

RESEARCH ARTICLE

The Asymptotic Statistics of Random Covering Surfaces

Michael Magee ¹ and Doron Puder ²

¹Department of Mathematical Sciences, Durham University, United Kingdom; E-mail: michael.r.magee@durham.ac.uk.

²School of Mathematical Sciences, Tel Aviv University, Tel Aviv, 6997801, Israel; E-mail: doronpuder@gmail.com.

Received: 2 March 2022; **Revised:** 9 April 2023; **Accepted:** 27 March 2023

2020 Mathematics Subject Classification: *Primary* – 20C15; *Secondary* – 20B30, 20P05, 20F34, 20F65, 20F70, 60B15

Abstract

Let Γ_g be the fundamental group of a closed connected orientable surface of genus $g \geq 2$. We develop a new method for integrating over the representation space $\mathbb{X}_{g,n} = \text{Hom}(\Gamma_g, S_n)$, where S_n is the symmetric group of permutations of $\{1, \dots, n\}$. Equivalently, this is the space of all vertex-labeled, n -sheeted covering spaces of the closed surface of genus g .

Given $\phi \in \mathbb{X}_{g,n}$ and $\gamma \in \Gamma_g$, we let $\text{fix}_\gamma(\phi)$ be the number of fixed points of the permutation $\phi(\gamma)$. The function fix_γ is a special case of a natural family of functions on $\mathbb{X}_{g,n}$ called Wilson loops. Our new methodology leads to an asymptotic formula, as $n \rightarrow \infty$, for the expectation of fix_γ with respect to the uniform probability measure on $\mathbb{X}_{g,n}$, which is denoted by $\mathbb{E}_{g,n}[\text{fix}_\gamma]$. We prove that if $\gamma \in \Gamma_g$ is not the identity and q is maximal such that γ is a q^{th} power in Γ_g , then

$$\mathbb{E}_{g,n}[\text{fix}_\gamma] = d(q) + O(n^{-1})$$

as $n \rightarrow \infty$, where $d(q)$ is the number of divisors of q . Even the weaker corollary that $\mathbb{E}_{g,n}[\text{fix}_\gamma] = o(n)$ as $n \rightarrow \infty$ is a new result of this paper. We also prove that $\mathbb{E}_{g,n}[\text{fix}_\gamma]$ can be approximated to any order $O(n^{-M})$ by a polynomial in n^{-1} .

Contents

1	Introduction	2
1.1	Related works I: Mirzakhani’s integral formulas	5
1.2	Related works II: Free groups	5
1.3	Related works III: Noncommutative probability	6
1.4	Related works IV: Residual finiteness	6
1.5	Related works V: Benjamini–Schramm convergence	7
1.6	Structure of the proofs and the issues that arise	7
1.7	Notation	9
2	Resolutions of core surfaces	9
2.1	Tiled surfaces and core surfaces	9
2.2	Expectations and probabilities of tiled surfaces	13
2.3	Resolutions	14
3	Background: representation theory of the symmetric group	19
3.1	Young diagrams	19

3.2	Young tableaux	20
3.3	Representations of symmetric groups	20
4	Preliminary representation theoretic results	22
4.1	Commutants	22
4.2	Bounds for the dimensions of irreducible representations	23
4.3	An estimate for matrix coefficients in skew modules	24
4.4	Families of YDs and zeta functions	25
5	The probability of an embedded tiled surface	28
5.1	Overview of this section	28
5.2	Tiled surfaces and cosets	30
5.3	Construction of auxiliary permutations	32
5.4	Integrating over cosets	33
5.5	A geometric bound for products of matrix coefficients	40
5.6	Construction of labelings of tiled surfaces from collections of tableaux	40
5.7	Stronger bounds for matrix coefficients	44
5.8	Approximating $\Xi_n(Y)$ by Laurent polynomials	45
5.9	Estimating $\Xi_n(Y)$	46
6	Proofs of main theorems	48
6.1	Proof of Theorem 1.1 and its extension to finitely generated subgroups	48
6.2	Proof of Theorems 1.2 and 1.3	49

1. Introduction

Let $g \geq 2$, and let Σ_g be a closed orientable topological surface of genus g . We fix a base point $o \in \Sigma_g$ and let

$$\Gamma_g \stackrel{\text{def}}{=} \pi_1(\Sigma_g, o) \cong \langle a_1, b_1, \dots, a_g, b_g \mid [a_1, b_1] \cdots [a_g, b_g] \rangle \tag{1.1}$$

be the fundamental group of Σ_g . Denote by

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

the *representation space* of all homomorphisms from Γ_g to S_n , where S_n is the symmetric group of permutations of $\{1, \dots, n\}$. From another point of view, the space $\mathbb{X}_{g,n}$ can be viewed as the space of degree- n covering maps of Σ_g . Indeed, for every not-necessarily-connected degree- n covering map

$$p: X \twoheadrightarrow \Sigma_g,$$

we may identify the fiber $p^{-1}(o)$ with $\{1, \dots, n\}$, and the monodromy action of $\pi_1(\Sigma_g, o)$ on the fiber then gives rise to a homomorphism $\phi \in \text{Hom}(\Gamma_g, S_n)$. This gives a one-to-one correspondence between $\mathbb{X}_{g,n}$ and degree- n covering maps with $p^{-1}(o) = \{1, \dots, n\}$. This correspondence is discussed in more detail in §§2.2.

The space $\mathbb{X}_{g,n}$ was studied by Liebeck and Shalev [LS04], who showed that a uniformly random homomorphism $\phi: \Gamma_g \rightarrow S_n$ satisfies $\phi(\Gamma_g) \supseteq A_n$ a.a.s. (asymptotically almost surely, namely, with probability tending to 1 as $n \rightarrow \infty$) [LS04, Thm. 1.12].¹ In particular the image is a.a.s. transitive, or, equivalently, the corresponding random degree- n covering space is a.a.s. connected. When Γ_g is replaced by a nonabelian free group, the analogous result holds by the famous theorem of Dixon [Dix69] that two random permutations in S_n a.a.s. generate S_n or A_n .

In the current work we address the problem of integration over the space $\mathbb{X}_{g,n}$. Namely, our goal is to analyze the expected value $\mathbb{E}_{g,n}[f]$ of functions f on $\mathbb{X}_{g,n}$ with respect to the uniform measure

¹The paper [LS04] considers, more generally, random homomorphisms from any Fuchsian group to S_n .

on $\mathbb{X}_{g,n}$. The functions on $\mathbb{X}_{g,n}$ that we consider are natural functions that arise from loops in Σ_g . Given an element $\gamma \in \Gamma_g$ and a character χ of S_n , we let

$$\chi_\gamma(\phi) \stackrel{\text{def}}{=} \chi(\phi(\gamma)), \quad \chi_\gamma : \mathbb{X}_{g,n} \rightarrow \mathbf{R}.$$

These functions are called *Wilson loops* in the physics literature [Lab13, Def. 6.4.1]. Our focus is on the character fix of S_n which assigns to every permutation its number of fixed points.

The main motivation behind this work is its relevance to the study of random covers of the closed surface Σ_g . Given some $1 \neq \gamma \in \Gamma_g$, consider the geodesic C_γ in Σ_g corresponding to the conjugacy class of γ . For every homomorphism $\phi \in \mathbb{X}_{g,n}$, the number of fixed points $\text{fix}_\gamma(\phi)$ is precisely the number of lifts of C_γ to a closed geodesic in the degree- n covering corresponding to ϕ . Indeed, the results of this paper are crucial ingredients in a subsequent work [MNP22] which gives new results on spectral gaps of random covers of a closed surface.

Another source of motivation is the rich theory that has been discovered around similar questions when surface groups are replaced by free groups (e.g., [Nic94, PP15, MP19, HP22] and see §§1.2 below). Expanding this theory to other groups is challenging, as the presence of a relation between the generators presents a fundamental difficulty that is not present for free groups. Surface groups, among the best understood and best behaved one-relator groups, are a natural starting point for this quest. To overcome the difficulty brought up by the existence of a relation, we develop in this work new machinery, both in representation theory of S_n and in combinatorial group theory.

Expected number of fixed points

Recall that the expectation $\mathbb{E}_{g,n} [\text{fix}_\gamma]$ is the average number of fixed points in $\phi(\gamma)$ where $\phi : \Gamma_g \rightarrow S_n$ is uniformly random. Our main results are the following two theorems.

Theorem 1.1. *Fix $g \geq 2$ and $\gamma \in \Gamma_g$. Then there is an infinite sequence of rational numbers*

$$a_1(\gamma), a_0(\gamma), a_{-1}(\gamma), a_{-2}(\gamma), \dots$$

such that for any $M \in \mathbf{N}$, as $n \rightarrow \infty$,

$$\mathbb{E}_{g,n} [\text{fix}_\gamma] = a_1(\gamma)n + a_0(\gamma) + \frac{a_{-1}(\gamma)}{n} + \dots + \frac{a_{-(M-1)}(\gamma)}{n^{M-1}} + O\left(\frac{1}{n^M}\right). \tag{1.2}$$

Theorem 1.2. *If $\gamma \in \Gamma_g$ is not the identity, then, as $n \rightarrow \infty$,*

$$\mathbb{E}_{g,n} [\text{fix}_\gamma] = O(1).$$

In fact, if $q \in \mathbf{N}$ is maximal such that $\gamma = \gamma_0^q$ for some $\gamma_0 \in \Gamma$, then, as $n \rightarrow \infty$,

$$\mathbb{E}_{g,n} [\text{fix}_\gamma] = d(q) + O\left(\frac{1}{n}\right),$$

where $d(q)$ is the number of divisors function. In other words, $a_1(\gamma) = 0$ and $a_0(\gamma) = d(q)$.

For example, consider the element a in $\Gamma_2 = \langle a, b, c, d \mid [a, b] [c, d] \rangle$. This element is not a proper power and so $\mathbb{E}_{2,n} [\text{fix}_a] = 1 + O(n^{-1})$ by Theorem 1.2. By Theorem 1.1, this average can be approximated to any order n^{-M} by a rational function in n . In this particular case, this rational function can be computed to get for, for example, $M = 5$,

$$\mathbb{E}_{2,n} [\text{fix}_a] = 1 + \frac{1}{n^2} + \frac{2}{n^3} + \frac{10}{n^4} + O\left(\frac{1}{n^5}\right).$$

Given a finite group G , the number of homomorphisms $\Gamma_g \rightarrow G$ is related to the Witten zeta function of G ,

$$\zeta^G(s) \stackrel{\text{def}}{=} \sum_{\chi \in \text{Irr}G} \chi(1)^{-s},$$

the summation being over the isomorphism classes of irreducible complex representations of G . These functions were introduced by Zagier [Zag94] after Witten’s work in [Wit91]. The connection is given by

$$|\text{Hom}(\Gamma_g, G)| = |G|^{2g-1} \zeta^G(2g-2). \tag{1.3}$$

This result goes back to Hurwitz [Hur02], who gave a more general formula for arbitrary Fuchsian groups (a proof in English is given in [LS04, Prop. 3.2]). It is also sometimes called ‘Mednykh’s formula’ in the literature after [Med78]. For the case $G = S_n$, the zeta function $\zeta^{S_n}(s)$ was studied in [Lul96, MP02, LS04, Gam06]. Inter alia, these works show that, for every $s > 0$,

$$\zeta^{S_n}(s) \xrightarrow{n \rightarrow \infty} 2.$$

Moreover, their results yield an asymptotic expansion in n which approximates $\zeta^{S_n}(s)$ as $n \rightarrow \infty$, in a similar manner to the one in Theorem 1.1. As such, their results can be thought of as the special case of $\gamma = 1$ of a version of Theorem 1.1. We elaborate more in §§5.1.

Common fixed points of subgroups

Our proof also yields the following more general result that concerns not only elements of Γ_g but also f.g. (finitely generated) subgroups. We write $J \leq_{\text{f.g.}} \Gamma_g$ to denote a f.g. subgroup J of Γ_g . Given $J \leq_{\text{f.g.}} \Gamma_g$ and $\phi \in \mathbb{X}_{g,n}$, we let $\text{fix}_J(\phi)$ denote the number of elements in $1, \dots, n$ that are fixed by all permutations in $\phi(J)$:

$$\text{fix}_J(\phi) \stackrel{\text{def}}{=} |\{i \in \{1, \dots, n\} \mid \sigma(i) = i \text{ for all } \sigma \in \phi(J)\}|.$$

In particular, $\text{fix}_{\langle \gamma \rangle} = \text{fix}_\gamma$ for all $\gamma \in \Gamma_g$. For $J \leq_{\text{f.g.}} \Gamma_g$, we let

$$\chi_{\max}(J) \stackrel{\text{def}}{=} \max \{ \chi(K) \mid J \leq K \leq_{\text{f.g.}} \Gamma_g \} \tag{1.4}$$

denote the largest Euler characteristic² of a f.g. subgroup $K \leq_{\text{f.g.}} \Gamma_g$ which contains J . Note that $\chi(\Gamma_g) = 2 - 2g \leq \chi_{\max}(J) \leq 1$ and that $\chi_{\max}(J) \geq \chi(J)$. It is also true that $\chi_{\max}(J) = 1$ if and only if $J = \{1\}$, and $\chi_{\max}(J) \geq 0$ if and only if J is cyclic. In addition, we let

$$\mathfrak{MOG}(J) \stackrel{\text{def}}{=} \{ K \leq_{\text{f.g.}} \Gamma_g \mid J \leq K \text{ and } \chi(K) = \chi_{\max}(J) \}$$

denote the set of ‘maximal overgroups’ – f.g. subgroups achieving the maximum from equation (1.4). This set is always finite – see Corollary 2.16.

Theorem 1.3. *Let $J \leq_{\text{f.g.}} \Gamma_g$ be a f.g. subgroup. Then*

$$\mathbb{E}_{g,n}[\text{fix}_J] = |\mathfrak{MOG}(J)| \cdot n^{\chi_{\max}(J)} + O\left(n^{\chi_{\max}(J)-1}\right).$$

Theorem 1.3 generalizes Theorem 1.2, as for $\gamma \neq 1$, $\chi_{\max}(\langle \gamma \rangle) = 0$ and

$$\mathfrak{MOG}(\langle \gamma \rangle) = \{ \langle \gamma_0^m \rangle \mid m|q \}.$$

²Every f.g. subgroup $K \leq \Gamma_g$ is either a free group, in which case $\chi(K) = 1 - \text{rank}(K)$, or a surface group of genus $h \geq g$, in which case $\chi(K) = 2 - 2h$.

The analog of Theorem 1.1 holds too for f.g. subgroups: There is an infinite sequence of rational numbers

$$a_1(J), a_0(J), a_{-1}(J), \dots$$

such that for any $M \in \mathbb{N}$, as $n \rightarrow \infty$,

$$\mathbb{E}_{g,n} [\text{fix}_J] = \sum_{i=-(M-1)}^1 a_i(J) n^i + O(n^{-M}),$$

and such that $a_1 = a_0 = \dots = a_{\chi_{\max}(J)+1} = 0$ and $a_{\chi_{\max}(J)} = |\mathfrak{MSG}(J)|$.

1.1. Related works I: Mirzakhani’s integral formulas

In [Mir07], Mirzakhani considered a similar problem to the one in this paper. Instead of integrating over the finite space $\text{Hom}(\Gamma_g, S_n)$, Mirzakhani obtained formulas for the integral of geometric functions over the the moduli space \mathcal{M}_g of complete hyperbolic surfaces of genus g , with respect to the Weil–Petersson volume form $d\text{Vol}_{\text{wp}}$.

The geometric functions that Mirzakhani considers are very much like our Wilson loops. Given any closed curve $\gamma \in \Sigma_g$, for any complete hyperbolic metric J on Σ_g there is a unique curve isotopic to γ that is shortest with respect to J , and the length of this curve is called the length of γ , denoted by $\ell_J([\gamma])$. Here, $[\gamma]$ is the isotopy class of γ .

Mirzakhani requires that γ be simple, meaning that it does not intersect itself. This condition is not present in the current paper and can be viewed as an advantage of our work. To obtain a function on \mathcal{M}_g , given a continuous function $f : \mathbf{R}_+ \rightarrow \mathbf{R}_+$, Mirzakhani considers the averaged function

$$f_\gamma(J) \stackrel{\text{def}}{=} \sum_{[\gamma'] \in \text{MCG}(\Sigma_g) \cdot [\gamma]} f(\ell_J([\gamma'])),$$

where $\text{MCG}(\Sigma_g)$ is the mapping class group of Σ_g . Because of the averaging over the mapping class group, f_γ descends to a function on \mathcal{M}_g . This type of averaging is not necessary in the current paper because $\mathbb{X}_{g,n} = \text{Hom}(\Gamma_g, S_n)$ is already finite; here, $\mathbb{X}_{g,n}$ is playing the role of the Teichmüller space and not the moduli space. In [Mir07, Thm. 8.1], Mirzakhani gives a formula for

$$\int_{\mathcal{M}_g} f_\gamma d\text{Vol}_{\text{wp}}$$

in terms of integrating f against Weil–Petersson volumes of moduli spaces. The power of this formula is that in the same paper [Mir07], Mirzakhani gives explicit recursive formulas for the calculations of Weil–Petersson volumes. For a more detailed discussion of these formulas, the reader should consult Wright’s survey of Mirzakhani’s work [Wri20, §4].

1.2. Related works II: Free groups

Let \mathbf{F}_r denote a free group on r generators. For $\gamma \in \mathbf{F}_r$, the problem of integrating the Wilson loop

$$\text{fix}_\gamma(\phi) \stackrel{\text{def}}{=} \text{fix}(\phi(\gamma)), \quad \text{fix}_\gamma : \text{Hom}(\mathbf{F}_r, S_n) \rightarrow \mathbf{R}$$

over $\text{Hom}(\mathbf{F}_r, S_n)$ with respect to the uniform probability measure is a basic problem that serves as a precursor to that of the current paper. As mentioned above, many of the considerations used with free groups no longer apply in the present paper. Indeed, $\text{Hom}(\mathbf{F}_r, S_n)$ can be identified with S_n^r and hence techniques for integrating over groups are relevant in a much more direct way than in the case of $\text{Hom}(\Gamma_g, S_n)$.

Despite being an easier problem, the theory is very rich. It was proved by Nica in [Nic94] that the analog of Theorem 1.2 holds for $\mathbb{E}_{\mathbf{F}_r, n}[\text{fix}_\gamma]$. A significantly sharper result was given by Puder and Parzanchevski in [PP15] where they proved that if $\gamma \in \mathbf{F}_r$, then as $n \rightarrow \infty$

$$\mathbb{E}_{\mathbf{F}_r, n}[\text{fix}_\gamma] = 1 + \frac{c(\gamma)}{n^{\pi(\gamma)-1}} + O\left(\frac{1}{n^{\pi(\gamma)}}\right),$$

where $\pi(\gamma) \in \{0, \dots, r\} \cup \{\infty\}$ is an algebraic invariant of γ called the *primitivity rank* and $c(\gamma) \in \mathbf{N}$ is explained in terms of the enumeration of special subgroups of \mathbf{F}_r determined by γ . Obtaining a similarly sharp result in the context of Γ_g is an interesting problem that should be taken up in the future.

Similar Laurent series expansions for the expected value of χ_γ on $\text{Hom}(\mathbf{F}_r, G(n))$ have been proved to exist, and studied, when $G(n)$ is one of the families of compact Lie groups $U(n), O(n), Sp(n)$ [MP19, MP22b], when $G(n)$ is a generalized symmetric group [MP21], and when $G(n) = GL_n(\mathbb{F}_q)$, where \mathbb{F}_q is a fixed finite field [EWPS21]. In all cases, χ is taken to be a natural character. For example, when $G(n) = U(n)$, one such χ is the trace of the matrix in the group. Moreover, for $G(n) = U(n), O(n), Sp(n)$ and χ the trace, all the coefficients of the Laurent series are understood [MP19, MP22b].

In works undertaken after the completion of this paper, the first named author has obtained analogs of Theorem 1.1 and the first part of Theorem 1.2 for³ $\text{Hom}(\Gamma_g, U(n))$ and the standard matrix trace [Mag22, Mag21]. The methods used in (*ibid.*) are inspired by those of the current work.

1.3. Related works III: Noncommutative probability

Theorem 1.2 has a direct consequence in the setting of Voiculescu’s noncommutative probability theory. Following [VDN92, Def. 2.2.2], a *C*-probability space* is a pair (\mathcal{B}, τ) , where \mathcal{B} is a unital C*-algebra and τ is a state⁴ on \mathcal{B} . We say that a sequence $\{(\mathcal{B}, \tau_n)\}_{n=1}^\infty$ of C*-probability spaces converges to (\mathcal{B}, τ) if for all elements $b \in \mathcal{B}$

$$\lim_{n \rightarrow \infty} \tau_n(b) = \tau(b).$$

The functions $\tau_n : \Gamma_g \rightarrow \mathbf{R}$ defined by $\tau_n(\gamma) \stackrel{\text{def}}{=} n^{-1} \mathbb{E}_{g, n}[\text{fix}_\gamma]$ extend to states on the full group C*-algebra $C^*(\Gamma_g)$ of Γ_g . There is also a unique state τ_{reg} on $C^*(\Gamma_g)$ that satisfies $\tau_{\text{reg}}(g) = 0$ for $g \neq 1$; we use the subscript *reg* because the GNS representation of τ_{reg} is the *left regular representation*. One has the following corollary of Theorem 1.2:

Corollary 1.4. *The C*-probability spaces $(C^*(\Gamma_g), \tau_n)$ converge to $(C^*(\Gamma_g), \tau_{\text{reg}})$ as $n \rightarrow \infty$.*

It is reasonable to hope that similar results can be obtained when Γ_g is replaced by any residually finite one-relator group (cf. §§1.4). We view Corollary 1.4 as an important first step in this program.

1.4. Related works IV: Residual finiteness

A f.g. discrete group Λ is *residually finite* if for any nonidentity $\lambda \in \Lambda$ there is a finite index subgroup $H \leq \Lambda$ such that $\lambda \notin H$. The residual finiteness of Γ_g has been known for a long time [Bau62, Hem72]. More recently, various quantifications of residual finiteness and of the related property of LERF⁵ have been proposed by various authors [BR10, LLM23]. Theorem 1.2 can serve as a strengthening of the residual finiteness of Γ_g , as we now explain.

³In this case, instead of the uniform measure on $\text{Hom}(\Gamma_g, S_n)$ that we use here, one should use a natural measure on $\text{Hom}(\Gamma_g, U(n))$ that arises from the Atiyah–Bott–Goldman symplectic form [Gol84, AB83] on (a nonsingular part of) the character variety $\text{Hom}(\Gamma_g, U(n))/U(n)$.

⁴A state on a unital C*-algebra is a positive linear functional such that $\tau(1) = 1$.

⁵Locally extended residual finiteness.

Note that residual finiteness of a group Λ is equivalent to, for all $e \neq \lambda \in \Lambda$, the existence of $n \in \mathbb{N}$ and $\phi \in \text{Hom}(\Lambda, S_n)$ such that $\phi(\lambda) \neq 1$. Theorem 1.2 combined with Markov’s inequality implies the following quantitative version of residual finiteness.

Corollary 1.5. *Given a nonidentity element $e \neq \gamma \in \Gamma_g$, for large enough n ,*

$$\frac{|\{\phi \in \text{Hom}(\Gamma_g, S_n) : \phi(\gamma) \neq 1\}|}{|\text{Hom}(\Gamma_g, S_n)|} \geq 1 - \frac{d(q)}{n} - O\left(\frac{1}{n^2}\right), \tag{1.5}$$

where q and $d(q)$ are as in Theorem 1.2, and the implied constant in the big- O term depends on γ .

In fact, the techniques of this paper can be used to show that, for example, for every $m \in \mathbb{N}$, the expected value of fix_γ^m is of the form $c(q) + O(n^{-1})$, where q is as in Theorem 1.2 and $c(q)$ is a positive integer. This would yield a probability bound similar to equation (1.5) but of the form $1 - \frac{c(q)}{n^m} + O(n^{-m-1})$. This is done explicitly in [PZ22, Corollary 1.8].

1.5. Related works V: Benjamini–Schramm convergence

In [BS01], Benjamini and Schramm introduced a notion of convergence of a sequence of finite graphs to a limiting graph, known now as *Benjamini–Schramm convergence*. This concept was extended to convergence of sequences of Riemannian manifolds in [ABB+11, ABB+17]. Theorem 1.2 has consequences for the Benjamini–Schramm convergence of random covers of Riemannian surfaces; there are various of these consequences but we present just one representative one here.⁶

Corollary 1.6. *Let X be a closed hyperbolic surface of genus ≥ 2 . Then as $n \rightarrow \infty$, uniformly random degree- n covering spaces of X converge in the sense of Benjamini–Schramm to the hyperbolic upper half plane \mathbb{H} .*

Concretely, this means that for any $L > 0$ and $\varepsilon > 0$, if X_n denotes a random degree- n cover of X (as above), then a.a.s. as $n \rightarrow \infty$,

$$\frac{\text{area}(X_n^{<L})}{\text{area}(X_n)} < \varepsilon,$$

where $X_n^{<L}$ is the points of X_n with local injectivity radius $< L$. To see how this follows from Theorem 1.2, viewing L as a constant, any point in $X_n^{<L}$ is in a neighborhood, with bounded area depending on L , of some simple closed geodesic of X_n with length $< 2L$ [Bus10, proof of Thm. 4.1.6]. Any such geodesic covers a closed (possibly nonprimitive) geodesic in X of length $< 2L$, and these in turn correspond to a finite list of conjugacy classes in Γ_g . Starting with a conjugacy class $[\gamma]$, the number of corresponding closed lifted geodesics in X_n is at most fix_γ . Using Markov’s inequality with Theorem 1.2 gives therefore a.a.s. that the number of simple closed geodesics of X_n with length $< 2L$ is bounded (depending on L). This means $\text{area}(X_n^{<L})$ is bounded a.a.s. and as $\text{area}(X_n)$ is linear in n , this completes the proof.

