*[12] $\varepsilon \leq \frac{1}{16}$.

Note that $\theta\left(Y'\right)$ is an integer, so any positive increase is an increase by at least one.

Proof. Suppose that in step **(b)** Y' is modified. Let Y'' denote the result of this modification and let $P = P(Y')$. Let k denote the number of new octagons added. First assume that P is a non-closed path, so $\chi\left(P\right) = 1$. We have $k \leq \mathfrak{he}\left(P\right) + 1$ because every hanging half-edge along P marks the passing from one new octagon to the next one. Every new octagon borders 8 edges in Z. For most new octagons, two of these edges contain hanging half-edges of P and are internal edges in Y'', so if j of the edges belong to P, the net contribution of the octagon to $\mathfrak{d}\left(Y''\right) - \mathfrak{d}\left(Y'\right)$ is at most $6-2j$. The exceptions are the two extreme octagons, which possibly meet only one hanging half-edge of P, and contribute a net of at most $7-2j$. The sum of the parameter j over all new octagons is exactly $\mathfrak{e}\left(P\right)$. In total, we obtain:

$$\begin{aligned}\mathfrak{d}\left(Y''\right) - \mathfrak{d}\left(Y'\right) &\leq 6k + 2 - 2\cdot\mathfrak{e}\left(P\right)\\ &\leq 6\left(\mathfrak{he}\left(P\right)+1\right) + 2 - 2\cdot\mathfrak{e}\left(P\right)\\ &= 2\left(3\cdot\mathfrak{he}\left(P\right) - \mathfrak{e}\left(P\right)\right) + 8\\ &< 2\left(\varepsilon\left|P\right| - 4\chi\left(P\right)\right) + 8 = 2\cdot\varepsilon\left|P\right|,\end{aligned}$$

where the last inequality comes from the definition of an ε-bad piece. If P is a whole boundary cycle of Y'_+, we have $k \leq \mathfrak{he}\left(P\right)$ and all octagons contribute at most $6-2j$ to $\mathfrak{d}\left(Y''\right) - \mathfrak{d}\left(Y'\right)$, so

$$\mathfrak{d}\left(Y''\right) - \mathfrak{d}\left(Y'\right) \leq 6k - 2\cdot\mathfrak{e}\left(P\right) \leq 6\cdot\mathfrak{he}\left(P\right) - 2\cdot\mathfrak{e}\left(P\right) < 2\left(\varepsilon\left|P\right| - 4\chi\left(P\right)\right) = 2\varepsilon\left|P\right|.$$

This proves Part 1.

There is a total of $8k$ directed edges at the boundaries of the new octagons. Of these, $\mathfrak{e}\left(P\right)$ are edges of P. Each of the remaining $8k - \mathfrak{e}\left(P\right)$ can 'host' two hanging half-edges of P, and each hanging half-edge appears in exactly 2 directed edges of new octagons. This gives

$$2\mathfrak{he}\left(P\right) \leq 2\left(8k - \mathfrak{e}\left(P\right)\right),$$

so $8k \geq \mathfrak{he}\left(P\right) + \mathfrak{e}\left(P\right) = \left|P\right|$. Hence

$$\theta\left(Y''\right) - \theta\left(Y'\right) = k - \left(\mathfrak{d}\left(Y''\right) - \mathfrak{d}\left(Y'\right)\right) > \frac{1}{8}\left|P\right| - 2\varepsilon\left|P\right| = \left(\frac{1}{8} - 2\varepsilon\right)\left|P\right|.$$

□

[12] For arbitrary $g \geq 2$, $\theta\left(Y'\right)$ increases by more than $\left(\frac{1}{4g} - 2\varepsilon\right)\left|P\left(Y'\right)\right|$, so we need $\varepsilon \leq \frac{1}{8g}$.

The upshot of the previous observations and Lemma 3.16 is that, provided $\varepsilon \leq \frac{1}{16}$, every time step **(a)** of the algorithm is reached, except for the first time, Y' has changed in step **(b)**, so $\theta(Y')$ has increased by at least one. Since

$$\theta(Y) = \mathfrak{f}(Y) - \mathfrak{d}(Y) \geq -\mathfrak{d}(Y),$$

and the algorithm halts at the latest after the first time that $\theta(Y)$ is positive, we deduce the following lemma:

LEMMA 3.17. *If $\varepsilon \leq \frac{1}{16}$, then during the octagons-vs-boundary algorithm, step **(a)** is reached at most $\mathfrak{d}(Y) + 2$ times. In particular, the algorithm always terminates.*

Now that we know the algorithm always terminates (assuming $\varepsilon \leq \frac{1}{16}$), and it clearly has deterministic output due to our a priori choices, if $Y \hookrightarrow Z$ is an embedding of a compact tiled surface Y into a tiled surface Z without boundary we write $\mathsf{OvB}_\varepsilon(Y \hookrightarrow Z)$ for the output of the OvB algorithm with parameter ε applied to $Y \hookrightarrow Z$. Thus $\mathsf{OvB}_\varepsilon(Y \hookrightarrow Z)$ is a tiled surface Y' with an attached embedding $Y \hookrightarrow Y'$. We can now make the following easy observation.

LEMMA 3.18. *Let $\varepsilon \leq \frac{1}{16}$, let $Y \hookrightarrow Z$ be an embedding of a compact tiled surface Y into a tiled surface Z without boundary, and let $Y' = \mathsf{OvB}_\varepsilon(Y \hookrightarrow Z)$. Then at least one of the following holds:*

- *Y' is ε-adapted.*
- *Y' is boundary reduced and $\mathfrak{f}(Y') > \mathfrak{d}(Y')$.*

We also want an upper bound on how $\mathfrak{d}(Y')$ and $\mathfrak{f}(Y')$ increase during the OvB algorithm.

LEMMA 3.19. *Assume*[13] *$\varepsilon \leq \frac{1}{32}$. Let Y be a compact tiled surface, Z be a boundaryless tiled surface and denote $\overline{Y} = \mathsf{OvB}_\varepsilon(Y \hookrightarrow Z)$. Then*

$$\mathfrak{d}(\overline{Y}) \leq 3\mathfrak{d}(Y), \tag{3.6}$$

$$\mathfrak{f}(\overline{Y}) \leq \mathfrak{f}(Y) + 4\mathfrak{d}(Y) + \mathfrak{d}(Y)^2. \tag{3.7}$$

Proof. If step **(a)** is only reached once, then the result of the algorithm, $\overline{Y}$, is equal to $\mathsf{BR}(Y \hookrightarrow Z)$. In this case we have $\mathfrak{d}(\overline{Y}) \leq \mathfrak{d}(Y)$ and $\mathfrak{f}(\overline{Y}) \leq \mathfrak{f}(Y) + \frac{\mathfrak{d}(Y)^2}{6}$ by Proposition 3.10 part 4, so the statement of the lemma holds. So from now on suppose step **(a)** is reached more than once.

Let $Y_1 = Y'$ at the penultimate time that step **(a)** is completed. Between the penultimate time that step **(a)** is completed and the algorithm terminates, step **(b)** takes place to form $Y_2 = Y'$, and then step **(a)** takes place one more time to form $Y_3 = \overline{Y}$ which is the output of the algorithm.

First we prove the bound on $\mathfrak{d}(Y_3)$. We have $\theta(Y_1) \leq 0$, so

$$\theta(Y_1) - \theta(Y) \leq 0 - (\mathfrak{f}(Y) - \mathfrak{d}(Y)) \leq \mathfrak{d}(Y).$$

[13] For arbitrary $g \geq 2$, we pick $\varepsilon \leq \frac{1}{16g}$. The statement of the lemma holds as is.

We claim that in every step of the OvB algorithm, the increase in θ is larger then the increase in $\mathfrak{d}$. Indeed, this is obviously true in step **(a)**, where θ does not decrease and $\mathfrak{d}$ does not increase. It is also true in step **(b)** by Lemma 3.16 and our assumption that $\varepsilon \leq \frac{1}{32}$. Therefore,

$$\mathfrak{d}(Y_1) - \mathfrak{d}(Y) \leq \theta(Y_1) - \theta(Y) \leq \mathfrak{d}(Y),$$

and we conclude that $\mathfrak{d}(Y_1) \leq 2\mathfrak{d}(Y)$.

Let $P = P(Y_1)$. By Lemma 3.16,

$$\begin{aligned}\mathfrak{d}(Y_2) &\leq \mathfrak{d}(Y_1) + 2\varepsilon|P| \overset{(3.4)}{\leq} \mathfrak{d}(Y_1) + 2\varepsilon \cdot \frac{4\mathfrak{d}(Y_1)}{3-\varepsilon} \\ &= \mathfrak{d}(Y_1)\left[1 + \frac{8\varepsilon}{3-\varepsilon}\right] \leq 1.1 \cdot \mathfrak{d}(Y_1) \leq 2.2 \cdot \mathfrak{d}(Y),\end{aligned}$$

where the penultimate inequality is based on that $\varepsilon \leq \frac{1}{32}$. Finally, $\mathfrak{d}(Y_3) \leq \mathfrak{d}(Y_2)$, so (3.6) is proven.

For the number of octagons, note first that

$$\mathfrak{f}(Y_1) = \theta(Y_1) + \mathfrak{d}(Y_1) \leq \mathfrak{d}(Y_1) \leq 2\mathfrak{d}(Y).$$

Let k denote the number of new octagons added in step **(b)** to form Y_2 from Y_1. As noted in the proof of Lemma 3.16, $k \leq \mathfrak{he}(P) + 1$. As $P = P(Y_1)$ is ε-bad, we have

$$\mathfrak{he}(P) \leq \frac{1}{3}(\mathfrak{e}(P) + \varepsilon|P|) \overset{(3.4)}{\leq} \frac{1}{3}\mathfrak{d}(Y_1)\left(1 + \frac{4\varepsilon}{3-\varepsilon}\right) < \mathfrak{d}(Y_1) \leq 2\mathfrak{d}(Y),$$

the penultimate inequality is based again on that $\varepsilon \leq \frac{1}{32}$. Thus $\mathfrak{f}(Y_2) - \mathfrak{f}(Y_1) \leq \mathfrak{he}(P) + 1 \leq 2\mathfrak{d}(Y)$.

Finally, by Proposition 3.10 part 4, $\mathfrak{f}(Y_3) - \mathfrak{f}(Y_2) \leq \frac{\mathfrak{d}(Y_2)^2}{6} \leq \mathfrak{d}(Y)^2$, and we conclude

$$\begin{aligned}\mathfrak{f}(Y_3) &= \mathfrak{f}(Y_1) + [\mathfrak{f}(Y_2) - \mathfrak{f}(Y_1)] + [\mathfrak{f}(Y_3) - \mathfrak{f}(Y_2)] \\ &\leq 2\mathfrak{d}(Y) + 2\mathfrak{d}(Y) + \mathfrak{d}(Y)^2 = 4\mathfrak{d}(Y) + \mathfrak{d}(Y)^2,\end{aligned}$$

which proves (3.7) in this case as well. □

3.6 Resolutions from the octagons-vs-boundary algorithm. Recall the definition of the tiled surface X_ϕ from Section 1 and Example 3.4. Given a tiled surface Y, we define

$$\mathbb{E}_n(Y) \overset{\text{def}}{=} \mathbb{E}_{\phi \in \mathbb{X}_n}[\#\text{morphisms } Y \to X_\phi].$$

This is the expected number of morphisms from Y to X_ϕ. Recall that we use the uniform probability measure on $\mathbb{X}_n$. We have the following result that relates this concept to Theorem 1.11.

Lemma 3.20. *Given $1 \neq \gamma \in \Gamma$, let $\mathcal{C}_\gamma$ be as in Example 3.5. Then*

$$\mathbb{E}_n[\mathsf{fix}_\gamma] = \mathbb{E}_n(\mathcal{C}_\gamma). \tag{3.8}$$

Proof. This is not hard to check but also follows from [MP20, Lem. 2.7]. □

We need to work not only with $\mathbb{E}_n(Y)$ for various tiled surfaces, but also with the expected number of times that Y *embeds* into X_ϕ. For a tiled surface Y, this is given by

$$\mathbb{E}_n^{\mathrm{emb}}(Y) \stackrel{\mathrm{def}}{=} \mathbb{E}_{\phi\in\mathbb{X}_n}[\#\text{embeddings } Y \hookrightarrow X_\phi].$$

We recall the following definition from [MP20, Def. 2.8].

Definition 3.21. *(Resolutions). A resolution $\mathcal{R}$ of a tiled surface Y is a collection of morphisms of tiled surfaces*

$$\mathcal{R} = \{f\colon Y \to W_f\},$$

such that every morphism $h\colon Y \to Z$ of Y into a tiled surface Z with no boundary decomposes <u>*uniquely*</u> *as $Y \xrightarrow{f} W_f \xhookrightarrow{\overline{h}} Z$, where $f \in \mathcal{R}$ and $\overline{h}$ is an* <u>*embedding*</u>*.*

The point of this definition is the following lemma also recorded in [MP20, Lem. 2.9].

Lemma 3.22. *If Y is a compact tiled surface and $\mathcal{R}$ is a finite resolution of Y, then*

$$\mathbb{E}_n(Y) = \sum_{f\in\mathcal{R}} \mathbb{E}_n^{\mathrm{emb}}(W_f). \tag{3.9}$$

The type of resolution we wish to use in this paper is the following.

Definition 3.23. *($\mathcal{R}_\varepsilon(Y)$). For a compact tiled surface Y, let $\mathcal{R}_\varepsilon(Y)$ denote the collection of all morphisms $Y \xrightarrow{f} W_f$ obtained as follows:*

- *$F : Y \to Z$ is a morphism of Y into a boundary-less tiled surface Z.*
- *U_F is the image of F in Z. Hence there is a given embedding $\iota_F : U_F \hookrightarrow Z$.*
- *W_f is given by $W_f = \mathsf{OvB}_\varepsilon(U_f \hookrightarrow Z)$ and $f = \iota_F \circ F\colon Y \to W_f$.*

Theorem 3.24. *Given a compact tiled surface Y amd $\varepsilon \leq \frac{1}{32}$ (or $\varepsilon \leq \frac{1}{16g}$ for arbitrary $g \geq 2$), the collection $\mathcal{R}_\varepsilon(Y)$ defined in Definition 3.23 is a finite resolution of Y.*

Proof. To see that $\mathcal{R}_\varepsilon(Y)$ is finite, note that there are finitely many options for U_F (this is a quotient of the compact complex Y). For any such U_F we have $\mathfrak{f}(U_F) \leq \mathfrak{f}(Y)$ and $\mathfrak{d}(U_F) \leq \mathfrak{d}(Y)$, and hence by Lemma 3.19 there is a bound on $\mathfrak{f}(W_f)$ depending only on Y. As we add a bounded number of octagons to obtain W_f, there is a bound also on $\mathfrak{v}(W_f)$ and on $\mathfrak{e}(W_f)$. This means that W_f is one of only finitely many tiled surfaces, and there are finitely many morphisms of Y to one of these.

Now we explain why $\mathcal{R}_\varepsilon(Y)$ is a resolution – this is essentially the same as [MP20, proof of Thm. 2.14]. Let $F\colon Y \to Z$ be a morphism with $\partial Z = \emptyset$. By the definition of $\mathcal{R}_\varepsilon(Y)$, it is clear that F decomposes as $Y \xrightarrow{f} W_f \hookrightarrow Z$ for the $f \in \mathcal{R}_\varepsilon(Y)$ that originates in F. To show uniqueness, assume that F decomposes in an additional way

$$Y \xrightarrow{f'} W_{f'} \hookrightarrow Z$$

where $W_{f'}$ is the result of the OvB algorithm for some $F'\colon Y \to Z'$ with $\partial Z' = \emptyset$. We claim that both decompositions are precisely the same decomposition of F (namely $W_{f'} = W_f$ and $f' = f$). First, $U_{F'} = F'(Y) \hookrightarrow W_{f'} \hookrightarrow Z$, so $U_{F'} = F'(Y) = F(Y) = U_F$. The OvB algorithm with input $F'(Y) \hookrightarrow Z'$ takes place entirely inside $W_{f'}$, and does not depend on the structure of $Z' \backslash W_{f'}$: the choices are made depending only on the structure of the boundary of Y' in step **(b)** of the OvB algorithm, as well as in every step of the procedure described in Proposition 3.10(3) to obtain $\mathsf{BR}(Y' \hookrightarrow Z)$ in step **(a)**. Moreover, the result of these steps depends only on the octagons of Z immediately adjacent to the boundary of Y'. But $W_{f'}$ is embedded in Z, and so it must be identical to W_f and f' identical to f. □

It is the following corollary of the previous results, applied to a tiled surface $\mathcal{C}_\gamma$ as in Example 3.5, that will be used in the rest of the paper. Recall that for $\gamma \in \Gamma$, $\ell_w(\gamma)$ denotes the word-length, with respect to the generators $\{a, b, c, d\}$, of a shortest representative of the conjugacy class of γ in Γ.

Corollary 3.25. *Let $1 \neq \gamma \in \Gamma$ and*[14] *$\varepsilon \leq \frac{1}{32}$. For any $f : \mathcal{C}_\gamma \to W_f$ in $\mathcal{R}_\varepsilon(\mathcal{C}_\gamma)$, either*

1. *W_f is boundary reduced, and $\chi(W_f) < -\mathfrak{f}(W_f) < -\mathfrak{d}(W_f)$, or*
2. *W_f is ε-adapted.*

Moreover, in either case,

$$\mathfrak{d}(W_f) \leq 6\ell_w(\gamma), \tag{3.10}$$

$$\mathfrak{f}(W_f) \leq 8\ell_w(\gamma) + 4(\ell_w(\gamma))^2. \tag{3.11}$$

Proof. The inequalities (3.10) and (3.11) are from Lemma 3.19 and the fact that $\mathfrak{d}(\mathcal{C}_\gamma) = 2\ell_w(\gamma)$ and $\mathfrak{f}(\mathcal{C}_\gamma) = 0$. It follows from the construction of $\mathcal{R}_\varepsilon(\mathcal{C}_\gamma)$ using the OvB algorithm that if $f \in \mathcal{R}_\varepsilon(Y)$ with $f : Y \to W_f$, and W_f is not ε-adapted, then W_f is boundary reduced and $\mathfrak{d}(W_f) < \mathfrak{f}(W_f)$. Combined with Lemma 3.7 this gives

$$\chi(W_f) \leq -2\mathfrak{f}(W_f) + \frac{1}{2}\mathfrak{d}(W_f) < -2\mathfrak{f}(W_f) + \frac{1}{2}\mathfrak{f}(W_f) \leq -\mathfrak{f}(W_f).$$ □

[14] For arbitrary $g \geq 2$, take $\varepsilon \leq \frac{1}{16g}$. The same result holds.

4 Representation Theory of Symmetric Groups

4.1 Background. We write S_n for the symmetric group of permutations of the set $[n]$. By convention S_0 is the trivial group with one element. If $m \leq n$, we always let $S_m \leq S_n$ be the subgroup of permutations fixing $[m+1, n]$ element-wise. For $k \leq n$, we will let $S'_k \leq S_n$ be our notation for the subgroup of permutations fixing $[n-k]$ element-wise. We write $\mathbf{C}[S_n]$ for the group algebra of S_n with complex coefficients.