1.6. Structure of the proofs and the issues that arise

The reader of the paper is advised to first read this §1.6, and then §6, where all the ideas of the paper are brought together to give concise proofs of Theorems 1.1, 1.2 and 1.3, before reading the other sections.

There are two main ideas of the paper that we will discuss momentarily. Here, we give a ‘high-level’ account of the strategy of proving our main theorems. At times, we oversimplify definitions to be more instructive. Let us fix $g = 2$ and discuss only Theorems 1.1 and 1.2. The extension of these results from cyclic groups to more general finitely generated subgroups is along the same lines. So fix $\gamma \in \Gamma_2$.

Firstly, we view $\mathbb{X}_n = \mathbb{X}_{2,n}$ as a space of random coverings of a fixed genus 2 surface Σ_2 . By fixing an octagonal fundamental domain of Σ_2 , each covering of Σ_2 is tiled by octagons. This leads us to the

⁶This consequence of Theorem 1.2 was first pointed out by Baker and Petri in [BP20].

notion of a *tilted surface*, defined precisely in Definition 2.1. A tiled surface involves not just a tiling but a labeling of the edges of the tiling by generators of the fundamental group of Σ_2 . Hence, all the main theorems can be reinterpreted in terms of random tiled surfaces that are called X_ϕ for $\phi \in \mathbb{X}_n$.

The first observation is that $\mathbb{E}_n [\text{fix}_\gamma] = \mathbb{E}_{2,n} [\text{fix}_\gamma]$, the expected number of fixed points of γ under $\phi \in \mathbb{X}_n$, is the expected number of times that we see a fixed annulus A , specified by γ , immersed in the random tiled surface X_ϕ . This annulus A may be the ‘core surface’ corresponding to $\langle \gamma \rangle$ – see Definition 2.6, the left part of Figure 2.3 and [MP22a, Lem. 5.1]. However, A needs not be embedded in X_ϕ . On the other hand, it is possible to produce a finite collection \mathcal{R} of tiled surfaces, each of which has an immersed copy of A , such that

$$\mathbb{E}_n [\text{fix}_\gamma] = \sum_{Y \in \mathcal{R}} \mathbb{E}_n^{\text{emb}}(Y), \quad (1.6)$$

where $\mathbb{E}_n^{\text{emb}}(Y)$ is the expected number of times that Y is embedded in the random X_ϕ .

We formalize types of collections \mathcal{R} that have the above property in Definition 2.8; we call them *resolutions* (of A). Of course, there is a great deal of flexibility in how \mathcal{R} is chosen; we will come back to this point shortly. The benefit to having equation (1.6) brings us to the first main idea of the paper:

We have a new method of calculating $\mathbb{E}_n^{\text{emb}}(Y)$, using the representation theory of symmetric groups S_n and more specifically, the approach to the representation theory of S_n developed by Vershik and Okounkov in [VO96].

This methodology is developed in §5. The necessary background on representation theory is given in §3, and in §4 we prove some preliminary representation theoretic results needed for §5. The reader may be interested to see that Theorem 1.1 has, at its source, Proposition 4.6. See also the overview of §5 in §§5.1.

This new methodology to calculate $\mathbb{E}_n^{\text{emb}}(Y)$ is sufficient to prove Theorem 1.1. However, in the proof of Theorem 1.2, a critical issue now intervenes. We expect, based on experience with similar projects (e.g., [PP15, MP19]) that

$$\mathbb{E}_n^{\text{emb}}(Y) \approx n^{\chi(Y)} \quad (1.7)$$

as $n \rightarrow \infty$. However, this cannot always be the case. For example, if, roughly speaking, it is possible to glue some octagons to Y to increase the Euler characteristic, forming Y' , then the observation that $\mathbb{E}_n^{\text{emb}}(Y) \geq \mathbb{E}_n^{\text{emb}}(Y')$ breaks equation (1.7). Then it is not unsurprising that the bounds we obtain for $\mathbb{E}_n^{\text{emb}}(Y)$ do not always agree with equation (1.7).

On the other hand, if Y has special properties that we call *boundary reduced* and *strongly boundary reduced*, then we can get appropriate bounds on $\mathbb{E}_n^{\text{emb}}(Y)$. We give the precise definitions of these properties in Definitions 2.4 and 2.5. They involve forbidding certain constellations from appearing in the boundary of Y . Even though these constellations are dictated by representation theory, forbidding them remarkably relates to natural geometric properties of Y . For example, if Y is not boundary reduced, then it is possible to add octagons to Y to decrease the number of edges in its boundary. To give some more intuition, being boundary reduced can be viewed as a discrete analog of a hyperbolic surface having geodesic boundary. This means that these properties are closely related with the problem of finding shortest representatives (with respect to word length) of elements of Γ_g that is addressed by Dehn’s algorithm [Deh12].

If Y is boundary reduced, then we can prove (Theorem 5.10 and Proposition 5.25)

$$\mathbb{E}_n^{\text{emb}}(Y) = O\left(n^{\chi(Y)}\right),$$

and if Y is strongly boundary reduced, we can prove (Theorem 5.10 and Proposition 5.26)

$$\mathbb{E}_n^{\text{emb}}(Y) = n^{\chi(Y)} \left(1 + O\left(n^{-1}\right)\right)$$

(see, again, Section 5.1 for a more detailed overview).⁷ Therefore, to prove Theorem 1.2, it suffices to produce resolutions of the annulus A where we can control which elements are (strongly) boundary reduced, control their Euler characteristics and count the number of elements with maximal Euler characteristic. The design of these resolutions is the second main theme of the paper.

For any tiled surface Z , we describe an algorithm to produce finite resolutions of Z with careful control on their properties as above. This is the main topic of §2. Precisely defining the annulus A that should be used as input, as well as its generalization for noncyclic subgroups $J \leq \Gamma$, and counting the outputs of our algorithm, requires introducing the concept of a *core surface of a subgroup* $J \leq \Gamma$. For example, above, A should be taken to be the core surface of $\langle \gamma \rangle$. The theory of core surfaces that we develop in a companion paper [MP22a] is analogous to that of Stallings’ core graphs for subgroups of free groups due to Stallings [Sta83], and we hope that the results therein may be of independent interest.

Remark 1.7. Another perspective on the value of $\mathbb{E}_n^{\text{emb}}(Y)$ for *arbitrary* tiles surfaces is given in [PZ22, Thm. 2.6]. Let Y be an arbitrary tiled surface, $p: Y \rightarrow \Sigma_g$ the restricted covering map, and $\chi^{\text{grp}}(Y)$ the Euler characteristic of the subgroup $p_*(\pi_1(Y)) \leq \pi_1(\Sigma_g) = \Gamma_g$. Then

$$\mathbb{E}_n^{\text{emb}}(Y) = n^{\chi^{\text{grp}}(Y)} \left(a_0 + O\left(n^{-1}\right) \right),$$

where a_0 is some positive integer. This theorem heavily relies on the results of the current paper.

1.7. Notation

Write \mathbf{N} for the natural numbers $1, 2, \dots$ and so on. For $n \in \mathbf{N}$, we use the notation $[n]$ for the set $\{1, \dots, n\}$. For $m \leq n, m, n \in \mathbf{N}$ we write $[m, n]$ for the set $\{m, m + 1, \dots, n\}$. If A and B are sets, we write $A \setminus B$ for the elements of A that are not in B . We write $(n)_\ell$ for the Pochhammer symbol

$$(n)_\ell \stackrel{\text{def}}{=} n(n - 1) \dots (n - \ell + 1).$$

If V is a vector space, we write $\text{End}(V)$ for the linear endomorphisms of V . If V is a unitary representation of some group, we write \check{V} for the dual representation. If P_1, \dots, P_k are a series of expressions we write $\mathbf{1}_{\{P_1, \dots, P_k\}}$ for a value which is 1 if all the statements P_i are true and 0 else. If V is a vector space, we write Id_V for the identity operator on that space. All integrals over finite sets are with respect to the uniform probability measure on the set. If X is a CW-complex, we write $X^{(i)}$ for its i -skeleton. If we use the symbol \pm more than once in the same expression or equation, we mean that the same sign is chosen each time. If implied constants in big-O notation depend on other constants, we indicate this by adding the constants as a subscript to the O , for example, $O_\varepsilon(f(n))$ means the implied constant depends on ε . We use Vinogradov notation $f(n) \ll g(n)$ to mean that there are constants $n_0 \geq 0$ and $C_0 > 0$ such that for $n > n_0, |f(n)| \leq C_0 g(n)$. We add subscripts to indicate dependence of the implied constants on other quantities or objects. If a, b are elements of the same group, we write $[a, b] \stackrel{\text{def}}{=} aba^{-1}b^{-1}$ for their commutator.

2. Resolutions of core surfaces

2.1. Tiled surfaces and core surfaces

In this §2.1, we summarize some definitions and results from [MP22a].⁸

⁷Example 3.16 and Figure 3.1 in [PZ22] illustrate that the coefficient of $n^{\chi(Y)}$ in $\mathbb{E}_n^{\text{emb}}(Y)$ for Y boundary reduced may indeed be strictly larger than 1. In fact, [PZ22] proves it is always a positive integer.

⁸A significant chunk of [MP22a] was part of the first version of the current paper. As we believe the theory of core surfaces is of independent interest and in order to keep the current paper to a manageable size, we decided to develop an expanded version of this theory in a separate paper.

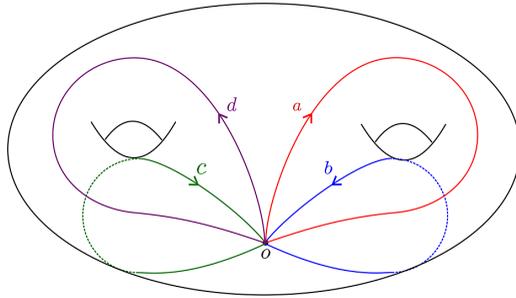


Figure 2.1. *The fixed CW-structure on Σ_2 .*

2.1.1. Tiled surfaces

Consider the construction of the surface Σ_g from a $4g$ -gon by identifying its edges in pairs according to the pattern $a_1 b_1 a_1^{-1} b_1^{-1} \dots a_g b_g a_g^{-1} b_g^{-1}$. This gives rise to a CW-structure on Σ_g consisting of one vertex (denoted o), $2g$ oriented 1-cells (denoted $a_1, b_1, \dots, a_g, b_g$) and one 2-cell which is the $4g$ -gon glued along $4g$ 1-cells.⁹ See Figure 2.1 (in our running examples with $g = 2$, we denote the generators of Γ_2 by a, b, c, d instead of a_1, b_1, a_2, b_2). We identify Γ_g with $\pi_1(\Sigma_g, o)$ so that in the presentation (1.1), words in the generators a_1, \dots, b_g correspond to the homotopy class of the corresponding closed paths based at o along the 1-skeleton of Σ_g .

Note that every covering space $p: Y \rightarrow \Sigma_g$ inherits a CW-structure from Σ_g : The vertices are the preimages of o , and the open 1-cells (2-cells) are the connected components of the preimages of the open 1-cells (2-cells, respectively) in Σ_g . In particular, this is true for the universal covering space $\widehat{\Sigma}_g$ of Σ_g , which we can now think of as a CW-complex. A subcomplex of a CW-complex is a subspace consisting of cells such that if some cell belongs to the subcomplex, then so do the cells of smaller dimension at its boundary.

Definition 2.1 (Tiled surface). [MP22a, Def. 3.1] A tiled surface Y is a subcomplex of a (not-necessarily-connected) covering space of Σ_g . In particular, a tiled surface is equipped with the restricted covering map $p: Y \rightarrow \Sigma_g$ which is an immersion. 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 $4g$ -gons.

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 $(a_1, b_1, \dots, a_g, b_g)$ are pulled back from Σ_g via p . These labels uniquely determine p as a combinatorial map between complexes. Figures 2.1 and 2.3 feature examples of tiled surfaces.

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

Definition 2.2 (Thick version of a tiled surface). [MP22a, Def. 3.2] Given a tiled surface Y which is a subcomplex of the covering space Y of Σ_g , adjoin to Y a small, closed, tubular neighborhood in Y around every edge and a small closed disc in Y around every vertex. The resulting closed surface, possibly with boundary, 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 $4g$ -gon, it is counted twice).

In particular, $\mathfrak{d}(Y) = 2\mathfrak{e}(Y) - 4g\mathfrak{f}(Y)$. We stress that we do not think of Y as a subcomplex 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_g . See [MP22a, §3] for a more detailed discussion.

⁹We use the terms vertices and edges interchangeably with 0-cells and 1-cells, respectively.

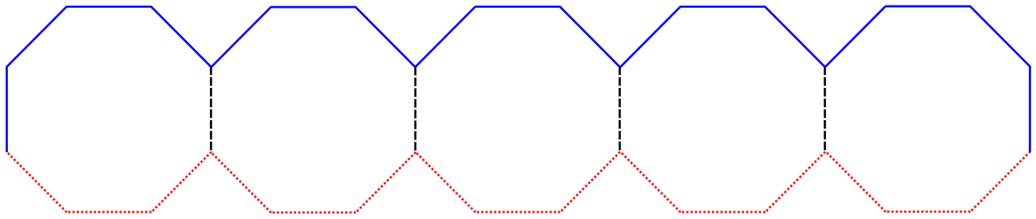


Figure 2.2. A long chain of total length 17 (blocks of sizes 4, 3, 3, 3, 4, in blue) and its complement of length 15 (in red).

It is occasionally useful, for example in Section 5, to augment the tiled surface Y by adding some new half-edges. Here, formally, a half-edge is a copy of the interval $[0, \frac{1}{2})$ which is an (open) half of an edge of a covering space of Σ_g .

Definition 2.3 (Tiled surface with hanging half-edges). [MP22a, §§3.2] Let Y be a tiled surface which is a subcomplex of the covering space $p: \Upsilon \rightarrow \Sigma_g$. We denote by Y_+ the tiled surface Y together with half-edges of Υ which do not belong to Y but are incident to vertices of Y . Every half-edge of Y_+ added to Y in this manner is called a *hanging half-edge*. The thick version of Y_+ is, as above, Y_+ together with a small, closed, tubular neighborhood in Υ around every edge or hanging half-edge, and a small closed disc in Υ around every vertex. We denote by ∂Y_+ the boundary of the thick version of Y_+ .

Note that there are exactly $4g$ half-edges incident to every vertex in Y_+ : Some of them originate from edges in Y and some are hanging half-edges.

Morphisms of tiled surfaces

If Y_1 and Y_2 are tiled surfaces, a *morphism* from Y_1 to Y_2 is a map of CW-complexes which maps i -cells to i -cells for $i = 0, 1, 2$ and respects the directions and labels of edges. Equivalently, this is a morphism of CW-complexes which commutes with the restricted covering maps $p_j: Y_j \rightarrow \Sigma_g$ ($j = 1, 2$). In particular, the restricted covering map from a tiled surface to Σ_g is itself a morphism of tiled surfaces. It is an easy observation that every morphism of tiled surfaces is an immersion (locally injective).

2.1.2. Blocks and chains

Some of the notions we use below are taken from [BS87]. See [MP22a, §§3.2] for a more detailed account.

Given a covering space Υ of Σ_g , every path in the 1-skeleton $\Upsilon^{(1)}$ corresponds to a word in $\{a_1^{\pm 1}, \dots, b_g^{\pm 1}\}$. A path that follows a (part of the) boundary of a $4g$ -gon is called a **block**. If it has length at least $2g + 1$ it is called a **long block**, and if it has length exactly $2g$, it is called a **half-block**. If a (noncyclic) block of length b sits along the boundary of a $4g$ -gon O , the **complement of the block** is the block of length $4g - b$ consisting of the complement set of edges along O , so the block and its complement share the same starting point and the same terminal point.

A **chain** is a path in $\Upsilon^{(1)}$ that consists of a sequence of blocks b_1, \dots, b_r such that if the last vertex of b_i and the first vertex of b_{i+1} is v , there is exactly one edge incident to v between the last edge of b_i and the first edges of b_{i+1} . In other words, if the $4g$ -gons corresponding to the blocks b_1, \dots, b_r are O_1, \dots, O_r , then O_i and O_{i+1} share an edge e with an endpoint v , and b_i ends at v and b_{i+1} starts at v . See Figure 2.2. A **long chain** is a chain with corresponding blocks of lengths

$$2g, 2g - 1, 2g - 1, \dots, 2g - 1, 2g.$$

A **half-chain** is a *cyclic* chain (so the corresponding path is closed) consisting of blocks each of which is of length $2g - 1$. The **complement of a long chain** is the chain with blocks of lengths $2g - 1, 2g - 1, \dots, 2g - 1$ which sits along the other side of the $4g$ -gons bordering the long chain and with the same starting point and endpoint. Note that the complement of a long chain is shorter by two

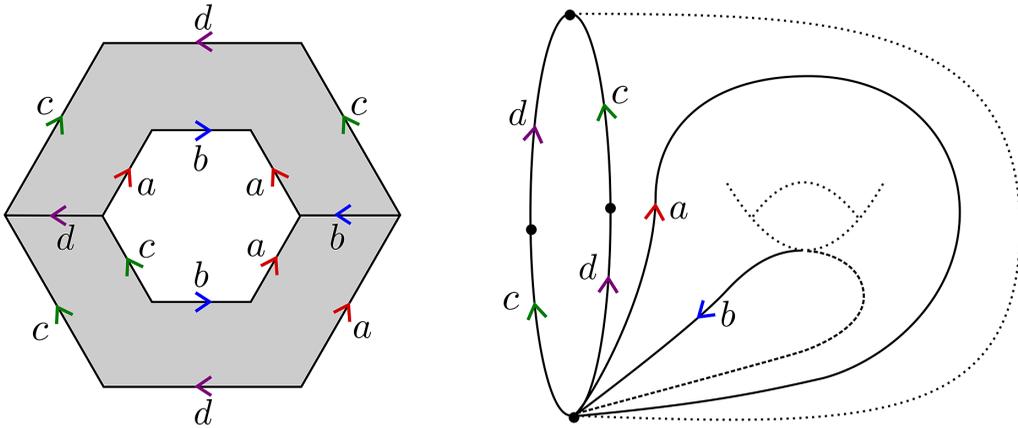


Figure 2.3. Fix $g = 2$, and let $\Gamma_2 = \langle a, b, c, d \mid [a, b][c, d] \rangle$. On the left is the core surface $\text{Core}(\langle aba^{-2}b^{-1}c \rangle)$. It consists of 12 vertices, 14 edges and two octagons and topologically it is an annulus. On the right is the core surface $\text{Core}(\langle a, b \rangle)$. It consists of four vertices, six edges and one octagon and topologically it is a genus-1 torus with one boundary component.

edges from the long chain (see Figure 2.2). The **complement of a half-chain** is defined as follows. If the half-chain sits along the boundary of the $4g$ -gons O_1, \dots, O_r , its complement is the half-chain sitting along the other sides of these $4g$ -gons: A block (of length $2g - 1$) of the half-chain along O_i is replaced by the path of length $2g - 1$ along O_i , with starting and terminal points one edge away from the starting and terminal points, respectively, of the block. The complement of a half-chain has the same length as the original half-chain. The left part of Figure 2.3 illustrates two complementing half-chains of length 6 each (with two octagons in between).

A **boundary cycle** of Y is a cycle in $Y^{(1)}$ corresponding to an oriented boundary component of the thick version of Y (see Definition 2.2). We always choose the orientation so that there are no $4g$ -gons to the immediate **left** of the boundary component as it is traversed. Therefore, boundary components of Y correspond to unique cycles. Note that $\mathfrak{d}(Y)$ is equal to the sum over boundary cycles of Y of the number of edges in each such cycle.

2.1.3. Boundary reduced and strongly boundary reduced tiled surfaces

The following definitions came up from our results in representation theory in §5, but they fit perfectly with classical results in combinatorial group theory [Deh12] and in particular with [BS87].

Definition 2.4 (Boundary reduced). A tiled surface Y is *boundary reduced* if no boundary cycle of Y contains a long block or a long chain.

In particular, if Y is boundary reduced, then every path that reads $[a_1, b_1] \dots [a_g, b_g]$ is not only closed, but there is also a $4g$ -gon attached to it. We also need a stronger version of this property.

Definition 2.5 (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.

Because a long block contains (at least two) half-blocks and a long chain contains (two) half-blocks, a strongly boundary reduced tiled surface is in particular boundary reduced. The relevance of the notions of being (strongly) boundary reduced is that our techniques for estimating $\mathbb{E}_n^{\text{emb}}(Y)$ for a tiled surface Y only give the right type of estimates when Y is boundary reduced – see Proposition 5.25. If Y is strongly boundary reduced we get even better estimates – see Proposition 5.26.

Let Y be a compact tiled surface which is a subcomplex of the covering space \mathbb{Y} of Σ_g . In [MP22a, §4], we describe the ‘boundary reduced closure’ $\text{BR}(Y \hookrightarrow \mathbb{Y})$ of Y in \mathbb{Y} which is the smallest intermediate tiled surface which is boundary reduced. By (ibid., Proposition 4.6), $\text{BR}(Y \hookrightarrow \mathbb{Y})$ is compact too. Likewise, $\text{SBR}(Y \hookrightarrow \mathbb{Y})$, the strongly boundary reduced closure, is the smallest intermediate tiled

surface which is strongly boundary reduced, but this one is not always compact. Our resolutions in Section 2.3 are based on a ‘compromise’ between these two types of closures.

2.1.4. Core surfaces

Finally, let us define the main object which was introduced and analyzed in [MP22a], with motivation coming from the current paper. In analogy to Stallings core graphs and their role in the study of free groups and their subgroups, we introduced the notion of *core surfaces* which relates to subgroups of Γ_g :

Definition 2.6 (Core surfaces). [MP22a, Def. 1.1] Given a subgroup $1 \neq J \leq \Gamma_g = \pi_1(\Sigma_g, o)$, consider the covering space $p: Y \rightarrow \Sigma_g$ corresponding to J , so $Y = J \backslash \widetilde{\Sigma}_g$. Define the **core surface of J** , denoted $\text{Core}(J)$, as the tiled surface which is a subcomplex of Y as follows: (i) take the union of all shortest-representative cycles in the 1-skeleton $Y^{(1)}$ of every free-homotopy class of essential closed curve in Y , and (ii) add every connected component of the complement which contains finitely many $4g$ -gons.

For completeness define the core surface of the trivial subgroup to be the zero-dimensional tiled surface consisting of a single vertex.

Note that the quotient $Y = J \backslash \widetilde{\Sigma}_g$ is invariant under conjugation of J , so $\text{Core}(J)$ depends only on the conjugacy class of J in Γ_g . Figure 2.3 gives two examples of core surfaces. As another example, if J is of finite index in Γ , $\text{Core}(J)$ is identical to $J \backslash \widetilde{\Sigma}_g$ and is a compact closed surface.

In [MP22a], we give an intrinsic definition of a core surface and show there is one-to-one correspondence between core surfaces (labeled by a_1, \dots, b_g) and conjugacy classes of subgroups of Γ_g , we provide a ‘folding process’ to construct $\text{Core}(J)$ from a finite generating set of J (provided, of course, that J is f.g.) and prove basic properties of core surfaces. In particular, $\text{Core}(J)$ is connected and strongly boundary reduced (ibid., Proposition 5.3), and whenever J is f.g., $\text{Core}(J)$ is compact (ibid., Proposition 5.8). We also show that whenever $H \leq J \leq \Gamma_g$, the natural morphism between the corresponding covering spaces $H \backslash \widetilde{\Sigma}_g \rightarrow J \backslash \widetilde{\Sigma}_g$, restricts to a morphism $\text{Core}(H) \rightarrow \text{Core}(J)$.

2.2. Expectations and probabilities of tiled surfaces

Correspondence between $\text{Hom}(\Gamma_g, S_n)$ and n -sheeted covering spaces of Σ_g

Let M be a connected topological space with basepoint m . Consider n -sheeted covering spaces of M with the fiber above m labeled by $[n]$ so that every point in the fiber has a different label. If M is ‘nice enough’, in particular if M is a surface, there is a one-to-one correspondence between these labeled n -sheeted covering spaces and the set of homomorphisms $\text{Hom}(\pi_1(M, m), S_n)$ (see, for instance, [Hat05, pp. 68-70]). If \hat{M} is a labeled n -sheeted covering space and $p: \hat{M} \rightarrow M$ the covering map, the corresponding homomorphism $\theta: \pi_1(M, m) \rightarrow S_n$ is given as follows: For $h \in \pi_1(M, m)$, consider γ , a closed path in M , based at m , which represents h . Then $\theta(h)(i) = j$ if and only if the lift of γ at the point i ends at the point j .¹⁰

In our case, this translates to a one-to-one correspondence between the representation space $\mathbb{X}_{g,n} = \text{Hom}(\Gamma_g, S_n)$ and labeled n -sheeted covering spaces of Σ_g (pointed at o). For $\phi: \Gamma_g \rightarrow S_n$, denote the corresponding covering space by $p_\phi: X_\phi \rightarrow \Sigma_g$. As explained above, X_ϕ inherits a CW-structure from Σ_g and is, therefore, a tiled surface. The fiber above o is precisely the vertices of X_ϕ , and they are labeled by $[n]$ in this construction.

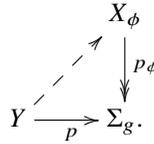
Expected number of fixed points as expected number of lifts

Given a compact tiled surface Y , we are interested in the expected number of morphisms from Y to a random n -covering of Σ_g , namely, in

$$\mathbb{E}_n(Y) \stackrel{\text{def}}{=} \mathbb{E}_{\phi \in \mathbb{X}_{g,n}} [\# \{ \text{morphisms } Y \rightarrow X_\phi \}],$$

¹⁰There is a subtle issue here with the direction in which permutations are multiplied in S_n . The map $\theta: \pi_1(M, m) \rightarrow S_n$ as defined here is a homomorphism only if permutations in S_n are composed from right to left. We refer to this issue in the case of $M = \Sigma_2$ in the beginning of Section 5.

where ϕ is sampled uniformly at random from $\mathbb{X}_{g,n}$. Equivalently, this is the expected number of lifts of the restricted covering map $p : Y \rightarrow \Sigma_g$ to the random n -covering X_ϕ :



Note that if Y is connected and $\phi \in \mathbb{X}_{g,n}$, the number of morphisms $Y \rightarrow X_\phi$ is at most n , as any vertex of Y can be lifted to one of the n vertices of X_ϕ , and each such lift can be extended in at most one way to a lift of the whole of Y . For suitable choices of Y , $\mathbb{E}_n(Y)$ is equal to the quantities $\mathbb{E}_{g,n}[\text{fix}_Y]$ and $\mathbb{E}_{g,n}[\text{fix}_J]$ that feature in our main theorems (Theorems 1.1, 1.2 and 1.3):

Lemma 2.7. *Let Y be a connected tiled surface and $p : Y \rightarrow \Sigma_g$ the restricted covering map. For arbitrary vertex $y \in Y$, assume that $p_*(\pi_1(Y, y))$ is conjugate to $J \leq_{\text{f.g.}} \Gamma_g$ (as a subgroup of $\pi_1(\Sigma_g, o) = \Gamma_g$). Then for all $n \in \mathbb{N}$,*

$$\mathbb{E}_n(Y) = \mathbb{E}_{g,n}[\text{fix}_J].$$

In particular, for $J \leq_{\text{f.g.}} \Gamma_g$,

$$\mathbb{E}_n(\text{Core}(J)) = \mathbb{E}_{g,n}[\text{fix}_J].$$

Proof. In fact, the equality holds at the level of the individual representation $\phi \in \mathbb{X}_n = \text{Hom}(\Gamma_g, S_n)$: The number of morphisms $Y \rightarrow X_\phi$ is equal to the number of common fixed points $\text{fix}_J(\phi)$. Indeed, because the number of common fixed points of $\phi(J)$ is the same as the number of fixed points of any conjugate, we may assume without loss of generality that $p_*(\pi_1(Y, y)) = J$. Now, $i \in [n]$ is a common fixed point of $\phi(J)$ if and only if $J \leq \pi_1(X_\phi, v_i)$, where v_i is the vertex of X_ϕ labeled i , and $\pi_1(X_\phi, v_i)$ is identified with the subgroup

$$(p_\phi)_*(\pi_1(X_\phi, v_i)) \leq \Gamma_g.$$

By standard facts from the theory of covering spaces [Hat05, Prop. 1.33 and 1.34], there is a lift of p to X_ϕ mapping the vertex y to v_i if and only if (the image in Γ_g of) $\pi_1(Y, y)$ is contained in (the image in Γ_g of) $\pi_1(X_\phi, v_i)$, and this lift, if exists, is unique.

The statement about core surfaces follows from the fact that (the image in Γ_g of) $\pi_1(\text{Core}(J))$ is conjugate to J [MP22a, Prop. 5.3]. □

Another type of expectation will also feature in this work. Given a compact tiled surface Y , denote

$$\mathbb{E}_n^{\text{emb}}(Y) \stackrel{\text{def}}{=} \mathbb{E}_{\phi \in \mathbb{X}_{g,n}} [\#\{\text{injective morphisms } Y \rightarrow X_\phi\}],$$

where the expectation is over a uniformly random $\phi \in \mathbb{X}_{g,n}$.

2.3. Resolutions

Definition 2.8 (Resolutions). A resolution \mathcal{R} of a tiled surface Y is a collection of morphisms of tiled surfaces

$$\mathcal{R} = \{f : Y \rightarrow W_f\},$$

such that every morphism $h : Y \rightarrow Z$ of Y into a tiled surface Z with no boundary decomposes uniquely as $Y \xrightarrow{f} W_f \xrightarrow{\bar{h}} Z$, where $f \in \mathcal{R}$ and \bar{h} is an embedding.

The purpose of introducing resolutions is the following obvious lemma. Recall the notation $\mathbb{E}_n(Y)$ and $\mathbb{E}_n^{\text{emb}}(Y)$ from Section 2.2.

Lemma 2.9. *If Y is a compact tiled surface and \mathcal{R} is a finite resolution of Y , then*

$$\mathbb{E}_n(Y) = \sum_{f \in \mathcal{R}} \mathbb{E}_n^{\text{emb}}(W_f). \tag{2.1}$$

Our main goal in the rest of this subsection is to prove the existence of a finite resolution whenever we are given a compact tiled surface Y – this is the content of Theorem 2.14 below. This resolution will consist strictly of boundary reduced tiled surfaces W_f , and some of these will even be strongly boundary reduced. We shall make use of Theorem 2.14 mainly for Y a core surface of a finitely generated subgroup of Γ . In this case, the resolution we construct has even nicer properties – see Proposition 2.15.

Ideally, we would have liked to get a resolution where all the elements are strongly boundary reduced. Unfortunately, such a resolution does not always exist. For example, when $g = 2$ and $\Gamma_2 = \langle a, b, c, d \mid [a, b][c, d] \rangle$, the core surface $Y = \text{Core}(\langle [a, b] \rangle)$ does not admit such a resolution as can be inferred from [MP22a, Fig. 4.2].

To prove the existence of a resolution with nice properties, we first define a process which outputs a ‘compromise’ between the BR-closure of a tiled surface and the SBR-closure, introduced in [MP22a, §4].

Definition 2.10. Fix $\chi_0 \in \mathbb{Z}$. Assume that $h: Y \rightarrow Z$ is a morphism between tiled surfaces where Y is compact and Z has no boundary. Let W_0 denote the h -image of Y in Z . Set $i = 0$. Perform the following algorithm we call the *growing process*:

1. If one of the following conditions holds:
 - (a) W_i is strongly boundary reduced, or
 - (b) W_i is boundary reduced and $\chi(W_i) < \chi_0$,
 terminate and return $h: Y \rightarrow W_i$.
2. Obtain W_{i+1} from W_i by adding to W_i (the closure of) every $4g$ -gon in $Z \setminus W_i$ which touches along its boundary an edge of ∂W_i which is part of a half-block (this includes the case of a long block), a long chain or a half-chain. Set $i := i + 1$ and return to item 1.

It is clear that every step of this process is deterministic. Note that if the process ends when W_i is strongly boundary reduced, then W_i is the unique smallest strongly boundary reduced tiled surface inside Z containing W_0 , denoted $\text{SBR}(Y \hookrightarrow Z)$ [MP22a, §4]. (In general, $\text{SBR}(Y \hookrightarrow Z)$ is not always compact, but in this case it is.) The growing process always terminates after finitely many steps.

Lemma 2.11. *The process described in Definition 2.10 always terminates.*

Proof. Let $\text{he}(W_i)$ denote the number of hanging half-edges along the boundary of $(W_i)_+$, and consider the triple

$$(\mathfrak{d}(W_i), \chi(W_i), -\text{he}(W_i)). \tag{2.2}$$

For every i , W_i is a compact subsurface of Z , and so the three quantities are well-defined integers. We claim that at every step in the growing process, the triple (2.2) strictly reduces with respect to the lexicographic order.

Indeed, assume we do not halt after i steps, and let O_1, \dots, O_k be the list of $4g$ -gons in $Z \setminus W_i$ which are added to W_i in order to obtain W_{i+1} . By the choice of $4g$ -gons, it is clear that $\mathfrak{d}(W_{i+1}) \leq \mathfrak{d}(W_i)$. If the inequality is strict, we are done. So assume $\mathfrak{d}(W_{i+1}) = \mathfrak{d}(W_i)$. This means that $\partial(W_i)$ contains no long blocks nor long chains, so it is boundary reduced, and that the edges in the complements in Z of the half-blocks and half-chains at $\partial(W_i)$ all belong to $\partial(W_{i+1})$. In other words, let p_1, \dots, p_m be these complements in Z : So p_j is either a half-block or a half-chain. The equality $\mathfrak{d}(W_{i+1}) = \mathfrak{d}(W_i)$ means that all the edges in p_1, \dots, p_m belong to ∂W_{i+1} .

It is easy to see that in this case $\chi(W_{i+1}) \leq \chi(W_i)$: The number of new $4g$ -gons and vertices in W_{i+1} at most balances the number of new edges. Let V denote the set of internal vertices in p_1, \dots, p_m (so not at their endpoints). We have strict inequality $\chi(W_{i+1}) < \chi(W_i)$ if and only if some $v \in V$ belongs to W_i or to two different complements from p_1, \dots, p_m .

Now, assume that $\mathfrak{d}(W_{i+1}) = \mathfrak{d}(W_i)$ and $\chi(W_{i+1}) = \chi(W_i)$. Then W_i is boundary reduced and each of the complements p_1, \dots, p_m is a connected piece of ∂W_{i+1} . If O_j touches a half-block of ∂W_i , its annexation adds a net of $(2g - 1)(4g - 2) - 2 = 8g(g - 1)$ hanging half-edges. Every $4g$ -gon along a half-chain of ∂W_i also adds on average a net of $8g(g - 1)$ hanging half-edges. So if we add at least one $4g$ -gon at the $(i + 1)$ st step, $-\mathfrak{h}e$ strictly decreases. So indeed the triple (2.2) strictly decreases lexicographically in every step.

Finally, there are at most finitely many steps in which $\mathfrak{d}(W_i)$ decreases because this is a nonnegative integer. So it is enough to show there cannot be infinitely many steps in which $\mathfrak{d}(W_i)$ is constant. If $\mathfrak{d}(W_{i+1}) = \mathfrak{d}(W_i)$, then W_i is boundary reduced. If $\chi(W_i)$ keeps decreasing, then eventually we hit the bound $\chi(W_i) < \chi_0$ and halt. If $\mathfrak{d}(W_i)$ and $\chi(W_i)$ are constant, then $\mathfrak{h}e(W_i)$ increases constantly, but in every tiled surface W , $\mathfrak{h}e(W) \leq 4g\mathfrak{d}(W)$, so there cannot be infinitely many steps of this type too. This proves the lemma. \square

Lemma 2.12. *There is a bound $B = B(Y)$, independent of $h: Y \rightarrow Z$, such that in the entire growing process, at most $B = B(Y)$ $4g$ -gons are added to W_0 .*

Proof. Note that every boundary edge of W_0 is necessarily an h -image of a boundary edge of Y so that $\mathfrak{d}(W_i) \leq \mathfrak{d}(W_0) \leq \mathfrak{d}(Y)$. In every step, we add at most $\frac{\mathfrak{d}(W_i)}{2g-1} \leq \frac{\mathfrak{d}(Y)}{2g-1}$ $4g$ -gons. So it is enough to bound the number of steps performed in the growing process until it terminates. There are at most $\mathfrak{d}(Y)$ steps in which $\mathfrak{d}(W_i)$ strictly decreases¹¹ ($\mathfrak{d}(W_{i+1}) < \mathfrak{d}(W_i)$), so there are at most $\mathfrak{d}(Y) + 1$ possible values of $\mathfrak{d}(W_i)$. In steps where $\mathfrak{d}(W_i)$ is unchanged, W_i is boundary reduced, so by definition $\chi(W_i) \geq \chi_0$ (otherwise, the process terminates). Let $\pi_0(Y)$ denote the number of connected components of Y . For all i , W_i is a subsurface of Z with at most $\pi_0(Y)$ connected components, and by the classification of surfaces, $\chi(W_i) \leq 2\pi_0(Y)$. There are at most $2\pi_0(Y) - \chi_0$ steps with $\mathfrak{d}(W_i)$ fixed and $\chi(W_i)$ strictly decreasing. Finally, when $\mathfrak{d}(W_i)$ is constant there are at most $2\pi_0(Y) - \chi_0 + 1$ possible values of $\chi(W_i)$, and if $\mathfrak{d}(W_{i+1}) = \mathfrak{d}(W_i)$ and $\chi(W_{i+1}) = \chi(W_i)$, then $\mathfrak{h}e(W_{i+1}) \geq \mathfrak{h}e(W_i) + 8g(g - 1)$ and $\mathfrak{h}e(W_i) \leq 4g\mathfrak{d}(W_i) \leq 4g\mathfrak{d}(Y)$, so there are at most $\frac{\mathfrak{d}(Y)}{2(g-1)}$ steps with the same value of $\mathfrak{d}(W_i)$ and $\chi(W_i)$. Overall there are at most

$$\mathfrak{d}(Y) + (\mathfrak{d}(Y) + 1) \left[(2\pi_0(Y) - \chi_0) + (2\pi_0(Y) - \chi_0 + 1) \cdot \frac{\mathfrak{d}(Y)}{2(g - 1)} \right] \tag{2.3}$$

steps in the growing process. Define $B(Y)$ to be $\frac{\mathfrak{d}(Y)}{2g-1}$ times equation (2.3). \square

We can now define the sought-after resolution for compact tiled surfaces.

Definition 2.13. Suppose that Y is a compact tiled surface and $\chi_0 \in \mathbb{Z}$ a fixed integer. Define the χ_0 -resolution of Y to be the collection

$$\mathcal{R} = \mathcal{R}(Y, \chi_0) = \{f: Y \rightarrow W_f\}$$

obtained from all possible morphisms $h: Y \rightarrow Z$ from Y to a tiled surface Z with no boundary via the growing process (the process applied with the parameter χ_0).

Theorem 2.14. *Suppose Y is a compact tiled surface and $\chi_0 \in \mathbb{Z}$ a fixed integer. The collection $\mathcal{R} = \mathcal{R}(Y, \chi_0)$ from Definition 2.13 is a finite resolution of Y which satisfies further*

- R1** for every $f \in \mathcal{R}$, the tiled surface W_f is compact and boundary reduced, and
- R2** for every $f \in \mathcal{R}$ with $\chi(W_f) \geq \chi_0$, the tiled surface W_f is strongly boundary reduced.

¹¹In fact, there are at most $\frac{\mathfrak{d}(Y)}{2}$ such steps as $\mathfrak{d}(W_i) = 2e(W_i) - 4gf(W_i)$ is always even.

Proof. By Lemma 2.11 and the halting conditions of the growing process, it is clear that every such morphism in \mathcal{R} satisfies **R1** and **R2**. Given a morphism $h: Y \rightarrow Z$ as in Definition 2.13, $h(Y)$ (also named W_0) is a quotient of Y and therefore the number of cells in $h(Y)$ is bounded. From Lemma 2.12, we now conclude that there is a bound on the number of cells in any W_f with $f \in \mathcal{R}$. This shows that \mathcal{R} is finite as there are finitely many tiled surfaces with given bounds on the number of cells and finitely many morphisms between two given compact tiled surfaces.

It remains to show that \mathcal{R} is a resolution. By the way it was constructed, it is clear that every morphism $h: Y \rightarrow Z$ with $\partial Z = \emptyset$ decomposes as

$$Y \xrightarrow{f} W_f \hookrightarrow Z \tag{2.4}$$

and that $f \in \mathcal{R}$. To show uniqueness, assume that h decomposes in an additional way

$$Y \xrightarrow{\varphi} W_\varphi \hookrightarrow Z, \tag{2.5}$$

where W_φ is the result of the growing process for some $h': Y \rightarrow Z'$ with $\partial Z' = \emptyset$. We claim that equations (2.4) and (2.5) are precisely the same decompositions of h . Indeed, the growing process defined by $h': Y \rightarrow Z'$ takes place entirely inside W_φ and does not depend on the structure of $Z' \setminus W_\varphi$: In the $(i + 1)$ st step of the growing process, the decision whether or not to annex more $4g$ -gons and where depends only on the structure and boundary of W_i . Consequently, the growing process defined by the morphism $h': Y \rightarrow Z'$ has the exact same output, in terms of the resulting element we add to \mathcal{R} , as the growing process defined by the composition $Y \xrightarrow{\varphi} W_\varphi \hookrightarrow Z$. But because the growing process is deterministic, the latter is identical to the growing process defined by $h: Y \rightarrow Z$. \square

As mentioned above, we will use Theorem 2.14 mainly with Y being a core surface. In this case, the theorem can be strengthened as follows. Recall from Section 1 that, given $J \leq_{f.g.} \Gamma_g$, we denote by $\mathfrak{MOG}(J)$ the set of f.g. overgroups of J with maximal Euler characteristic, and by $\chi_{\max}(J)$ this maximal Euler characteristic.

Proposition 2.15 (Addendum to Theorem 2.14). *Let $J \leq_{f.g.} \Gamma_g$, and let $\chi_0 \in \mathbb{Z}$. Let $\mathcal{R}_{J, \chi_0} = \{f: \text{Core}(J) \rightarrow W_f\}$ be the resolution $\mathcal{R}(\text{Core}(J), \chi_0)$ from Definition 2.13. Then \mathcal{R}_{J, χ_0} satisfies further the following two properties.*

- R2** *For every $f \in \mathcal{R}_{J, \chi_0}$ with $\chi(W_f) \geq \chi_0$, the tiled surface W_f is the core surface of some $K \leq_{f.g.} \Gamma_g$ with $J \leq K$ and f is the natural morphism between the two core surfaces (the restriction of $J \setminus \widetilde{\Sigma}_g \rightarrow K \setminus \widetilde{\Sigma}_g$).*
- R4** *Assume that $\chi_0 \leq \chi_{\max}(J)$. Then for every $K \in \mathfrak{MOG}(J)$, the natural morphism $\text{Core}(J) \rightarrow \text{Core}(K)$ belongs to \mathcal{R}_{J, χ_0} .*

For $K \leq_{f.g.} \Gamma_g$, we have $\chi(K) = \chi(\text{Core}(K))$ [MP22a, Prop. 5.3]. Proposition 2.15 thus shows that, as long as $\chi_0 \leq \chi_{\max}(J)$, there is a bijection between the elements of $\mathfrak{MOG}(J)$ and the elements in the resolution with maximal Euler characteristic.

Corollary 2.16. *For every $J \leq_{f.g.} \Gamma_g$, the set $\mathfrak{MOG}(J)$ of f.g. overgroups of maximal Euler characteristic is finite.*

Proof of Proposition 2.15. Suppose that $f: \text{Core}(J) \rightarrow W_f$ satisfies $\chi(W_f) \geq \chi_0$. In particular, W_f is strongly boundary reduced by **R2**. Let $j \in \text{Core}(J)$ be a vertex and assume without loss of generality that $J = p_*(\pi_1(\text{Core}(J), j))$ (the fact that $p_*(\pi_1(\text{Core}(J), j)) \leq \pi_1(\Sigma_g, o) = \Gamma_g$ is conjugate to J follows from [MP22a, Prop. 5.3 and Cor. 4.11]). Define $K \stackrel{\text{def}}{=} \pi_1(W_f, f(j))$. As W_f is compact, K is finitely generated. Let $p_K: (K \setminus \widetilde{\Sigma}_g, k) \rightarrow \Sigma_g$ be the pointed covering space with $\pi_1(K \setminus \widetilde{\Sigma}_g, k) = K$. By the unique lifting property from the theory of covering spaces [Hat05, Prop. 1.33 and 1.34], as W_f

is connected, there is a unique lift α of the restricted covering map $p_{W_f} : W_f \rightarrow \Sigma_g$ to $K \backslash \widetilde{\Sigma}_g$, which maps $f(j)$ to k . We will show that α gives an isomorphism between W_f and $\text{Core}(K)$.

$$\begin{array}{ccc}
 (\text{Core}(K), k) \xrightarrow{\hookrightarrow} (K \backslash \widetilde{\Sigma}_g, k) & & (\text{Core}(K), k) \xrightarrow{\hookrightarrow} (K \backslash \widetilde{\Sigma}_g, k) \\
 \nearrow & \exists! \alpha \nearrow & \downarrow p_K \\
 (\text{Core}(J), j) \xrightarrow{f} (W_f, f(j)) \xrightarrow{p_{W_f}} (\Sigma_g, o) & & (W_f, f(j)) \xrightarrow{\hookrightarrow} (Z, f(j)) \\
 & & \downarrow m
 \end{array}$$

First, we show that $\alpha(W_f) \subseteq \text{Core}(K)$. Recall that W_f is the result of the growing process for some $h : \text{Core}(J) \rightarrow Z$ with $\partial Z = \emptyset$. Consider $W_0 \stackrel{\text{def}}{=} h(\text{Core}(J)) \subseteq Z$. Recall that $W_f = \text{SBR}(W_0 \hookrightarrow Z)$. By the unique lifting property, $\alpha \circ f$ is the natural morphism $\text{Core}(J) \rightarrow K \backslash \widetilde{\Sigma}_g$, which, by [MP22a, Lem. 5.4], has an image contained in $\text{Core}(K)$. Hence, $\alpha(W_0) \subseteq \text{Core}(K)$. As $\text{Core}(K)$ is strongly boundary reduced [MP22a, Prop. 5.3], we have that $\text{SBR}(\alpha(W_0) \hookrightarrow K \backslash \widetilde{\Sigma}_g)$ is contained in $\text{Core}(K)$. By [MP22a, Lem. 4.7],

$$\alpha(W_f) = \alpha(\text{SBR}(W_0 \hookrightarrow Z)) \subseteq \text{SBR}(\alpha(W_0) \hookrightarrow K \backslash \widetilde{\Sigma}_g) \subseteq \text{Core}(K).$$

Now, Z is a covering space of Σ_g and we may assume it is connected (because $\text{Core}(J)$ is). Thus, Z is identical to $L \backslash \widetilde{\Sigma}_g$ for some $L = \pi_1(Z, h(j))$. By property **R2**, W_f is strongly boundary reduced, and so by [MP22a, Cor. 4.11] its embedding in Z is π_1 -injective. In other words, $K \leq L$ and, therefore, there is a morphism $m : (K \backslash \widetilde{\Sigma}_g, k) \rightarrow (Z, f(j))$. By the unique lifting property, the composition $m \circ \alpha : (W_f, f(j)) \rightarrow (Z, f(j))$ must be identical to the embedding $(W_f, f(j)) \hookrightarrow (Z, f(j))$ and therefore α is injective. So $\alpha(W_f)$ is a strongly boundary reduced subtiled surface of $\text{Core}(K)$ with fundamental group K . By [MP22a, Lem. 5.7], it follows that $\alpha(W_f) \supseteq \text{Core}(K)$. We conclude that $\alpha : W_f \rightarrow \text{Core}(K)$ is an isomorphism, and **R3** is proven.

To prove **R4**, suppose that $K \in \mathfrak{M}\mathfrak{D}\mathfrak{G}(J)$. Let $h : \text{Core}(J) \rightarrow K \backslash \widetilde{\Sigma}_g$ be the natural morphism. By the definition of the resolution \mathcal{R}_J , the morphism h factors as $\text{Core}(J) \xrightarrow{f} W_f \hookrightarrow K \backslash \widetilde{\Sigma}_g$ for some $f \in \mathcal{R}_J$. Because $h(\text{Core}(J)) \subseteq \text{Core}(K)$ (by [MP22a, Lem. 5.4]) and because $\text{Core}(K)$ is strongly boundary reduced, we have $W_f \subseteq \text{Core}(K)$.

Let C be a connected component of the difference between the thick version of $\text{Core}(K)$ and the thick version of W_f . As $\text{Core}(K)$ is compact, \overline{C} is compact. As $\text{Core}(K)$ is connected, \overline{C} must intersect ∂W_f and in particular has at least one boundary component. Since W_f is boundary reduced, \overline{C} is not homeomorphic to a disc, and so $\chi(\overline{C}) \leq 0$. Now,

$$\chi(K) = \chi(\text{Core}(K)) = \chi(W_f) + \sum_C \chi(\overline{C}),$$

the sum being over all connected components as above. We conclude that $\chi(W_f) \geq \chi(K) = \chi_{\max}(J) \geq \chi_0$. By **R2**, W_f is strongly boundary reduced and by **R3**, $W_f = \text{Core}(M)$ for some subgroup M . But then $M \in \mathfrak{M}\mathfrak{D}\mathfrak{G}(J)$, $\chi(\text{Core}(M)) = \chi(K)$, and every connected component C as above satisfies $\chi(\overline{C}) = 0$. As \overline{C} has at least one boundary component, it must be an annulus. But then $\text{Core}(M)$ is a deformation retract of $\text{Core}(K)$, so $M = K$ up to conjugation and so $W_f = \text{Core}(M) = \text{Core}(K)$. \square

Corollary 2.17. *Suppose $1 \neq \gamma \in \Gamma_g$ is a nontrivial element. Let q be the maximal natural number such that $\gamma = \gamma_0^q$ for some $\gamma_0 \in \Gamma_g$, and $d(q)$ the number of positive divisors of q . Then $\text{Core}(\langle \gamma \rangle)$ has a*

finite resolution $\mathcal{R}_\gamma = \{f: \text{Core}(\langle\gamma\rangle) \rightarrow W_f\}$ with W_f boundary reduced for every $f \in \mathcal{R}_\gamma$ and with exactly $d(q)$ elements $f \in \mathcal{R}_\gamma$ with $\chi(W_f) \geq 0$. Moreover, these $d(q)$ elements are precisely

$$f_m: \text{Core}(\langle\gamma\rangle) \rightarrow \text{Core}(\langle\gamma_0^m\rangle) \tag{2.6}$$

for $m \mid q$, where f_m is the natural morphism between the core surfaces.