Young diagrams. A Young diagram (YD) of size n is a collection of n boxes, arranged in left-aligned rows in the plane, such that the number of boxes in each row is non-increasing from top to bottom. A Young diagram is uniquely specified by the sequence $\lambda_1, \lambda_2, \ldots, \lambda_r$ where λ_i is the number of boxes in the ith row (and there are r rows). We have $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_r > 0$; such a sequence of integers is called a *partition.* We view YDs and partitions interchangeably in this paper. If $\sum_i \lambda_i = n$ we write $\lambda \vdash n$. Two important examples of partitions are (n), with all boxes of the corresponding YD in the first row, and $(1)^n \stackrel{\text{def}}{=} (\underbrace{1, \ldots, 1}_{n})$, with all boxes of the corresponding YD in the first column. If μ, λ are YDs, we write $\mu \subset \lambda$ if all boxes of μ are contained in λ (when both are aligned to the same top-left borders). We say $\mu \subset_k \lambda$ if $\mu \subset \lambda$ and there are k boxes of λ that are not in μ. We write $\emptyset$ for the empty YD with no boxes. If λ is a YD, $\check{\lambda}$ is the conjugate YD obtained by reflecting λ in the diagonal (switching rows and columns).

A skew Young diagram (SYD) is a pair of Young diagrams μ and λ with $\mu \subset \lambda$. This pair is denoted λ/μ and thought of as the collection of boxes of λ that are not in μ. We identify a YD λ with the SYD $\lambda/\emptyset$ so that YDs are special cases of SYDs. The size of a SYD λ/μ is the number of boxes it contains; i.e. the number of boxes of λ that are not in μ. The size is denoted by $|\lambda/\mu|$, or if λ is a YD, $|\lambda|$.

Young tableaux. Let λ/μ be a SYD, with $\lambda \vdash n$ and $\mu \vdash k$. A standard Young tableau of shape λ/μ is a filling of the boxes of λ/μ with the numbers $[k+1, n]$ such that each number appears in exactly one box and the numbers in each row (resp. column) are strictly increasing from left to right (resp. top to bottom). We refer to standard Young tableaux just as *tableaux* in this paper. We write $\mathrm{Tab}(\lambda/\mu)$ for the collection of tableaux of shape λ/μ. Given a tableau T, we denote by $T|_{\leq m}$ (resp. $T|_{>m}$) the tableau formed by the numbers-in-boxes of T with numbers in the set $[m]$ (resp. $[m+1, n]$). The shape of $T|_{\leq m}$ and of $T|_{>m}$ is a SYD in general. If T is a tableau and the shape of T is a YD we let $\mu_m(T)$ be the YD that is the shape of $T|_{\leq m}$. If $\nu \subset \mu \subset \lambda$, $T \in \mathrm{Tab}(\mu/\nu)$ and $R \in \mathrm{Tab}(\lambda/\mu)$, then we write $T \sqcup R$ for the tableau in $\mathrm{Tab}(\lambda/\nu)$ obtained by adjoining R to T in the obvious way.

Irreducible representations. The equivalence classes of irreducible unitary representations of S_n are in one-to-one correspondence with Young diagrams of size n. Given a YD $\lambda \vdash n$, we write V^λ for the corresponding irreducible representation of S_n; each V^λ is a finite dimensional Hermitian complex vector space with an action of S_n by

unitary linear automorphisms. Hence V^λ can also be thought of as a module for $\mathbf{C}[S_n]$. We write $d_\lambda \stackrel{\text{def}}{=} \dim V^\lambda$. It is well-known, and also follows from the discussion of the next paragraphs, that $d_\lambda = |\mathrm{Tab}(\lambda)|$. Note that $d_\lambda = d_{\check{\lambda}}$ since reflection in the diagonal gives a bijection between $\mathrm{Tab}(\lambda)$ and $\mathrm{Tab}(\check{\lambda})$.

We now give an account of the Vershik–Okounkov approach to the representation theory of symmetric groups from [VO04]. According to the usual ordering of $[n]$ there is a filtration of subgroups

$$S_0 \leq S_1 \leq S_2 \leq \cdots \leq S_n.$$

If W is any unitary representation of S_n, $m \in [n]$ and $\mu \vdash m$, we write W_μ for the span of vectors in copies of V^μ in the restriction of W to S_m; we call W_μ the μ-isotypic subspace of W.

It follows from the branching law for restriction of representations between S_m and S_{m-1} that for $\lambda \vdash n$ and $T \in \mathrm{Tab}(\lambda)$ the intersection

$$\left(V^\lambda\right)_{\mu_1(T)} \cap \left(V^\lambda\right)_{\mu_2(T)} \cap \cdots \cap \left(V^\lambda\right)_{\mu_{n-1}(T)}$$

is one-dimensional. Vershik–Okounkov specify a unit vector v_T in this intersection. The collection

$$\{\, v_T \,:\, T \in \mathrm{Tab}(\lambda) \,\}$$

is an orthonormal basis for V^λ called a Gelfand-Tsetlin basis.

Modules from SYDs. If $m, n \in \mathbf{N}$, $\lambda \vdash n$, $\mu \vdash m$ and $\mu \subset \lambda$, then

$$V^{\lambda/\mu} \stackrel{\text{def}}{=} \mathrm{Hom}_{S_m}(V^\mu, V^\lambda)$$

is a unitary representation of S'_{n-m} as S'_{n-m} is in the centralizer of S_m in S_n. We write $d_{\lambda/\mu}$ for the dimension of this representation. There is also an analogous Gelfand-Tsetlin orthonormal basis of $V^{\lambda/\mu}$ indexed by $T \in \mathrm{Tab}(\lambda/\mu)$; the basis element corresponding to a skew tableau T will be denoted w_T. It follows that $d_{\lambda/\mu} = |\mathrm{Tab}(\lambda/\mu)|$. Note that when $\mu = \lambda$, $\mathrm{Tab}\,(\lambda/\mu) = \{\emptyset\}$ ($\emptyset$ the empty tableau), and the representation $V^{\lambda/\mu}$ is one-dimensional with basis $w_\emptyset$.

One has the following consequence of Frobenius reciprocity (cf. e.g. [MP20, Lem. 3.1]).

Lemma 4.1. *Let $n \in \mathbf{N}$, $m \in [n]$ and $\mu \vdash m$. Then*

$$\sum_{\lambda \vdash n\colon \mu \subset \lambda} d_{\lambda/\mu} d_\lambda = \frac{n!}{m!} d_\mu.$$

4.2 Effective bounds for dimensions. Throughout the paper, we will write b_λ for the number of boxes outside the first row of a YD λ, and write $\check{b}_\lambda$ for the number of boxes outside the first column of λ. More generally, we write $b_{\lambda/\nu}$ (resp. $\check{b}_{\lambda/\nu}$) for the number of boxes outside the first row (resp. column) of the SYD λ/ν, so $b_{\lambda/\nu} = b_\lambda - b_\nu$ and $\check{b}_{\lambda/\nu} = \check{b}_\lambda - \check{b}_\nu$. We need the following bounds on dimensions of representations.

LEMMA 4.2. [MP20, Lem. 4.3]. *If $n \in \mathbf{N}$, $m \in [n]$, $\lambda \vdash n$, $\nu \vdash m$, $\nu \subset \lambda$ and $m \geq 2b_\lambda$, then*

$$\frac{(n-b_\lambda)^{b_\lambda}}{b_\lambda^{b_\lambda} m^{b_\nu}} \leq \frac{d_\lambda}{d_\nu} \leq \frac{b_\nu^{b_\nu} n^{b_\lambda}}{(m-b_\nu)^{b_\nu}}. \tag{4.1}$$

The condition $m \geq 2b_\lambda$ ensures that both ν and λ have most boxes in their first row. This is an important and recurring theme of the paper (see e.g. Proposition 4.6).

LEMMA 4.3. *Let λ/ν be a skew Young diagram of size n. Then*

$$d_{\lambda/\nu} \leq (n)_{b_{\lambda/\nu}} \quad \text{and} \quad d_{\lambda/\nu} \leq (n)_{\check{b}_{\lambda/\nu}}.$$

Proof. There are at most $\binom{n}{b_{\lambda/\nu}}$ options for the set of $b_{\lambda/\nu}$ elements outside the first row. Given these, there are at most $b_{\lambda/\nu}!$ choices for how to place them outside the first row. The proof of the second inequality is analogous. □

4.3 Effective bounds for the zeta function of the symmetric group. The Witten zeta function of the symmetric group S_n is defined for a real parameter s as

$$\zeta^{S_n}(s) \stackrel{\text{def}}{=} \sum_{\lambda \vdash n} \frac{1}{d_\lambda^{\,s}}. \tag{4.2}$$

This function, and various closely related functions, play a major role in this paper. One main reason for its appearance is due to a formula going back to Hurwitz [Hur02] that states

$$|\mathbb{X}_{g,n}| = |\mathrm{Hom}(\Gamma_g, S_n)| = |S_n|^{2g-1}\zeta^{S_n}(2g-2). \tag{4.3}$$

This is also sometimes called Mednykh's formula [Med78]. We first give the following result due to Liebeck and Shalev [LS04, Thm. 1.1] and independently, Gamburd [Gam06, Prop. 4.2]. We refer the reader to Section 1.4 for the definition of notations (e.g. $O, \ll$) that we use in this Section 4.

Theorem 4.4. [LS04, Gam06]. *For any $s > 0$, as $n \to \infty$*

$$\zeta^{S_n}(s) = 2 + O\left(n^{-s}\right).$$

This has the following corollary when combined with (4.3).

Corollary 4.5. *For any $g \in \mathbf{N}$ with $g \geq 2$, we have*

$$\frac{|\mathbb{X}_{g,n}|}{(n!)^{2g-1}} = 2 + O(n^{-2}).$$

As well as the previous results, we also need to know how well $\zeta^{S_n}(2g-2)$ is approximated by restricting the summation in (4.2) to λ with a bounded number of boxes either outside the first row or outside the first column. We let $\Lambda(n,b)$ denote the collection of $\lambda \vdash n$ such that $\lambda_1 \leq n-b$ and $\check{\lambda}_1 \leq n-b$. In other words, $\Lambda(n,b)$ is the collection of YDs $\lambda \vdash n$ with both $b_\lambda \geq b$ and $\check{b}_\lambda \geq b$. A version of the next proposition, when b is fixed and $n \to \infty$, is due independently to Liebeck and Shalev [LS04, Prop. 2.5] and Gamburd [Gam06, Prop. 4.2]. Here, we need a version that holds uniformly over b that is not too large compared to n.

Proposition 4.6. *Fix $s > 0$. There exists a constant $\kappa = \kappa(s) > 1$ such that when $b^2 \leq \frac{n}{3}$,*

$$\sum_{\lambda \in \Lambda(n,b)} \frac{1}{d_\lambda{}^s} \ll_s \left(\frac{\kappa b^{2s}}{(n-b^2)^s} \right)^b. \tag{4.4}$$

Proof. Here we follow Liebeck and Shalev [LS04, proof of Prop. 2.5] and make the proof uniform over b. Let $\Lambda_0(n,b)$ denote the collection of $\lambda \vdash n$ with $\check{\lambda}_1 \leq \lambda_1 \leq n-b$. Since $d_\lambda = d_{\check{\lambda}}$,

$$\sum_{\lambda \in \Lambda(n,b)} \frac{1}{d_\lambda{}^s} \leq 2 \sum_{\lambda \in \Lambda_0(n,b)} \frac{1}{d_\lambda{}^s},$$

so it suffices to prove a bound for $\sum_{\lambda \in \Lambda_0(n,b)} \frac{1}{d_\lambda{}^s}$. Let $\Lambda_1(n,b)$ denote the elements λ of $\Lambda_0(n,b)$ with $\lambda_1 \geq \frac{2n}{3}$. We write

$$\sum_{\lambda \in \Lambda_0(n,b)} \frac{1}{d_\lambda{}^s} = \Sigma_1 + \Sigma_2$$

where

$$\Sigma_1 \stackrel{\text{def}}{=} \sum_{\lambda \in \Lambda_1(n,b)} \frac{1}{d_\lambda{}^s}, \quad \Sigma_2 \stackrel{\text{def}}{=} \sum_{\lambda \in \Lambda_0(n,b) - \Lambda_1(n,b)} \frac{1}{d_\lambda{}^s}.$$

Bound for Σ_1. By [LS04, Lem. 2.1] if $\lambda \in \Lambda_1(n,b)$ then since $\lambda_1 \geq \frac{n}{2}$, $d_\lambda \geq \binom{\lambda_1}{n-\lambda_1}$. Indeed, for completeness, following [LS04, Proof of Lem. 2.1] we can find many tableaux of shape λ as follows. Put the numbers $1, \ldots, n-\lambda_1$ in the left most entries of the first row of λ. Then for any of the $\binom{\lambda_1}{n-\lambda_1}$ choices of size $n-\lambda_1$ subsets of $[n-\lambda_1+1, n]$, there is obviously a tableau of shape λ with those numbers outside the first row.

Let $p(m)$ denote the number of $\mu \vdash m$. The number of $\lambda \in \Lambda_1(n,b)$ with a valid fixed value of λ_1 is $p(n-\lambda_1)$ since $\lambda_1 \geq \frac{n}{2}$ and hence any YD with $n-\lambda_1$ boxes can be added below the fixed first row of λ_1 boxes to form λ. Therefore

$$\Sigma_1 \leq \sum_{\lambda_1=\lceil \frac{2n}{3} \rceil}^{n-b} \frac{p\left(n-\lambda_1\right)}{\binom{\lambda_1}{n-\lambda_1}^s} = \sum_{\ell=b}^{\lfloor \frac{n}{3} \rfloor} \frac{p(\ell)}{\binom{n-\ell}{\ell}^s}.$$

We now split the sum into two ranges to estimate $\Sigma_1 \leq \Sigma_1' + \Sigma_1''$ where

$$\Sigma_1' = \sum_{\ell=b}^{b^2} \frac{p(\ell)}{\binom{n-\ell}{\ell}^s}, \quad \Sigma_1'' = \sum_{\ell=b^2+1}^{\lfloor \frac{n}{3} \rfloor} \frac{p(\ell)}{\binom{n-\ell}{\ell}^s}.$$

First we deal with Σ_1'. We have $p(\ell) \leq c_1^{\sqrt{\ell}}$ for some $c_1 > 1$ [Apo76, Thm. 14.5]. As $\ell \leq n-\ell$,

$$\binom{n-\ell}{\ell} \geq \frac{(n-\ell)^\ell}{\ell^\ell}.$$

This gives

$$\Sigma_1' \leq \sum_{\ell=b}^{b^2} c_1^{\sqrt{\ell}} \left(\frac{\ell}{n-\ell}\right)^{s\ell} \leq {c_1}^b \sum_{\ell=b}^{b^2} \left(\frac{b^2}{n-b^2}\right)^{s\ell} \ll_s {c_1}^b \left(\frac{b^2}{n-b^2}\right)^{sb}, \tag{4.5}$$

where the last inequality used that $\frac{b^2}{(n-b^2)} \leq \frac{1}{2}$ as we assume $b^2 \leq \frac{n}{3}$.

To deal with Σ_1'' we make the following claim.

Claim. *There is $n_{00} > 0$ such that when $n \geq n_{00}$ and $\ell \leq \frac{n}{3}$*

$$\binom{n-\ell}{\ell} \geq \left(\frac{2n}{3}\right)^{\sqrt{\ell}}. \tag{4.6}$$

Proof of claim. Observe that when $\ell \leq \frac{n}{3}$

$$\begin{aligned}\binom{n-\ell}{\ell} &\geq \frac{(n-\ell)^\ell}{\ell^\ell} = (n-\ell)^{\sqrt{\ell}} (n-\ell)^{\ell-\sqrt{\ell}} \ell^{-\ell} \\ &\geq \left(\frac{2n}{3}\right)^{\sqrt{\ell}} (2\ell)^{\ell-\sqrt{\ell}} \ell^{-\ell} = \left(\frac{2n}{3}\right)^{\sqrt{\ell}} \left(\frac{2^{\sqrt{\ell}-1}}{\ell}\right)^{\sqrt{\ell}}.\end{aligned}$$

We have $2^{\sqrt{\ell}-1} \geq \ell$ when $\ell \geq 49$ which proves the claim in this case. On the other hand, it is easy to see that there is a $n_{00} > 0$ such that (4.6) holds when $n \geq n_{00}$ and $1 \leq \ell < 49$. *This proves the claim.* □

The claim gives

$$\Sigma_1'' \leq \sum_{\ell=b^2+1}^{\lfloor \frac{n}{3} \rfloor} \left(\frac{c_2}{n^s}\right)^{\sqrt{\ell}}$$

for some $c_2 = c_2(s) > 1$ when $n \geq n_{00}$. Let $n_0 = n_0(s) \geq n_{00}$ be such that when $n \geq n_0$, $\frac{c_2}{n^s} < e^{-1}$. Let $q = q\,(n) \stackrel{\text{def}}{=} \frac{c_2}{n^s}$. Then when $n \geq n_0$, $\log(q) \leq -1$ and

$$\Sigma_1'' \leq \int_{b^2}^{\infty} q^{\sqrt{x}} dx = \frac{2q^b}{\log q}\left(\frac{1}{\log q} - b\right).$$

We obtain

$$\Sigma_1'' \leq 2(b+1)q^b \leq \frac{2(b+1){c_2}^b}{n^{sb}}. \tag{4.7}$$

Together with (4.5) this yields:

$$\Sigma_1 \ll_s {c_1}^b \left(\frac{b^2}{n-b^2}\right)^{sb} + \frac{2(b+1){c_2}^b}{n^{sb}} \ll_s \left(\frac{\kappa b^{2s}}{(n-b^2)^s}\right)^b \tag{4.8}$$

with $\kappa = \kappa\,(s) = \max\,(c_1, c_2)$.

Bound for Σ_2. If $\lambda \in \Lambda_0(n,b) - \Lambda_1(n,b)$ then $\check{\lambda}_1 \leq \lambda_1 < \frac{2n}{3}$ and [LS04, Prop. 2.4] gives the existence of an absolute $c_0 > 1$ such that

$$d_\lambda \geq {c_0}^n.$$

Thus for large enough n and $b^2 \leq \frac{n}{3}$

$$\Sigma_2 \leq \sum_{\lambda \in \Lambda_0(n,b) - \Lambda_1(n,b)} c_0^{-ns} \leq p(n) c_0^{-ns} \leq c_1^{\sqrt{n}} c_0^{-ns} \ll_s n^{-bs}. \tag{4.9}$$

Putting (4.8) and (4.9) together proves the proposition. □

5 Estimates for the Probabilities of Tiled Surfaces

Before reading this Section 5, we recommend the reader to have read Section 1.3 for context and motivation.

5.1 Prior results. The aim of this Section 5.1 is to introduce an already known formula (Theorem 5.1) for the quantities $\mathbb{E}_n^{\mathrm{emb}}(Y)$ that are essential to this paper, and to give some known first estimates for the quantities appearing therein (Lemma 5.2). To better understand their source and logic, the reader is advised to look at [MP20, Section 5].

We continue to assume $g = 2$. Throughout this entire Section 5 we will assume that Y is a fixed compact tiled surface. We let $\mathfrak{v} = \mathfrak{v}(Y)$, $\mathfrak{e} = \mathfrak{e}(Y)$, $\mathfrak{f} = \mathfrak{f}(Y)$ denote the number of vertices, edges, and octagons of Y, respectively. Throughout this section, f will stand for one of the letters a, b, c, d. For each letter $f \in \{a, b, c, d\}$, let $\mathfrak{e}_f$ denote the number of f-labeled edges of Y.

In [MP20, Section 5.3] we constructed permutations

$$\sigma_f^+, \sigma_f^-, \tau_f^+, \tau_f^- \in S'_{\mathfrak{v}} \subset S_n$$

for each $f \in \{a, b, c, d\}$ satisfying certain five properties named **P1**, **P2**, **P3**, **P4**, and **P5** that are essential to the development of the theory, but not illuminating to state precisely here. We henceforth view these permutations as fixed, given Y.