Proof. Construct $\mathcal{R}_\gamma = \{f: \text{Core}(\langle\gamma\rangle) \rightarrow W_f\}$ as $\mathcal{R}(\text{Core}(\langle\gamma\rangle), 0)$ from Definition 2.13. By Theorem 2.14 and Proposition 2.15, the elements in \mathcal{R}_γ with $\chi(W_f) = 0$ are precisely the core surfaces of the subgroups in $\mathfrak{M}\mathfrak{D}\mathfrak{G}(\langle\gamma\rangle)$. So it only remains to show that $\mathfrak{M}\mathfrak{D}\mathfrak{G}(\langle\gamma\rangle)$ are precisely $\langle\gamma_0^m\rangle$ with $m \mid q$.

But f.g. subgroups $K \leq \Gamma_g$ with $\chi(K) = 0$ are necessarily cyclic. Assume $K = \langle\delta\rangle \in \mathfrak{M}\mathfrak{D}\mathfrak{G}(\langle\gamma\rangle)$, so $\gamma \in \langle\delta\rangle$, and we may assume that γ is a positive power of δ (otherwise switch to δ^{-1}). Every finitely generated subgroup of Γ of infinite index is free (e.g., [Sco78]), and a subgroup of finite index of Γ_g is isomorphic to Γ_h for some $h \geq g$ and so cannot be generated by less than $2g$ elements. We conclude that the subgroup $\langle\delta, \gamma_0\rangle \leq \Gamma$ is free. Because there is a relation $\gamma_0^q = \delta^k$ for some $k \in \mathbb{N}$, it must be a cyclic subgroup. By definition, γ_0 is not a proper power, and so δ must be a positive power of γ_0 , and hence $\delta = \gamma_0^m$ for some $m \mid q$. \square

3. Background: representation theory of the symmetric group

In this section, we give background on the complex representation theory of S_n that will be used in the sequel. We follow the Vershik–Okounkov approach to the representation theory of S_n developed in [VO96].

3.1. Young diagrams

A *partition* is a sequence $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_\ell)$ with each $\lambda_i \in \mathbb{N}$ and $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_\ell$. If $\sum \lambda_i = n$, we write this as $\lambda \vdash n$. Such partitions are in one-to-one correspondence with *Young diagrams* (YD). A YD consists of a collection of left-aligned rows of identical square boxes, where the number of boxes in each row is nonincreasing from top to bottom. Given a partition $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_\ell)$, the corresponding YD has ℓ rows and λ_i boxes in the i^{th} row, where i increases from top to bottom. We think of partitions as YDs, and vice versa, freely throughout the sequel. If λ and μ are two YDs, we say $\mu \subset \lambda$ if every box of μ is a box of λ . We say $\mu \subset_k \lambda$ if $\mu \subset \lambda$ and μ and λ differ by k boxes.

A *skew YD* (SYD) is formally a pair of YDs μ and λ with $\mu \subset \lambda$ and is denoted by λ/μ . We also think of λ/μ as a diagram consisting of the boxes of λ that are not in μ . We can think of a YD λ also as a skew diagram $\lambda = \lambda/\emptyset$, where \emptyset is the empty diagram with no boxes. Therefore, statements that we make about SYDs apply in this way also to YDs.

The *size* $|\lambda/\mu|$ of an SYD λ/μ is the number of boxes that it contains, or $\sum \lambda_i - \sum \mu_i$. If \square is a particular box appearing in an SYD, we let $i(\square)$ be the row number (starting at 1, counting from top to bottom) of the box and $j(\square)$ the column number (starting at 1, counting from left to right) of the box. The *content* of a box \square in an SYD is

$$c(\square) \stackrel{\text{def}}{=} j(\square) - i(\square).$$

If \square_1 and \square_2 are two boxes in an SYD, we let

$$\text{ax}(\square_1, \square_2) \stackrel{\text{def}}{=} c(\square_1) - c(\square_2);$$

this is called the *axial distance* between \square_1 and \square_2 .

If λ is a YD, we write $\check{\lambda}$ for the YD obtained from λ by swapping rows and columns, namely, by transposing. This $\check{\lambda}$ is called the *conjugate* of λ .

3.2. Young tableaux

Let λ/μ be an SYD with $\lambda \vdash n$ and $\mu \vdash m$. A *standard Young tableau* of shape λ/μ is a filling of the boxes of λ/μ with the numbers $m + 1, \dots, n$ such that

- each number appears in exactly one box of λ/μ , and
- the numbers in the boxes are strictly increasing from left to right and from top to bottom.

In the sequel, we will refer to standard Young tableaux simply as *tableaux*. For λ/μ an SYD, we write $\text{Tab}(\lambda/\mu)$ for the collection of tableaux of shape λ/μ . If $\lambda \vdash n$, $T \in \text{Tab}(\lambda)$ and $m \in [n]$, we write $\mu_m(T)$ for the YD obtained by deleting the boxes containing $m + 1, \dots, n$ from T , so $\mu_m(T) \vdash m$. We also write $T|_{\leq m} \in \text{Tab}(\mu_m(T))$ for the tableau formed by the numbers-in-boxes of T that are $\leq m$, and $T|_{> m}$ for the tableau formed by the numbers-in-boxes of T that are $> m$. In general, the shape of $T|_{> m}$ will be an SYD. If T is a tableau of shape λ/μ , where $\lambda \vdash n$ and $\mu \vdash m$, and $m < i \leq n$, we write \boxed{i}_T for the box containing i in T .

If $\lambda \vdash n$ and $\mu \subset \lambda$, then we have a concatenation between $\text{Tab}(\mu)$ and $\text{Tab}(\lambda/\mu)$: If $T \in \text{Tab}(\mu)$ and $R \in \text{Tab}(\lambda/\mu)$, let $T \sqcup R$ be the tableau obtained by adjoining R to T .

3.3. Representations of symmetric groups

The irreducible unitary representations of S_n are parameterized, up to unitary equivalence, by YDs of size n . This correspondence between YDs and representations is denoted by

$$\lambda \mapsto V^\lambda.$$

Each V^λ is a finite-dimensional complex vector space with a unitary action of S_n and is also a module for the group algebra $\mathbf{C}[S_n]$. Let $d_\lambda \stackrel{\text{def}}{=} \dim V^\lambda$. It is known that $d_\lambda = |\text{Tab}(\lambda)|$.

We now follow Vershik–Okounkov [VO96]. The natural ordering of $[n]$ induces a filtration

$$S_1 \subset S_2 \subset \dots \subset S_{n-1} \subset S_n$$

of S_n , where S_m is the subgroup of S_n fixing each of the numbers in $[m + 1, n]$. If W is any unitary representation of S_n , for $m \in [n]$ and μ a YD of size m , we write W_μ for the linear span in W of all elements in the image of $\text{Hom}_{S_m}(V^\mu, W)$. In other words, W_μ is the span of copies of V^μ in the restriction of W to S_m . This W_μ is called the μ -isotypic subspace of W .

Vershik and Okounkov describe a specific orthonormal basis of V^λ , called a *Gelfand–Tsetlin basis*, that will be useful to us here. The basis is indexed by $T \in \text{Tab}(\lambda)$; each such T gives a basis vector v_T . The vectors v_T can be characterized up to multiplication by complex scalars of modulus 1 in the following way. The intersection of subspaces

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

is one-dimensional and contains the unit vector v_T [VO96, §1]. One important corollary of this is that if $\mu \vdash m \in [n]$, then $(V^\lambda)_\mu \neq \{0\}$ if and only if $\mu \subset \lambda$. Also, note that if $\mu_1, \mu_2 \subset \lambda$, $\mu_1, \mu_2 \vdash m \in [n]$, and $\mu_1 \neq \mu_2$, then $(V^\lambda)_{\mu_1}$ is orthogonal to $(V^\lambda)_{\mu_2}$.

More generally, if λ/μ is an SYD with $\lambda \vdash n$ and $\mu \vdash m$, then there is a *skew module* $V^{\lambda/\mu}$ that is a unitary representation of S'_{n-m} where we write S'_{n-m} for the copy of S_{n-m} in S_n that fixes the elements $[m]$. Formally,

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

where the action of S'_{n-m} is by left multiplication: for $\varphi \in \text{Hom}_{S_m}(V^\mu, V^\lambda)$, $\tau \in S'_{n-m}$ and $v \in V^\mu$, $(\tau \cdot \varphi)(v) \stackrel{\text{def}}{=} \tau \cdot (\varphi(v))$. This action preserves $V^{\lambda/\mu}$ as S'_{n-m} is in the centralizer of S_m in $\mathbf{C}[S_n]$. We write $d_{\lambda/\mu}$ for the dimension of $V^{\lambda/\mu}$. Since $d_{\lambda/\mu}$ is the multiplicity of V^μ in the restriction of V^λ to S_m , by Frobenius reciprocity, it is also the multiplicity of V^λ in the induced representation $\text{Ind}_{S_m}^{S_n} V^\mu$. By calculating the dimension of $\text{Ind}_{S_m}^{S_n} V^\mu$ in two ways, we obtain the following result that will be useful later.

Lemma 3.1. *Let $n \in \mathbf{N}$, $m \in [n]$ and $\mu \vdash m$. Then,*

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

The module $V^{\lambda/\mu}$ has an orthonormal basis w_T indexed by $T \in \text{Tab}(\lambda/\mu)$ [VO96, §7]. One also has the following property that we will use later [CSST10, Eq. (3.65)].

Lemma 3.2. *Let $n \in \mathbf{N}$, $m \in [n]$, $\lambda \vdash n$ and $\mu \vdash m$ and assume that $\mu \subset \lambda$. Then the map*

$$v_T \otimes w_R \mapsto v_{T \sqcup R}, \quad T \in \text{Tab}(\mu), R \in \text{Tab}(\lambda/\mu)$$

linearly extends to an isomorphism of unitary $(S_m \times S'_{n-m})$ -representations $V^\mu \otimes V^{\lambda/\mu} \cong (V^\lambda)_\mu$.

There is also an explicit formula for the action of S'_{n-m} on $V^{\lambda/\mu}$. A full exposition of this formula can be found in [VO96, §6]. Recall that S'_{n-m} is generated by the *Coxeter generators*

$$s_i \stackrel{\text{def}}{=} (i \ i + 1)$$

for $m < i < n$, where $(i \ i + 1)$ is our notation for a transposition switching i and $i + 1$. Therefore, it is sufficient to describe how the s_i act on $V^{\lambda/\mu}$. Say that T is admissible for s_i if the boxes containing i and $i + 1$ in T are neither in the same row nor the same column.

For $T \in \text{Tab}(\lambda/\mu)$, let

$$s_i T = \begin{cases} T & \text{if } T \text{ is not admissible for } s_i \\ T' & \text{if } T \text{ is admissible for } s_i, \end{cases}$$

where T' is the tableaux obtained from T by swapping i and $i + 1$. The admissibility condition ensures T' is a valid standard Young tableau. Then one has *Young's orthogonal form*

$$s_i w_T = \frac{1}{\text{ax}(\boxed{i+1}_T, \boxed{i}_T)} w_T + \sqrt{1 - \frac{1}{\text{ax}(\boxed{i+1}_T, \boxed{i}_T)^2}} w_{s_i T}. \tag{3.1}$$

Note that as a special case of this formula, if T is not admissible for s_i , then $\text{ax}(\boxed{i+1}_T, \boxed{i}_T) = \pm 1$ and

$$s_i w_T = \frac{1}{\text{ax}(\boxed{i+1}_T, \boxed{i}_T)} w_T = \begin{cases} w_T & \text{if } i \text{ and } i + 1 \text{ are in the same row,} \\ -w_T & \text{if } i \text{ and } i + 1 \text{ are in the same column.} \end{cases} \tag{3.2}$$

Remark 3.3. For completeness of some of our statements, we need to define the notions above also for S_0 , the symmetric group of the empty set. This is the trivial group. Whenever $\mu = \lambda$, we have $\text{Tab}(\lambda/\mu) = \{\emptyset\}$, and the representation $V^{\lambda/\mu}$ is one-dimensional with basis w_T , for T the empty tableau.

4. Preliminary representation theoretic results

In this section, we give some preliminary results on representation theory that will be used in the rest of the paper. Although some results here seem to be novel (in particular Proposition 4.4), this section plays only a supporting role in the paper.

4.1. Commutants

Recall that if V is a finite-dimensional vector space and \mathcal{A} is a subalgebra of $\text{End}(V)$, then the *commutant* of \mathcal{A} in $\text{End}(V)$ is the algebra of elements $b \in \text{End}(V)$ such that

$$ba = ab$$

for all $a \in \mathcal{A}$. For $m \in [n]$ and $\lambda \vdash n$, let $Z(\lambda, m, n)$ denote the commutant of the image of $\mathbf{C}[S_m]$ in $\text{End}(V^\lambda)$. We identify

$$\text{End}(V^\lambda) \cong V^\lambda \otimes \check{V}^\lambda \tag{4.1}$$

and give $\text{End}(V^\lambda)$ the Hermitian inner product induced from V^λ .

Lemma 4.1. *Let $m \in [n]$ and $\lambda \vdash n$. The algebra $Z(\lambda, m, n)$ has an orthonormal basis given by*

$$\left\{ \mathcal{E}_{\mu, R_1, R_2}^\lambda \stackrel{\text{def}}{=} \frac{1}{\sqrt{d_\mu}} \sum_{T \in \text{Tab}(\mu)} v_{T \sqcup R_1} \otimes \check{v}_{T \sqcup R_2} : \mu \vdash m, \mu \subset \lambda, R_1, R_2 \in \text{Tab}(\lambda/\mu) \right\}. \tag{4.2}$$

Proof. Let $\mathcal{A} \subseteq \text{End}(V^\lambda)$ be the algebra generated by the $\mathcal{E}_{\mu, R_1, R_2}^\lambda$ (over all μ) and $\mathcal{A}_\mu \subseteq \mathcal{A}$ be the algebra generated by the $\mathcal{E}_{\mu, R_1, R_2}^\lambda$ with a fixed value of μ . Suppose that $Q \in \text{Tab}(\lambda)$. The formula for the action of $\mathcal{E}_{\mu, R_1, R_2}^\lambda$ on v_Q is

$$\mathcal{E}_{\mu, R_1, R_2}^\lambda (v_Q) = \mathbf{1} \{ \mu_m(Q) = \mu, Q|_{>m} = R_2 \} \frac{1}{\sqrt{d_\mu}} v_{Q|_{\leq m} \sqcup R_1}. \tag{4.3}$$

It is clear that the $\mathcal{E}_{\mu, R_1, R_2}^\lambda$ are an orthonormal set of elements in $\text{End}(V^\lambda)$. It follows from equation (4.3) that

$$\mathcal{E}_{\mu_1, R_1, R_2}^\lambda \mathcal{E}_{\mu_2, R_3, R_4}^\lambda = \mathbf{1} \{ \mu_1 = \mu_2, R_3 = R_2 \} \frac{1}{\sqrt{d_{\mu_1}}} \mathcal{E}_{\mu_1, R_1, R_4}^\lambda,$$

so the $\mathcal{E}_{\mu, R_1, R_2}^\lambda$ are an orthonormal basis for the algebra \mathcal{A} , and those with a fixed μ are an orthonormal basis for \mathcal{A}_μ . Furthermore, we have

$$\mathcal{A} = \bigoplus_{\mu \subset \lambda, \mu \vdash m} \mathcal{A}_\mu.$$

For $\mu \vdash m$, let $p_\mu \in \mathbf{C}[S_m]$ be the central idempotent projection onto the μ -isotypic component of $\mathbf{C}[S_m]$ and let P_μ^λ be the image of p_μ in $\text{End}(V^\lambda)$. The element P_μ^λ is the orthogonal projection onto $(V^\lambda)_\mu$, and P_μ^λ is in the center of $Z(\lambda, m, n)$. Hence, for every $z \in Z(\lambda, m, n)$, we can write

$$z = \bigoplus_{\mu \vdash m, \mu \subset \lambda} z^{(\mu)},$$

where $z^{(\mu)} \stackrel{\text{def}}{=} P_\mu^\lambda z P_\mu^\lambda$. Moreover, if we let \mathcal{B}_μ be the algebra generated by the image of $\mathbf{C}[S_m]$ in $\text{End}(V_\mu^\lambda)$, each $z^{(\mu)}$ must be in the commutant \mathcal{B}'_μ of \mathcal{B}_μ in $\text{End}(V_\mu^\lambda)$. On the other hand, if $z = \bigoplus_{\mu \vdash m, \mu \subset \lambda} z^{(\mu)}$ and each $z^{(\mu)} \in \mathcal{B}'_\mu$, then $z \in Z(\lambda, m, n)$. This shows that

$$Z(\lambda, m, n) = \bigoplus_{\mu \vdash m, \mu \subset \lambda} \mathcal{B}'_\mu. \tag{4.4}$$

Since V^μ is an irreducible module for $\mathbf{C}[S_m]$, the algebra generated by $\mathbf{C}[S_m]$ in $\text{End}(V^\mu)$ is the whole of $\text{End}(V^\mu)$. Hence, under the isomorphism of Lemma 3.2, the algebra \mathcal{B}_μ is identified with $\text{End}(V^\mu) \otimes \text{CId}_{V_\lambda/\mu}$. By a classical theorem, due to Tomita [Tom67] in the generality of von Neumann algebras,¹² the commutant of a tensor product is the tensor product of the two commutants. Therefore, still using the isomorphism of Lemma 3.2, we have

$$\mathcal{B}'_\mu \cong \text{CId}_{V^\mu} \otimes \text{End}(V_\lambda/\mu).$$

This space is the algebra \mathcal{A}_μ , so $\mathcal{B}'_\mu = \mathcal{A}_\mu$ and

$$Z(\lambda, m, n) = \bigoplus_{\mu \vdash m, \mu \subset \lambda} \mathcal{B}'_\mu = \bigoplus_{\mu \vdash m, \mu \subset \lambda} \mathcal{A}_\mu = \mathcal{A}$$

as required. □

4.2. Bounds for the dimensions of irreducible representations

In this section, we give bounds related to the dimensions of irreducible representations that we use later. We first note a very simple bound for the dimensions of irreducible representations of S_n . For a YD, denote by $b_\lambda \stackrel{\text{def}}{=} |\lambda| - \lambda_1$ the number of boxes outside the first row.

Lemma 4.2. *Let $\lambda \vdash n$. Suppose that $\lambda_1 = n - b_\lambda \geq \frac{n}{2}$. Then*

$$\binom{\lambda_1}{b_\lambda} \leq d_\lambda \leq n^{b_\lambda}.$$

Proof. The first inequality is given by Liebeck and Shalev in [LS04, Lem. 2.1]. To bound d_λ from above, note that the standard tableaux of shape λ can be obtained by choosing which λ_1 elements of $[n]$ are in the first row (of which there are at most $\binom{n}{b_\lambda}$ choices) and choosing the remaining b_λ numbers' locations outside the first row, of which there are at most $b_\lambda!$ choices. Hence,

$$d_\lambda \leq b_\lambda! \binom{n}{b_\lambda} \leq n^{b_\lambda}. \tag{4.5}$$

Lemma 4.3. *Let $\lambda \vdash n, \nu \subset_k \lambda$. If $n \geq k + 2b_\lambda$, then*

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

Proof. By assumption, $n \geq k + 2b_\lambda \geq k + 2b_\nu$, so $n - b_\lambda \geq \frac{n}{2}$ and $n - k - b_\nu \geq \frac{n-k}{2}$ and Lemma 4.2 applies. The statement now follows from Lemma 4.2 together with the inequality $\binom{p}{q} \geq \left(\frac{p}{q}\right)^q$, holding for $p, q \in \mathbf{N}$ with $p \geq q \geq 1$. □

¹²This is, however, easy to prove in the special case here that we use it.

4.3. An estimate for matrix coefficients in skew modules

Recall that the Coxeter generators of S_n are $s_i = (i \ i + 1)$ for $i \in [n - 1]$. If $\tau \in S_n$, we write $\ell_{\text{cox}}(\tau)$ for the minimal length of a word of Coxeter generators that equals τ . Assume that $\lambda \vdash n$, $m \in [n]$ and $\nu \vdash m$. For $T \in \text{Tab}(\lambda/\nu)$, we write $\text{top}(T) \subseteq [m + 1, n]$ for the set of elements in the top row of T (which may be empty: It is of size $\lambda_1 - \nu_1$).

For any two subsets A, B of $[n]$, we define $d(A, B) = |A \setminus B|$. When restricted to subsets of $[n]$ with exactly p elements, for some $p \in [0, n]$, this function is a metric. Moreover, the function d is clearly invariant under S_n , that is, if $\sigma \in S_n$ and $A, B \subseteq [n]$, then $d(\sigma(A), \sigma(B)) = d(A, B)$.

Proposition 4.4. *Suppose $m \leq n$, $\lambda \vdash n$, $\nu \vdash m$ and $\nu \subset \lambda$, and write $k = n - m$. If $\lambda_1 + \nu_1 > n + k^2$, then for any $T, T' \in \text{Tab}(\lambda/\nu)$ and $\sigma \in S'_k$ we have*

$$|\langle \sigma w_T, w_{T'} \rangle| \leq \left(\frac{k^2}{\lambda_1 + \nu_1 - n} \right)^{d(\sigma \text{top}(T), \text{top}(T'))} \tag{4.6}$$

Note that if the top row of λ/ν is empty, namely, if $\nu_1 = \lambda_1$, then $\text{top}(T) = \emptyset$ for every $T \in \text{Tab}(\lambda/\nu)$ and the upper bound in equation (4.6) is trivial: $(k^2/(\lambda_1 + \nu_1 - n))^0 = 1$. In particular, this is the case if $m = n$, in which case $k = 0$, the bound is $0^0 = 1$ and we have an action of the trivial group on a one-dimensional space spanned by w_T for T the empty tableau (see Remark 3.3).

Proof. If $k = 0$ the statement is trivial, so we may assume $k \geq 1$. We prove equation (4.6) as a consequence of the following slightly stronger statement:

(S) If $\lambda_1 + \nu_1 \geq n + \ell_{\text{cox}}(\sigma)$, then for any $T \in \text{Tab}(\lambda/\nu)$, $A_0 \subseteq [m + 1, n]$ of size $\lambda_1 - \nu_1$ and any unit vector u in

$$W_{A_0} \stackrel{\text{def}}{=} \text{span}(\{w_{T'} \mid T' \in \text{Tab}(\lambda/\nu), \text{top}(T') = A_0\})$$

we have

$$|\langle \sigma w_T, u \rangle| \leq \left(\frac{\ell_{\text{cox}}(\sigma)}{\nu_1 + \lambda_1 - n} \right)^{d(\sigma \text{top}(T), A_0)} \tag{4.7}$$

The proposition follows from (S) by using the bound $\ell_{\text{cox}}(\sigma) \leq k^2$ and setting $A_0 = \text{top}(T')$, $u = w_{T'}$.

Let $D \stackrel{\text{def}}{=} d(\sigma \text{top}(T), A_0)$. We prove (S) by induction on $\ell \stackrel{\text{def}}{=} \ell_{\text{cox}}(\sigma)$. The base case of the induction is $\ell = 0$. Then $\sigma = \text{id}$ and

$$|\langle \sigma w_T, u \rangle| = |\langle w_T, u \rangle| = 0$$

unless $\text{top}(T) = A_0$, meaning $D = 0$. On the other hand, if $D = 0$, then

$$|\langle \sigma w_T, u \rangle| \leq 1 = 0^0 = \left(\frac{\ell}{\nu_1 + \lambda_1 - n} \right)^D$$

as required.

For the inductive step, for $\ell \geq 1$ we write $\sigma = s_j \sigma'$, where $\ell_{\text{cox}}(\sigma') = \ell - 1$ and $j \in [m + 1, n - 1]$. Two scenarios can occur.

(i) Suppose $s_j A_0 = A_0$. In this case, by the definition of the action of the Coxeter generators in equation (3.1), $s_j u$ is a unit vector in W_{A_0} . Also, by the invariance of the distance function under s_j ,

$$d(\sigma' \text{top}(T), A_0) = d(\sigma \text{top}(T), s_j A_0) = d(\sigma \text{top}(T), A_0) = D.$$

The inductive hypothesis then yields

$$|\langle \sigma w_T, u \rangle| = |\langle \sigma' w_T, s_j u \rangle| \leq \left(\frac{\ell - 1}{\nu_1 + \lambda_1 - n} \right)^D \leq \left(\frac{\ell}{\nu_1 + \lambda_1 - n} \right)^D,$$

as required.

(ii) Suppose otherwise that $s_j A_0 \neq A_0$. This means that exactly one of j and $j + 1$ are in A_0 . We write

$$u = \sum_{T' \in \text{Tab}(\lambda/\nu) : \text{top}(T')=A_0} \beta_{T'} w_{T'}. \tag{4.8}$$

For each T' with $\text{top}(T') = A_0$, we have

$$|\text{ax}(\boxed{j+1}_{T'}, \boxed{j}_{T'})| \geq \nu_1 + \lambda_1 - n.$$

From equation (4.8) and the formula for the action of Coxeter generators (3.1), we can therefore write $s_j u = w_1 + w_2$, where $w_1 \in W_{s_j A_0}$ and $w_2 \in W_{A_0}$ are orthogonal vectors with $\|w_1\| \leq 1$ and $\|w_2\| \leq (\nu_1 + \lambda_1 - n)^{-1}$. Hence,

$$|\langle \sigma w_T, u \rangle| = |\langle \sigma' w_T, s_j u \rangle| \leq |\langle \sigma' w_T, w_1 \rangle| + |\langle \sigma' w_T, w_2 \rangle|.$$

Note that $d(\sigma' \text{top}(T), s_j A_0) = d(\sigma \text{top}(T), A_0) = D$, so by the inductive hypothesis

$$|\langle \sigma w_T, u \rangle| \leq \left(\frac{\ell - 1}{\nu_1 + \lambda_1 - n} \right)^D + \frac{1}{\nu_1 + \lambda_1 - n} \left(\frac{\ell - 1}{\nu_1 + \lambda_1 - n} \right)^{d(\sigma' \text{top}(T), A_0)}. \tag{4.9}$$

By the triangle inequality,

$$D - 1 = d(\sigma' \text{top}(T), s_j A_0) - d(s_j A_0, A_0) \leq d(\sigma' \text{top}(T), A_0),$$

so using $\nu_1 + \lambda_1 \geq n + \ell$ we obtain from equation (4.9)

$$\begin{aligned} |\langle \sigma w_T, u \rangle| &\leq \left(\frac{\ell - 1}{\nu_1 + \lambda_1 - n} \right)^D + \frac{1}{\nu_1 + \lambda_1 - n} \left(\frac{\ell - 1}{\nu_1 + \lambda_1 - n} \right)^{D-1} \\ &= \frac{(\ell - 1)^{D-1} \ell}{(\nu_1 + \lambda_1 - n)^D} \leq \left(\frac{\ell}{\nu_1 + \lambda_1 - n} \right)^D, \end{aligned}$$

as required. □

4.4. Families of YDs and zeta functions

Recall from §1 that the zeta function of S_n is defined by

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

and that for $g \geq 2$

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

Let $\Lambda(n, b)$ denote the collection of $\lambda \vdash n$ such that $b_\lambda \geq b$ and $b_{\bar{\lambda}} \geq b$. In other words, $\Lambda(n, b)$ is the collection of YDs of size n with at least b boxes outside the first row and at least b boxes outside

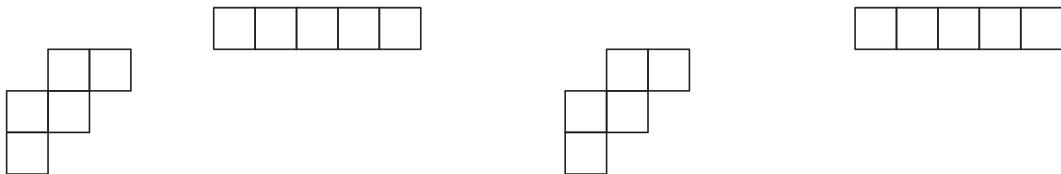


Figure 4.1. This figure depicts two elements of a family of SYDs $\lambda(n)/\nu(n - 10)$ for $n = 16$ and $n = 18$. Here, we can take $\lambda = (6, 3, 2, 1)$ and $\nu = (1, 1)$.

the first column. One has the following useful result of Liebeck and Shalev [LS04, Prop. 2.5] and, independently, Gamburd [Gam06, Prop. 4.2]:

Proposition 4.5. For fixed $b \geq 0$ and real $s > 0$, as $n \rightarrow \infty$

$$\sum_{\lambda \in \Lambda(n,b)} \frac{1}{d_\lambda^s} = O_b(n^{-sb}).$$

The proof of Theorem 1.1 will crucially depend on certain families of YDs that interact nicely with the skew modules $V^{\lambda/\nu}$. Given a YD λ , we will write $\lambda(n)$ for the unique YD $\lambda(n) \vdash n$ which is obtained from λ by either deleting boxes from or adding boxes to the first row of λ , if it exists. To be precise, $\lambda(n)$ exists if and only if $n \geq |\lambda| - (\lambda_1 - \lambda_2)$, interpreting $\lambda_2 = 0$ if λ only has one row.

Now, given $k \in \mathbb{N}$, and YDs $\nu \subset_k \lambda$, assume that n_1 and n_2 are large enough so that $\lambda(n_i)$ and $\nu(n_i - k)$ both exist and so that the first row (of length $\lambda_1 - \nu_1$, which could be zero) of the SYD $\lambda(n_i)/\nu(n_i - k)$ does not border the second row, namely, $\nu(n_i - k)_1 \geq \lambda_2$. Then there is a natural way to identify $\text{Tab}(\lambda(n_1)/\nu(n_1 - k))$ with $\text{Tab}(\lambda(n_2)/\nu(n_2 - k))$ by simply adding $n_2 - n_1$ to all numbers in boxes of a tableau in $\text{Tab}(\lambda(n_1)/\nu(n_1 - k))$ and shifting the first row right or left as needed. If $\nu_1 \geq \lambda_2$ and $T \in \text{Tab}(\lambda/\nu)$, we write $T(n)$ for the resulting tableau in $\text{Tab}(\lambda(n)/\nu(n - k))$.

Given $n \in \mathbb{N}$ and $k \in [n]$, recall that we write S'_k for the subgroup of S_n that acts as the identity on $[n - k]$. Throughout the paper, we fix isomorphisms

$$S_k \xrightarrow{\cong} S'_k, \quad \sigma \mapsto \rho_k^{-1} \circ \sigma \circ \rho_k, \tag{4.10}$$

where we view $S_k \leq S_n$ in the usual way and

$$\rho_k(i) = \begin{cases} i + k & \text{if } i \in [n - k] \\ i - n + k & \text{if } i \in [n - k + 1, n]. \end{cases}$$

Using these isomorphisms allows us to identify the different subgroups S'_k as n varies: This will recur at several points of the sequel. It also allows us to note in the following proposition that matrix coefficients of skew modules $V^{\lambda(n)/\nu(n-k)}$ are holomorphic functions of n^{-1} , for sufficiently large n . Recall that $w_{T_i(n)}$ are elements of the Gelfand–Tsetlin basis for $V^{\lambda(n)/\nu(n-k)}$.

Proposition 4.6. Let $k \in \mathbb{N}$, $\sigma \in S_k$, and $\nu \subset_k \lambda$ be two YDs that differ by k boxes. Suppose that $\nu_1 \geq \lambda_2$. Given $T_1, T_2 \in \text{Tab}(\lambda/\nu)$, there is a function $F = F_{\sigma, \lambda, \nu, T_1, T_2}$ that is holomorphic in the ball of radius $|\lambda|^{-1}$ around zero, has Taylor expansion around 0 with rational coefficients and such that for all $n \geq |\lambda|$, viewing σ as an element of $S'_k \leq S_n$ via the isomorphism (4.10),

$$\langle \sigma w_{T_1(n)}, w_{T_2(n)} \rangle = F(n^{-1}).$$

Proof. Since $V^{\lambda(n)/\nu(n-k)}$ is finite-dimensional with dimension independent of n (as long as $n \geq |\lambda|$), it suffices to prove the result in the case that $\sigma \in S_k$ is a Coxeter generator s_i with $i \in [k - 1]$.

Interpreted as an element of $S'_k \leq S_n$ via equation (4.10), σ corresponds to the Coxeter generator $s_{i+n-k} \in S_n$.

Let a be the axial distance between $j \stackrel{\text{def}}{=} i + |\lambda| - k$ and $j + 1 = i + 1 + |\lambda| - k$ in T_1 . Note that

$$\text{ax} \left(\boxed{i+n-k}_{T_1(n)}, \boxed{i+1+n-k}_{T_1(n)} \right) = \begin{cases} a+n-|\lambda| & \text{if } j \text{ in the first row of } T_1 \text{ and } j+1 \text{ not,} \\ a-(n-|\lambda|) & \text{if } j+1 \text{ in the first row of } T_1 \text{ and } j \text{ not,} \\ a & \text{otherwise.} \end{cases}$$

In the first, case $a > 0$, and in the second case, $a < 0$. By the description in equation (3.1) of how the Coxeter generators act, $\langle s_{i+n-k} w_{T_1(n)}, w_{T_2(n)} \rangle$ is therefore one of the following functions of n :

$$0, \frac{1}{a-n+|\lambda|}, \frac{1}{a+n-|\lambda|}, \sqrt{1 - \frac{1}{(a-n+|\lambda|)^2}}, \sqrt{1 - \frac{1}{(a+n-|\lambda|)^2}}.$$

If one replaces n by z^{-1} , each of these yields a holomorphic function of z when $|z|$ is sufficiently small. □

The dimensions of representations in a family $\lambda(n)$ are polynomials in n .

Lemma 4.7. *Given a YD λ , consider the family of YDs $\lambda(n)$. There is a polynomial $G = G_\lambda \in \mathbb{Q}[t]$ of degree b_λ with rational coefficients such that for every n such that $\lambda(n)$ exists,*

$$d_{\lambda(n)} = G(n).$$

Furthermore, the complex zeros of G are integers n with $n \in [0, |\lambda|]$, and the leading coefficient is $\frac{1}{m}$ for some integer m .

For example, if $\lambda(n) = (n - 4, 3, 1)$, then $d_{\lambda(n)} = \frac{n(n-1)(n-3)(n-6)}{8}$ for every $n \geq 7$.

Proof. This easily follows from the hook-length formula for the dimension d_λ [FRT54]. □

Lemma 4.7 together with Proposition 4.5 have the following nice consequence for the zeta function ζ^{S_n} that will be crucial in proving Theorem 1.1.

Proposition 4.8. *For any $s \in \mathbb{N}$ and $M \in \mathbb{N}$, there is a polynomial $P_{s,M} \in \mathbb{Z}[t]$ with integer coefficients of degree $< M$ such that*

$$\zeta^{S_n}(s) = 2 \cdot P_{s,M}(n^{-1}) + O(n^{-M})$$

as $n \rightarrow \infty$. The constant coefficient of $P_{s,M}$ is equal to 1.

For example, for $s = 2$ and $M = 5$ we have

$$\zeta^{S_n}(2) = 2 \left(1 + \frac{1}{n^2} + \frac{2}{n^3} + \frac{11}{n^4} \right) + O\left(\frac{1}{n^5}\right).$$

Proof. Fix $s \in \mathbb{N}$ and $M \in \mathbb{N}$ as in the statement of the proposition. Let $b = \lceil \frac{M}{s} \rceil$. Proposition 4.5 implies that

$$\zeta^{S_n}(s) = \sum_{\substack{\lambda \vdash n \\ \lambda \notin \Lambda(n,b)}} \frac{1}{d_\lambda^s} + O(n^{-M})$$

as $n \rightarrow \infty$. The λ in the sum above have either $< b$ boxes outside their first row or $< b$ boxes outside their first column. For $n > 2b$, these options are mutually exclusive. Moreover, the map $\lambda \mapsto \check{\lambda}$ maps YDs of the first kind to YDs of the second kind bijectively and vice versa. Hence, if we let $\Lambda_{b,\lambda < b}(n)$ be the collection of $\lambda \vdash n$ with $< b$ boxes outside their first row, then as $n \rightarrow \infty$,

$$\begin{aligned} \zeta^{S_n}(s) &= \sum_{\lambda \in \Lambda_{b,\lambda < b}(n)} \left(\frac{1}{d_\lambda^s} + \frac{1}{d_{\check{\lambda}}^s} \right) + O(n^{-M}) \\ &= 2 \sum_{\lambda \in \Lambda_{b,\lambda < b}(n)} \frac{1}{d_\lambda^s} + O(n^{-M}). \end{aligned}$$

Now, if $n > 2b$, there is a finite collection of YDs $\{\mu_1, \dots, \mu_\ell\}$, depending on b , with $|\mu_i| < 2b$ for all i , such that for each $n > 2b$

$$\Lambda_{b,\lambda < b}(n) = \{\mu_1(n), \mu_2(n), \dots, \mu_\ell(n)\}.$$

For each of these μ_i , let G_{μ_i} be the polynomial provided by Lemma 4.7. No G_{μ_i} has any zero z with $|z| > 2b$. Hence,

$$\zeta^{S_n}(s) = 2 \sum_{i=1}^{\ell} \frac{1}{(G_{\mu_i}(n))^s} + O(n^{-M})$$

as $n \rightarrow \infty$. Because of the special structure of G_{μ_i} , as elaborated in Lemma 4.7, $(G_{\mu_i}(n))^{-1}$ is equal to a power series in n^{-1} with integer coefficients. Since $s \in \mathbf{N}$, $(G_{\mu_i}(n))^{-s}$ is too equal to a power series in n^{-1} with integer coefficients. This proves the first statement. Because the degree of $G_{\mu_i}(n)$ is positive unless $\mu_i(n) = (n)$ in which case $G_{\mu_i}(n) = 1$, the constant coefficient of $P_{s,M}$ must be 1. \square

In fact, it is the following direct corollary of Proposition 4.8 that we will need.

Corollary 4.9. *For any $s \in \mathbf{N}$ and $M \in \mathbf{N}$, there is a polynomial $Q_{s,M} \in \mathbf{Z}[t]$ of degree $< M$ and constant coefficient 1 such that as $n \rightarrow \infty$,*

$$\frac{1}{\zeta^{S_n}(s)} = \frac{1}{2} Q_{s,M}(n^{-1}) + O(n^{-M}).$$

5. The probability of an embedded tiled surface

5.1. Overview of this section

This short overview is meant to make the results of this section more transparent and to stress an analogy with known results about the zeta function of S_n . For simplicity, we assume $g = 2$ throughout this Section §5 and denote $\Gamma = \Gamma_2 = \langle a, b, c, d \mid [a, b] [c, d] \rangle$.

As explained in Section §1,

$$|\mathbb{X}_n| = |\mathbb{X}_{2,n}| = (n!)^3 \cdot \sum_{\lambda \vdash n} \frac{1}{d_\lambda^2}. \tag{5.1}$$

If $\{\lambda(n)\}_{n \geq n_0}$ is a family of YDs obtained by extending the first row, as in Section 4.4, then $d_{\lambda(n)}$ is a polynomial in n of degree b_λ (Lemma 4.7), and so the contribution of $\lambda(n)$ and of $\check{\lambda}(n)$ to equation (5.1) is a rational function in n for every $n \geq n_0$. Proposition 4.5 (due to [LS04, Gam06]) states that up to order $O(n^{-2b})$, the zeta function in equation (5.1) is determined by those families of YDs with $b_\lambda < b$ and their transpose.

In this Section 5, we prove something analogous for $\mathbb{E}_n^{\text{emb}}(Y)$, where Y is a compact tiled surface. We write $\mathfrak{v} = \mathfrak{v}(Y)$, $\mathfrak{e} = \mathfrak{e}(Y)$, $\mathfrak{f} = \mathfrak{f}(Y)$ for the number of vertices, edges and octagons of Y , respectively. We will use the letter f for an element of $\{a, b, c, d\}$ and write $e_f = e_f(Y)$ for the number of f -labeled edges of Y . The first major result is that, depending on a choice of four constant permutations $\sigma_f^\pm, \tau_f^\pm \in S'_v$ per letter (defined in Section 5.3), we have

$$\mathbb{E}_n^{\text{emb}}(Y) = \frac{(n!)^3}{|\mathbb{X}_n|} \cdot \frac{(n)_v (n)_f}{\prod_f (n)_{e_f}} \cdot \sum_{\nu \vdash n-v} H_Y(\nu), \tag{5.2}$$

where $H_Y(\nu)$ is some explicit function. This follows from Theorem 5.10. Notice that as $n \rightarrow \infty$, the first fraction in equation (5.2) is $\frac{(n!)^3}{|\mathbb{X}_n|} = \frac{1}{2} + O(n^{-2})$ by Proposition 4.5, and the second one is $\frac{(n)_v (n)_f}{\prod_f (n)_{e_f}} = n^{\chi(Y)} (1 + O(n^{-1}))$. So equation (5.2) gives that

$$\mathbb{E}_n^{\text{emb}}(Y) = \left(\frac{1}{2} + O(n^{-1})\right) n^{\chi(Y)} \cdot \sum_{\nu \vdash n-v} H_Y(\nu). \tag{5.3}$$

Next, our analysis shows that by considering, as above, families of YDs $\{\nu(n)\}_{n \geq n_0}$, then for large enough n , $H_Y(\nu(n))$ is equal to a converging series $\sum_{j=-\infty}^K \beta_j n^j$, with $K = K(Y, \nu)$ some integer. Section 5.8 then shows that, for any given M , there is finite set of families $\nu(n)$, with b_ν and $b_{\bar{\nu}}$ bounded, such that all remaining summands in $\sum_{\nu \vdash n-v} H_Y(\nu)$ outside these families contribute jointly $O(n^{-M})$ – this is analogous to Proposition 4.5. Because every tiled surface admits finite resolutions as in Section 2.3, this quickly leads to the proof of Theorem 1.1 in Section 6.

In fact, the analysis so far could have been carried out with graphs (core graphs à la Stallings) rather than with tiled surfaces. The importance of tiled surfaces and, moreover, of (strongly) boundary reduced tiled surfaces, is in our ability to determine the order of magnitude of $H_Y(\nu)$. Our analysis here culminates in Proposition 5.21 and Section 5.9, from which it follows that when Y is boundary reduced,

$$H_Y(\nu(n)) = \frac{1}{d_\nu^2} \cdot O(1)$$

as $n \rightarrow \infty$, and when Y is strongly boundary reduced,

$$H_Y(\nu(n)) = \frac{1}{d_\nu^2} \left(1 + O\left(\frac{1}{n}\right)\right). \tag{5.4}$$

This shows that the analysis of the zeta function in equation (5.1) can be viewed as a special case of our results. Indeed, when $Y = Y_\emptyset$ is the empty tiled surface (which is, in particular, strongly boundary reduced), equation (5.2) together with equation (5.4) yields that

$$|\mathbb{X}_n| = |\mathbb{X}_n| \cdot \mathbb{E}_n^{\text{emb}}(Y_\emptyset) = (n!)^3 \cdot \sum_{\nu \vdash n} H_{Y_\emptyset}(\nu) = (n!)^3 \cdot \sum_{\nu \vdash n} \frac{1}{d_\nu^2} \left(1 + O\left(\frac{1}{n}\right)\right).$$

What we achieve here is the extension of this result to general strongly boundary reduced tiled surfaces, with an extra factor of $n^{\chi(Y)}$ appearing. If Y is merely boundary reduced, we obtain the same result up to multiplicative constants.

A remark about composing permutations

A technical but important remark is due. The bijection

$$\begin{aligned} \phi &\mapsto X_\phi \\ \text{Hom}(\Gamma, S_n) &\rightarrow \{\text{degree-}n \text{ covers of } \Sigma_2\} \end{aligned}$$

described previously involves the version of S_n where permutations are composed with the left-most permutation acting first. On the other hand, since in the rest of the paper, we work with permutations in detail and specifically with the representation theory of S_n , it is more standard to view permutations as functions from $[n]$ to $[n]$ and hence to multiply according to composition of functions (functions acting from the left). *So in the rest of the paper, permutations will be composed with the right-most permutation acting first.* These two versions of S_n are isomorphic, of course, and one isomorphism is given by $\text{inv}: S_n \rightarrow S_n$ defined by $\sigma \mapsto \sigma^{-1}$.

Accordingly, by postmultiplication with inv , we turn a homomorphism

$$\phi \in \mathbb{X}_n = \text{Hom}(\Gamma_2, S_n \text{ (left-to-right version)})$$

into a homomorphism

$$\bar{\phi} \stackrel{\text{def}}{=} \text{inv} \circ \phi \in \text{Hom}(\Gamma_2, S_n \text{ (right-to-left version)}).$$

The homomorphism $\bar{\phi}$ satisfies $\bar{\phi}(\gamma) = \phi(\gamma)^{-1}$ for every $\gamma \in \Gamma_2$. In particular, with composition of permutations from right to left, the four permutations $\phi(a), \phi(b), \phi(c), \phi(d) \in S_n$ satisfy

$$[\phi(a)^{-1}, \phi(b)^{-1}] [\phi(c)^{-1}, \phi(d)^{-1}] = \left[\overline{\phi(a)}, \overline{\phi(b)} \right] \left[\overline{\phi(c)}, \overline{\phi(d)} \right] = 1,$$

or, equivalently (taking the inverse of the resulting permutation),

$$[\phi(d)^{-1}, \phi(c)^{-1}] [\phi(b)^{-1}, \phi(a)^{-1}] = 1.$$

This means that the word $[d^{-1}, c^{-1}] [b^{-1}, a^{-1}]$ will appear below at several points. Note that the image of $\gamma \in \Gamma$ under $\bar{\phi}$ is the inverse of $\phi(\gamma)$. But since a permutation and its inverse have the same cycle-structure in S_n , this does not affect the statistics we study in this paper.

5.2. Tiled surfaces and cosets

We assume that Y is a compact tiled surface. In this section, we assume $n \in \mathbb{N}$ with $n \geq \mathbf{v}$. We fix an arbitrary bijection $\mathcal{J} : Y^{(0)} \rightarrow [\mathbf{v}]$ and view (Y, \mathcal{J}) as fixed in this §5. We modify \mathcal{J} slightly for technical reasons¹³ by letting

$$\mathcal{J}_n : Y^{(0)} \rightarrow [n - \mathbf{v} + 1, n], \quad \mathcal{J}_n(\mathbf{v}) \stackrel{\text{def}}{=} \mathcal{J}(\mathbf{v}) + n - \mathbf{v}. \tag{5.5}$$

Then (Y, \mathcal{J}_n) is a vertex-labeled tiled surface for each n . We are interested in the quantity $\mathbb{E}_n^{\text{emb}}(Y)$, but because the uniform measure on \mathbb{X}_n is invariant under conjugation by S_n and S_n acts transitively on ordered tuples of size \mathbf{v} in $[n]$, we have

$$\mathbb{E}_n^{\text{emb}}(Y) = \frac{n!}{(n - \mathbf{v})!} \frac{|\mathbb{X}_n(Y, \mathcal{J}_n)|}{|\mathbb{X}_n|}, \tag{5.6}$$

where

$$\mathbb{X}_n(Y, \mathcal{J}_n) \stackrel{\text{def}}{=} \{ \phi \in \mathbb{X}_n : \text{there is an embedding } Y \hookrightarrow X_\phi \text{ inducing } \mathcal{J}_n \}.$$

(Recall from §§2.2 that the vertices of X_ϕ are labeled by $[n]$. Also, recall that an embedding $Y \hookrightarrow X_\phi$ inducing \mathcal{J}_n , if it exists, is unique.) Hence, we are interested in the size of the set $\mathbb{X}_n(Y, \mathcal{J}_n)$. *Henceforth, we use the map \mathcal{J}_n and the previous labelings of the vertices of X_ϕ to identify the vertex sets of Y and X_ϕ with subsets of \mathbb{N} .*

¹³The reason for using this modification comes from a convention in the representation theoretic methods we use below.

For each letter $f \in \{a, b, c, d\}$, let $\mathcal{V}_f^- = \mathcal{V}_f^-(Y) \subset [n - \mathbf{v} + 1, n]$ be the subset of vertices of Y with outgoing f -labeled edges, and $\mathcal{V}_f^+ \subset [n - \mathbf{v} + 1, n]$ those vertices of Y with incoming f -labeled edges. Note that $e_f = |\mathcal{V}_f^\pm|$. We let G_f denote the subgroup of S_n that fixes \mathcal{V}_f^- and write

$$G \stackrel{\text{def}}{=} G_a \times G_b \times G_c \times G_d \leq S_n^4.$$

For each $f \in \{a, b, c, d\}$, we let $g_f^0 \in S_n$ be a fixed element with the property that for every pair of vertices i, j of Y (so $i, j \in [n - \mathbf{v} + 1, n]$) with a directed f -labeled edge from i to j , we have $g_f^0(i) = j$. Recall the notation $S'_\mathbf{v}$ for the subgroup of S_n fixing $[n - \mathbf{v}]$ pointwise. We choose the g_f^0 consistently for each n in the sense that g_f^0 is chosen when $n = \mathbf{v}$ and then defined for arbitrary n by the isomorphisms $S_\mathbf{v} \cong S'_\mathbf{v}$ given in equation (4.10). We write $g^0 \stackrel{\text{def}}{=} (g_a^0, g_b^0, g_c^0, g_d^0)$. Notice that $g_f^0(\mathcal{V}_f^-) = \mathcal{V}_f^+$.

In the rest of the paper, whenever we write an integral over a group, it is performed with respect to the uniform probability measure. Let

$$\Theta_\lambda(Y, \mathcal{J}_n) \stackrel{\text{def}}{=} \int_{h_f \in G_f} \chi_\lambda \left(\left[\left(g_d^0 h_d \right)^{-1}, \left(g_c^0 h_c \right)^{-1} \right] \left[\left(g_b^0 h_b \right)^{-1}, \left(g_a^0 h_a \right)^{-1} \right] \right), \tag{5.7}$$

where χ_λ is the character of S_n corresponding to the irreducible representation V^λ . We will calculate the size of $\mathbb{X}_n(Y, \mathcal{J}_n)$ using the following result.

Proposition 5.1. *We have*

$$|\mathbb{X}_n(Y, \mathcal{J}_n)| = \frac{\prod_{f \in a,b,c,d} (n - e_f)!}{n!} \sum_{\lambda \vdash n} d_\lambda \Theta_\lambda(Y, \mathcal{J}_n).$$

Proof. We begin by observing that with $g_a^0, g_b^0, g_c^0, g_d^0$ as above, the map

$$\mathbb{X}_n \rightarrow S_n^4, \quad \phi \mapsto (\phi(a), \phi(b), \phi(c), \phi(d))$$

restricts to a bijection between $\mathbb{X}_n(Y, \mathcal{J}_n)$ and the tuples $(g_a, g_b, g_c, g_d) \in S_n^4$ such that both (g_a, g_b, g_c, g_d) is in the coset $(g_a^0, g_b^0, g_c^0, g_d^0)G$ and $[g_d^{-1}, g_c^{-1}] [g_b^{-1}, g_a^{-1}] = 1$.