Recall from Section 4.1 that for YDs $\mu \subset \lambda$ we say $\mu \subset_k \lambda$ if λ has k more boxes than μ. Also recall from Section 4.1 that $\sqcup$ denotes concatenation of Young tableaux, and for a SYD λ/ν, if T is a (standard) tableau of shape λ/ν, w_T denotes a Gelfand-Tsetlin basis vector in $V^{\lambda/\nu}$ associated to T. In the same situation, we write $\langle \bullet, \bullet \rangle$ for the inner product in the unitary representation $V^{\lambda/\nu}$. In the prequel paper [MP20, Thm. 5.10] the following theorem was proved.

Theorem 5.1. *For $n \geq \mathfrak{v}$ we have*

$$\mathbb{E}_n^{\mathrm{emb}}(Y) = \frac{(n!)^3}{|\mathbb{X}_n|} \cdot \frac{(n)_{\mathfrak{v}}\,(n)_{\mathfrak{f}}}{\prod_f (n)_{\mathfrak{e}_f}} \cdot \Xi_n(Y) \tag{5.1}$$

where

$$\Xi_n(Y) \stackrel{\mathrm{def}}{=} \sum_{\substack{\lambda,\nu:\\ \nu \subset_{\mathfrak{v}-\mathfrak{f}} \lambda \vdash n-\mathfrak{f}}} d_\lambda d_\nu \sum_{\substack{\mu_a,\mu_b,\mu_c,\mu_d\\ \forall f,\, \nu \subset \mu_f \subset_{\mathfrak{e}_f - \mathfrak{f}} \lambda}} \frac{1}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \Upsilon_n\left(\left\{\sigma_f^\pm, \tau_f^\pm\right\}, \nu, \{\mu_f\}, \lambda\right), \tag{5.2}$$

$$\Upsilon_n\left(\left\{\sigma_f^\pm, \tau_f^\pm\right\}, \nu, \{\mu_f\}, \lambda\right) \stackrel{\mathrm{def}}{=} \sum_{\substack{r_f^+, r_f^- \in \mathrm{Tab}(\mu_f/\nu)\\ s_f, t_f \in \mathrm{Tab}(\lambda/\mu_f)}} \mathcal{M}\left(\left\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\right\}\right) \tag{5.3}$$

and $\mathcal{M}(\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\})$ is the following product of matrix coefficients:

$$\begin{aligned}\mathcal{M}\left(\left\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\right\}\right) \stackrel{\mathrm{def}}{=} &\left\langle \sigma_b^- \left(\sigma_a^+\right)^{-1} w_{r_a^+ \sqcup s_a}, w_{r_b^- \sqcup s_b} \right\rangle \left\langle \tau_a^+ \left(\sigma_b^+\right)^{-1} w_{r_b^+ \sqcup s_b}, w_{r_a^+ \sqcup t_a} \right\rangle \\ &\cdot \left\langle \tau_b^+ \left(\tau_a^-\right)^{-1} w_{r_a^- \sqcup t_a}, w_{r_b^+ \sqcup t_b} \right\rangle \left\langle \sigma_c^- \left(\tau_b^-\right)^{-1} w_{r_b^- \sqcup t_b}, w_{r_c^- \sqcup s_c} \right\rangle\end{aligned}$$

$$\cdot \left\langle \sigma_d^- \left(\sigma_c^+\right)^{-1} w_{r_c^+ \sqcup s_c}, w_{r_d^- \sqcup s_d} \right\rangle \left\langle \tau_c^+ \left(\sigma_d^+\right)^{-1} w_{r_d^+ \sqcup s_d}, w_{r_c^+ \sqcup t_c} \right\rangle$$
$$\cdot \left\langle \tau_d^+ \left(\tau_c^-\right)^{-1} w_{r_c^- \sqcup t_c}, w_{r_d^+ \sqcup t_d} \right\rangle \left\langle \sigma_a^- \left(\tau_d^-\right)^{-1} w_{r_d^- \sqcup t_d}, w_{r_a^- \sqcup s_a} \right\rangle. \tag{5.4}$$

Note that $\frac{(n!)^3}{|\mathbb{X}_n|} \overset{n\to\infty}{\to} 2$ by (4.3) and Theorem 4.4, and that $\frac{(n)_{\mathfrak{v}}(n)_{\mathfrak{f}}}{\prod_f (n)_{\mathfrak{e}_f}} = n^{\chi(Y)} \left(1 + O\left(n^{-1}\right)\right)$, so the more mysterious term in (5.1) is $\Xi_n(Y)$. In light of Theorem 5.1, we will repeatedly discuss $\nu, \{\mu_f\}, \lambda$ satisfying

$$\nu \subset_{\mathfrak{v}-\mathfrak{e}_f} \mu_f \subset_{\mathfrak{e}_f-\mathfrak{f}} \lambda \vdash n - \mathfrak{f} \quad \forall f \in \{a,b,c,d\} \tag{5.5}$$

and $\{r_f^\pm, s_f, t_f\}$ satisfying

$$r_f^+, r_f^- \in \mathrm{Tab}(\mu_f/\nu), \quad s_f, t_f \in \mathrm{Tab}(\lambda/\mu_f) \quad \forall f \in \{a,b,c,d\}. \tag{5.6}$$

To give good estimates for $\Xi_n(Y)$, we need an effective bound for the quantities $\mathcal{M}(\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\})$ that was obtained in [MP20]. Before giving this bound, we recall some notation. For $T \in \mathrm{Tab}(\lambda/\nu)$, we write $\mathrm{top}(T)$ for the set of elements in the top row of T (the row of length $\lambda_1 - \nu_1$ which may be empty). For any two sets A, B in $[n]$, we define $d(A,B) = |A \backslash B|$. Given $\{r_f^\pm, s_f, t_f\}$ as in (5.6), we define

$$\begin{aligned}
&D_{\mathrm{top}}\left(\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\}\right) \\
&\overset{\mathrm{def}}{=} d\left(\sigma_b^- \left(\sigma_a^+\right)^{-1} \mathrm{top}(r_a^+ \sqcup s_a), \mathrm{top}(r_b^- \sqcup s_b)\right) + d\left(\tau_a^+ \left(\sigma_b^+\right)^{-1} \mathrm{top}(r_b^+ \sqcup s_b), \mathrm{top}(r_a^+ \sqcup t_a)\right) \\
&\quad + d\left(\tau_b^+ \left(\tau_a^-\right)^{-1} \mathrm{top}(r_a^- \sqcup t_a), \mathrm{top}(r_b^+ \sqcup t_b)\right) + d\left(\sigma_c^- \left(\tau_b^-\right)^{-1} \mathrm{top}(r_b^- \sqcup t_b), \mathrm{top}(r_c^- \sqcup s_c)\right) \\
&\quad + d\left(\sigma_d^- \left(\sigma_c^+\right)^{-1} \mathrm{top}(r_c^+ \sqcup s_c), \mathrm{top}(r_d^- \sqcup s_d)\right) + d\left(\tau_c^+ \left(\sigma_d^+\right)^{-1} \mathrm{top}(r_d^+ \sqcup s_d), \mathrm{top}(r_c^+ \sqcup t_c)\right) \\
&\quad + d\left(\tau_d^+ \left(\tau_c^-\right)^{-1} \mathrm{top}(r_c^- \sqcup t_c), \mathrm{top}(r_d^+ \sqcup t_d)\right) + d\left(\sigma_a^- \left(\tau_d^-\right)^{-1} \mathrm{top}(r_d^- \sqcup t_d), \mathrm{top}(r_a^- \sqcup s_a)\right).
\end{aligned} \tag{5.7}$$

Lemma 5.2. [MP20, Lem. 5.14]. *Let $\nu, \{\mu_f\}, \lambda$ be as in (5.5) and $\{r_f^\pm, s_f, t_f\}$ be as in (5.6). If $\lambda_1 + \nu_1 > n - \mathfrak{f} + (\mathfrak{v} - \mathfrak{f})^2$, then*

$$\left|\mathcal{M}\left(\left\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\right\}\right)\right| \leq \left(\frac{(\mathfrak{v}-\mathfrak{f})^2}{\lambda_1 + \nu_1 - (n - \mathfrak{f})}\right)^{D_{\mathrm{top}}\left(\left\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\right\}\right)}.$$

The condition $\lambda_1 + \nu_1 > n - \mathfrak{f} + (\mathfrak{v} - \mathfrak{f})^2$ corresponds to the bound given by Lemma 5.2 being non-trivial, and we will be applying Lemma 5.2 when both λ and ν have $O(n^{1/4})$ boxes outside their first rows and $\mathfrak{v}, \mathfrak{f} \ll n^{1/4}$. In particular, $\lambda_1 + \nu_1$ is of order $2n$, while $\mathfrak{f}$ and $(\mathfrak{v} - \mathfrak{f})^2$ are of much smaller order. Hence the condition will be met for sufficiently large n (Fig. 5).

Recall from Section 4.2 that b_ν is the number of boxes of a Young diagram ν outside the first row, and $\check{b}_\nu$ is the number of boxes outside the first column. We have the following trivial upper bound for $D_{\mathrm{top}}(\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\})$:

$$D_{\mathrm{top}}\left(\left\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\right\}\right) \leq 8\left(b_\lambda - b_\nu\right). \tag{5.8}$$

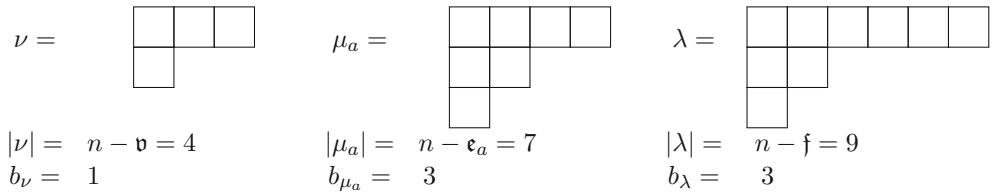

Figure 5: An example of possible ν, μ_a, λ appearing in Theorem 5.1 (supposing e.g. $n = 10, \mathfrak{v} = 6, \mathfrak{e}_a = 3, \mathfrak{f} = 1$).

We recall the following estimate obtained in [MP20, Prop. 5.22].

PROPOSITION 5.3. *Let $\varepsilon \geq 0$. Suppose that $\nu, \{\mu_f\}, \lambda$ are as in (5.5) and $\{r_f^\pm, s_f, t_f\}$ are as in (5.6). If Y is ε-adapted then*

$$D_{\text{top}}\left(\left\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\right\}\right) \geq b_\lambda + 3b_\nu - b_{\mu_a} - b_{\mu_b} - b_{\mu_c} - b_{\mu_d} + \varepsilon b_{\lambda/\nu}. \tag{5.9}$$

5.2 Partitioning Ξ_n and preliminary estimates. In this Section 5.2 we show how the condition that Y is ε-adapted leads to bounds on Ξ_n. We continue to view Y as fixed and hence suppress dependence of quantities on Y. We write $\mathfrak{D} = \mathfrak{D}(Y) \stackrel{\text{def}}{=} \mathfrak{v} - \mathfrak{f}$. Note that $\mathfrak{D} \geq 0$ by (3.2), with equality if and only if Y has no boundary. By Lemma 3.6, $\mathfrak{D} \leq \mathfrak{d}(Y)$. So $\mathfrak{D}$ is another measure of the size of the boundary of Y, and it plays an important role in some of our bounds below. We will use the notation $\Xi_n^{P(\nu)}$ where P is a proposition concerning ν to mean

$$\Xi_n^{P(\nu)} \stackrel{\text{def}}{=} \sum_{\substack{\nu \subset_{\mathfrak{v}-\mathfrak{f}} \lambda \vdash n-\mathfrak{f} \\ P(\nu) \text{ holds true}}} d_\lambda d_\nu \sum_{\nu \subset \mu_f \subset_{\mathfrak{e}_f - \mathfrak{f}} \lambda} \frac{1}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \Upsilon_n\left(\left\{\sigma_f^\pm, \tau_f^\pm\right\}, \nu, \{\mu_f\}, \lambda\right).$$

We will continue to use this notation, for various propositions P, throughout the rest of the paper. We want to give bounds for various $\Xi_n^{P(\nu)}$ under the condition that Y is either boundary reduced (namely, 0-adapted) or, moreover, ε-adapted for some $\varepsilon > 0$. We will always assume $\mathfrak{v} \leq n^{1/4}$ and so also $\mathfrak{D} = \mathfrak{v} - \mathfrak{f} \leq n^{1/4}$. Note that $b_\nu \leq \mathfrak{D}$ and $\check{b}_\nu \leq \mathfrak{D}$ cannot hold simultaneously as $\mathfrak{v}, \mathfrak{D} \leq n^{1/4}$, and as all but one box of $\nu \vdash n - \mathfrak{v}$ is either outside the first row or first column, one has the simple inequality $b_\nu + \check{b}_\nu + 1 \geq n - \mathfrak{v}$.

Then for $n \gg 1$ we have

$$\Xi_n = \Xi_n^{\nu=(n-\mathfrak{v})} + \Xi_n^{\nu=(1)^{n-\mathfrak{v}}} + \Xi_n^{0<b_\nu \leq \mathfrak{D}; \check{b}_\nu > 0} + \Xi_n^{0<\check{b}_\nu \leq \mathfrak{D}; b_\nu > 0} + \Xi_n^{b_\nu, \check{b}_\nu > \mathfrak{D}}.$$

Moreover by [MP20, Lem. 5.9] we have

$$\Xi_n^{\nu=(n-\mathfrak{v})} = \Xi_n^{\nu=(1)^{n-\mathfrak{v}}}, \quad \Xi_n^{0<b_\nu \leq \mathfrak{D}; \check{b}_\nu > 0} = \Xi_n^{0<\check{b}_\nu \leq \mathfrak{D}; b_\nu > 0}$$

hence

$$\Xi_n = 2\Xi_n^{\nu=(n-\mathfrak{v})} + 2\Xi_n^{0<b_\nu \leq \mathfrak{D}; \check{b}_\nu > 0} + \Xi_n^{b_\nu, \check{b}_\nu > \mathfrak{D}}. \tag{5.10}$$

This is according to three regimes for b_ν and $\check{b}_\nu$:

- The *zero regime*: when b_ν or $\check{b}_\nu$ equal 0. The contribution from here is $2\Xi_n^{\nu=(n-\nu)}$.
- The *intermediate regime:* when $b_\nu, \check{b}_\nu > 0$ but one of them is at most $\mathfrak{D}$. The contribution from this regime is $2\Xi_n^{0<b_\nu\leq\mathfrak{D};\check{b}_\nu>0}$.
- The *large regime:* when both $b_\nu, \check{b}_\nu > \mathfrak{D}$. The contribution from this regime is $\Xi_n^{b_\nu,\check{b}_\nu>\mathfrak{D}}$.

The strategy for bounding these different contributions is to further partition the tuples $(\nu, \{\mu_f\}, \lambda)$ according to the data $b_\lambda, \{b_{\mu_f}\}, b_\nu, \check{b}_\lambda, \{\check{b}_{\mu_f}\}, \check{b}_\nu$.

Definition 5.4. *For* $\overline{B} = \left(B_\lambda, \{B_{\mu_f}\}, B_\nu, \check{B}_\lambda, \{\check{B}_{\mu_f}\}, \check{B}_\nu\right)$ *we write*

$$(\nu, \{\mu_f\}, \lambda) \vdash \overline{B}$$

if (5.5) holds, and ν, $\{\mu_f\}$ *and* λ *have the prescribed number of blocks outside the first row and outside the first column, namely,*

$$b_\lambda = B_\lambda, \check{b}_\lambda = \check{B}_\lambda,\ b_\nu = B_\nu, \check{b}_\nu = \check{B}_\nu \quad \text{and} \quad \forall f \in \{a,b,c,d\} \quad b_{\mu_f} = B_{\mu_f}, \check{b}_{\mu_f} = \check{B}_{\mu_f}.$$

We denote by $\mathcal{B}_n(Y)$ *the collection of tuples* $\overline{B}$ *which admit at least one tuple of YDs* $(\nu, \{\mu_f\}, \lambda)$. *Finally, we let*

$$\Xi_n^{\overline{B}} = \Xi_n^{\overline{B}}(Y) \stackrel{\text{def}}{=} \sum_{(\nu,\{\mu_f\},\lambda)\vdash\overline{B}} \frac{d_\lambda d_\nu}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \Upsilon_n\left(\left\{\sigma_f^\pm, \tau_f^\pm\right\}, \nu, \{\mu_f\}, \lambda\right). \tag{5.11}$$

Note that $\Xi_n(Y) = \sum_{\overline{B}\in\mathcal{B}_n(Y)} \Xi_n^{\overline{B}}$. Also, note that $\overline{B} \in \mathcal{B}_n(Y)$ imposes restrictions on the possible values of $B_\lambda, \{B_{\mu_f}\}, B_\nu, \check{B}_\lambda, \{\check{B}_{\mu_f}\}, \check{B}_\nu$. For example, for every $f \in \{a,b,c,d\}$, $0 \leq B_{\mu_f} - B_\nu \leq \mathfrak{v} - \mathfrak{e}_f$ and $0 \leq B_\lambda - B_{\mu_f} \leq \mathfrak{e}_f - \mathfrak{f}$, and likewise for the $\check{B}$'s. In addition, $B_\nu + \check{B}_\nu + 1 \geq n - \nu$, and so on.

We first give a general estimate for the quotient of dimensions in the summands in (5.11).

Lemma 5.5. *Suppose that* $\mathfrak{v} \leq n^{1/4}$ *and that* $(\nu, \{\mu_f\}, \lambda)$ *satisfy (5.5). If* $b_\nu \leq \mathfrak{D}$ *then*

$$\frac{d_\lambda d_\nu}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \ll \frac{1}{{d_\nu}^2} b_\lambda^{5b_\lambda} n^{\left(b_\lambda + 3b_\nu - \sum_f b_{\mu_f}\right)}. \tag{5.12}$$

Proof. By Lemma 4.2,

$$\frac{d_\nu}{d_{\mu_f}} \leq \frac{{b_{\mu_f}}^{b_{\mu_f}} (n-\mathfrak{v})^{b_\nu}}{\left(n - \mathfrak{e}_f - b_{\mu_f}\right)^{b_{\mu_f}}} \leq \frac{{b_\lambda}^{b_\lambda} n^{b_\nu}}{\left(n - 2n^{1/4}\right)^{b_{\mu_f}}},$$

where the second inequality is based on that $\mathfrak{e}_f + b_{\mu_f} \leq \mathfrak{e}_f + (b_\nu + \mathfrak{v} - \mathfrak{e}_f) = b_\nu + \mathfrak{v} \leq 2n^{1/4}$. The hypotheses of Lemma 4.2 are met here since

$$2b_{\mu_f} - |\nu| \leq 2(b_\nu + \mathfrak{v} - \mathfrak{e}_f) - (n - \mathfrak{v}) \leq 5\mathfrak{v} - n \leq 5n^{\frac{1}{4}} - n \leq 0$$

for $n \gg 1$. Similarly, since $2b_\lambda - |\nu| \le 2(b_\nu + \mathfrak{D}) - (n - \mathfrak{v}) \le 5n^{\frac{1}{4}} - n \le 0$ for $n \gg 1$, Lemma 4.2 gives $\frac{d_\lambda}{d_\nu} \le \frac{b_\nu^{b_\nu}(n-\mathfrak{f})^{b_\lambda}}{(n-\mathfrak{v}-b_\nu)^{b_\nu}} \le \frac{b_\lambda^{b_\lambda} n^{b_\lambda}}{(n-2n^{1/4})^{b_\nu}}$. Altogether,

$$\frac{d_\lambda d_\nu^3}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \le \frac{b_\lambda^{5b_\lambda} n^{(b_\lambda+4b_\nu)}}{\left(n-2n^{1/4}\right)^{b_\nu+\sum_f b_{\mu_f}}} = b_\lambda^{5b_\lambda} n^{\left(b_\lambda+3b_\nu-\sum_f b_{\mu_f}\right)} \left(\frac{1}{1-2n^{-3/4}}\right)^{b_\nu+\sum_f b_{\mu_f}}$$
$$\le b_\lambda^{5b_\lambda} n^{\left(b_\lambda+3b_\nu-\sum_f b_{\mu_f}\right)} \cdot \left(\frac{1}{1-2n^{-3/4}}\right)^{9n^{1/4}}.$$

As $\left(\frac{1}{1-2n^{-3/4}}\right)^{9n^{1/4}} \overset{n\to\infty}{\to} 1$, the right hand side of the last inequality is at most $2b_\lambda^{5b_\lambda} n^{\left(b_\lambda+3b_\nu-\sum_f b_{\mu_f}\right)}$ for large enough n. □

We next give bounds for the individual $\Xi_n^{\overline{B}}$.

Lemma 5.6. *There is $\kappa > 1$ such that if Y is ε-adapted for $\varepsilon \ge 0$, $\mathfrak{v} \le n^{1/4}$ and $B_\nu \le \mathfrak{D}$, then*

$$\left|\Xi_n^{\overline{B}}\right| \ll B_\lambda^{10B_\lambda} \left(\mathfrak{D}^{24} n^{-\varepsilon}\right)^{B_\lambda - B_\nu} \left(\frac{\kappa \mathfrak{D}^4}{(n-\mathfrak{v}-\mathfrak{D}^2)^2}\right)^{B_\nu}.$$

Proof. By assumption, $B_\nu \le \mathfrak{D} \le \mathfrak{v} \le n^{\frac{1}{4}}$. So for every $(\nu, \{\mu_f\}, \lambda) \vdash \overline{B}$,

$$\begin{aligned}\lambda_1 + \nu_1 - (n-\mathfrak{f}) &= (n - \mathfrak{f} - B_\lambda) + (n - \mathfrak{v} - B_\nu) - (n - \mathfrak{f}) \\ &\ge n - \mathfrak{v} - B_\nu - (B_\nu + \mathfrak{D}) \ge n - 4\mathfrak{v},\end{aligned}$$

and Lemma 5.2 gives that whenever $\left\{r_f^\pm, s_f, t_f\right\}$ satisfy (5.6),

$$\left|\mathcal{M}\left(\left\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\right\}\right)\right| \le \left(\frac{\mathfrak{D}^2}{n-4\mathfrak{v}}\right)^{D_{\mathrm{top}}\left(\left\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\right\}\right)}. \tag{5.13}$$

Proposition 5.3 gives

$$B_\lambda + 3B_\nu - \sum_f B_{\mu_f} \le D_{\mathrm{top}}\left(\left\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\right\}\right) - \varepsilon\left(B_\lambda - B_\nu\right),$$

so by Lemma 5.5

$$\frac{d_\lambda d_\nu^3}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \left|\mathcal{M}\left(\left\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\right\}\right)\right|$$
$$\ll B_\lambda^{5B_\lambda} n^{-\varepsilon(B_\lambda - B_\nu)} \left(\frac{n\mathfrak{D}^2}{n-4\mathfrak{v}}\right)^{D_{\mathrm{top}}\left(\left\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\right\}\right)}.$$

Now using the trivial upper bound $D_{\text{top}}\left(\left\{\sigma_f^{\pm}, \tau_f^{\pm}, r_f^{\pm}, s_f, t_f\right\}\right) \leq 8(B_\lambda - B_\nu)$ in (5.8) and $B_\lambda - B_\nu \leq \mathfrak{v} - \mathfrak{f} \leq \mathfrak{v} \leq n^{\frac{1}{4}}$, we obtain that for large enough n,

$$\left(\frac{n\mathfrak{D}^2}{n-4\mathfrak{v}}\right)^{D_{\text{top}}\left(\left\{\sigma_f^{\pm},\tau_f^{\pm},r_f^{\pm},s_f,t_f\right\}\right)} \leq \mathfrak{D}^{16(B_\lambda - B_\nu)}\left(\frac{1}{1-4n^{-3/4}}\right)^{8n^{1/4}} \leq 2\mathfrak{D}^{16(B_\lambda - B_\nu)}.$$

Therefore,

$$\frac{d_\lambda d_\nu^3}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}}\left|\mathcal{M}\left(\left\{\sigma_f^{\pm}, \tau_f^{\pm}, r_f^{\pm}, s_f, t_f\right\}\right)\right| \ll {B_\lambda}^{5B_\lambda}\left(\mathfrak{D}^{16} n^{-\varepsilon}\right)^{B_\lambda - B_\nu}.$$

From this we obtain

$$\begin{aligned}\left|\Xi_n^{\overline{B}}\right| &\ll {B_\lambda}^{5B_\lambda}\left(\mathfrak{D}^{16} n^{-\varepsilon}\right)^{B_\lambda - B_\nu} \sum_{(\nu,\{\mu_f\},\lambda)\vdash\overline{B}} \frac{1}{d_\nu^2} \sum_{\substack{r_f^+, r_f^- \in \operatorname{Tab}(\mu_f/\nu) \\ s_f, t_f \in \operatorname{Tab}(\lambda/\mu_f)}} 1 \\ &\leq {B_\lambda}^{5B_\lambda}\left(\mathfrak{D}^{24} n^{-\varepsilon}\right)^{B_\lambda - B_\nu} \sum_{(\nu,\{\mu_f\},\lambda)\vdash\overline{B}} \frac{1}{d_\nu^2}\end{aligned}$$

since there are at most $(\mathfrak{D})_{(B_\lambda - B_\nu)} \leq \mathfrak{D}^{(B_\lambda - B_\nu)}$ choices of $r_f^+ \sqcup s_f$ or of $r_f^- \sqcup t_f$ for all f, by Lemma 4.3. For fixed ν above, there are at most $B_\lambda^{5B_\lambda}$ choices of $\{\mu_f\}$ and λ such that $(\nu, \{\mu_f\}, \lambda) \vdash \overline{B}$. For example, the boxes outside the first row of λ uniquely determine λ and form a YD of size B_λ; there are at most $B_\lambda! \leq B_\lambda^{B_\lambda}$ of these. Hence

$$\left|\Xi_n^{\overline{B}}\right| \ll {B_\lambda}^{10B_\lambda}\left(\mathfrak{D}^{24} n^{-\varepsilon}\right)^{B_\lambda - B_\nu} \sum_{\nu\vdash n-\mathfrak{v}\,:\, b_\nu = B_\nu} \frac{1}{d_\nu^2}.$$

Note that above, we have $\nu_1 = n - \mathfrak{v} - B_\nu \geq n - 2n^{\frac{1}{4}}$, so $\check{b}_\nu \geq n - 2n^{\frac{1}{4}} - 1 \geq n^{\frac{1}{4}} \geq B_\nu$ for $n \gg 1$, and in this case $\nu \in \Lambda(n - \mathfrak{v}, B_\nu)$. Moreover, for $n \gg 1$, $B_\nu^2 \leq n^{\frac{1}{2}} \leq \frac{n - n^{\frac{1}{4}}}{3} \leq \frac{n-\mathfrak{v}}{3}$ and so we can finally apply Proposition 4.6 to obtain for the same $\kappa = \kappa(2) > 1$ from Proposition 4.6 that

$$\begin{aligned}\left|\Xi_n^{\overline{B}}\right| &\ll {B_\lambda}^{10B_\lambda}\left(\mathfrak{D}^{24} n^{-\varepsilon}\right)^{B_\lambda - B_\nu}\left(\frac{\kappa {B_\nu}^4}{\left(n - \mathfrak{v} - {B_\nu}^2\right)^2}\right)^{B_\nu} \\ &\leq {B_\lambda}^{10B_\lambda}\left(\mathfrak{D}^{24} n^{-\varepsilon}\right)^{B_\lambda - B_\nu}\left(\frac{\kappa \mathfrak{D}^4}{\left(n - \mathfrak{v} - \mathfrak{D}^2\right)^2}\right)^{B_\nu}.\end{aligned}$$

□

Since Lemma 5.6 is only useful for B_ν or $\check{B}_\nu$ small compared to n we have to supplement it with the following weaker bound.

Lemma 5.7. *If Y is any tiled surface and $\overline{B} \in \mathcal{B}_n(Y)$ then*

$$\left|\Xi_n^{\overline{B}}\right| \leq (\mathfrak{D}!)^8 \sum_{(\nu,\{\mu_f\},\lambda)\vdash\overline{B}} \frac{d_\lambda}{d_\nu^3}.$$

Proof. Since $\mathcal{M}(\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm s_f, t_f\})$ is a product of matrix coefficients of unit vectors in unitary representations, we obtain $|\mathcal{M}(\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm s_f, t_f\})| \leq 1$. Therefore, with assumptions as in the lemma, and arguing similarly as in the proof of Lemma 5.6, we obtain

$$\begin{aligned}\left|\Xi_n^{\overline{B}}\right| &\leq \sum_{(\nu,\{\mu_f\},\lambda)\vdash\overline{B}} \frac{d_\lambda d_\nu}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \sum_{\substack{r_f^+, r_f^- \in \mathrm{Tab}(\mu_f/\nu) \\ s_f, t_f \in \mathrm{Tab}(\lambda/\mu_f)}} 1 \\ &\overset{(*)}{\leq} (\mathfrak{D}!)^8 \sum_{(\nu,\{\mu_f\},\lambda)\vdash\overline{B}} \frac{d_\lambda d_\nu}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \\ &\leq (\mathfrak{D}!)^8 \sum_{(\nu,\{\mu_f\},\lambda)\vdash\overline{B}} \frac{d_\lambda}{d_\nu^3},\end{aligned}$$

where in $(*)$ we used the fact there are at most $|\lambda/\nu|! = (\mathfrak{v}-\mathfrak{f})!$ choices of $r_f^+ \sqcup s_f$ and of $r_f^- \sqcup t_f$. □

5.3 The zero regime of b_ν. We only need analytic estimates for $\Xi_n^{\nu=(n-\mathfrak{v})}$ when Y is boundary reduced (so 0-adapted); when Y is ε-adapted for $\varepsilon > 0$ we will take a different, more algebraic approach to $\Xi_n^{\nu=(n-\mathfrak{v})}$ in Section 5.7.

Lemma 5.8. *If Y is boundary reduced and $\mathfrak{v} \leq n^{1/4}$ then*

$$\left|\Xi_n^{\nu=(n-\mathfrak{v})}\right| \ll (\mathfrak{D}+1)^9 \mathfrak{D}^{34\mathfrak{D}}.$$

Proof. If $\nu = (n-\mathfrak{v})$ then $B_\nu = 0$. Inserting the bounds from Lemma 5.6 with $\varepsilon = 0$ (since Y is boundary reduced, see Lemma 3.13) and $B_\nu = 0$ gives

$$\left|\Xi_n^{\nu=(n-\mathfrak{v})}\right| \ll \sum_{\overline{B}\in\mathcal{B}_n(Y):\ B_\nu=0} B_\lambda^{10B_\lambda} \mathfrak{D}^{24B_\lambda}.$$

Because $\overline{B} \in \mathcal{B}_n(Y)$, there exist some $(\nu, \{\mu_f\}, \lambda) \vdash \overline{B}$ and satisfying (5.5). We then have since $\nu \subset_{\mathfrak{v}-\mathfrak{f}} \lambda$, and $b_\nu = 0$, $B_\lambda = b_\lambda \leq b_\nu + \mathfrak{v} - \mathfrak{f} = \mathfrak{v} - \mathfrak{f} = \mathfrak{D}$. In $\mathcal{B}_n(Y)$, the set of $\overline{B}'s$ with $B_\nu = 0$ and a fixed value of B_λ is of size at most $(\mathfrak{D}+1)^9$. Indeed, there are at most $B_\lambda + 1 \leq \mathfrak{D}+1$ options for B_{μ_f} for each f. Since $n - \mathfrak{v} - 1 = \check{B}_\nu \leq \check{B}_{\mu_f} \leq \check{B}_\lambda \leq n - \mathfrak{f} - 1$, there are at most $\mathfrak{v} - \mathfrak{f} + 1 = \mathfrak{D}+1$ possible values of each of $\check{B}_{\mu_f}$ and $\check{B}_\lambda$. In total then there are at most $(\mathfrak{D}+1)^9$ choices. Hence

$$\left|\Xi_n^{\nu=(n-\mathfrak{v})}\right| \ll (\mathfrak{D}+1)^9 \sum_{B_\lambda=0}^{\mathfrak{D}} \left({B_\lambda}^{10}\mathfrak{D}^{24}\right)^{B_\lambda} \leq (\mathfrak{D}+1)^9$$
$$\sum_{B_\lambda=0}^{\mathfrak{D}} (\mathfrak{D}^{34})^{B_\lambda} \ll (\mathfrak{D}+1)^9\mathfrak{D}^{34\mathfrak{D}}.$$

□

5.4 The intermediate regime of b_ν.

LEMMA 5.9. *Assume that* $\mathfrak{v} \leq n^{1/4}$.

1. *If* Y *is boundary reduced with* $\mathfrak{D} \leq n^{1/10}$ *then*
$$\left|\Xi_n^{0<b_\nu\leq\mathfrak{D};\check{b}_\nu>0}\right| \ll \frac{\left(\mathfrak{D}^{34}2^{10}\right)^{\mathfrak{D}+1}}{\left(n-\mathfrak{v}-\mathfrak{D}^2\right)^2}. \tag{5.14}$$
2. *For any* $\varepsilon \in (0,1)$, *there is* $\eta = \eta(\varepsilon) \in (0, \frac{1}{100})$ *such that if* Y *is* ε*-adapted, with* $\mathfrak{D} \leq n^\eta$ *then*
$$\left|\Xi_n^{0<b_\nu\leq\mathfrak{D};\check{b}_\nu>0}\right| \ll_\varepsilon \frac{1}{n}. \tag{5.15}$$

Proof. When $\mathfrak{D} = 0$, the inequality $0 < b_\nu \leq \mathfrak{D}$ cannot hold, and so $\Xi_n^{0<b_\nu\leq\mathfrak{D};\check{b}_\nu>0} = 0$ by definition, and both statements hold. So assume $\mathfrak{D} \geq 1$. We can also assume that $\mathfrak{D} \leq n^{1/10}$.

For any $\varepsilon \geq 0$, the bounds from Lemma 5.6 give

$$\left|\Xi_n^{0<b_\nu\leq\mathfrak{D};\check{b}_\nu>0}\right| \ll \sum_{\substack{\overline{B}\in\mathcal{B}_n(Y):\\ 0<B_\nu\leq\mathfrak{D};\check{B}_\nu>0}} {B_\lambda}^{10B_\lambda} \left(\mathfrak{D}^{24}n^{-\varepsilon}\right)^{B_\lambda-B_\nu} \left(\frac{\kappa\mathfrak{D}^4}{\left(n-\mathfrak{v}-\mathfrak{D}^2\right)^2}\right)^{B_\nu}.$$

Arguing similarly as in the proof of Lemma 5.8, the number of $\overline{B}$'s in the sum above with a fixed value of B_ν and B_λ is $\ll \mathfrak{D}^{10}$. Also note that $B_\lambda \leq B_\nu + \mathfrak{D} \leq 2\mathfrak{D}$. We obtain

$$\begin{aligned}\left|\Xi_n^{0<b_\nu\leq\mathfrak{D};\check{b}_\nu>0}\right| &\ll \mathfrak{D}^{10} \sum_{\substack{0<B_\nu\leq\mathfrak{D}\\ B_\nu\leq B_\lambda\leq B_\nu+\mathfrak{D}}} {B_\lambda}^{10B_\lambda} \left(\mathfrak{D}^{24}n^{-\varepsilon}\right)^{B_\lambda-B_\nu} \left(\frac{\kappa\mathfrak{D}^4}{\left(n-\mathfrak{v}-\mathfrak{D}^2\right)^2}\right)^{B_\nu}\\ &\leq \mathfrak{D}^{10} \sum_{B_\nu=1}^{\mathfrak{D}} \left(\frac{\kappa(2\mathfrak{D})^{10}\mathfrak{D}^4}{\left(n-\mathfrak{v}-\mathfrak{D}^2\right)^2}\right)^{B_\nu} \sum_{B_\lambda=B_\nu}^{B_\nu+\mathfrak{D}} \left(\mathfrak{D}^{24}{B_\lambda}^{10}n^{-\varepsilon}\right)^{B_\lambda-B_\nu}.\end{aligned}$$

As $B_\lambda \leq 2\mathfrak{D}$, we bound the second summation by $\sum_{t=0}^{\mathfrak{D}} \left(\mathfrak{D}^{34}2^{10}n^{-\varepsilon}\right)^t$. By our assumption that $\mathfrak{D} \leq n^{1/10}$ and $\mathfrak{v} \leq n^{1/4}$, we have $\frac{\kappa(2\mathfrak{D})^{10}\mathfrak{D}^4}{(n-\mathfrak{v}-\mathfrak{D}^2)^2} \leq \frac{1}{2}$ for large enough n. Hence

$$\left|\Xi_n^{0<b_\nu\leq\mathfrak{D};\check{b}_\nu>0}\right| \ll \frac{\mathfrak{D}^{10}\cdot\kappa(2\mathfrak{D})^{10}\mathfrak{D}^4}{\left(n-\mathfrak{v}-\mathfrak{D}^2\right)^2}\sum_{t=0}^{\mathfrak{D}}\left(\mathfrak{D}^{34}2^{10}n^{-\varepsilon}\right)^t$$
$$\ll \frac{\mathfrak{D}^{24}}{\left(n-\mathfrak{v}-\mathfrak{D}^2\right)^2}\sum_{t=0}^{\mathfrak{D}}\left(\mathfrak{D}^{34}2^{10}n^{-\varepsilon}\right)^t. \tag{5.16}$$

If Y is boundary reduced, it is 0-adapted (Lemma 3.13), so (5.16) yields

$$\left|\Xi_n^{0<b_\nu\leq\mathfrak{D};\check{b}_\nu>0}\right| \ll \frac{\mathfrak{D}^{24}}{\left(n-\mathfrak{v}-\mathfrak{D}^2\right)^2}\cdot\left(\mathfrak{D}^{34}2^{10}\right)^{\mathfrak{D}} \leq \frac{\left(\mathfrak{D}^{34}2^{10}\right)^{\mathfrak{D}+1}}{\left(n-\mathfrak{v}-\mathfrak{D}^2\right)^2}$$

proving the first statement.

For the second statement, given $\varepsilon>0$, let $\eta=\frac{\varepsilon}{100}$ and assume $1\leq\mathfrak{D}\leq n^\eta$. The choice of η implies that for $n\gg_\varepsilon 1$, $\mathfrak{D}^{34}2^{10}n^{-\varepsilon}\leq\frac{1}{2}$, so (5.16) gives

$$\left|\Xi_n^{0<b_\nu\leq\mathfrak{D};\check{b}_\nu>0}\right| \ll_\varepsilon \frac{\mathfrak{D}^{24}}{\left(n-\mathfrak{v}-\mathfrak{D}^2\right)^2} \ll_\varepsilon \frac{1}{n}.$$

□

5.5 The large regime of $b_\nu, \check{b}_\nu$. In the large regime of b_ν and $\check{b}_\nu$ we use the same estimate for any type of tiled surface.