Now, let

$$I = \int_{h_f \in G_f} \mathbf{1} \left\{ \left[\left(g_d^0 h_d \right)^{-1}, \left(g_c^0 h_c \right)^{-1} \right] \left[\left(g_b^0 h_b \right)^{-1}, \left(g_a^0 h_a \right)^{-1} \right] = 1 \right\}.$$

Then it is immediate that

$$|\mathbb{X}_n(Y, \mathcal{J}_n)| = |G| \cdot I = \prod_{f \in a,b,c,d} (n - e_f)! \cdot I.$$

Finally, use Schur orthogonality to write, as functions on S_n ,

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

insert this into the definition of I and interchange summation and integration to complete the proof. \square

In the next sections, we will focus our attention on the quantities $\Theta_\lambda(Y, \mathcal{J}_n)$.

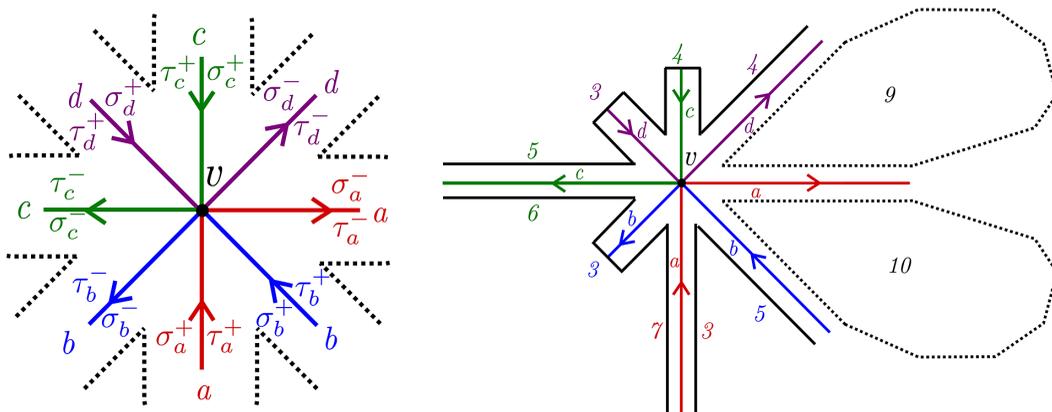


Figure 5.1. The figure on the left shows a local picture of a vertex v in the thick version of some tiled surface with hanging half-edges Y_+ , and the correspondence between the 16 maps σ_f^\pm, τ_f^\pm and the 16 sides of half-edges incident to v . The figure on the right illustrates how numbering of octagons, of exposed sides of full-edges and of hanging half-edges determines the values of σ_f^\pm, τ_f^\pm at v . In this figure, continuous black lines mark pieces of the boundary of the thick version of Y_+ , whereas dotted black lines mark boundary pieces of $Y_+^{(1)}$ to which octagons are glued in Y_+ . The vertex v in the center of the figure is incident with two octagons, numbered 9 and 10; with three hanging half-edges numbered 3 (outgoing b and incoming d) and 4 (incoming c) and with five half-edges belonging to full-edges, with a total of six exposed sides, the numbering of which is described in the figure. The images of this vertex under σ_f^\pm and τ_f^\pm are listed in page 53.

5.3. Construction of auxiliary permutations

In order to obtain an expression for $\Theta_\lambda(Y, \mathcal{J}_n)$ that leads to good analytic estimates, we introduce further maps

$$\sigma_f^+, \sigma_f^-, \tau_f^+, \tau_f^- \in S'_v \subset S_n$$

for each $f \in \{a, b, c, d\}$. These should be thought of as orderings of the vertices with indices from $[n - \mathbf{v} + 1, n]$, other than the one fixed by \mathcal{J}_n . We will first describe the construction of these maps and then note their properties.

Recall from §2.1 that $Y^{(1)}$ carries the structure of a ribbon graph and this gives us a way to thicken it up to an oriented surface with boundary with an embedded copy of the graph $Y^{(1)}$. Also recall, from §2.1, that we constructed a larger object Y_+ by adding extra hanging half-edges to the vertices. The one-skeleton $Y_+^{(1)}$ also has a cyclic ordering of the half-edges (hanging or otherwise) at each vertex and so $Y_+^{(1)}$ can be thickened up to a ‘cut’ ribbon graph with some half-ribbon edges. In this picture, every edge is thickened to a thin rectangle, and every hanging half-edge is thickened up to a thin half-rectangle.

So every vertex of Y has eight incident half-edges (hanging or otherwise), and each of these half-edges has two sides. The 16 maps $\{\sigma_f^\pm, \tau_f^\pm\}_{f \in \{a, b, c, d\}}$ correspond to these 16 sides of half-edges at each vertex: σ_f^- and τ_f^- correspond to the sides of the outgoing f -half-edge, while σ_f^+ and τ_f^+ correspond to the sides of the incoming f -half-edge. Finally, σ_f^\pm correspond to the left side of the outgoing and incoming f -half-edges, while τ_f^\pm correspond to the right side of these f -half-edges, where ‘left’ and ‘right’ here are with respect to the direction of the half-edge. (We keep our convention from §2.1 that boundary cycles are oriented so that the object lies to the right. In particular, the boundary of an octagon is $[a, b] [c, d]$ when followed in counterclockwise direction.) See the left-hand side of Figure 5.1.

The definition of σ_f^\pm, τ_f^\pm is based on the following choices:

- **Numbering octagons:** Number the \mathfrak{f} octagons of Y by distinct elements in $[\mathfrak{v} - \mathfrak{f} + 1, \mathfrak{v}]$.
- **Numbering full edges at ∂Y :** For every $f \in \{a, b, c, d\}$, there are $e_f - \mathfrak{f}$ left-sides of full f -edges that belong to the boundary ∂Y (as compared with \mathfrak{f} left sides of full f -edges that meet octagons of Y). Number them by distinct values in $[\mathfrak{v} - e_f + 1, \mathfrak{v} - \mathfrak{f}]$. Similarly, number the $e_f - \mathfrak{f}$ right sides of full f -edges belonging to ∂Y by distinct values in the same range $[\mathfrak{v} - e_f + 1, \mathfrak{v} - \mathfrak{f}]$.
- **Numbering hanging half-edges:** For each $f \in \{a, b, c, d\}$, there are precisely $\mathfrak{v} - e_f$ outgoing f -labeled hanging half-edges, and we number them by distinct numbers in $[\mathfrak{v} - e_f]$. Using the matching determined by g_f^0 between outgoing and incoming f -labeled hanging half-edges, the numbering we have just chosen induces a numbering also of the incoming f -labeled hanging half-edges by numbers in $[\mathfrak{v} - e_f]$.

We now define σ_f^\pm and τ_f^\pm as follows. For every vertex v of Y and a side of an incident half-edge (hanging or otherwise), we need to determine the image of v under the permutation among σ_f^\pm, τ_f^\pm corresponding to this side-of-half-edge.

- If the half-edge is part of a full-edge of Y , then
 - if the side in question meets an octagon numbered i , we map $v \mapsto n - \mathfrak{v} + i$, and
 - if the side in question belongs to ∂Y and the full-edge is numbered j , map $v \mapsto n - \mathfrak{v} + j$.
- If this is a hanging half-edge numbered k , we map $v \mapsto n - \mathfrak{v} + k$.

This is illustrated in the right-hand side of Figure 5.1, which shows some vertex v of some tiled surface Y , and the numbering of incident octagons, of exposed sides of full-edges and of hanging half-edges. In that case, the images of v under σ_f^\pm and τ_f^\pm are the following:

$$\begin{aligned} \sigma_a^-(v) &= n - \mathfrak{v} + 9 & \sigma_b^-(v) &= n - \mathfrak{v} + 3 & \sigma_c^-(v) &= n - \mathfrak{v} + 6 & \sigma_d^-(v) &= n - \mathfrak{v} + 4 \\ \sigma_a^+(v) &= n - \mathfrak{v} + 7 & \sigma_b^+(v) &= n - \mathfrak{v} + 5 & \sigma_c^+(v) &= n - \mathfrak{v} + 4 & \sigma_d^+(v) &= n - \mathfrak{v} + 3 \\ \tau_a^-(v) &= n - \mathfrak{v} + 10 & \tau_b^-(v) &= n - \mathfrak{v} + 3 & \tau_c^-(v) &= n - \mathfrak{v} + 5 & \tau_d^-(v) &= n - \mathfrak{v} + 9 \\ \tau_a^+(v) &= n - \mathfrak{v} + 3 & \tau_b^+(v) &= n - \mathfrak{v} + 10 & \tau_c^+(v) &= n - \mathfrak{v} + 4 & \tau_d^+(v) &= n - \mathfrak{v} + 3. \end{aligned}$$

The following properties of the maps we defined are all evident from the construction.

Lemma 5.2. *When the vertices of Y are identified with $[n - \mathfrak{v} + 1, n]$ according to \mathcal{J}_n , the 16 maps $\sigma_f^\pm, \sigma_f^\mp, \tau_f^\pm, \tau_f^\mp$ we defined indeed belong to $S'_\mathfrak{v} \subset S_n$. Moreover, they satisfy the following properties:*

P1 For all $f \in \{a, b, c, d\}$, $\sigma_f^\pm(\mathcal{V}_f^\pm) = \tau_f^\pm(\mathcal{V}_f^\pm) = [n - e_f + 1, n]$.

P2 For all $f \in \{a, b, c, d\}$, $(\sigma_f^+)^{-1}\sigma_f^- = (\tau_f^+)^{-1}\tau_f^- = g_f^0$.

P3 For all $f \in \{a, b, c, d\}$, $\sigma_f^\pm|_{[n] \setminus \mathcal{V}_f^\pm} = \tau_f^\pm|_{[n] \setminus \mathcal{V}_f^\pm}$.

P4 Each of the following permutations fixes every element of $[n - \mathfrak{f} + 1, n]$:

$$\sigma_b^-(\sigma_a^+)^{-1}, \tau_a^+(\sigma_b^+)^{-1}, \tau_b^+(\tau_a^-)^{-1}, \sigma_c^-(\tau_b^-)^{-1}, \sigma_d^-(\sigma_c^+)^{-1}, \tau_c^+(\sigma_d^+)^{-1}, \tau_d^+(\tau_c^-)^{-1}, \sigma_a^-(\tau_d^-)^{-1}.$$

P5 The permutations σ_f^\pm, τ_f^\pm are the same for each n in the sense that they change with n via the fixed isomorphisms between $S_\mathfrak{v}$ and $S'_\mathfrak{v} \leq S_n$ in equation (4.10).

From now on, assume that we have fixed σ_f^\pm, τ_f^\pm with properties **P1–P5**. We do this once and for all for every tiled surface Y (including the choice of \mathcal{J}).

5.4. Integrating over cosets

We briefly review some linear algebra. Recall that \check{V}^λ is the vector space of complex linear functionals on V^λ . If V^λ has orthonormal basis $\{v_i\}$, then \check{V}^λ has a dual basis $\{\check{v}_i\}$ defined by $\check{v}_i(v) \stackrel{\text{def}}{=} \langle v, v_i \rangle$. Requiring the \check{v}_i to be orthonormal defines a Hermitian inner product on \check{V}^λ . The action of S_n on \check{V}^λ is by $g[\phi](v) \stackrel{\text{def}}{=} \phi(g^{-1}v)$. If $A_{ji} \stackrel{\text{def}}{=} \langle gv_i, v_j \rangle$ so that g acts by the matrix $A = (A_{ij})$ on V^λ in this basis, then

$$g[\check{v}_i](v_k) = \check{v}_i \left[\sum_{\ell} (A^{-1})_{\ell k} v_{\ell} \right] = (A^{-1})_{ik},$$

so

$$g[\check{v}_i] = \sum_k (A^{-1})_{ik} \check{v}_k.$$

We now give some motivation for what follows. We wish to integrate the function

$$S_n^4 \rightarrow \mathbf{R}, \quad (g_a, g_b, g_c, g_d) \mapsto \chi_{\lambda} \left([g_d^{-1}, g_c^{-1}] [g_b^{-1}, g_a^{-1}] \right),$$

over a coset in S_n^4 . This function can clearly be written as a finite sum of finite products of matrix coefficients of the g_f and g_f^{-1} in V^{λ} . However, this is not the route we wish to take. Instead, following a philosophy similar to that used in the development of the Weingarten calculus (see, for example, [CS06]), we aim to write this function more holistically as (what is essentially) one single matrix coefficient in one single representation. To this end, consider the vector space

$$W^{\lambda} \stackrel{\text{def}}{=} V_a^{\lambda} \otimes \check{V}_a^{\lambda} \otimes V_b^{\lambda} \otimes \check{V}_b^{\lambda} \otimes V_c^{\lambda} \otimes \check{V}_c^{\lambda} \otimes V_d^{\lambda} \otimes \check{V}_d^{\lambda} \tag{5.8}$$

as a unitary representation of S_n^4 . We write an element of S_n^4 as (g_a, g_b, g_c, g_d) and the subscripts above indicate which coordinate acts on which factor.

Let $B_{\lambda} \in \text{End}(W^{\lambda})$ be defined via matrix coefficients by the formula

$$\langle B_{\lambda} (v_1 \otimes \check{v}_2 \otimes v_3 \otimes \check{v}_4 \otimes v_5 \otimes \check{v}_6 \otimes v_7 \otimes \check{v}_8), w_1 \otimes \check{w}_2 \otimes w_3 \otimes \check{w}_4 \otimes w_5 \otimes \check{w}_6 \otimes w_7 \otimes \check{w}_8 \rangle \stackrel{\text{def}}{=} \langle v_1, w_3 \rangle \langle v_3, v_2 \rangle \langle w_2, v_4 \rangle \langle w_4, w_5 \rangle \langle v_5, w_7 \rangle \langle v_7, v_6 \rangle \langle w_6, v_8 \rangle \langle w_8, w_1 \rangle. \tag{5.9}$$

Remark 5.3. One could also order the tensor factors in equation (5.8) according to the order specified by the word $[a, b] [c, d]$, namely $V_a^{\lambda} \otimes V_b^{\lambda} \otimes \check{V}_a^{\lambda} \otimes \check{V}_b^{\lambda} \otimes V_c^{\lambda} \otimes V_d^{\lambda} \otimes \check{V}_c^{\lambda} \otimes \check{V}_d^{\lambda}$. In this case, the definition of B_{λ} would be more natural: $\langle v_1, w_2 \rangle \langle v_2, v_3 \rangle \langle w_3, v_4 \rangle \langle w_4, w_5 \rangle \langle v_5, w_6 \rangle \langle v_6, v_7 \rangle \langle w_7, v_8 \rangle \langle w_8, w_1 \rangle$ and easily generalizable to arbitrary words. We chose to stick with the order in equation (5.8) for ease of notation in the sequel, for example, in Lemma 5.6.

Lemma 5.4. For any $(g_a, g_b, g_c, g_d) \in S_n^4$, we have

$$\text{tr}_{W^{\lambda}} (B_{\lambda} \circ (g_a, g_b, g_c, g_d)) = \chi_{\lambda} \left([g_d^{-1}, g_c^{-1}] [g_b^{-1}, g_a^{-1}] \right).$$

Proof. Let v_i be any orthonormal basis of V^{λ} . Let $a_{ji} \stackrel{\text{def}}{=} \langle g_a v_i, v_j \rangle$ be the matrix coefficients of the matrix $a = (a_{ij})$ by which g acts on V^{λ} with respect to $\{v_i\}$. Similarly, define matrices b, c, d for g_b, g_c, g_d in V^{λ} . We have

$$\begin{aligned} \text{tr}_{W^{\lambda}} (B_{\lambda} \circ (g_a, g_b, g_c, g_d)) &= \sum_{i_1, \dots, i_8} \langle B_{\lambda} \circ (g_a, g_b, g_c, g_d) v_{i_1} \otimes \check{v}_{i_2} \otimes v_{i_3} \otimes \check{v}_{i_4} \otimes v_{i_5} \otimes \check{v}_{i_6} \otimes v_{i_7} \otimes \check{v}_{i_8}, \\ &\quad v_{i_1} \otimes \check{v}_{i_2} \otimes v_{i_3} \otimes \check{v}_{i_4} \otimes v_{i_5} \otimes \check{v}_{i_6} \otimes v_{i_7} \otimes \check{v}_{i_8} \rangle \\ &= \sum_{\substack{i_1, \dots, i_8 \\ j_1, \dots, j_8}} a_{j_1 i_1} (a^{-1})_{i_2 j_2} b_{j_3 i_3} (b^{-1})_{i_4 j_4} c_{j_5 i_5} (c^{-1})_{i_6 j_6} d_{j_7 i_7} (d^{-1})_{i_8 j_8} \cdot \\ &\quad \langle B_{\lambda} v_{j_1} \otimes \check{v}_{j_2} \otimes v_{j_3} \otimes \check{v}_{j_4} \otimes v_{j_5} \otimes \check{v}_{j_6} \otimes v_{j_7} \otimes \check{v}_{j_8}, \\ &\quad v_{i_1} \otimes \check{v}_{i_2} \otimes v_{i_3} \otimes \check{v}_{i_4} \otimes v_{i_5} \otimes \check{v}_{i_6} \otimes v_{i_7} \otimes \check{v}_{i_8} \rangle \end{aligned}$$

which equals

$$\begin{aligned}
 &= \sum_{j_1, j_2, j_4, j_5, j_6, j_8, i_1, i_4} a_{j_1 i_1} (a^{-1})_{j_4 j_2} b_{j_2 j_1} (b^{-1})_{i_4 j_4} c_{j_5 i_4} (c^{-1})_{j_8 j_6} d_{j_6 j_5} (d^{-1})_{i_1 j_8} \\
 &= \sum_{j_1, j_2, j_4, j_5, j_6, j_8, i_1, i_4} (d^{-1})_{i_1 j_8} (c^{-1})_{j_8 j_6} d_{j_6 j_5} c_{j_5 i_4} (b^{-1})_{i_4 j_4} (a^{-1})_{j_4 j_2} b_{j_2 j_1} a_{j_1 i_1} \\
 &= \chi_\lambda \left(\left[g_d^{-1}, g_c^{-1} \right] \left[g_b^{-1}, g_a^{-1} \right] \right).
 \end{aligned}$$

The third equality used equation (5.9). □

Using Lemma 5.4 allows us to relate $\Theta_\lambda(Y, \mathcal{J}_n)$ to orthogonal projections in the space W^λ . For each $f \in \{a, b, c, d\}$, let P_f be the orthogonal projection in W^λ onto the vectors that are invariant by G_f . We let $Q \stackrel{\text{def}}{=} P_a P_b P_c P_d$.

Lemma 5.5. *We have $\Theta_\lambda(Y, \mathcal{J}_n) = \text{tr}_{W^\lambda} (B_\lambda g^0 Q)$.*

Proof. Using Lemma 5.4, we can write

$$\begin{aligned}
 \Theta_\lambda(Y, \mathcal{J}_n) &= \int_{h_f \in G_f} \chi_\lambda \left(\left[\left(g_d^0 h_d \right)^{-1}, \left(g_c^0 h_c \right)^{-1} \right] \left[\left(g_b^0 h_b \right)^{-1}, \left(g_a^0 h_a \right)^{-1} \right] \right) \\
 &= \text{tr}_{W^\lambda} \left(B_\lambda g^0 P_a P_b P_c P_d \right) = \text{tr}_{W^\lambda} \left(B_\lambda g^0 Q \right).
 \end{aligned}$$
□

Hence, we now wish to calculate $\text{tr}_{W^\lambda} (B_\lambda g^0 Q)$. For each $f \in \{a, b, c, d\}$ and $T \in \text{Tab}(\lambda)$, let

$$v_T^{\sigma_f^\pm} \stackrel{\text{def}}{=} \left(\sigma_f^\pm \right)^{-1} (v_T), \quad v_T^{\tau_f^\pm} \stackrel{\text{def}}{=} \left(\tau_f^\pm \right)^{-1} (v_T). \tag{5.10}$$

Similarly, if $\nu \subset_b \lambda$, recalling that if $T \in \text{Tab}(\lambda/\nu)$, w_T denotes the corresponding Gelfand–Tsetlin basis element of $V^{\lambda/\nu}$, we define

$$w_T^{\sigma_f^\pm} \stackrel{\text{def}}{=} \left(\sigma_f^\pm \right)^{-1} (w_T), \quad w_T^{\tau_f^\pm} \stackrel{\text{def}}{=} \left(\tau_f^\pm \right)^{-1} (w_T);$$

this makes sense as σ_f^\pm and τ_f^\pm are in S'_b . Recalling the notation $\mathcal{E}_{\mu, R_1, R_2}^\lambda$ from Lemma 4.1 (where $\mu \subset \lambda$), we define

$$\mathcal{E}_{\mu, R_1, R_2}^{\lambda, f, \pm} \stackrel{\text{def}}{=} \left(\sigma_f^\pm \right)^{-1} \otimes \left(\tau_f^\pm \right)^{-1} \left(\mathcal{E}_{\mu, R_1, R_2}^\lambda \right). \tag{5.11}$$

Lemma 5.6. *For $\lambda \vdash n$, the elements*

$$\left\{ \mathcal{E}_{\mu_a, S_a, T_a}^{\lambda, a, -} \otimes \mathcal{E}_{\mu_b, S_b, T_b}^{\lambda, b, -} \otimes \mathcal{E}_{\mu_c, S_c, T_c}^{\lambda, c, -} \otimes \mathcal{E}_{\mu_d, S_d, T_d}^{\lambda, d, -} : \mu_f \subset_{e_f} \lambda, S_f, T_f \in \text{Tab}(\lambda/\mu_f) \right\} \tag{5.12}$$

are an orthonormal basis for the G -invariant vectors in W^λ .

Proof. Consider W^λ as a module for $\mathbf{G} \stackrel{\text{def}}{=} S_a^{(1)} \times S_a^{(2)} \times S_b^{(1)} \times S_b^{(2)} \times S_c^{(1)} \times S_c^{(2)} \times S_d^{(1)} \times S_d^{(2)}$, where all the $S_f^{(i)}$ are isomorphic copies of S_n , $S_f^{(1)}$ acts on the V_f^λ factor and $S_f^{(2)}$ acts on the \check{V}_f^λ factor of W^λ . Given subgroups H_f of S_n for each $f \in \{a, b, c, d\}$, write $\Delta(H_a, H_b, H_c, H_d)$ for the subgroup consisting of tuples of the form $(g_a, g_a, g_b, g_b, g_c, g_c, g_d, g_d)$ with each $g_f \in H_f$. The statement of Lemma 5.6 is equivalent to the statement that the set given in equation (5.12) is an orthonormal basis for the $\Delta(G_a, G_b, G_c, G_d)$ -invariant elements.

We have

$$\begin{aligned} & \mathcal{E}_{\mu_a, S_a, T_a}^{\lambda, a, -} \otimes \mathcal{E}_{\mu_b, S_b, T_b}^{\lambda, b, -} \otimes \mathcal{E}_{\mu_c, S_c, T_c}^{\lambda, c, -} \otimes \mathcal{E}_{\mu_d, S_d, T_d}^{\lambda, d, -} = \\ & (\sigma_a^-, \tau_a^-, \sigma_b^-, \tau_b^-, \sigma_c^-, \tau_c^-, \sigma_d^-, \tau_d^-)^{-1} \mathcal{E}_{\mu_a, S_a, T_a}^{\lambda} \otimes \mathcal{E}_{\mu_b, S_b, T_b}^{\lambda} \otimes \mathcal{E}_{\mu_c, S_c, T_c}^{\lambda} \otimes \mathcal{E}_{\mu_d, S_d, T_d}^{\lambda}. \end{aligned}$$

We note that $(\sigma_a^-, \tau_a^-, \sigma_b^-, \tau_b^-, \sigma_c^-, \tau_c^-, \sigma_d^-, \tau_d^-)$ acts unitarily on W^λ , and by Lemma 4.1, the vectors $\mathcal{E}_{\mu_a, S_a, T_a}^{\lambda} \otimes \mathcal{E}_{\mu_b, S_b, T_b}^{\lambda} \otimes \mathcal{E}_{\mu_c, S_c, T_c}^{\lambda} \otimes \mathcal{E}_{\mu_d, S_d, T_d}^{\lambda}$ are an orthonormal basis for the $\Delta(S_{n-e_a}, S_{n-e_b}, S_{n-e_c}, S_{n-e_d})$ -invariant vectors in W^λ . Therefore, the set given in equation (5.12) is an orthonormal basis of invariant vectors for the group

$$(\sigma_a^-, \tau_a^-, \sigma_b^-, \tau_b^-, \sigma_c^-, \tau_c^-, \sigma_d^-, \tau_d^-)^{-1} \Delta(S_{n-e_a}, S_{n-e_b}, S_{n-e_c}, S_{n-e_d}) (\sigma_a^-, \tau_a^-, \sigma_b^-, \tau_b^-, \sigma_c^-, \tau_c^-, \sigma_d^-, \tau_d^-).$$

It remains to prove that this group is $\Delta(G_a, G_b, G_c, G_d)$. By property P1, this group is contained in $G_a \times G_a \times G_b \times G_b \times G_c \times G_c \times G_d \times G_d$. Combining this with property P3, the group displayed above is equal to $\Delta(G_a, G_b, G_c, G_d)$, as required. \square

Lemma 5.7. *We have*

$$g^0 \left(\mathcal{E}_{\mu_a, S_a, T_a}^{\lambda, a, -} \otimes \mathcal{E}_{\mu_b, S_b, T_b}^{\lambda, b, -} \otimes \mathcal{E}_{\mu_c, S_c, T_c}^{\lambda, c, -} \otimes \mathcal{E}_{\mu_d, S_d, T_d}^{\lambda, d, -} \right) = \mathcal{E}_{\mu_a, S_a, T_a}^{\lambda, a, +} \otimes \mathcal{E}_{\mu_b, S_b, T_b}^{\lambda, b, +} \otimes \mathcal{E}_{\mu_c, S_c, T_c}^{\lambda, c, +} \otimes \mathcal{E}_{\mu_d, S_d, T_d}^{\lambda, d, +}.$$

Proof. This follows from property P2 together with the definitions of $\mathcal{E}_{\mu_f, S_f, T_f}^{\lambda, f, \pm}$ in equation (5.11). \square

Proposition 5.8. *Recalling the definition of $\Theta_\lambda(Y, \mathcal{J}_n)$ from equation (5.7), we have*

$$\Theta_\lambda(Y, \mathcal{J}_n) = \sum_{\nu \subset_{\mathfrak{v}-\mathfrak{f}} \lambda' \subset_{\mathfrak{f}} \lambda} d_{\lambda/\lambda'} d_\nu \sum_{\nu \subset_{\mu_f} \subset_{\epsilon_f} \nu \subset_{\mathfrak{f}} \lambda'} \frac{1}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} Y_n \left(\left\{ \sigma_f^\pm, \tau_f^\pm \right\}, \nu, \{\mu_f\}, \lambda' \right), \quad (5.13)$$

where

$$Y_n \left(\left\{ \sigma_f^\pm, \tau_f^\pm \right\}, \nu, \{\mu_f\}, \lambda' \right) \stackrel{\text{def}}{=} \sum_{\substack{r_f^+, r_f^- \in \text{Tab}(\mu_f/\nu) \\ s_f, t_f \in \text{Tab}(\lambda'/\mu_f)}} \mathcal{M} \left(\left\{ \sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f \right\} \right) \quad (5.14)$$

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

$$\begin{aligned} \mathcal{M} \left(\left\{ \sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f \right\} \right) & \stackrel{\text{def}}{=} \left\langle \sigma_b^- (\sigma_a^+)^{-1} w_{r_a^+ \sqcup s_a}, w_{r_b^- \sqcup s_b} \right\rangle \left\langle \tau_a^+ (\sigma_b^+)^{-1} w_{r_b^+ \sqcup s_b}, w_{r_a^+ \sqcup t_a} \right\rangle \cdot \\ & \left\langle \tau_b^+ (\tau_a^-)^{-1} w_{r_a^- \sqcup t_a}, w_{r_b^+ \sqcup t_b} \right\rangle \left\langle \sigma_c^- (\tau_b^-)^{-1} w_{r_b^- \sqcup t_b}, w_{r_c^- \sqcup s_c} \right\rangle \cdot \\ & \left\langle \sigma_d^- (\sigma_c^+)^{-1} w_{r_c^+ \sqcup s_c}, w_{r_d^- \sqcup s_d} \right\rangle \left\langle \tau_c^+ (\sigma_d^+)^{-1} w_{r_d^+ \sqcup s_d}, w_{r_c^+ \sqcup t_c} \right\rangle \cdot \\ & \left\langle \tau_d^+ (\tau_c^-)^{-1} w_{r_c^- \sqcup t_c}, w_{r_d^+ \sqcup t_d} \right\rangle \left\langle \sigma_a^- (\tau_d^-)^{-1} w_{r_d^- \sqcup t_d}, w_{r_a^- \sqcup s_a} \right\rangle. \quad (5.15) \end{aligned}$$

Before proving Proposition 5.8, we say a word about the interpretation of the formula. Recall that the permutations σ_f^\pm and τ_f^\pm all belong to $S'_v \leq S_n$. But by property P4, the eight permutations appearing in equation (5.15) all restrict to the identity on $[n - \mathfrak{f} + 1, n]$, and so can be seen as permutations on $[n - \mathfrak{v} + 1, n - \mathfrak{f}]$. For every $\pi \in \left\{ \sigma_f^\pm, \tau_f^\pm \right\}$, $\pi^{-1}([n - \mathfrak{v} + 1, n - \mathfrak{f}])$ correspond to vertices where the corresponding side-of-half- f -edge belongs either to a hanging half-edge, or to an exposed side of a full-edge. One should think of r_f^- as the skew Young tableau consisting of indices of outgoing hanging half-edges labeled f , of r_f^+ as the tableau of incoming hanging half-edges labeled f , of s_f as the tableau

of exposed left-sides of full f -edges and of t_f as the tableau of exposed right-sides of full f -edges. Then indeed, for example, the indices corresponding to σ_f^- are $r_f^- \sqcup s_f$, those corresponding to σ_f^+ are $r_f^+ \sqcup s_f$, those corresponding to τ_f^- are $r_f^- \sqcup t_f$, and those corresponding to τ_f^+ are $r_f^+ \sqcup t_f$.

Proof of Proposition 5.8. By Lemmas 5.5, 5.6 and 5.7, we have

$$\begin{aligned} \Theta_\lambda(Y, \mathcal{J}_n) &= \text{tr}_{W^\lambda} \left(B_\lambda g^0 Q \right) \\ &= \sum_{\mu_f \subset_{e_f} \lambda, S_f, T_f \in \text{Tab}(\lambda/\mu_f)} \langle B_\lambda g^0 \left[\mathcal{E}_{\mu_a, S_a, T_a}^{\lambda, a, -} \otimes \mathcal{E}_{\mu_b, S_b, T_b}^{\lambda, b, -} \otimes \mathcal{E}_{\mu_c, S_c, T_c}^{\lambda, c, -} \otimes \mathcal{E}_{\mu_d, S_d, T_d}^{\lambda, d, -} \right], \\ &\quad \mathcal{E}_{\mu_a, S_a, T_a}^{\lambda, a, -} \otimes \mathcal{E}_{\mu_b, S_b, T_b}^{\lambda, b, -} \otimes \mathcal{E}_{\mu_c, S_c, T_c}^{\lambda, c, -} \otimes \mathcal{E}_{\mu_d, S_d, T_d}^{\lambda, d, -} \rangle \\ &= \sum_{\mu_f \subset_{e_f} \lambda, S_f, T_f \in \text{Tab}(\lambda/\mu_f)} \langle B_\lambda \left[\mathcal{E}_{\mu_a, S_a, T_a}^{\lambda, a, +} \otimes \mathcal{E}_{\mu_b, S_b, T_b}^{\lambda, b, +} \otimes \mathcal{E}_{\mu_c, S_c, T_c}^{\lambda, c, +} \otimes \mathcal{E}_{\mu_d, S_d, T_d}^{\lambda, d, +} \right], \\ &\quad \mathcal{E}_{\mu_a, S_a, T_a}^{\lambda, a, -} \otimes \mathcal{E}_{\mu_b, S_b, T_b}^{\lambda, b, -} \otimes \mathcal{E}_{\mu_c, S_c, T_c}^{\lambda, c, -} \otimes \mathcal{E}_{\mu_d, S_d, T_d}^{\lambda, d, -} \rangle. \end{aligned} \tag{5.16}$$

Using equations (4.2), (5.9) and (5.11), we obtain

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

Since $\sigma_f^\pm, \tau_f^\pm \in S'_0$ for all $f \in \{a, b, c, d\}$, the only way the product of matrix coefficients above can be nonzero is if there is $\nu \vdash n - \mathfrak{v}$ such that $\nu \subset \mu_f$ for all $f \in \{a, b, c, d\}$, and all $R_f^+|_{\leq n-\mathfrak{v}}, R_f^-|_{\leq n-\mathfrak{v}}$ are equal and of shape ν . Also, recall from Section 3.3 that the action of $\sigma \in S'_0$ on a tableau of shape $\lambda \vdash n$ depends only on the boxes with numbers from $[n - \mathfrak{v} + 1, n]$. This gives

$$\begin{aligned} &\left\langle B_\lambda \left[\mathcal{E}_{\mu_a, S_a, T_a}^{\lambda, a, +} \otimes \mathcal{E}_{\mu_b, S_b, T_b}^{\lambda, b, +} \otimes \mathcal{E}_{\mu_c, S_c, T_c}^{\lambda, c, +} \otimes \mathcal{E}_{\mu_d, S_d, T_d}^{\lambda, d, +} \right], \mathcal{E}_{\mu_a, S_a, T_a}^{\lambda, a, -} \otimes \mathcal{E}_{\mu_b, S_b, T_b}^{\lambda, b, -} \otimes \mathcal{E}_{\mu_c, S_c, T_c}^{\lambda, c, -} \otimes \mathcal{E}_{\mu_d, S_d, T_d}^{\lambda, d, -} \right\rangle \\ &= \sum_{\nu \subset_{\mathfrak{v}} \lambda} \frac{d_\nu}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \sum_{r_f^+, r_f^- \in \text{Tab}(\mu_f/\nu)} \\ &\quad \left\langle w_{r_a^+ \sqcup S_a}^{\sigma_a^+}, w_{r_b^- \sqcup S_b}^{\sigma_b^-} \right\rangle \cdot \left\langle w_{r_b^+ \sqcup S_b}^{\sigma_b^+}, w_{r_a^+ \sqcup T_a}^{\tau_a^+} \right\rangle \cdot \left\langle w_{r_a^- \sqcup T_a}^{\tau_a^-}, w_{r_b^+ \sqcup T_b}^{\tau_b^+} \right\rangle \cdot \left\langle w_{r_b^- \sqcup T_b}^{\tau_b^-}, w_{r_c^- \sqcup S_c}^{\sigma_c^-} \right\rangle \cdot \\ &\quad \left\langle w_{r_c^+ \sqcup S_c}^{\sigma_c^+}, w_{r_d^- \sqcup S_d}^{\sigma_d^-} \right\rangle \cdot \left\langle w_{r_d^+ \sqcup S_d}^{\sigma_d^+}, w_{r_c^+ \sqcup T_c}^{\tau_c^+} \right\rangle \cdot \left\langle w_{r_c^- \sqcup T_c}^{\tau_c^-}, w_{r_d^+ \sqcup T_d}^{\tau_d^+} \right\rangle \cdot \left\langle w_{r_d^- \sqcup T_d}^{\tau_d^-}, w_{r_a^- \sqcup S_a}^{\sigma_a^-} \right\rangle. \end{aligned} \tag{5.17}$$

Putting equations (5.16) and (5.17) together yields

$$\Theta_\lambda(Y, \mathcal{J}_n) = \sum_{\nu \subset \mathfrak{v}\lambda} d_\nu \sum_{\nu \subset \mu_f \subset \mathfrak{c}_{e_f}\lambda} \frac{1}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \sum_{r_f^+, r_f^- \in \text{Tab}(\mu_f/\nu)} \sum_{s_f, t_f \in \text{Tab}(\lambda/\mu_f)} \left\langle w_{r_a^+ \sqcup s_a}^{\sigma_a^+}, w_{r_b^- \sqcup s_b}^{\sigma_b^-} \right\rangle \left\langle w_{r_b^+ \sqcup s_b}^{\sigma_b^+}, w_{r_a^+ \sqcup t_a}^{\tau_a^+} \right\rangle \left\langle w_{r_a^- \sqcup t_a}^{\tau_a^-}, w_{r_b^+ \sqcup t_b}^{\tau_b^+} \right\rangle \left\langle w_{r_b^- \sqcup t_b}^{\tau_b^-}, w_{r_c^- \sqcup s_c}^{\sigma_c^-} \right\rangle \left\langle w_{r_c^+ \sqcup s_c}^{\sigma_c^+}, w_{r_d^- \sqcup s_d}^{\sigma_d^-} \right\rangle \left\langle w_{r_d^+ \sqcup s_d}^{\sigma_d^+}, w_{r_c^+ \sqcup t_c}^{\tau_c^+} \right\rangle \left\langle w_{r_c^- \sqcup t_c}^{\tau_c^-}, w_{r_d^+ \sqcup t_d}^{\tau_d^+} \right\rangle \left\langle w_{r_d^- \sqcup t_d}^{\tau_d^-}, w_{r_a^- \sqcup s_a}^{\sigma_a^-} \right\rangle.$$

Now, $\langle w_{r_a^+ \sqcup s_a}^{\sigma_a^+}, w_{r_b^- \sqcup s_b}^{\sigma_b^-} \rangle = \langle \sigma_b^- (\sigma_a^+)^{-1} w_{r_a^+ \sqcup s_a}, w_{r_b^- \sqcup s_b} \rangle$ and so on, and property **P4** implies that each pair of skew Young tableaux occurring in the same matrix coefficient above have the elements $[n - \mathfrak{f} + 1, n]$ in the same boxes, if the matrix coefficient is nonzero. This implies that if the product of matrix coefficients is nonzero then all the skew Young tableaux above have the elements $[n - \mathfrak{f} + 1, n]$ in the same boxes and there is $\lambda' \subset \lambda$ such that $\mu_f \subset \mathfrak{c}_{e_f - \mathfrak{f}} \lambda'$ for all f . Therefore, the above is equal to

$$\Theta_\lambda(Y, \mathcal{J}_n) = \sum_{\nu \subset \mathfrak{v} - \mathfrak{f} \lambda' \subset \mathfrak{f} \lambda} d_{\lambda/\lambda'} d_\nu \sum_{\nu \subset \mu_f \subset \mathfrak{c}_{e_f - \mathfrak{f}} \lambda'} \frac{1}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \sum_{r_f^+, r_f^- \in \text{Tab}(\mu_f/\nu)} \sum_{s_f, t_f \in \text{Tab}(\lambda'/\mu_f)} \left\langle \sigma_b^- (\sigma_a^+)^{-1} w_{r_a^+ \sqcup s_a}, w_{r_b^- \sqcup s_b} \right\rangle \cdot \left\langle \tau_a^+ (\sigma_b^+)^{-1} w_{r_b^+ \sqcup s_b}, w_{r_a^+ \sqcup t_a} \right\rangle \cdot \left\langle \tau_b^+ (\tau_a^-)^{-1} w_{r_a^- \sqcup t_a}, w_{r_b^+ \sqcup t_b} \right\rangle \cdot \left\langle \sigma_c^- (\tau_b^-)^{-1} w_{r_b^- \sqcup t_b}, w_{r_c^- \sqcup s_c} \right\rangle \cdot \left\langle \sigma_d^- (\sigma_c^+)^{-1} w_{r_c^+ \sqcup s_c}, w_{r_d^- \sqcup s_d} \right\rangle \cdot \left\langle \tau_c^+ (\sigma_d^+)^{-1} w_{r_d^+ \sqcup s_d}, w_{r_c^+ \sqcup t_c} \right\rangle \cdot \left\langle \tau_d^+ (\tau_c^-)^{-1} w_{r_c^- \sqcup t_c}, w_{r_d^+ \sqcup t_d} \right\rangle \cdot \left\langle \sigma_a^- (\tau_d^-)^{-1} w_{r_d^- \sqcup t_d}, w_{r_a^- \sqcup s_a} \right\rangle.$$

This finishes the proof. □

It is also useful to know the following.

Lemma 5.9. We have $\Upsilon_n \left(\left\{ \sigma_f^\pm, \tau_f^\pm \right\}, \nu, \{\mu_f\}, \lambda' \right) = \Upsilon_n \left(\left\{ \sigma_f^\pm, \tau_f^\pm \right\}, \check{\nu}, \{\check{\mu}_f\}, \check{\lambda}' \right)$.

Proof. This uses that as $S_{n-\mathfrak{f}}$ modules, $V^{\lambda'}$ and $V^{\check{\lambda}'} \otimes \text{sign}$ are isomorphic by the map $w_T \mapsto w_{\check{T}}$. This gives

$$\begin{aligned} \mathcal{M} \left(\left\{ \sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f \right\} \right) &= \text{sign} \left(\sigma_d^- (\sigma_c^+)^{-1} \right) \text{sign} \left(\tau_c^+ (\sigma_d^+)^{-1} \right) \text{sign} \left(\tau_d^+ (\tau_c^-)^{-1} \right) \text{sign} \left(\sigma_a^- (\tau_d^-)^{-1} \right) \\ &\quad \cdot \text{sign} \left(\sigma_b^- (\sigma_a^+)^{-1} \right) \text{sign} \left(\tau_a^+ (\sigma_b^+)^{-1} \right) \text{sign} \left(\tau_b^+ (\tau_a^-)^{-1} \right) \text{sign} \left(\sigma_c^- (\tau_b^-)^{-1} \right) \cdot \\ &\quad \cdot \mathcal{M} \left(\left\{ \sigma_f^\pm, \tau_f^\pm, \check{r}_f^\pm, \check{s}_f, \check{t}_f \right\} \right) \\ &= \mathcal{M} \left(\left\{ \sigma_f^\pm, \tau_f^\pm, \check{r}_f^\pm, \check{s}_f, \check{t}_f \right\} \right), \end{aligned}$$

where the last line used **P2** to get $(\tau_f^+)^{-1} \tau_f^- = (\sigma_f^+)^{-1} \sigma_f^- = g_0^f$. Using this identity gives the result. □

We are now ready to give an exact expression for $\mathbb{E}_n^{\text{emb}}(Y)$, which is the main result of this §5.4.

Theorem 5.10. For $n \geq \mathfrak{v}$, we have¹⁴

$$\mathbb{E}_n^{\text{emb}}(Y) = \frac{(n!)^3}{|\mathbb{X}_n|} \cdot \frac{(n)_{\mathfrak{v}} (n)_{\mathfrak{f}}}{\prod_f (n)_{e_f}} \mathbb{E}_n(Y), \tag{5.18}$$

¹⁴For general $g \geq 2$, the exponent of $(n!)$ in equation (5.18) is $2g - 1$.

where

$$\Xi_n(Y) \stackrel{\text{def}}{=} \sum_{\nu \subset_{v-\bar{f}} \lambda' \vdash n-\bar{f}} d\lambda d\nu \sum_{\nu \subset \mu_f \subset_{e_f-\bar{f}} \lambda'} \frac{1}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \Upsilon_n \left(\left\{ \sigma_f^\pm, \tau_f^\pm \right\}, \nu, \{\mu_f\}, \lambda' \right). \tag{5.19}$$

Remark 5.11. Although the expression (5.19) seems to depend on the choices of σ_f^\pm etc. that we have made already, the relation (5.18) shows that it only depends on Y .

Proof of Theorem 5.10. Recall from equation (5.6) that $\mathbb{E}_n^{\text{emb}}(Y) = \frac{n!}{(n-\mathbf{v})!} \frac{|\mathbb{X}_n(Y, \mathcal{J}_n)|}{|\mathbb{X}_n|}$. Combining this with Propositions 5.1 and 5.8 gives

$$\mathbb{E}_n^{\text{emb}}(Y) = \frac{\prod_{f \in a,b,c,d} (n - e_f)!}{(n - \mathbf{v})! |\mathbb{X}_n|} \sum_{\nu \subset_{v-\bar{f}} \lambda' \subset_{\bar{f}} \lambda \vdash n} d\lambda d_{\lambda'} d\nu \sum_{\nu \subset \mu_f \subset_{e_f-\bar{f}} \lambda'} \frac{1}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \Upsilon_n \left(\left\{ \sigma_f^\pm, \tau_f^\pm \right\}, \nu, \{\mu_f\}, \lambda' \right).$$

Applying Lemma 3.1 to the summation over λ , for fixed λ' , yields

$$\mathbb{E}_n^{\text{emb}}(Y) = \frac{n! \prod_{f \in a,b,c,d} (n - e_f)!}{(n - \mathbf{v})! (n - \bar{f})! |\mathbb{X}_n|} \sum_{\nu \subset_{v-\bar{f}} \lambda' \vdash n-\bar{f}} d\lambda d\nu \sum_{\nu \subset \mu_f \subset_{e_f-\bar{f}} \lambda'} \frac{1}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \Upsilon_n \left(\left\{ \sigma_f^\pm, \tau_f^\pm \right\}, \nu, \{\mu_f\}, \lambda' \right).$$

□

In view of Theorem 5.10, from now on we will not need to refer to the partition $\lambda \vdash n$, but only to the partitions $\nu \subset_{v-e_f} \mu_f \subset_{e_f-\bar{f}} \lambda' \vdash n - \bar{f}$. For ease of notation, from now on we shall abuse notation and **write λ instead of λ'** .

Before moving on, we prove that the recently defined functions $\Upsilon_n(\{\sigma_f^\pm, \tau_f^\pm\}, \nu, \{\mu_f\}, \lambda)$ are analytic functions of n^{-1} when ν, μ_f, λ each vary in a family of YDs. Recall the notation $\lambda(n)$ and $T(n)$ from Section 4.4.

Lemma 5.12. *Still assume that (Y, \mathcal{J}) , σ_f^\pm and τ_f^\pm are all fixed in the sense of Sections 5.4 and 5.6. Suppose that we are given YDs $\nu \subset_{v-\bar{f}} \lambda$ and for each $f \in \{a, b, c, d\}$ a YD μ_f with*

$$\nu \subset_{v-e_f} \mu_f \subset_{e_f-\bar{f}} \lambda.$$

There is a function $Y^(\nu, \{\mu_f\}, \lambda, \bullet)$ that is holomorphic in some open disc in \mathbf{C} with center 0 such that for all n sufficiently large (depending on Y), $\lambda(n - \bar{f})$, $\mu_f(n - e_f)$ and $\nu(n - \mathbf{v})$ all exist and*

$$\Upsilon_n \left(\left\{ \sigma_f^\pm, \tau_f^\pm \right\}, \nu(n - \mathbf{v}), \{\mu_f(n - e_f)\}, \lambda(n - \bar{f}) \right) = Y^* \left(\nu, \{\mu_f\}, \lambda, n^{-1} \right).$$

In addition, the coefficients of the Taylor series of $Y^(\nu, \{\mu_f\}, \lambda, \bullet)$ are in \mathbf{Q} .*

Proof. The proof relies crucially on property **P5** of the permutations σ_f^\pm, τ_f^\pm stating that they are obtained from fixed permutations in S_ν . This means that each of the summands

$$\mathcal{M} \left(\left\{ \sigma_f^\pm, \tau_f^\pm, s_f(n - \bar{f}), t_f(n - \bar{f}), r_f^-(n - e_f), r_f^+(n - e_f) \right\} \right)$$

of

$$\Upsilon_n \left(\left\{ \sigma_f^\pm, \tau_f^\pm \right\}, \nu(n - \mathbf{v}), \{\mu_f(n - e_f)\}, \lambda(n - \bar{f}) \right)$$

(cf. equation (5.14)) agrees with a function of n^{-1} that is holomorphic in an open disc with center 0 and with rational coefficients of its Taylor series by Proposition 4.6. Since there are only finitely many summands in equation (5.14), this proves the lemma. □

We also give a coarse bound for the quantities $\Upsilon_n(\{\sigma_f^\pm, \tau_f^\pm\}, \nu, \{\mu_f\}, \lambda)$; this will be improved later in Proposition 5.21.

Lemma 5.13. *We have*

$$\left| \Upsilon_n \left(\left\{ \sigma_f^\pm, \tau_f^\pm \right\}, \nu, \{\mu_f\}, \lambda \right) \right| \leq (d_{\lambda/\nu})^8 \leq ((\mathbf{v} - \mathbf{f})!)^8.$$

Proof. For fixed $\nu, \{\mu_f\}, \lambda$, the range of summation in equation (5.14) is parameterized 1 : 1 by the eight tableaux of the form

$$r_f^+ \sqcup s_f, r_f^- \sqcup t_f \in \text{Tab}(\lambda/\nu).$$

Also, since the matrix coefficients in equation (5.15) involve unit vectors in a unitary representation, each summand in equation (5.15) is ≤ 1 in absolute value. Hence, the lemma follows. \square

5.5. A geometric bound for products of matrix coefficients

We continue to keep all the notations and assumptions of §§5.4. We will show that we can give improved bounds for the product of matrix coefficients $\mathcal{M}(\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\})$ defined in equation (5.15) in terms of geometric properties of Y . Recall the definitions of the functions top and d from §§4.3. We define

$$D_{\text{top}} \left(\left\{ \sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f \right\} \right) \stackrel{\text{def}}{=} \tag{5.20}$$

$$d \left(\sigma_b^- (\sigma_a^+)^{-1} \text{top}(r_a^+ \sqcup s_a), \text{top}(r_b^- \sqcup s_b) \right) + d \left(\tau_a^+ (\sigma_b^+)^{-1} \text{top}(r_b^+ \sqcup s_b), \text{top}(r_a^+ \sqcup t_a) \right) +$$

$$d \left(\tau_b^+ (\tau_a^-)^{-1} \text{top}(r_a^- \sqcup t_a), \text{top}(r_b^+ \sqcup t_b) \right) + d \left(\sigma_c^- (\tau_b^-)^{-1} \text{top}(r_b^- \sqcup t_b), \text{top}(r_c^- \sqcup s_c) \right) +$$

$$d \left(\sigma_d^- (\sigma_c^+)^{-1} \text{top}(r_c^+ \sqcup s_c), \text{top}(r_d^- \sqcup s_d) \right) + d \left(\tau_c^+ (\sigma_d^+)^{-1} \text{top}(r_d^+ \sqcup s_d), \text{top}(r_c^+ \sqcup t_c) \right) +$$

$$d \left(\tau_d^+ (\tau_c^-)^{-1} \text{top}(r_c^- \sqcup t_c), \text{top}(r_d^+ \sqcup t_d) \right) + d \left(\sigma_a^- (\tau_d^-)^{-1} \text{top}(r_d^- \sqcup t_d), \text{top}(r_a^- \sqcup s_a) \right).$$

Proposition 4.4 directly implies the following result.

Lemma 5.14. *If $\lambda_1 + \nu_1 > n - \mathbf{f} + (\mathbf{v} - \mathbf{f})^2$, then*

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

Remark 5.15. If ν has a fixed bound on the number of boxes outside its first row and Y is fixed, then the hypothesis of Lemma 5.14 is satisfied for sufficiently large n .

The quantities $D_{\text{top}}(\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\})$ have a useful interpretation in terms of the combinatorics of the boundary cycles of Y . To explain this, we construct from the data $\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\}$ a labeling of the sides/half-sides of rectangles/half-rectangles in $Y_+^{(1)}$.