LEMMA 5.10. *If $\mathfrak{v}\leq n^{1/4}$ and $\mathfrak{D}\leq n^{1/24}$ then*

$$\left|\Xi_n^{b_\nu,\check{b}_\nu>\mathfrak{D}}\right| \ll \frac{(\mathfrak{D}+1)^4}{\left(n-\mathfrak{v}-\mathfrak{D}^2\right)^2}.$$

Proof. Using the bound from Lemma 5.7 gives

$$\begin{aligned}
\left|\Xi_n^{b_\nu,\check{b}_\nu>\mathfrak{D}}\right| &\leq \sum_{\overline{B}\in\mathcal{B}_n(Y):\ B_\nu,\check{B}_\nu>\mathfrak{D}}(\mathfrak{D}!)^8\sum_{(\nu,\{\mu_f\},\lambda)\vdash\overline{B}}\frac{d_\lambda}{d_\nu^3}\\
&\leq (\mathfrak{D}!)^8\sum_{\nu\vdash n-\mathfrak{v},b_\nu>\mathfrak{D},\check{b}_\nu>\mathfrak{D}}d_\nu^{-3}\sum_{\nu\subset_{\mathfrak{v}-\mathfrak{f}}\lambda}d_\lambda\sum_{\nu\subset\mu_f\subset_{\mathfrak{e}_f-\mathfrak{f}}\lambda}1\\
&\leq (\mathfrak{D}!)^{12}\sum_{\nu\vdash n-\mathfrak{v},b_\nu>\mathfrak{D},\check{b}_\nu>\mathfrak{D}}d_\nu^{-3}\sum_{\nu\subset_{\mathfrak{v}-\mathfrak{f}}\lambda}d_\lambda \leq \mathfrak{D}^{12\mathfrak{D}}\frac{(n-\mathfrak{f})!}{(n-\mathfrak{v})!}\sum_{\nu\vdash n-\mathfrak{v},b_\nu>\mathfrak{D},\check{b}_\nu>\mathfrak{D}}d_\nu^{-2}\\
&\ll \mathfrak{D}^{12\mathfrak{D}}n^{\mathfrak{D}}\left(\frac{\kappa\,(\mathfrak{D}+1)^4}{\left(n-\mathfrak{v}-(\mathfrak{D}+1)^2\right)^2}\right)^{\mathfrak{D}+1}\\
&= \left(\frac{\kappa n\mathfrak{D}^{12}\,(\mathfrak{D}+1)^4}{\left(n-\mathfrak{v}-(\mathfrak{D}+1)^2\right)^2}\right)^{\mathfrak{D}}\frac{\kappa\,(\mathfrak{D}+1)^4}{\left(n-\mathfrak{v}-(\mathfrak{D}+1)^2\right)^2}.
\end{aligned}$$

The second-last inequality used Lemma 4.1 and the final inequality used Proposition 4.6. Since we assume $\mathfrak{D}\leq n^{1/24}$ and $\mathfrak{v}\leq n^{1/4}$ we obtain the stated result. □

5.6 Assembly of analytic estimates for Ξ_n. Now we combine the estimates obtained in Sections 5.3, 5.4, 5.5. First we give the culmination of our previous estimates when Y is boundary reduced.

PROPOSITION 5.11. *There is $A_0 > 0$ such that if Y is boundary reduced, $\mathfrak{v} \leq n^{1/4}$, and $\mathfrak{D} \leq n^{1/24}$, then*

$$|\Xi_n| \ll (A_0\mathfrak{D})^{A_0\mathfrak{D}}.$$

Proof. With assumptions as in the proposition, splitting Ξ_n as in (5.10) and using Lemmas 5.8, 5.9(1), and 5.10 gives

$$|\Xi_n| \ll (\mathfrak{D}+1)^9\mathfrak{D}^{34\mathfrak{D}} + \frac{\left(\mathfrak{D}^{34}2^{10}\right)^{\mathfrak{D}+1}}{\left(n-\mathfrak{v}-\mathfrak{D}^2\right)^2} + \frac{(\mathfrak{D}+1)^4}{\left(n-\mathfrak{v}-\mathfrak{D}^2\right)^2}.$$

If $\mathfrak{D} = 0$ this gives $|\Xi_n| \ll 1$ which proves the result. If $1 \leq \mathfrak{D} \leq n^{1/24}$ we obtain $|\Xi_n| \ll (A_0\mathfrak{D})^{A_0\mathfrak{D}}$ as required. □

Next we show that if Y is ε-adapted, then $\mathfrak{D}$ can be as large as a fractional power of n while Ξ_n is still very well approximated by $2\Xi_n^{\nu=(n-\mathfrak{v})}$.

PROPOSITION 5.12. *For any $\varepsilon \in (0,1)$, there is $\eta = \eta(\varepsilon) \in (0, \frac{1}{100})$ such that if Y is ε-adapted with $\mathfrak{D} \leq n^\eta$ and $\mathfrak{v} \leq n^{1/4}$, then*

$$\left|\Xi_n - 2\Xi_n^{\nu=(n-\mathfrak{v})}\right| \ll_\varepsilon \frac{1}{n}.$$

Proof. Lemmas 5.9(2) and 5.10 yield that given $\varepsilon \in (0,1)$, there is $\eta = \eta(\varepsilon) \in (0, \frac{1}{100})$, such that if $\mathfrak{D} \leq n^\eta$, $\mathfrak{v} \leq n^{1/4}$ and Y is ε-adapted, then

$$\begin{aligned}\left|\Xi_n - 2\Xi_n^{\nu=(n-\mathfrak{v})}\right| &= \left|2\Xi_n^{0<b_\nu\leq\mathfrak{D};\check{b}_\nu>0} + \Xi_n^{b_\nu,\check{b}_\nu>\mathfrak{D}}\right| \\ &\ll_\varepsilon \frac{1}{n} + \frac{(\mathfrak{D}+1)^4}{\left(n-\mathfrak{v}-\mathfrak{D}^2\right)^2} \ll \frac{1}{n}.\end{aligned}$$

□

REMARK. For general g, the condition $\eta(\epsilon) < \frac{1}{100}$ of Proposition 5.12 should be replaced by $\eta(\epsilon) < \frac{1}{Cg}$ for some universal $C \geq 100$.

5.7 A new expression for $\Xi_n^{\nu=(n-\mathfrak{v})}$. We continue to fix a compact tiled surface Y. The goal of this section is to give a formula for $\Xi_n^{\nu=(n-\mathfrak{v})}$ that is more precise than is possible to obtain with the methods of the previous section. This will be done by refining the methods of [MP20, Section 5].

We will assume throughout that $n \geq \mathfrak{v}$. We fix a bijective map $\mathcal{J} : Y^{(0)} \to [\mathfrak{v}]$, and as in [MP20, Section 5] for each $n \in \mathbf{N}$ we modify $\mathcal{J}$ by letting

$$\mathcal{J}_n : Y^{(0)} \to [n-\mathfrak{v}+1, n], \quad \mathcal{J}_n(v) \stackrel{\text{def}}{=} \mathcal{J}(v) + n - \mathfrak{v}. \tag{5.17}$$

We use the map $\mathcal{J}_n$ to identify the vertex set of Y with $[n-\mathfrak{v}+1,n]$. Let $\mathcal{V}_f^- = \mathcal{V}_f^-(Y) \subset [n-\mathfrak{v}+1,n]$ be the subset of vertices of Y with outgoing f-labeled edges, and $\mathcal{V}_f^+ \subset [n-\mathfrak{v}+1,n]$ those vertices of Y with incoming f-labeled edges. Note that $\mathfrak{e}_f = |\mathcal{V}_f^-| = |\mathcal{V}_f^+|$. Recall that $S'_{\mathfrak{v}} \leq S_n$ is the subgroup of permutations fixing $[n-\mathfrak{v}]$ element-wise. For each $f \in \{a,b,c,d\}$ we fix $g_f^0 \in S'_{\mathfrak{v}}$ such that for every pair of vertices i,j of Y in $[n-\mathfrak{v}+1,n]$ with a directed f-labeled edge from i to j, we have $g_f^0(i) = j$. Note that $g_f^0(\mathcal{V}_f^-) = \mathcal{V}_f^+$. We let $g^0 \stackrel{\text{def}}{=} (g_a^0, g_b^0, g_c^0, g_d^0) \in S_n^4$. For each $f \in \{a,b,c,d\}$ let G_f be the subgroup of S_n fixing pointwise $\mathcal{V}_f^-$. Let $G \stackrel{\text{def}}{=} G_a \times G_b \times G_c \times G_d \leq S_n^4$.

Our formula for $\Xi_n^{\nu=(n-\mathfrak{v})}$ will involve the size of the set

$$\mathbb{X}_n^*(Y,\mathcal{J}) \stackrel{\text{def}}{=} \left\{(\alpha_a,\alpha_b,\alpha_c,\alpha_d) \in g^0 G \,|\, W(\alpha_a,\alpha_b,\alpha_c,\alpha_d) \in S_{n-\mathfrak{v}}\right\} \tag{5.18}$$

where[15] $W(g_a,g_b,g_c,g_d) \stackrel{\text{def}}{=} g_d^{-1}g_c^{-1}g_dg_cg_b^{-1}g_a^{-1}g_bg_a$. Note that a similar set, denoted $\mathbb{X}_n(Y,\mathcal{J})$ in [MP20, Section 5.2], is the set in which the condition is that $W(\alpha_a,\alpha_b,\alpha_c,\alpha_d) = 1$ rather than the identity only when restricted to $[n-\mathfrak{v}+1,n]$, as in (5.18). This smaller set $\mathbb{X}_n(Y,\mathcal{J})$ counts the number of covers $\phi \in \mathrm{Hom}(\Gamma_2, S_n)$ in which $(Y,\mathcal{J})$ embeds.

The main result of this Section 5.7 is the following.

Proposition 5.13. *With notations as above,*

$$\Xi_n^{\nu=(n-\mathfrak{v})} = \frac{(n)_{\mathfrak{v}}\,|\mathbb{X}_n^*(Y,\mathcal{J})|}{(n)_{\mathfrak{f}}\prod_{f\in a,b,c,d}(n-\mathfrak{e}_f)!}.$$

Recall that $(n)_q$ is the Pochhammer symbol as defined in Section 1.4. *In the rest of the paper, whenever we write an integral over a group, it is performed with respect to the uniform measure on the relevant group.* Let

$$I \stackrel{\text{def}}{=} \int_{h_f\in G_f}\int_{\pi\in S_{n-\mathfrak{v}}} \mathbf{1}\left\{W\left(g_a^0h_a, g_b^0h_b, g_c^0h_c, g_d^0h_d\right)\pi = 1\right\}.$$

The following lemma is immediate as a result of relating sums to normalized integrals.

Lemma 5.14. *We have* $|\mathbb{X}_n^*(Y,\mathcal{J})| = |S_{n-\mathfrak{v}}|\cdot|G|\cdot I$.

For a Young diagram λ of size m, we write χ_λ for the trace of the irreducible representation of S_m on V^λ.

[15] The reason we use this word instead of the relator $[g_a,g_b][g_c,g_d]$ of Γ_2 is the same as in [MP20]: the one-to-one correspondence between $\mathbb{X}_n$ and degree-n covers of a genus 2 surface uses the version of the symmetric group where permutations are multiplied as functions acting from the *right*, whereas in this section we want to multiply permutations as functions on $[n]$ acting from the *left*.

Corollary 5.15. *We have*

$$|\mathbb{X}_n^*(Y,\mathcal{J})| = \frac{\prod_{f\in a,b,c,d}(n-\mathfrak{e}_f)!}{(n)_{\mathfrak{v}}} \sum_{\lambda\vdash n} d_\lambda \Theta_\lambda^{(n-\mathfrak{v})}(Y,\mathcal{J})$$

where

$$\Theta_\lambda^{(n-\mathfrak{v})}(Y,\mathcal{J}) \stackrel{\text{def}}{=} \int_{h_f\in G_f}\int_{\pi\in S_{n-\mathfrak{v}}} \chi_\lambda\left(W\left(g_a^0h_a, g_b^0h_b, g_c^0h_c, g_d^0h_d\right)\pi\right). \tag{5.19}$$

Proof. Using Schur orthogonality, write

$$\mathbf{1}\{g=1\} = \frac{1}{n!}\sum_{\lambda\vdash n} d_\lambda \chi_\lambda(g),$$

hence

$$I = \frac{1}{n!}\sum_{\lambda\vdash n} d_\lambda \Theta_\lambda^{(n-\mathfrak{v})}(Y,\mathcal{J}).$$

We have $|G| = \prod_{f\in\{a,b,c,d\}}(n-\mathfrak{e}_f)!$, hence by Lemma 5.14

$$\begin{aligned}|\mathbb{X}_n^*(Y,\mathcal{J})| &= (n-\mathfrak{v})! \prod_{f\in\{a,b,c,d\}}(n-\mathfrak{e}_f)! \cdot \frac{1}{n!}\sum_{\lambda\vdash n} d_\lambda \Theta_\lambda^{(n-\mathfrak{v})}(Y,\mathcal{J}) \\ &= \frac{\prod_{f\in a,b,c,d}(n-\mathfrak{e}_f)!}{(n)_{\mathfrak{v}}}\sum_{\lambda\vdash n} d_\lambda \Theta_\lambda^{(n-\mathfrak{v})}(Y,\mathcal{J}).\end{aligned}$$

□

Consider the vector space

$$W^\lambda \stackrel{\text{def}}{=} V^\lambda \otimes \check{V}^\lambda \otimes V^\lambda \otimes \check{V}^\lambda \otimes V^\lambda \otimes \check{V}^\lambda \otimes V^\lambda \otimes \check{V}^\lambda$$

as a unitary representation of S_n^8. This is a departure from [MP20, Section 5] where W^λ was thought of as a representation of S_n^4; we take a more flexible setup here. The reader may find it useful to see [MP20, Section 5.4] for extra background on representation theory. The inner product on V^λ gives an isomorphism $V^\lambda \cong \check{V}^\lambda$, $v\mapsto \check{v}$. Let $B_\lambda \in \mathrm{End}(W^\lambda)$ be defined as in [MP20, Equation (5.9)] by the formula

$$\begin{aligned}&\langle B_\lambda\left(v_1\otimes\check{v}_2\otimes v_3\otimes\check{v}_4\otimes v_5\otimes\check{v}_6\otimes v_7\otimes\check{v}_8\right), w_1\otimes\check{w}_2\otimes w_3\otimes\check{w}_4\otimes w_5\otimes\check{w}_6\otimes w_7\otimes\check{w}_8\rangle \\ &\stackrel{\text{def}}{=} \langle v_1,w_3\rangle\langle v_3,v_2\rangle\langle w_2,v_4\rangle\langle w_4,w_5\rangle\langle v_5,w_7\rangle\langle v_7,v_6\rangle\langle w_6,v_8\rangle\langle w_8,w_1\rangle.\end{aligned} \tag{5.20}$$

We note the following, extending [MP20, Lem. 5.4].

Lemma 5.16. *For any* $(g_1,g_2,g_3,g_4,g_5,g_6,g_7,g_8)\in S_n^8$, *we have*

$$\mathrm{tr}_{W^\lambda}(B_\lambda\circ(g_1,g_2,g_3,g_4,g_5,g_6,g_7,g_8)) = \chi_\lambda(g_8^{-1}g_6^{-1}g_7g_5g_4^{-1}g_2^{-1}g_3g_1).$$

Proof. The proof is a direct calculation directly generalizing [MP20, Lem. 5.4]. □

Let Q be the orthogonal projection in W^λ onto the vectors that are invariant by G acting on W^λ by the map

$$(g_a, g_b, g_c, g_d) \in G \mapsto (g_a, g_a, g_b, g_b, g_c, g_c, g_d, g_d) \in S_n^8.$$

This projection appeared also in [MP20, Section 5.4].

Lemma 5.17. *We have* $\Theta_\lambda^{(n-\mathfrak{v})}(Y, \mathcal{J}) = \mathrm{tr}_{W^\lambda}(\mathfrak{p} B_\lambda g^0 Q)$ *where* $\mathfrak{p}$ *denotes the operator*

$$\mathfrak{p} \stackrel{\text{def}}{=} \int_{\pi \in S_{n-\mathfrak{v}}} (\pi, 1, 1, 1, 1, 1, 1, 1) \in \mathrm{End}\left(W^\lambda\right).$$

Remark 5.18. Note that $\mathfrak{p}$ is the projection in $\mathrm{End}(W^\lambda)$ onto the triv-isotypic subspace for the action of $S_{n-\mathfrak{v}}$ on the first factor of W^λ (while being the identity on the remaining seven factors). This is a self-adjoint operator.

Proof. Recall the definition of $\Theta_\lambda^{(n-\mathfrak{v})}(Y, \mathcal{J})$ in (5.19). Using Lemma 5.16, for every set of fixed values of the h_f and π, we have

$$\begin{aligned}
&\chi_\lambda\left(W\left(g_a^0 h_a, g_b^0 h_b, g_c^0 h_c, g_d^0 h_d\right)\pi\right) \\
&\quad = \mathrm{tr}_{W^\lambda}\left(B_\lambda \circ \left(g_a^0 h_a \pi, g_a^0 h_a, g_b^0 h_b, g_b^0 h_b, g_c^0 h_c, g_c^0 h_c, g_d^0 h_d, g_d^0 h_d\right)\right)
\end{aligned}$$

Therefore,

$$\Theta_\lambda^{(n-\mathfrak{v})}(Y, \mathcal{J}) = \mathrm{tr}_{W^\lambda}(B_\lambda g^0 Q \mathfrak{p}) = \mathrm{tr}_{W^\lambda}(\mathfrak{p} B_\lambda g^0 Q). \qquad \square$$

Using Lemma 5.17, we now find a new expression for $\Theta_\lambda^{(n-\mathfrak{v})}(Y, \mathcal{J})$ by calculating $\mathrm{tr}_{W^\lambda}(\mathfrak{p} B_\lambda g^0 Q)$.

Proposition 5.19. *We have*

$$\Theta_\lambda^{(n-\mathfrak{v})}(Y, \mathcal{J}) = \sum_{(n-\mathfrak{v}) \subset \mu_f \subset_{\mathfrak{e}_f - \mathfrak{f}} \lambda' \subset_{\mathfrak{f}} \lambda} \frac{d_{\lambda/\lambda'}}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \Upsilon_n\left(\left\{\sigma_f^\pm, \tau_f^\pm\right\}, (n-\mathfrak{v}), \{\mu_f\}, \lambda'\right). \tag{5.21}$$

Proof. This calculation is very similar to the proof of [MP20, Prop. 5.8] where $\mathrm{tr}_{W^\lambda}(B_\lambda g^0 Q)$ was calculated. The only difference here is the presence of the additional operator $\mathfrak{p}$. Therefore we will not give all the details. The proof follows [MP20, proof of Prop. 5.8] using properties **P1**–**P4** of $\sigma_f^\pm, \tau_f^\pm$. One also uses that $\mathfrak{p}$ is a self-adjoint projection. The role that $\mathfrak{p}$ plays in the proof is that instead of obtaining a summation over all $\nu \subset_{\mathfrak{v}} \lambda$, the projection $\mathfrak{p}$ forces only the relevant $\nu = (n - \mathfrak{v})$ to appear.

Indeed, the calculation leading to [MP20, Equation (5.17)] is replaced by

$$
\begin{aligned}
&\Big\langle \mathfrak{p} B_\lambda \left[\mathcal{E}^{\lambda,a,+}_{\mu_a,S_a,T_a} \otimes \mathcal{E}^{\lambda,b,+}_{\mu_b,S_b,T_b} \otimes \mathcal{E}^{\lambda,c,+}_{\mu_c,S_c,T_c} \otimes \mathcal{E}^{\lambda,d,+}_{\mu_d,S_d,T_d}\right], \\
&\quad \mathcal{E}^{\lambda,a,-}_{\mu_a,S_a,T_a} \otimes \mathcal{E}^{\lambda,b,-}_{\mu_b,S_b,T_b} \otimes \mathcal{E}^{\lambda,c,-}_{\mu_c,S_c,T_c} \otimes \mathcal{E}^{\lambda,d,-}_{\mu_d,S_d,T_d}\Big\rangle \\
&\quad = \Big\langle B_\lambda \left[\mathcal{E}^{\lambda,a,+}_{\mu_a,S_a,T_a} \otimes \mathcal{E}^{\lambda,b,+}_{\mu_b,S_b,T_b} \otimes \mathcal{E}^{\lambda,c,+}_{\mu_c,S_c,T_c} \otimes \mathcal{E}^{\lambda,d,+}_{\mu_d,S_d,T_d}\right], \\
&\quad \mathfrak{p}\left(\mathcal{E}^{\lambda,a,-}_{\mu_a,S_a,T_a} \otimes \mathcal{E}^{\lambda,b,-}_{\mu_b,S_b,T_b} \otimes \mathcal{E}^{\lambda,c,-}_{\mu_c,S_c,T_c} \otimes \mathcal{E}^{\lambda,d,-}_{\mu_d,S_d,T_d}\right)\Big\rangle \\
&\quad = \frac{1}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \sum_{R^\pm_f \in \mathrm{Tab}(\mu_f)} \left\langle v^{\sigma_a^+}_{R_a^+ \sqcup S_a}, v^{\sigma_b^-}_{R_b^- \sqcup S_b}\right\rangle \left\langle v^{\sigma_b^+}_{R_b^+ \sqcup S_b}, v^{\tau_a^+}_{R_a^+ \sqcup T_a}\right\rangle \left\langle v^{\tau_a^-}_{R_a^- \sqcup T_a}, v^{\tau_b^+}_{R_b^+ \sqcup T_b}\right\rangle \\
&\qquad \cdot \left\langle v^{\tau_b^-}_{R_b^- \sqcup T_b}, v^{\sigma_c^-}_{R_c^- \sqcup S_c}\right\rangle \left\langle v^{\sigma_c^+}_{R_c^+ \sqcup S_c}, v^{\sigma_d^-}_{R_d^- \sqcup S_d}\right\rangle \left\langle v^{\sigma_d^+}_{R_d^+ \sqcup S_d}, v^{\tau_c^+}_{R_c^+ \sqcup T_c}\right\rangle \\
&\qquad\quad \left\langle v^{\tau_c^-}_{R_c^- \sqcup T_c}, v^{\tau_d^+}_{R_d^+ \sqcup T_d}\right\rangle \left\langle v^{\tau_d^-}_{R_d^- \sqcup T_d}, \mathfrak{p}_0 v^{\sigma_a^-}_{R_a^- \sqcup S_a}\right\rangle
\end{aligned}
$$

where $\mathfrak{p}_0$ is orthogonal projection to the $S_{n-\mathfrak{v}}$-invariant vectors in V^λ. Then the same discussion as precedes [MP20, Equation (5.17)] applies now to show that the above is zero unless there is $\nu \vdash n-\mathfrak{v}$ such that $\nu \subset \mu_f$ for all $f \in \{a,b,c,d\}$, and all $R_f^+|_{\le n-\mathfrak{v}}$, $R_f^-|_{\le n-\mathfrak{v}}$ are equal and of shape ν, except now, the presence of $\mathfrak{p}_0$ forces $\nu = (n-\mathfrak{v})$. Then the rest of the proof is the same. □

Proof of Proposition 5.13. Combining Corollary 5.15 and Proposition 5.19 we obtain

$$
\begin{aligned}
&|\mathbb{X}^*_n(Y,\mathcal{J})| \\
&\quad = \frac{\prod_{f\in\{a,b,c,d\}}(n-\mathfrak{e}_f)!}{(n)_\mathfrak{v}} \sum_{\lambda\vdash n} d_\lambda \sum_{(n-\mathfrak{v})\subset\mu_f\subset_{\mathfrak{e}_f-\mathfrak{f}}\lambda'\subset_\mathfrak{f}\lambda} \frac{d_{\lambda/\lambda'}}{d_{\mu_a}d_{\mu_b}d_{\mu_c}d_{\mu_d}} \Upsilon_n \\
&\qquad \left(\left\{\sigma_f^\pm, \tau_f^\pm\right\}, (n-\mathfrak{v}), \{\mu_f\}, \lambda'\right) \\
&\quad = \frac{\prod_{f\in a,b,c,d}(n-\mathfrak{e}_f)!(n)_\mathfrak{f}}{(n)_\mathfrak{v}} \sum_{(n-\mathfrak{v})\subset\mu_f\subset_{\mathfrak{e}_f-\mathfrak{f}}\lambda'\vdash n-\mathfrak{f}} \frac{d_{\lambda'}}{d_{\mu_a}d_{\mu_b}d_{\mu_c}d_{\mu_d}} \Upsilon_n \\
&\qquad \left(\left\{\sigma_f^\pm, \tau_f^\pm\right\}, (n-\mathfrak{v}), \{\mu_f\}, \lambda'\right) \\
&\quad = \frac{\prod_{f\in a,b,c,d}(n-\mathfrak{e}_f)!(n)_\mathfrak{f}}{(n)_\mathfrak{v}} \Xi_n^{\nu=(n-\mathfrak{v})},
\end{aligned}
$$

where the second equality used Lemma 4.1 and the third used $d_{(n-\mathfrak{v})} = 1$. This gives the result. □

5.8 Understanding $|\mathbb{X}^*_n(Y,\mathcal{J})|$. Recall the definition of $\mathbb{X}^*_n(Y,\mathcal{J})$ in (5.18). Because these 4-tuples of permutations generally do not correspond to covers of the surface Σ_2, they are better analyzed as n-degree covers of the bouquet of four loops, namely, as graphs on n vertices labeled by $[n]$ with directed edges labeled by a,b,c,d, and exactly one incoming f-edge and one outgoing f-edge in every vertex and every

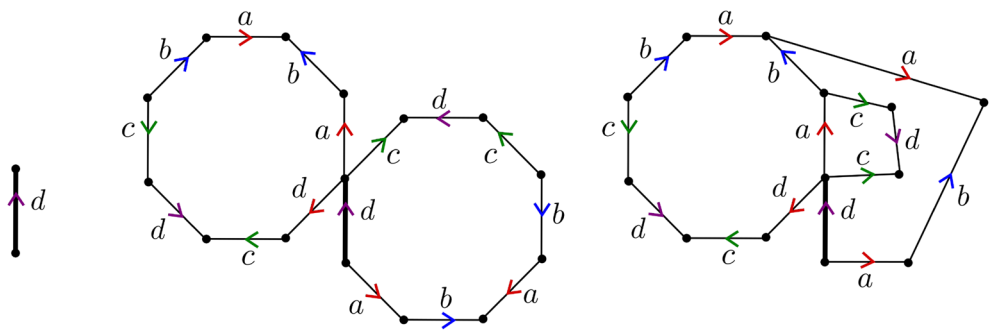

Figure 6: On the left is a tiled surface Y consisting of two vertices and a single d-edge between them. The middle part shows $\hat{Y}$: where we "grow" an octagon from every vertex of Y, and in which $Y^{(1)}$ is embedded. The figure on the right shows another element of $\mathcal{Q}(Y)$: a folded quotient of $\hat{Y}$ where $Y^{(1)}$ is still embedded.

$f \in \{a, b, c, d\}$. Equivalently, these graphs are the Schreier graphs depicting the action of S_n on $[n]$ with respect to the four permutations $\alpha_a, \alpha_b, \alpha_c, \alpha_d$.

Such a Schreier graph $\mathcal{G}$ corresponds to some 4-tuple $(\alpha_a, \alpha_b, \alpha_c, \alpha_d) \in \mathbb{X}_n^*(Y, \mathcal{J})$ if and only if the following two conditions are satisfied. The assumption that $(\alpha_a, \alpha_b, \alpha_c, \alpha_d) \in g^0 G$ means that $Y^{(1)}$, the 1-skeleton of Y, is embedded in $\mathcal{G}$, in an embedding that extends $\mathcal{J}_n$ on the vertices. The condition that $W(\alpha_a, \alpha_b, \alpha_c, \alpha_d) \in S_{n-\mathfrak{v}}$, means that at every vertex of $\mathcal{G}$ with label in $[n - \mathfrak{v} + 1, n]$, there is a closed path of length 8 that spells out the word $[a, b][c, d]$.

In Lemma 5.20 below we show that the number of such graphs (equal to $|\mathbb{X}_n^*(Y, \mathcal{J})|$) is rational in n. To this end, we apply techniques based on Stallings core graphs, in a similar fashion to the techniques applied in [Pud14, PP15].

Construct a finite graph $\hat{Y}$ as follows. Start with $Y^{(1)}$, the 1-skeleton of Y. At every vertex attach a closed cycle of length 8 spelling out $[a, b]\,[c, d]$. Then fold the resulting graph, in the sense of Stallings,[16] to obtain $\hat{Y}$. In other words, at each vertex v of $Y^{(1)}$, if there is a closed path at v spelling $[a, b]\,[c, d]$, do nothing. Otherwise, find the largest prefix of $[a, b]\,[c, d]$ that can be read on a path p starting at v and the largest suffix of $[a, b]\,[c, d]$ that can be read on a path s terminating at v. Because Y is a tiled surface, $|p| + |s| < 8$. Attach a path of length $8 - |p| - |s|$ between the endpoint of p and the beginning of s which spells out the missing part of the word $[a, b]\,[c, d]$. In this description, no folding is required. Note, in particular, that $Y^{(1)}$ is embedded in $\hat{Y}$.

By the discussion above, the Schreier graphs $\mathcal{G}$ corresponding to $\mathbb{X}_n^*(Y, \mathcal{J})$ are the graphs in which there is an embedding of $Y^{(1)}$ which *extends to a morphism of directed edge-labeled graphs* of $\hat{Y}$. We group these $\mathcal{G}$ according to the image of $\hat{Y}$. So

[16] Folding a graph with directed and labeled edges means that as long as there is a vertex with two incoming edges with the same label, or two outgoing edges with the same label, these two edges are merged, and so are their other endpoints. It is well known that this process has a unique outcome [Sta83, Section 3].

denote by $\mathcal{Q}(Y)$ the possible images of $\hat{Y}$ in the graphs $\mathcal{G}$: these are precisely the folded quotients of $\hat{Y}$ (edges can only be merged with equally-labeled other edges) which restrict to a bijection on $Y^{(1)}$. In particular, $\hat{Y} \in \mathcal{Q}(Y)$. We illustrate these concepts in Figure 6. As $\hat{Y}$ is a finite graph, the set $\mathcal{Q}(Y)$ is finite.

LEMMA 5.20. *For every $n \geq 8\mathfrak{v}(Y)$,*

$$|\mathbb{X}_n^*(Y,\mathcal{J})| = \frac{(n!)^4}{(n)_{\mathfrak{v}(Y)}} \sum_{H\in\mathcal{Q}(Y)} \frac{(n)_{\mathfrak{v}(H)}}{\prod_{f\in\{a,b,c,d\}} (n)_{\mathfrak{e}_f(H)}}. \tag{5.22}$$

Proof. By the discussion above it is enough to show that for every $H \in \mathcal{Q}(Y)$ and $n \geq 8\mathfrak{v}(Y)$, the number of Schreier graphs $\mathcal{G}$ on n vertices where the image of $\hat{Y}$ is H, is precisely

$$\frac{(n!)^4}{(n)_{\mathfrak{v}(Y)}} \cdot \frac{(n)_{\mathfrak{v}(H)}}{\prod_{f\in\{a,b,c,d\}} (n)_{\mathfrak{e}_f(H)}}.$$

First, note that $\mathfrak{v}(H) \leq \mathfrak{v}\left(\hat{Y}\right) \leq 8\mathfrak{v}(Y)$, so under the assumption that $n \geq 8\mathfrak{v}(Y)$, H can indeed be embedded in Schreier graphs on n vertices. The number of possible labelings of the vertices of H, which must extend the labeling of the vertices of $Y^{(1)}$, is

$$(n-\mathfrak{v}(Y))(n-\mathfrak{v}(Y)-1)\cdots(n-\mathfrak{v}(H)+1) = \frac{(n)_{\mathfrak{v}(H)}}{(n)_{\mathfrak{v}(Y)}}.$$

There are exactly $\mathfrak{e}_a(H)$ constraints on the permutation α_a for it to agree with the data in the vertex-labeled H, so there are $(n-\mathfrak{e}_a(H))! = \frac{n!}{(n)_{\mathfrak{e}_a(H)}}$ such permutations. The same logic applied to the other letters gives the required result. □

Combining Lemma 5.20 with Proposition 5.13 gives the following corollary.

COROLLARY 5.21. *For $n \geq 8\mathfrak{v}(Y)$ we have*

$$\Xi_n^{\nu=(n-\mathfrak{v})}(Y) = \frac{\prod_{f\in\{a,b,c,d\}}(n)_{\mathfrak{e}_f(Y)}}{(n)_{\mathfrak{f}(Y)}} \sum_{H\in\mathcal{Q}(Y)} \frac{(n)_{\mathfrak{v}(H)}}{\prod_{f\in\{a,b,c,d\}}(n)_{\mathfrak{e}_f(H)}}.$$

In particular, if Y is fixed and $n \to \infty$, we have

$$\Xi_n^{\nu=(n-\mathfrak{v})}(Y) = \sum_{H\in\mathcal{Q}(Y)} n^{\mathfrak{e}(Y)-\mathfrak{f}(Y)+\chi(H)}\left(1+O_Y\left(\frac{1}{n}\right)\right). \tag{5.23}$$

Proof. The first statement follows directly from Lemma 5.20 and Proposition 5.13. To obtain the second statement from the first, we use that all Pochammer symbols $(n)_q$ appearing therein have q bounded depending on Y and hence $(n)_q = n^q + O_Y(n^{q-1})$. □

Note that in the construction of $\hat{Y}$ from $Y^{(1)}$, we add a "handle" (a sequence of edges) to the graph for every vertex of Y that does not admit a closed cycle spelling $[a,b]\,[c,d]$. Hence the Euler characteristic of $\hat{Y}$ is equal to that of $Y^{(1)}$ minus the number of such vertices in Y. If Y has an octagon attached along every closed cycle spelling $[a,b]\,[c,d]$, there are $\mathfrak{v}\,(Y)-\mathfrak{f}\,(Y)$ such vertices, so

$$\chi\left(\hat{Y}\right)=\chi\left(Y^{(1)}\right)-\left(\mathfrak{v}\,(Y)-\mathfrak{f}\,(Y)\right)=\mathfrak{f}\,(Y)-\mathfrak{e}\,(Y)\,. \tag{5.24}$$

In particular, this is the case when Y is (strongly) boundary reduced. This is important because of the role of $\chi\,(H)$ in (5.23) for $H\in\mathcal{Q}(Y)$. It turns out that when Y is strongly boundary reduced, $\hat{Y}$ has Euler characteristic strictly larger than all other graphs in $\mathcal{Q}\,(Y)$:

Lemma 5.22. *If Y is strongly boundary reduced, then for every $H\in\mathcal{Q}(Y)\backslash\{\hat{Y}\}$,*

$$\chi\,(H)<\chi\left(\hat{Y}\right).$$

Proof. We use [MP20, Prop. 5.26] that states that if Y is strongly boundary reduced, then as $n\to\infty$,

$$\Xi_n(Y)=2+O_Y\left(n^{-1}\right). \tag{5.25}$$

When Y is fixed and $n\to\infty$, it follows from Lemmas 5.9(1) and 5.10 that

$$\Xi_n(Y)=2\Xi_n^{(n-\mathfrak{v})}(Y)+O_Y\left(n^{-2}\right). \tag{5.26}$$

Combining (5.25) and (5.26) gives

$$\Xi_n^{(n-\mathfrak{v})}(Y)=1+O_Y\left(n^{-1}\right). \tag{5.27}$$

Comparing (5.27) with (5.23) shows that there is exactly one $H\in\mathcal{Q}(Y)$ with $\chi\,(H)=\mathfrak{f}(Y)-\mathfrak{e}(Y)$, and all remaining graphs in $\mathcal{Q}(Y)$ have strictly smaller Euler characteristic. Finally, (5.24) shows this H must be $\hat{Y}$ itself. □

5.9 Bounds on $\mathbb{E}_n^{\mathrm{emb}}(Y)$ for ε-adapted Y. In this section we give the final implications of the previous sections for $\mathbb{E}_n^{\mathrm{emb}}(Y)$ for ε-adapted Y. Recall the definition of $\mathcal{Q}\,(Y)$ from Section 5.8. We will need the following easy bound for Pochhammer symbols.

Lemma 5.23. *Let $n\in\mathbf{N}$ and $q\in\mathbf{N}\cup\{0\}$ with $q\le\frac{1}{2}n$. Then*

$$n^q\left(1-\frac{q^2}{n}\right)\le n^q\exp\left(\frac{-q^2}{n}\right)\le(n)_q\le n^q.$$

Proof. The first inequality is based on $1-x\le e^{-x}$. The second one is based on writing $(n)_q=n^q\left(1-\frac{1}{n}\right)\cdots\left(1-\frac{q-1}{n}\right)$ and using $e^{-2x}\le 1-x$ which holds for $x\in\left[0,\frac{1}{2}\right]$. The third inequality is obvious. □

PROPOSITION 5.24. *Let $\varepsilon \in (0,1)$ and $\eta = \eta(\varepsilon) \in (0, \frac{1}{100})$ be the parameter provided by Proposition 5.12 for this ε. Let $n \in \mathbf{N}$ and $M = M(n) \geq 1$. Let Y be ε-adapted with $\mathfrak{D}(Y) \leq n^{\eta}$ and $\mathfrak{v}(Y), \mathfrak{e}(Y), \mathfrak{f}(Y) \leq M \leq n^{1/4}$. Then*

$$\frac{\mathbb{E}_n^{\mathrm{emb}}(Y)}{n^{\chi(Y)}} = \left(1 + O_{\varepsilon}\left(\frac{M^2}{n}\right)\right)\left(1 + \sum_{H \in \mathcal{Q}(Y)\setminus\{\hat{Y}\}} n^{\chi(H)+\mathfrak{e}(Y)-\mathfrak{f}(Y)}\right). \tag{5.28}$$

Proof. Assume all parameters are as in the statement of the proposition. By Theorem 5.1 and Proposition 5.12 we have

$$\frac{\mathbb{E}_n^{\mathrm{emb}}(Y)}{n^{\chi(Y)}} = \frac{(n!)^3}{|\mathbb{X}_n|} \cdot \frac{(n)_{\mathfrak{v}(Y)}(n)_{\mathfrak{f}(Y)}}{\prod_f (n)_{\mathfrak{e}_f(Y)} n^{\chi(Y)}}\left[2\Xi_n^{\nu=(n-\mathfrak{v})}(Y) + O_{\varepsilon}\left(\frac{1}{n}\right)\right].$$

By Lemma 5.23, $\frac{(n)_{\mathfrak{v}(Y)}(n)_{\mathfrak{f}(Y)}}{\prod_f (n)_{\mathfrak{e}_f(Y)} n^{\chi(Y)}} = 1 + O\left(\frac{M^2}{n}\right)$. By Corollary 4.5, $\frac{(n!)^3}{|\mathbb{X}_n|} = \frac{1}{2} + O\left(\frac{1}{n^2}\right)$. With Corollary 5.21, this gives

$$\begin{aligned}\frac{\mathbb{E}_n^{\mathrm{emb}}(Y)}{n^{\chi(Y)}} &= \left[\frac{1}{2} + O\left(\frac{M^2}{n}\right)\right]\left[2\frac{\prod_f (n)_{\mathfrak{e}_f(Y)}}{(n)_{\mathfrak{f}(Y)}} \sum_{H \in \mathcal{Q}(Y)} \frac{(n)_{\mathfrak{v}(H)}}{\prod_f (n)_{\mathfrak{e}_f(H)}}\right] + O_{\varepsilon}\left(\frac{1}{n}\right) \\ &\overset{\text{Lem. 5.23}}{=} \left[1 + O\left(\frac{M^2}{n}\right)\right] \sum_{H \in \mathcal{Q}(Y)} n^{\mathfrak{e}(Y)-\mathfrak{f}(Y)+\chi(H)} + O_{\varepsilon}\left(\frac{1}{n}\right),\end{aligned} \tag{5.29}$$

where the use of Lemma 5.23 is justified since for every $H \in \mathcal{Q}(Y)$, $\mathfrak{v}(H) \leq \mathfrak{v}(\hat{Y}) \leq 8\mathfrak{v}(Y) \leq 8M$, and $\mathfrak{e}(H) \leq \mathfrak{e}(\hat{Y}) \leq \mathfrak{e}(Y) + 8\mathfrak{v}(Y) \leq 9M$. In the summation in (5.29), the top power of n is realized by $\hat{Y}$ and is equal to zero (by (5.24) and Lemma 5.22), so we obtain

$$\frac{\mathbb{E}_n^{\mathrm{emb}}(Y)}{n^{\chi(Y)}} = \left[1 + O\left(\frac{M^2}{n}\right)\right]\left(1 + \sum_{H \in \mathcal{Q}(Y)\setminus\{\hat{Y}\}} n^{\chi(H)+\mathfrak{e}(Y)-\mathfrak{f}(Y)}\right) + O_{\varepsilon}\left(\frac{1}{n}\right),$$

which yields (5.28). □

The drawback of Proposition 5.24 is that we do not know how to directly estimate the sum over $H \in \mathcal{Q}(Y)\setminus\{\hat{Y}\}$ that appears therein. Because we can not directly deal with this sum, we instead use Proposition 5.24 to deduce in the remaining results of this section that for ε-adapted Y we can control $\mathbb{E}_n^{\mathrm{emb}}(Y)$ using $\mathbb{E}_m^{\mathrm{emb}}(Y)$ with m much smaller than n.

COROLLARY 5.25. *Let $\varepsilon \in (0,1)$, and $\eta = \eta(\varepsilon) \in (0, \frac{1}{100})$ be the parameter provided by Proposition 5.12 for this ε. Let $m \in \mathbf{N}$. Let Y be ε-adapted with $\mathfrak{D}(Y) \leq m^{\eta}$ and $\mathfrak{v}(Y), \mathfrak{e}(Y), \mathfrak{f}(Y) \leq m^{1/4}$. Then*

$$\frac{\mathbb{E}_m^{\mathrm{emb}}(Y)}{m^{\chi(Y)}} \gg_{\varepsilon} 1 + \sum_{H \in \mathcal{Q}(Y)\setminus\{\hat{Y}\}} m^{\chi(H)+\mathfrak{e}(Y)-\mathfrak{f}(Y)}.$$

In particular, $\mathbb{E}_m^{\mathrm{emb}}(Y) \gg_{\epsilon} m^{\chi(Y)}$.

Remark 5.26. As a direct consequence of Corollary 5.25, under the same conditions, if $\chi(Y) < 0$ and $n \geq m$ we obtain

$$\frac{m}{n}\mathbb{E}_m^{\mathrm{emb}}(Y) \gg \frac{m^{\chi(Y)+1}}{n} \gg n^{\chi(Y)}.$$

While Corollary 5.25 is a direct consequence of Proposition 5.24, we can get more information by combining Proposition 5.24 with what we already know about $\mathcal{Q}(Y)$.

Proposition 5.27. *Let $\varepsilon \in (0,1)$, η be as in Proposition 5.12 and $K > 1$. Let $n \in \mathbf{N}$ and $m = m(n) \in \mathbf{N}$ with $m < n$ and $m \overset{n\to\infty}{\to} \infty$. Let Y be ε-adapted and suppose that $\mathfrak{v}(Y), \mathfrak{e}(Y), \mathfrak{f}(Y) \leq (K\log n)^2 \leq m^{1/4}$ and that $\mathfrak{D}(Y) \leq K\log n \leq m^{\eta}$. Then*

$$\frac{\mathbb{E}_n^{\mathrm{emb}}(Y)}{n^{\chi(Y)}} = 1 + O_{\varepsilon,K}\left(\frac{(\log n)^4}{n}\right) + O_{\varepsilon,K}\left(\frac{m}{n}\frac{\mathbb{E}_m^{\mathrm{emb}}(Y)}{m^{\chi(Y)}}\right). \tag{5.30}$$

Proof. With assumptions as in the proposition, Proposition 5.24 gives

$$\begin{aligned}\frac{\mathbb{E}_n^{\mathrm{emb}}(Y)}{n^{\chi(Y)}} &= \left(1 + O_{\varepsilon,K}\left(\frac{(\log n)^4}{n}\right)\right)\left(1 + \sum_{H\in\mathcal{Q}(Y)\setminus\{\hat{Y}\}} n^{\chi(H)+\mathfrak{e}(Y)-\mathfrak{f}(Y)}\right)\\ &= 1 + O_{\varepsilon,K}\left(\frac{(\log n)^4}{n}\right) + O_{\varepsilon,K}\left(\sum_{H\in\mathcal{Q}(Y)\setminus\{\hat{Y}\}} n^{\chi(H)+\mathfrak{e}(Y)-\mathfrak{f}(Y)}\right).\end{aligned}$$

Finally, because for every $H \in \mathcal{Q}(Y)\setminus\{\hat{Y}\}$ we have $\chi(H) + \mathfrak{e}(Y) - \mathfrak{f}(Y) \leq -1$ and $m < n$,

$$\begin{aligned}\sum_{H\in\mathcal{Q}(Y)\setminus\{\hat{Y}\}} n^{\chi(H)+\mathfrak{e}(Y)-\mathfrak{f}(Y)} &= \sum_{H\in\mathcal{Q}(Y)\setminus\{\hat{Y}\}} \left(\frac{n}{m}\right)^{\chi(H)+\mathfrak{e}(Y)-\mathfrak{f}(Y)} m^{\chi(H)+\mathfrak{e}(Y)-\mathfrak{f}(Y)}\\ &\leq \frac{m}{n}\sum_{H\in\mathcal{Q}(Y)\setminus\{\hat{Y}\}} m^{\chi(H)+\mathfrak{e}(Y)-\mathfrak{f}(Y)} \overset{\text{Cor. 5.25}}{\ll_{\varepsilon}} \frac{m}{n}\frac{\mathbb{E}_m^{\mathrm{emb}}(Y)}{m^{\chi(Y)}},\end{aligned}$$

concluding the proof of the proposition. □

6 Proof of Theorem 1.11

The reader is suggested to have read the overview in Section 1.3 before attempting to read this section of the paper.

6.1 Setup. We remind the reader that $g = 2$. We are given $c > 0$, and an element $\gamma \in \Gamma$ of cyclic word length $\ell_w(\gamma) \leq c \log n$. We assume that γ is not a proper power of another element of Γ. We remind the reader that $\mathcal{C}_\gamma$ is an annular tiled surface associated to γ as in Example 3.5. By Lemma 3.20,

$$\mathbb{E}_n\left[\mathsf{fix}_\gamma\right] = \mathbb{E}_n\left(\mathcal{C}_\gamma\right),$$

where $\mathbb{E}_n(\mathcal{C}_\gamma)$ is the expected number of morphisms from $\mathcal{C}_\gamma$ to the random surface X_ϕ. Let $\varepsilon = \frac{1}{32}$ (for general g, $\varepsilon = \frac{1}{16g}$) and let $\mathcal{R}_\varepsilon(\mathcal{C}_\gamma)$ be the finite resolution of $\mathcal{C}_\gamma$ provided by Definition 3.23 and Theorem 3.24. Each element of this resolution is a morphism $h : \mathcal{C}_\gamma \to W_h$ where W_h is a tiled surface. By Lemma 3.22 we have for any $n \geq 1$

$$\mathbb{E}_n\left[\mathsf{fix}_\gamma\right] = \sum_{h \in \mathcal{R}_\varepsilon(\mathcal{C}_\gamma)} \mathbb{E}_n^{\mathrm{emb}}\left(W_h\right), \tag{6.1}$$

where $\mathbb{E}_n^{\mathrm{emb}}\left(W_h\right)$ is the expected number of embeddings of W_h into the random tiled surface X_ϕ. Associated to each W_h here, $\mathfrak{v}(W_h), \mathfrak{e}(W_h)$, and $\mathfrak{f}(W_h)$ are the number of vertices, edges, and faces of W_h. Also associated to W_h are $\mathfrak{d}(W_h)$, the number of edges in the boundary of W_h, $\chi(W_h)$, the topological Euler characteristic of W_h, and $\mathfrak{D}(W_h) = \mathfrak{v}(W_h) - \mathfrak{f}(W_h)$.

By Corollary 3.25, there is a constant $K = K(c) > 0$, such that for each $h \in \mathcal{R}_\varepsilon(\mathcal{C}_\gamma)$, and for $n \geq 3$, we have

$$\begin{aligned} \mathfrak{d}(W_h) &\leq K \log n, \\ \mathfrak{f}(W_h) &\leq K (\log n)^2. \end{aligned}$$

By Lemma 3.6 we have $\mathfrak{v}(W_h) \leq \mathfrak{d}(W_h) + \mathfrak{f}(W_h)$, so $\mathfrak{D}\left(W_h\right) \leq \mathfrak{d}\left(W_h\right)$ and

$$\mathfrak{D}(W_h) \leq K \log n.$$

We also have $\mathfrak{e}(W_h) \leq 4\mathfrak{v}(W_h)$ by (3.2). Hence by increasing K if necessary we can also ensure

$$\mathfrak{v}(W_h), \mathfrak{e}(W_h) \leq K (\log n)^2.$$

6.2 Part I: the contribution from non-ε-adapted surfaces. Our first goal is to control the contribution to $\mathbb{E}_n[\mathsf{fix}_\gamma]$ in (6.1) from non-ε-adapted surfaces. Let $\mathcal{R}_\varepsilon{}^{(\text{non-}\varepsilon\text{-ad})}\left(\mathcal{C}_\gamma\right)$ denote the set of morphisms $h : \mathcal{C}_\gamma \to W_h$ in $\mathcal{R}_\varepsilon(\mathcal{C}_\gamma)$ such that W_h is not ε-adapted. In particular, such W_h is boundary reduced and $\mathfrak{f}\left(W_h\right) > \mathfrak{d}\left(W_h\right)$.

Proposition 6.1. *There is a constant $A > 0$ such that for any $c > 0$, if $\ell_w(\gamma) \leq c \log n$, then*

$$\sum_{h \in \mathcal{R}_\varepsilon{}^{(\text{non-}\varepsilon\text{-ad})}(\mathcal{C}_\gamma)} \mathbb{E}_n^{\mathrm{emb}}\left(W_h\right) \ll_c \frac{(\log n)^A}{n}.$$

Proof. We first do some counting. Let us count $h \in \mathcal{R}_\varepsilon^{(\text{non-}\varepsilon\text{-ad})}(\mathcal{C}_\gamma)$ by their value of $\mathfrak{D}(W_h)$ and $\mathfrak{f}(W_h)$. By Corollary 3.25 every $h \in \mathcal{R}_\varepsilon^{(\text{non-}\varepsilon\text{-ad})}(\mathcal{C}_\gamma)$ has

$$\chi(W_h) < -\mathfrak{f}(W_h) < -\mathfrak{d}(W_h). \tag{6.2}$$

Combining (6.2) with Lemma 3.6 yields

$$0 \le \mathfrak{D}(W_h) \le \mathfrak{d}(W_h) < \mathfrak{f}(W_h). \tag{6.3}$$

Notice that (6.3) implies $\mathfrak{f}(W_h) \ge 1$. First we bound the number of possible W_h with $\mathfrak{D}(W_h) = \mathfrak{D}_0$ and $\mathfrak{f}(W_h) = \mathfrak{f}_0$ for fixed $\mathfrak{D}_0 < \mathfrak{f}_0$. Note that in this case $\mathfrak{v}(W_h) = \mathfrak{v}_0 \overset{\text{def}}{=} \mathfrak{D}_0 + \mathfrak{f}_0$. We may over-count the number of W_h with $\mathfrak{v}_0$ vertices by counting the number of W_h together with a labeling of their vertices by $[\mathfrak{v}_0]$. We first construct the one-skeleton of such a tiled surface: there are at most ${\mathfrak{v}_0}^{\mathfrak{v}_0}$ choices for the a-labeled edges, and also for the b-labeled edges etc. Because W_h are all boundary reduced, there is an octagon attached to any closed $[a,b]\,[c,d]$ path, so the one-skeleton completely determines the entire tiled surface. Hence there are at most ${\mathfrak{v}_0}^{4\mathfrak{v}_0}$ choices for W_h with $\mathfrak{v}(W_h) = \mathfrak{v}_0$.

We also have to estimate how many ways there are to map $\mathcal{C}_\gamma$ into such a W_h. Fixing arbitrarily a vertex v of $\mathcal{C}_\gamma$, any morphism $\mathcal{C}_\gamma \to W_h$ is uniquely determined by where v goes; hence there are at most $\mathfrak{v}_0$ morphisms and so in total there are at most

$$\mathfrak{v}_0^{4\mathfrak{v}_0+1} \le \mathfrak{v}_0^{5\mathfrak{v}_0} = (\mathfrak{D}_0 + \mathfrak{f}_0)^{5(\mathfrak{D}_0+\mathfrak{f}_0)} \le (2\mathfrak{f}_0)^{10\mathfrak{f}_0}$$

elements $h \in \mathcal{R}_\varepsilon^{(\text{non-}\varepsilon\text{-ad})}(\mathcal{C}_\gamma)$ with $\mathfrak{D}(W_h) = \mathfrak{D}_0$ and $\mathfrak{f}(W_h) = \mathfrak{f}_0$. Hence there are at most $K \log n \cdot (2\mathfrak{f}_0)^{10\mathfrak{f}_0}$ elements $h \in \mathcal{R}_\varepsilon^{(\text{non-}\varepsilon\text{-ad})}(\mathcal{C}_\gamma)$ with $\mathfrak{f}(W_h) = \mathfrak{f}_0$.

We are going to use Theorem 5.1 that relates $\mathbb{E}_n^{\text{emb}}(W_h)$ to a certain quantity $\Xi_n(W_h)$. By Proposition 5.11 there is $A_0 > 1$ such that for $h \in \mathcal{R}_\varepsilon^{(\text{non-}\varepsilon\text{-ad})}(\mathcal{C}_\gamma)$

$$|\Xi_n(W_h)| \ll_K (A_0\mathfrak{D}(W_h))^{A_0\mathfrak{D}(W_h)} \le (A_0\mathfrak{f}(W_h))^{A_0\mathfrak{f}(W_h)}, \tag{6.4}$$

so by Theorem 5.1, Corollary 4.5, and Lemma 5.23 we get

$$\begin{aligned}\mathbb{E}_n^{\text{emb}}(W_h) &\overset{\text{Thm 5.1}}{=} \frac{n!^3}{|\mathbb{X}_n|}\frac{(n)_{\mathfrak{v}(W_h)}(n)_{\mathfrak{f}(W_h)}}{\prod_f (n)_{\mathfrak{e}_f(W_h)}}\Xi_n(W_h) \overset{\text{Cor. 4.5}}{\ll} \frac{(n)_{\mathfrak{v}(W_h)}(n)_{\mathfrak{f}(W_h)}}{\prod_f (n)_{\mathfrak{e}_f(W_h)}}\Xi_n(W_h)\\ &\overset{\text{Lemma 5.23}}{\ll_K} n^{\chi(W_h)}\Xi_n(W_h) \overset{(6.4)}{\ll_K} n^{\chi(W_h)}(A_0\mathfrak{f}(W_h))^{A_0\mathfrak{f}(W_h)}.\end{aligned} \tag{6.5}$$

Therefore, for every $1 \le f_0 \le K(\log n)^2$,

$$\begin{aligned}\sum_{\substack{h\in\mathcal{R}_\varepsilon^{(\text{non-}\varepsilon\text{-ad})}(\mathcal{C}_\gamma)\\ \mathfrak{f}(W_h)=\mathfrak{f}_0}} \mathbb{E}_n^{\text{emb}}(W_h) &\ll_K (A_0\mathfrak{f}_0)^{A_0\mathfrak{f}_0} \sum_{\substack{h\in\mathcal{R}_\varepsilon^{(\text{non-}\varepsilon\text{-ad})}(\mathcal{C}_\gamma)\\ \mathfrak{f}(W_h)=\mathfrak{f}_0}} n^{\chi(W_h)}\\ &\overset{(6.2)}{\le} (A_0\mathfrak{f}_0)^{A_0\mathfrak{f}_0} \sum_{\substack{h\in\mathcal{R}_\varepsilon^{(\text{non-}\varepsilon\text{-ad})}(\mathcal{C}_\gamma)\\ \mathfrak{f}(W_h)=\mathfrak{f}_0}} n^{-\mathfrak{f}_0}\end{aligned}$$

$$\leq K\log n\left(\frac{(A_0\mathfrak{f}_0)^{A_0}\,(2\mathfrak{f}_0)^{10}}{n}\right)^{\mathfrak{f}_0}$$
$$\leq K\log n\cdot\left(\frac{A_0^{A_0}2^{10}\left(K(\log n)^2\right)^{A_0+10}}{n}\right)^{\mathfrak{f}_0}.$$

So

$$\sum_{h\in\mathcal{R}_\varepsilon{}^{(\text{non-}\varepsilon\text{-ad})}(\mathcal{C}_\gamma)}\mathbb{E}_n^{\text{emb}}\left(W_h\right)=\sum_{\mathfrak{f}_0=1}^{K(\log n)^2}\sum_{\substack{h\in\mathcal{R}_\varepsilon{}^{(\text{non-}\varepsilon\text{-ad})}(\mathcal{C}_\gamma)\\ \mathfrak{f}(W_h)=\mathfrak{f}_0}}\mathbb{E}_n^{\text{emb}}\left(W_h\right)$$
$$\ll_K K\log n\cdot\sum_{\mathfrak{f}_0=1}^{K(\log n)^2}\left(\frac{A_0^{A_0}2^{10}\left(K(\log n)^2\right)^{A_0+10}}{n}\right)^{\mathfrak{f}_0}$$
$$\ll_K\frac{(\log n)^{2A_0+21}}{n},$$

where the last inequality is based on that $\frac{A_0^{A_0}2^{10}(K(\log n)^2)^{A_0+10}}{n}\leq\frac{1}{2}$ for $n\gg_K 1$. □

6.3 Part II: the contribution from ε-adapted surfaces. Write $\mathcal{R}_\varepsilon{}^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)\subset\mathcal{R}_\varepsilon(\mathcal{C}_\gamma)$ for the collection of morphisms $h:\mathcal{C}_\gamma\to W_h$ in $\mathcal{R}_\varepsilon(\mathcal{C}_\gamma)$ such that W_h is ε-adapted. In light of Proposition 6.1 it remains to deal with the contributions to $\mathbb{E}_n[\mathsf{fix}_\gamma]$ from $\mathcal{R}_\varepsilon{}^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)$. Indeed we have by Proposition 6.1 and (6.1)

$$\mathbb{E}_n[\mathsf{fix}_\gamma]=\sum_{h\in\mathcal{R}_\varepsilon{}^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)}\mathbb{E}_n^{\text{emb}}\left(W_h\right)+O_c\left(\frac{(\log n)^A}{n}\right).\tag{6.6}$$

Recall that if W_h is ε-adapted, it is, in particular, strongly boundary reduced, and so by [MP20, Section 1.6], $\mathbb{E}_n^{\text{emb}}\left(W_h\right)=n^{\chi(W_h)}\left[1+O\left(n^{-1}\right)\right]$ as $n\to\infty$. By Theorem 1.10, $\mathbb{E}_n[\mathsf{fix}_\gamma]=1+O\left(n^{-1}\right)$. Comparing this with (6.6), we conclude that there is exactly one $h_0\in\mathcal{R}_\varepsilon(\mathcal{C}_\gamma)$ with $\chi\left(W_{h_0}\right)=0$. This h_0 also satisfies that W_{h_0} is ε-adapted.[17]

Still, we are missing some information about $\mathcal{R}_\varepsilon{}^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)$ that we will need: for example, the ability to count how many $h:\mathcal{C}_\gamma\to W_h$ there are in $\mathcal{R}_\varepsilon{}^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)$ with different orders of contributions (i.e. $n^{\chi(W_h)}$) to (6.6). We are going to use a trick to get around this missing information.

Let $\eta\in(0,\frac{1}{100})$ be the parameter provided by Proposition 5.12 for the current $\varepsilon=\frac{1}{32}$ (the reason for choosing η like this now is just so that we can momentarily apply Corollary 5.25 and Proposition 5.27). Let m be an auxiliary parameter given by

$$m=\left\lceil(K\log n)^{1/\eta}\right\rceil$$

[17] It can be shown that h_0 is the result of the OvB algorithm when applied to the embedding $\mathcal{C}_\gamma\hookrightarrow\langle\gamma\rangle\backslash\widetilde{\Sigma_2}$ with $\widetilde{\Sigma_2}$ the universal cover of Σ_2—see [MP20, Section 2].

so that when $n \gg_c 1$, for all $h \in \mathcal{R}_\varepsilon^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)$, $\mathfrak{D}(W_h) \leq K \log n \leq m^\eta$ and $\mathfrak{v}(W_h), \mathfrak{e}(W_h), \mathfrak{f}(W_h) \leq K(\log n)^2 \leq m^{1/4}$. Moreover, $(\log n)^{100} \ll_c m \ll_c (\log n)^{\frac{1}{\eta}}$. To exploit the fact that each $\mathbb{E}_n^{\text{emb}}(W_h)$ is controlled by $\mathbb{E}_m^{\text{emb}}(W_h)$ (Corollary 5.25 and Proposition 5.27), we will at two points use the inequality

$$m \geq \mathbb{E}_m[\mathsf{fix}_\gamma] \overset{(6.1)}{=} \sum_{h \in \mathcal{R}_\varepsilon(\mathcal{C}_\gamma)} \mathbb{E}_m^{\text{emb}}(W_h) \geq \sum_{h \in \mathcal{R}_\varepsilon^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)} \mathbb{E}_m^{\text{emb}}(W_h). \tag{6.7}$$

We begin with

$$\begin{aligned}
\sum_{h \in \mathcal{R}_\varepsilon^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)} \mathbb{E}_n^{\text{emb}}(W_h) &\overset{\text{Prop. 5.27}}{=} \sum_{h \in \mathcal{R}_\varepsilon^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)} n^{\chi(W_h)} \left[1 + O_c\left(\frac{(\log n)^4}{n}\right)\right. \\
&\qquad \left. + O_c\left(\frac{m}{n} \frac{\mathbb{E}_m^{\text{emb}}(W_h)}{m^{\chi(W_h)}}\right)\right] \\
&= \sum_{h \in \mathcal{R}_\varepsilon^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)} n^{\chi(W_h)} \left[1 + O_c\left(\frac{(\log n)^4}{n}\right)\right] \\
&\quad + O_c\left(\frac{m}{n} \sum_{h \in \mathcal{R}_\varepsilon^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)} \mathbb{E}_m^{\text{emb}}(W_h)\right) \\
&\overset{(6.7)}{=} \sum_{h \in \mathcal{R}_\varepsilon^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)} n^{\chi(W_h)} \left(1 + O_c\left(\frac{(\log n)^4}{n}\right)\right) + O_c\left(\frac{m^2}{n}\right).
\end{aligned} \tag{6.8}$$

The middle estimate above used that $\chi(W_h) \leq 0$ for all $h \in \mathcal{R}_\varepsilon^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)$, and so $\left(\frac{n}{m}\right)^{\chi(W_h)} \leq 1$. The contribution to (6.8) from h_0 is $1 + O_c\left(\frac{(\log n)^4}{n}\right)$. So we obtain

$$\sum_{h \in \mathcal{R}_\varepsilon^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)} \mathbb{E}_n^{\text{emb}}(W_h) = 1 + O_c\left(\frac{(\log n)^4}{n}\right) \tag{6.9}$$

$$+ O\left(\frac{m^2}{n}\right) + O\left(\sum_{\substack{h \in \mathcal{R}_\varepsilon^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma) \\ \chi(W_h) < 0}} n^{\chi(W_h)}\right). \tag{6.10}$$

To deal with the last error term, we relate it to the expectations at level m. Indeed,

$$\left|\sum_{\substack{h\in\mathcal{R}_\varepsilon{}^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)\\ \chi(W_h)<0}} n^{\chi(W_h)}\right| = \sum_{\substack{h\in\mathcal{R}_\varepsilon{}^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)\\ \chi(W_h)<0}} \left(\frac{n}{m}\right)^{\chi(W_h)} m^{\chi(W_h)} \leq \frac{m}{n} \sum_{h\in\mathcal{R}_\varepsilon{}^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)} m^{\chi(W_h)}$$
$$\overset{\text{Cor. 5.25}}{\ll} \frac{m}{n} \sum_{h\in\mathcal{R}_\varepsilon{}^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)} \mathbb{E}_m^{\text{emb}}(W_h) \overset{(6.7)}{\leq} \frac{m^2}{n}.$$

Incorporating this estimate into (6.10) gives

$$\sum_{h\in\mathcal{R}_\varepsilon{}^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)} \mathbb{E}_n^{\text{emb}}(W_h) = 1 + O_c\left(\frac{(\log n)^4}{n}\right) + O\left(\frac{m^2}{n}\right) = 1 + O_c\left(\frac{(\log n)^A}{n}\right),$$

where $A = \frac{2}{\eta}$. Combining this with (6.6) and increasing A if necessary we obtain

$$\mathbb{E}_n[\mathsf{fix}_\gamma] = 1 + O_c\left(\frac{(\log n)^A}{n}\right)$$

as required. This concludes the proof of Theorem 1.11. □

Remark 6.2. The arguments above show that

$$\sum_{h\in\mathcal{R}_\varepsilon{}^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)} m^{\chi(W_h)} \ll \sum_{h\in\mathcal{R}_\varepsilon{}^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)} \mathbb{E}_m^{\text{emb}}(W_h) \ll m,$$

hence the number of elements of $h \in \mathcal{R}_\varepsilon{}^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma)$ with $\chi(W_h) = \chi$ is $\ll m^{1-\chi}$. In general, given arbitrary $\gamma \in \Gamma$, and $\epsilon > 0$ we obtain by the same argument that for some $\eta = \eta(\epsilon) > 0$ we have

$$\#\{\, h \in \mathcal{R}_\varepsilon{}^{(\varepsilon\text{-ad})}(\mathcal{C}_\gamma) \,:\, \chi(W_h) = \chi \,\} \ll_\epsilon \left(\ell(\gamma)^{\frac{1}{\eta}}\right)^{1-\chi}.$$

We mention this side-effect of our proof in case it is of independent interest.

Acknowledgments

We thank Amitay Kamber for the observation that our methods prove Theorem 1.7. We thank the anonymous referees for their careful readings and insightful comments. Frédéric Naud is supported by Institut Universitaire de France. Doron Puder was supported by the Israel Science Foundation: ISF Grant 1071/16. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant Agreement No. 850956 and Grant Agreement No. 949143).

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

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

References

[Alo86] N. Alon. Eigenvalues and expanders. *Combinatorica*, (2)6 (1986), 83–96. Theory of computing (Singer Island, Fla., 1984).

[AM85] N. Alon and V.D. Milman. λ_1, isoperimetric inequalities for graphs, and superconcentrators. *J. Combin. Theory Ser. B*, (1)38 (1985), 73–88.

[Apo76] T.M. Apostol. *Introduction to Analytic Number Theory. Undergraduate Texts in Mathematics*. Springer-Verlag, New York-Heidelberg (1976).

[BBD88] P. Buser, M. Burger, and J. Dodziuk. *Riemann Surfaces of Large Genus and Large λ_1, Geometry and Analysis on Manifolds (Katata/Kyoto, 1987), Lecture Notes in Mathematics*, Vol. 1339. Springer, Berlin (1988), pp. 54–63.

[BC19] C. Bordenave and B. Collins. Eigenvalues of random lifts and polynomials of random permutation matrices. *Ann. of Math. (2)*, (3)190 (2019), 811–875.

[BCP21] T. Budzinski, N. Curien, and B. Petri. The diameter of random Belyĭ surfaces. *Algebraic & Geometric Topology*, (6)21 (2021), 2929–2957.

[Bea84] A.F. Beardon. *A Primer on Riemann Surfaces. London Mathematical Society Lecture Note Series*, Vol. 78. Cambridge University Press, Cambridge (1984).

[BH99] M.R. Bridson and A. Haefliger. *Metric Spaces of Non-positive Curvature. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*, Vol. 319. Springer-Verlag, Berlin (1999).

[BK93] Ya. M. Barzdin and A.N. Kolmogorov. *On the Realization of Networks in Three-Dimensional Space*. Springer Netherlands, Dordrecht (1993), pp. 194–202.

[BM04] R. Brooks and E. Makover. Random construction of Riemann surfaces. *J. Differential Geom.*, (1)68 (2004), 121–157.

[Bol88] B. Bollobás. The isoperimetric number of random regular graphs. *European J. Combin.*, (3)9 (1988), 241–244.

[Bor] C. Bordenave. A new proof of Friedman's second eigenvalue theorem and its extension to random lifts. *Annales Scientifiques de l'Ecole Normale Supérieure*, to appear, available at 1502.04482.

[Bor16] D. Borthwick. *Spectral Theory of Infinite-Area Hyperbolic Surfaces, Progress in Mathematics*, second ed., Vol. 318. Birkhäuser/Springer, [Cham] (2016).

[BS87a] J.S. Birman and C. Series. Dehn's algorithm revisited, with applications to simple curves on surfaces. In: *Combinatorial Group Theory and Topology (Alta, Utah, 1984), Annals of Mathematics Studies*, Vol. 111. Princeton University Press, Princeton, NJ (1987), pp. 451–478.

[BS87b] A. Broder and E. Shamir. On the second eigenvalue of random regular graphs. In: *The 28th Annual Symposium on Foundations of Computer Science* (1987), pp. 286–294.

[Bus10] P. Buser. *Geometry and Spectra of Compact Riemann Surfaces*, Modern Birkhäuser Classics (2010). Reprint of the 1992 edition.

[FP22] J. Friedman and D. Puder. A note on the trace method for random regular graphs. *Israel J. Math.*, (2022+), to appear, 2006.13605.

[Fri03] J. Friedman. Relative expanders or weakly relatively Ramanujan graphs. *Duke Math. J.*, (1)118 (2003), 19–35.

[Fri08] J. Friedman. A proof of Alon's second eigenvalue conjecture and related problems. *Mem. Amer. Math. Soc.*, (910)195 (2008), viii+100.

[Gam06] A. Gamburd. Poisson–Dirichlet distribution for random Belyi surfaces. *The Annals of Probability*, (5)34 (2006), 1827–1848.

[GJ78] S. Gelbart and H. Jacquet. A relation between automorphic representations of GL(2) and GL(3). *Ann. Sci. École Norm. Sup. (4)*, (4)11 (1978), 471–542.

[GK19] K. Golubev and A. Kamber. Cutoff on hyperbolic surfaces. *Geom. Dedicata*, 203 (2019), 225–255.

[Hej76] D.A. Hejhal. *The Selberg Trace Formula for* $\mathrm{PSL}(2,\mathbf{R})$. *Vol. I, Lecture Notes in Mathematics*, Vol. 548. Springer-Verlag, Berlin-New York (1976).

[Hej83] D.A. Hejhal. *The Selberg Trace Formula for* $\mathrm{PSL}(2,\mathbf{R})$. *Vol. 2, Lecture Notes in Mathematics*, Vol. 1001. Springer-Verlag, Berlin (1983).

[Hid21] W. Hide. *Spectral gap for Weil–Petersson random surfaces with cusps*, preprint, 2107.14555 (2021).

[HM21] W. Hide and M. Magee. *Near optimal spectral gaps for hyperbolic surfaces*, preprint, 2107.05292 (2021).

[Hub74] H. Huber. Über den ersten Eigenwert des Laplace-Operators auf kompakten Riemannschen Flächen. *Comment. Math. Helv.*, 49 (1974), 251–259.

[Hum18] P. Humphries. Density theorems for exceptional eigenvalues for congruence subgroups. *Algebra Number Theory*, (7)12 (2018), 1581–1610.

[Hur02] A. Hurwitz. Ueber die anzahl der Riemann'schen flächen mit gegebenen verzweigungspunkten. *Mathematische Annalen*, (1)55 (1902), 53–66.

[Hux86] M.N. Huxley. *Exceptional Eigenvalues and Congruence Subgroups, The Selberg Trace Formula and Related Topics (Brunswick, Maine, 1984), Contemporary Mathematics*, Vol. 53. American Mathematical Society, Providence, RI (1986), pp. 341–349.

[Iwa89] H. Iwaniec. *Selberg's Lower Bound of the First Eigenvalue for Congruence Groups, Number Theory, Trace Formulas and Discrete Groups (Oslo, 1987).* Academic Press, Boston, MA (1989), pp. 371–375.

[Iwa96] H. Iwaniec. The lowest eigenvalue for congruence groups. In: *Topics in Geometry, Progr. Nonlinear Differential Equations Appl.*, Vol. 20. Birkhäuser Boston, Boston, MA (1996), pp. 203–212.

[Iwa02] H. Iwaniec. *Spectral Methods of Automorphic Forms, Graduate Studies in Mathematics*, second ed., Vol. 53. American Mathematical Society, Providence, RI; Revista Matemática Iberoamericana, Madrid (2002).

[Jen84] F. Jenni. Über den ersten Eigenwert des Laplace-Operators auf ausgewählten Beispielen kompakter Riemannscher Flächen. *Comment. Math. Helv.*, (2)59 (1984), 193–203.

[JL70] H. Jacquet and R.P. Langlands. *Automorphic Forms on* GL(2), *Lecture Notes in Mathematics*, Vol. 114. Springer-Verlag, Berlin-New York (1970)

[Kim03] H.H. Kim. Functoriality for the exterior square of GL_4 and the symmetric fourth of GL_2. *J. Amer. Math. Soc.*, (1)16 (2003), 139–183. With appendix 1 by D. Ramakrishnan and appendix 2 by H. H. Kim and P. Sarnak

[KS02] H.H. Kim and F. Shahidi. Functorial products for $\mathrm{GL}_2 \times \mathrm{GL}_3$ and the symmetric cube for GL_2. *Ann. of Math. (2)* (3)155 (2002), 837–893, With an appendix by Colin J. Bushnell and Guy Henniart.

[LP81] P.D. Lax and R.S. Phillips. The asymptotic distribution of lattice points in Euclidean and non-Euclidean spaces. In: *Functional Analysis and Approximation (Oberwolfach, 1980), International Series of Numerical Mathematics*, Vol. 60. Birkhäuser (1981), pp. 373–383.

[LPS88] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, (3)8 (1988), 261–277.

[LRS95] W. Luo, Z. Rudnick, and P. Sarnak. On Selberg's eigenvalue conjecture. *Geom. Funct. Anal.*, (2)5 (1995), 387–401.

[LS04] M.W. Liebeck and A. Shalev. Fuchsian groups, coverings of Riemann surfaces, subgroup growth, random quotients and random walks. *J. Algebra* (2)276, (2004), 552–601.

[LW21] M. Lipnowski and A. Wright. *Towards optimal spectral gaps in large genus*, preprint, 2103.07496 (2021).

[Med78] A.D. Mednyhk. Determination of the number of nonequivalent coverings over a compact Riemann surface. *Dokl. Akad. Nauk SSSR*, (2)239 (1978), 269–271.

[Mir13] M. Mirzakhani. Growth of Weil-Petersson volumes and random hyperbolic surfaces of large genus. *J. Differential Geom.*, (2)94 (2013), 267–300.

[MN20] M. Magee and F. Naud. Explicit spectral gaps for random covers of Riemann surfaces. *Publ. Math. Inst. Hautes Études Sci.*, 132 (2020), 137–179.

[MN21] M. Magee and F. Naud. *Extension of Alon's and Friedman's conjectures to Schottky surfaces*, preprint, 2106.02555 (2021)

[Mon22] L. Monk. Benjamini–Schramm convergence and spectrum of random hyperbolic surfaces of high genus. *Analysis & PDE* (2022), to appear, 2002.00869.

[MP20] M. Magee and D. Puder. *The asymptotic statistics of random covering surfaces*, preprint, 2003.05892v5 (2020).

[MP21] M. Magee and D. Puder. *Core surfaces*, preprint, 2108.00717 (2021).

[Nil91] A. Nilli, On the second eigenvalue of a graph. *Discrete Math.*, (2)91 (1991), 207–210.

[Pat76] S.J. Patterson. The limit set of a Fuchsian group. *Acta Math.* (3–4)136 (1976), 241–273.

[PP15] D. Puder and O. Parzanchevski. Measure preserving words are primitive. *Journal of the American Mathematical Society*, (1)28 (2015), 63–97.

[Pud14] D. Puder. Primitive words, free factors and measure preservation. *Israel Journal of Mathematics*, (1)201 (2014), 25–73.

[Pud15] D. Puder. Expansion of random graphs: new proofs, new results. *Invent. Math.*, (3)201 (2015), 845–908.

[Sar87] P. Sarnak. Statistical properties of eigenvalues of the Hecke operators. In: *Analytic Number Theory and Diophantine Problems (Stillwater, OK, 1984), Progress in Mathematics*, Vol. 70. Birkhäuser Boston, Boston, MA (1987), pp. 321–331.

[Sel56] A. Selberg. Harmonic analysis and discontinuous groups in weakly symmetric Riemannian spaces with applications to Dirichlet series. *J. Indian Math. Soc. (N.S.)*, 20 (1956), 47–87.

[Sel65] A. Selberg. On the estimation of Fourier coefficients of modular forms. In: *Proceedings of Symposia in Pure Mathematics*, Vol. VIII. American Mathematical Society, Providence, RI (1965), pp. 1–15.

[Sta83] J.R. Stallings. Topology of finite graphs. *Inventiones Mathematicae*, (3)71 (1983), 551–565.

[VO04] A.M. Vershik and A. Yu. Okun'kov. A new approach to representation theory of symmetric groups. II. *Zap. Nauchn. Sem. S.-Peterburg. Otdel. Mat. Inst. Steklov. (POMI)* 307 (2004), no. Teor. Predst. Din. Sist. Komb. i Algoritm. Metody. 10, 57–98, 281.

[Wei48] A. Weil. On some exponential sums. *Proc. Nat. Acad. Sci. U.S.A.*, 34 (1948), 204–207.

[Wri20] A. Wright. A tour through Mirzakhani's work on moduli spaces of Riemann surfaces. *Bulletin of the American Mathematical Society*, (3)57 (2020), 359–408.

[WX21] Y. Wu and Y. Xue. *Random hyperbolic surfaces of large genus have first eigenvalues greater than* $\frac{3}{16} - \epsilon$, preprint, 2102.05581 (2021). https://link.springer.com/article/10.1007/s00039-022-00595-7.

M. Magee
Department of Mathematical Sciences, Durham University, Lower Mountjoy DH1 3LE, Durham, UK. michael.r.magee@durham.ac.uk

F. Naud
IMJ-PRG - UMR7586, Université Pierre et Marie Curie, 4 place Jussieu, 75252 Paris Cedex 05, France. frederic.naud@imj-prg.fr

D. Puder
School of Mathematical Sciences, Tel Aviv University, 6997801 Tel Aviv, Israel.
doronpuder@gmail.com

Received: September 7, 2021
Revised: February 17, 2022
Accepted: March 15, 2022