5.6. Construction of labelings of tiled surfaces from collections of tableaux

In this section, we keep all of the notation from the previous sections. In particular, we have a fixed vertex-labeled compact tiled surface (Y, \mathcal{J}_n) with \mathbf{v} vertices, e_f f -labeled edges for each $f \in \{a, b, c, d\}$, and \mathbf{f} octagons. We fix the data

$$\begin{aligned} \mathbf{v} \vdash n - \mathbf{v}, \lambda \vdash n - \mathbf{f} \\ \mathbf{v} \subset \mu_f \subset_{e_f - \mathbf{f}} \lambda \quad \forall f \in \{a, b, c, d\} \\ r_f^+, r_f^- \in \text{Tab}(\mu_f/\nu), s_f, t_f \in \text{Tab}(\lambda/\mu_f) \quad \forall f \in \{a, b, c, d\}. \end{aligned} \tag{5.21}$$

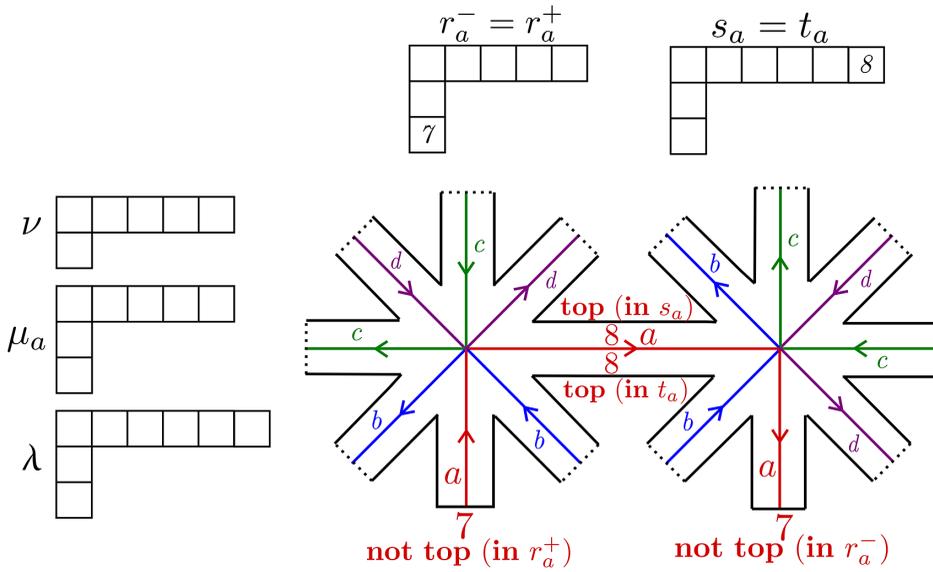


Figure 5.2. Illustration of how the tableaux r_f^-, r_f^+, s_f, t_f induce the ‘top’ labeling.

All this data uniquely determine one summand of $\Upsilon_n(\{\sigma_f^\pm, \tau_f^\pm\}, \nu, \{\mu_f\}, \lambda)$ as in equation (5.15) and hence also of $\Xi_n(Y)$ as in equation (5.19).

Recall that in §§5.3 we constructed the maps σ_f^\pm and τ_f^\pm according to numbering in $[\mathbf{v} - \mathbf{f}]$ of octagons, of exposed sides of full-edges and of hanging half-edges of Y_+ . By adding $n - \mathbf{v}$, these numbers are in $[n - \mathbf{v} + 1, n - \mathbf{f}]$, give rise to the images of the corresponding vertices of Y through σ_f^\pm, τ_f^\pm , and are the elements in the tableaux r_f^\pm, s_f, t_f . Given these tableaux, we assign a ‘top’ label to the hanging half-edges and exposed sides of full-edges which appear in the top row of the corresponding tableau. Namely,

- Every exposed left side (resp. right side) of an f -full-edge is labeled ‘top’ if the corresponding element in s_f (resp. t_f) lies in the top row.¹⁵
- Every outgoing (resp. incoming) hanging f -half-edge is labeled ‘top’ if the corresponding element in r_f^- (resp. r_f^+) lies in the top row.

This labeling scheme is illustrated in Figure 5.2.

The purpose of introducing these labelings is the following diagrammatic interpretation of D_{top} . We view the boundary ∂Y_+ of (the thick version of) Y_+ to consist of hanging half-edges and of exposed sides of full-edges.

Lemma 5.16. *The quantity $D_{\text{top}}(\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\})$ is half the number of incidences between two consecutive parts of ∂Y_+ among which one is labeled ‘top’ and other one is not.*

Proof. This follows simply by careful consideration of the definition (5.20) of D_{top} , as the permutations appearing in that definition map the index of one part of ∂Y_+ to a neighboring part. For example, $\sigma_b^-(\sigma_a^+)^{-1}$ maps the index on an exposed left side of an a -full-edge to the neighboring index which either belongs to an exposed left side of a b -full-edge or to a hanging outgoing b -half-edge.

The reason for the factor $\frac{1}{2}$ is that for A, B of the same size, $d(A, B)$ is half the size of the symmetric difference of A and B . While in the definition of D_{top} we count differences $d(A, B)$, in counting switches between ‘top’ parts to ‘nontop’ parts of ∂Y_+ and vice versa, we refer to symmetric differences. \square

¹⁵To be sure, the top row in this case is row number one of λ/μ_f , which may be empty: its length is $\lambda_1 - (\mu_f)_1$.

The benefit of Lemma 5.16 is that it allows us to connect the possible properties of Y being boundary reduced or strongly boundary reduced to nontrivial bounds for matrix coefficients.

A *piece* of ∂Y_+ is a contiguous collection of exposed sides of full-edges and of hanging half-edges. Given a piece P , we write $e(P)$ for the number of exposed sides-of-full-edges in P , $h_e(P)$ for the number of hanging half-edges in P and $\chi(P)$ for the Euler characteristic of P , which may be 0 (if P is a circle) or 1 (if P is topologically a path). A *piece collection* \mathcal{P} of ∂Y_+ is a collection of *disjoint* pieces of ∂Y_+ (without intersection even of endpoints).

Definition 5.17. Define

$$\text{Defect}(P) \stackrel{\text{def}}{=} e(P) - 3h_e(P)$$

and¹⁶

$$\max \text{Defect}(Y) \stackrel{\text{def}}{=} \max_{\mathcal{P} \neq \emptyset} \sum_{P \in \mathcal{P}} \text{Defect}(P) - 8\chi(P),$$

where the maximum is over all nonempty piece collections of ∂Y_+ .

Lemma 5.18. *If Y is boundary reduced, then*

$$\max \text{Defect}(Y) \leq 0. \tag{5.22}$$

If Y is strongly boundary reduced, then

$$\max \text{Defect}(Y) \leq -2. \tag{5.23}$$

Proof. Assume first that Y is boundary reduced. Recall that Y has no long blocks, so no blocks of size > 4 and no long chains, so between every two blocks of size 4 in the same piece, there must be either two consecutive hanging edges or one block of size ≤ 2 . As a result, for every piece P of ∂Y_+ that is a circle, we have $\text{Defect}(P) \leq 0$ so $\text{Defect}(P) - 8\chi(P) \leq 0$. For every piece P of ∂Y_+ that is a path, we have $\text{Defect}(P) \leq 4$: This bound is attained for example when P corresponds to a block of size 4 or to two consecutive blocks of size 4 and 3. Therefore, $\text{Defect}(P) - 8\chi(P) \leq -4$. Hence, all contributions to $\max \text{Defect}(Y)$ are nonpositive, and we obtain equation (5.22).

If Y is strongly boundary reduced, there are no blocks in ∂Y of length > 3 , and so every piece P of ∂Y_+ which is a path that satisfies $\text{Defect}(P) \leq 3$ and $\text{Defect}(P) - 8\chi(P) \leq -5$: This bound is attained when P corresponds to a chain of consecutive blocks of size 3 each. If P is a piece of ∂Y_+ that is a circle but not a cyclic chain, then there are two consecutive hanging half-edges and $\text{Defect}(P) = \text{Defect}(P) - \chi(P) \leq -3$. If P is a cyclic chain, then by [MP22a, Lem. 3.6] it cannot have only one block of size 2, so $\text{Defect}(P) = \text{Defect}(P) - \chi(P) \leq -2$. This proves equation (5.23). \square

The following result relates the structure of the pieces of ∂Y to the quantities D_{top} and D_{left} appearing in our previous bound (Lemma 5.14) for matrix coefficients. Given an SYD λ/ν , we write $b_{\lambda/\nu}$ for the number of boxes of λ/ν outside the first row.

Proposition 5.19. *Suppose we are given $\nu, \mu_f, \lambda, r_f^\pm, s_f, t_f$ as in equation (5.21). If $b_{\lambda/\nu} > 0$, then¹⁷*

$$b_{\lambda/\nu} - b_{\mu_a/\nu} - b_{\mu_b/\nu} - b_{\mu_c/\nu} - b_{\mu_d/\nu} - D_{\text{top}} \left(\left\{ \left\{ \sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f \right\} \right\} \right) \leq \frac{1}{8} \max \text{Defect}(Y).$$

Proof. We define a collection of pieces \mathcal{P} of ∂Y_+ according to the ‘top’ labels: Pieces are contiguous segments of ∂Y_+ (hanging half-edges or exposed sides of full edges) which are **not** labeled ‘top’. This

¹⁶For general genus $g \geq 2$, the definitions are $\text{Defect}(P) = e(P) - (2g - 1)h_e(P)$ and $\max \text{Defect}(Y) = \max_{\mathcal{P}} \text{Defect}(P) - 4g\chi(P)$.

¹⁷For general $g \geq 2$ the bound is $\frac{1}{4g} \max \text{Defect}(Y)$.

collection is nonempty if and only if $b_{\lambda/\nu} > 0$, which holds by assumption. Let \mathcal{P}_0 denote the collection of such pieces that are circles, \mathcal{P}_1 denote the collection of such pieces that are paths and $\mathcal{P} = \mathcal{P}_0 \sqcup \mathcal{P}_1$. It follows from Lemma 5.16 that

$$D_{\text{top}} \stackrel{\text{def}}{=} D_{\text{top}} \left(\left\{ \sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f \right\} \right) = |\mathcal{P}_1|.$$

Now, $b_{\lambda/\nu} = \frac{1}{8} \sum_{P \in \mathcal{P}} [\mathfrak{e}(P) + \mathfrak{h}\mathfrak{e}(P)]$ because for every $f \in \{a, b, c, d\}$, every square outside the top row of λ/ν corresponds either to two hanging half-edges (in case this square belongs to μ_f/ν) or to two exposed-sides-of-full-edges (in case this square belongs to λ/μ_f). Similarly, $b_{\mu_a/\nu} + b_{\mu_b/\nu} + b_{\mu_c/\nu} + b_{\mu_d/\nu} = \frac{1}{2} \sum_{P \in \mathcal{P}} \mathfrak{h}\mathfrak{e}(P)$. Thus,

$$\begin{aligned} b_{\lambda/\nu} - b_{\mu_a/\nu} - b_{\mu_b/\nu} - b_{\mu_c/\nu} - b_{\mu_d/\nu} - D_{\text{top}} &= \sum_{P \in \mathcal{P}} \left(\frac{1}{8} (\mathfrak{e}(P) + \mathfrak{h}\mathfrak{e}(P)) - \frac{1}{2} \mathfrak{h}\mathfrak{e}(P) \right) - \sum_{P \in \mathcal{P}_1} 1 \\ &= \sum_{P \in \mathcal{P}_0} \frac{1}{8} \text{Defect}(P) + \sum_{P \in \mathcal{P}_1} \frac{1}{8} (\text{Defect}(P) - 8) \\ &= \sum_{P \in \mathcal{P}} \frac{1}{8} (\text{Defect}(P) - 8\chi(P)) \leq \frac{1}{8} \max \text{Defect}(Y). \end{aligned}$$

□

Proposition 5.20. *Suppose we are given $\nu, \mu_f, \lambda, r_f^-, r_f^+, s_f, t_f$ as in equation (5.21).*

1. *If Y is boundary reduced,*

$$D_{\text{top}} \left(\left\{ \sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f \right\} \right) \geq b_{\lambda/\nu} - b_{\mu_a/\nu} - b_{\mu_b/\nu} - b_{\mu_c/\nu} - b_{\mu_d/\nu}. \tag{5.24}$$

2. *If Y is strongly boundary reduced, then equation (5.24) becomes an equality if and only if*

$$D_{\text{top}} \left(\left\{ \sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f \right\} \right) = b_{\lambda/\nu} = b_{\mu_a/\nu} = b_{\mu_b/\nu} = b_{\mu_c/\nu} = b_{\mu_d/\nu} = 0. \tag{5.25}$$

Otherwise,

$$D_{\text{top}} \left(\left\{ \sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f \right\} \right) \geq b_{\lambda/\nu} - b_{\mu_a/\nu} - b_{\mu_b/\nu} - b_{\mu_c/\nu} - b_{\mu_d/\nu} + 1. \tag{5.26}$$

Proof. Note that, for any Y , if $b_{\lambda/\nu} = 0$, then $D_{\text{top}}(\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\}) = b_{\mu_a/\nu} = b_{\mu_b/\nu} = b_{\mu_c/\nu} = b_{\mu_d/\nu} = 0$. Otherwise, Proposition 5.19 applies, and we obtain the statement of the proposition by combining Proposition 5.19 with the inequalities (5.22) and (5.23). To obtain equation (5.26) from equation (5.23) (instead of a bound featuring $\frac{1}{4}$), one uses that all b_\bullet quantities and $D_{\text{top}}(\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\})$ are integer valued. □

Proposition 5.20 together with Lemma 5.16 have the following important consequence for the quantities $Y_n(\{\sigma_f^\pm, \tau_f^\pm\}, \nu, \{\mu_f\}, \lambda)$.

Proposition 5.21. *Suppose that $\nu, \{\mu_f\}, \lambda$ are as in equation (5.21) and that $\lambda_1 + \nu_1 > n - \mathfrak{f} + (\nu - \mathfrak{f})^2$.*

1. *If Y is boundary reduced, then*

$$\left| Y_n \left(\left\{ \sigma_f^\pm, \tau_f^\pm \right\}, \nu, \{\mu_f\}, \lambda \right) \right| \leq ((\nu - \mathfrak{f})!)^8 \left(\frac{(\nu - \mathfrak{f})^2}{\nu_1 + \lambda_1 - (n - \mathfrak{f})} \right)^{b_{\lambda/\nu} - b_{\mu_a/\nu} - b_{\mu_b/\nu} - b_{\mu_c/\nu} - b_{\mu_d/\nu}}. \tag{5.27}$$

2. If Y is strongly boundary reduced and $b_{\lambda/\nu} > 0$, then

$$\left| Y_n \left(\left\{ \sigma_f^\pm, \tau_f^\pm \right\}, \nu, \{\mu_f\}, \lambda \right) \right| \leq ((\mathbf{v} - \mathbf{f})!)^8 \left(\frac{(\mathbf{v} - \mathbf{f})^2}{\nu_1 + \lambda_1 - (n - \mathbf{f})} \right)^{1+b_{\lambda/\nu}-b_{\mu_a/\nu}-b_{\mu_b/\nu}-b_{\mu_c/\nu}-b_{\mu_d/\nu}} \quad (5.28)$$

3. For any Y , if $b_{\lambda/\nu} = 0$, then

$$Y_n \left(\left\{ \sigma_f^\pm, \tau_f^\pm \right\}, \nu, \{\mu_f\}, \lambda \right) = 1.$$

Proof. Part 1. Suppose that Y is boundary reduced. As argued in the proof of 5.13, there are at most $(d_{\lambda/\nu})^8 \leq ((\mathbf{v} - \mathbf{f})!)^8$ summands in the definition (5.14) of $Y_n(\{\sigma_f^\pm, \tau_f^\pm\}, \nu, \{\mu_f\}, \lambda)$. Each summand is some $\mathcal{M}(\{\sigma_f^\pm, \tau_f^\pm, r_f^\pm, s_f, t_f\})$, so by Lemma 5.14 and Proposition 5.20 Part 1, we get that each summand of equation (5.14) has absolute value

$$\leq \left(\frac{(\mathbf{v} - \mathbf{f})^2}{\nu_1 + \lambda_1 - (n - \mathbf{f})} \right)^{b_{\lambda/\nu}-b_{\mu_a/\nu}-b_{\mu_b/\nu}-b_{\mu_c/\nu}-b_{\mu_d/\nu}}.$$

This proves equation (5.27).

Part 2. Suppose now that Y is strongly boundary reduced and that $b_{\lambda/\nu} > 0$. This time, Lemma 5.14 and Proposition 5.20 give that each summand of $Y_n(\{\sigma_f^\pm, \tau_f^\pm\}, \nu, \{\mu_f\}, \lambda)$ in equation (5.14) has absolute value

$$\leq \left(\frac{(\mathbf{v} - \mathbf{f})^2}{\nu_1 + \lambda_1 - (n - \mathbf{f})} \right)^{1+b_{\lambda/\nu}-b_{\mu_a/\nu}-b_{\mu_b/\nu}-b_{\mu_c/\nu}-b_{\mu_d/\nu}}.$$

As in Part 1, there are $\leq ((\mathbf{v} - \mathbf{f})!)^8$ summands of equation (5.14), so this proves equation (5.28).

Part 3. Suppose that $b_{\lambda/\nu} = 0$. Then there is only one possible choice for the tableaux r_f^-, r_f^+, s_f, t_f in equation (5.14). Moreover, $V^{\lambda/\nu}$ is the trivial module for the relevant copy of $S_{\mathbf{v}-\mathbf{f}}$, hence the product of matrix coefficients appearing in equation (5.15) is equal to 1, and $Y_n(\{\sigma_f^\pm, \tau_f^\pm\}, \nu, \{\mu_f\}, \lambda) = 1$. \square

5.7. Stronger bounds for matrix coefficients

In this section, we give a strengthening of Proposition 5.19 that is used in a sequel to this paper [MNP22]. A reader who is only interested in the current paper may skip this short §5.7.

Proposition 5.22. *Suppose we are given YDs $\nu \vdash n - \mathbf{v}$, $\lambda \vdash n - \mathbf{f}$ and μ_f such that*

$$\nu \subset \mu_f \subset_{e_f - \mathbf{f}} \lambda, \quad \forall f \in \{a, b, c, d\},$$

and tableaux $r_f^+, r_f^- \in \text{Tab}(\mu_f/\nu)$ and $s_f, t_f \in \text{Tab}(\lambda/\mu_f)$. Fix $\varepsilon \geq 0$, and suppose in addition that for every piece P of ∂Y we have

$$\text{Defect}(P) - 4\chi(P) \leq -\varepsilon (e(P) + \mathfrak{h}e(P)).$$

Then

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

Proof. We follow the same construction of a piece collection \mathcal{P} as in the proof of Proposition 5.19. This leads to

$$\begin{aligned}
 b_{\lambda/\nu} - b_{\mu_a/\nu} - b_{\mu_b/\nu} - b_{\mu_c/\nu} - b_{\mu_d/\nu} - D_{\text{top}} &= \frac{1}{8} \sum_{P \in \mathcal{P}} \text{Defect}(P) - 8\chi(P) \leq \frac{1}{8} \sum_{P \in \mathcal{P}} \text{Defect}(P) - 4\chi(P) \\
 &\leq -\frac{\varepsilon}{8} \sum_{P \in \mathcal{P}} (\mathfrak{e}(P) + \mathfrak{h}\mathfrak{e}(P)) = -\varepsilon b_{\lambda/\nu},
 \end{aligned}$$

as required. □

5.8. Approximating $\Xi_n(Y)$ by Laurent polynomials

In this section, we keep all the notations and assumptions of §§5.4. We want to show that we can replace the summation over ν, μ_f , and λ given the definition of $\Xi_n(Y)$ in equation (5.19) by a sum of finite size, independent of n , at the cost of a controllable error term. To state this precisely, recalling the definition of $\Lambda(n, b)$ from §§4.4, and letting $b \in \mathbb{N}$, we introduce the quantity

$$\Xi_n^{(b)}(Y) \stackrel{\text{def}}{=} \sum_{\substack{\nu \subset_{\mathfrak{v}-\mathfrak{f}} \lambda \vdash n-\mathfrak{f} \\ \nu \notin \Lambda(n-\mathfrak{v}, b)}} d_\lambda d_\nu \sum_{\nu \subset \mu_f \subset_{\mathfrak{e}_f-\mathfrak{f}} \lambda} \frac{1}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \Upsilon_n \left(\left\{ \sigma_f^\pm, \tau_f^\pm \right\}, \nu, \{\mu_f\}, \lambda \right). \tag{5.29}$$

In this summation, we restrict to ν that have less than b boxes either outside the first row or outside the first column. Note that whereas $\Xi_n(Y)$ does not depend on any of the choices of g^0 and numberings we made in Sections 5.2 and 5.3 (see Remark 5.11), $\Xi_n^{(b)}(Y)$ may depend on these choices.

Lemma 5.23. *For a fixed tiled surface Y and $b \in \mathbb{N}$, for any vertex-ordering \mathcal{J} of Y as above, we have*

$$\Xi_n(Y) = \Xi_n^{(b)}(Y) + O\left(n^{\mathfrak{v}-\mathfrak{f}-2b}\right)$$

as $n \rightarrow \infty$. The implied constant depends on b, \mathfrak{v} and \mathfrak{f} .

Proof. Note that $d_{\mu_f} \geq d_\nu$, and for fixed ν and λ the number of μ_f with $\nu \subset \mu_f \subset_{\mathfrak{e}_f-\mathfrak{f}} \lambda$ is $\leq (\mathfrak{v} - \mathfrak{f})!$ (there is an injection from the collection of such μ_f to $\text{Tab}(\lambda/\nu)$ by filling the boxes of μ_f/ν with $[n - \mathfrak{v} + 1, n - \mathfrak{e}_f]$ and the other boxes of λ/ν arbitrarily). Using these observations together with Lemma 5.13, we obtain

$$\begin{aligned}
 \left| \Xi_n(Y) - \Xi_n^{(b)}(Y) \right| &= \left| \sum_{\substack{\nu \subset_{\mathfrak{v}-\mathfrak{f}} \lambda \vdash n-\mathfrak{f} \\ \nu \in \Lambda(n-\mathfrak{v}, b)}} d_\lambda d_\nu \sum_{\nu \subset \mu_f \subset_{\mathfrak{e}_f-\mathfrak{f}} \lambda} \frac{1}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \Upsilon \left(\left\{ \sigma_f^\pm, \tau_f^\pm \right\}, \nu, \{\mu_f\}, \lambda \right) \right| \\
 &\leq ((\mathfrak{v} - \mathfrak{f})!)^{12} \sum_{\substack{\nu \subset_{\mathfrak{v}-\mathfrak{f}} \lambda \vdash n-\mathfrak{f} \\ \nu \in \Lambda(n-\mathfrak{v}, b)}} \frac{d_\lambda}{d_\nu^3}.
 \end{aligned}$$

By Lemma 3.1, for a fixed $\nu \vdash n - \mathfrak{v}$, we have

$$\sum_{\lambda: \nu \subset_{\mathfrak{v}-\mathfrak{f}} \lambda} d_\lambda \leq \sum_{\lambda: \nu \subset_{\mathfrak{v}-\mathfrak{f}} \lambda} d_\lambda d_{\lambda/\nu} = \frac{(n - \mathfrak{f})!}{(n - \mathfrak{v})!} d_\nu.$$

Therefore, by Proposition 4.5,

$$\begin{aligned}
 \left| \Xi_n(Y) - \Xi_n^{(b)}(Y) \right| &\leq ((\mathfrak{v} - \mathfrak{f})!)^{12} \frac{(n - \mathfrak{f})!}{(n - \mathfrak{v})!} \sum_{\nu \in \Lambda(n-\mathfrak{v}, b)} \frac{1}{d_\nu^2} \\
 &= O_{b, \mathfrak{v}, \mathfrak{f}} \left(n^{\mathfrak{v}-\mathfrak{f}-2b} \right).
 \end{aligned}$$

□

Proposition 5.24. *For any $M \in \mathbf{N}$, there is a Laurent polynomial $\Xi_M^*(Y) \in \mathbf{Q}[t, t^{-1}]$ such that as $n \rightarrow \infty$*

$$\Xi_n(Y) = \Xi_M^*(Y)[n] + O\left(n^{-M}\right).$$

Proof. Let $b = \left\lceil \frac{v-f+M}{2} \right\rceil$. Then Lemma 5.23 yields that as $n \rightarrow \infty$,

$$\Xi_n(Y) = \Xi_n^{(b)}(Y) + O\left(n^{-M}\right). \tag{5.30}$$

Similarly to the proof of Proposition 4.8, we note that for $n - v > 2b$, the collection of $v \vdash n - v$ such that $v \notin \Lambda(n - v, b)$ is the disjoint union of $\Lambda_{b, \lambda < b}(n - v) = \{v \vdash n - v \mid v_1 > n - v - b\}$ and the dual partitions $\{\check{v} \mid v \in \Lambda_{b, \lambda < b}(n - v)\}$. Because each μ_f and λ in equation (5.29) extend v by a fixed number of boxes, there is a finite number ℓ of tuples of YDs

$$(v^i, \mu_a^i, \mu_b^i, \mu_c^i, \mu_d^i, \lambda^i), \quad 1 \leq i \leq \ell$$

with $v^i \in \Lambda_{b, \lambda < b}(2b)$ such that for each $1 \leq i \leq \ell$ and $f \in \{a, b, c, d\}$

$$v^i \subset_{v-e_f} \mu_f^i \subset_{e_f-f} \lambda^i.$$

Thus, equation (5.29) can be rewritten as

$$\begin{aligned} \Xi_n^{(b)}(Y) &= \sum_{i=1}^{\ell} \frac{d_{\lambda^i(n-f)} d_{v^i(n-v)}}{d_{\mu_a^i(n-e_a)} d_{\mu_b^i(n-e_b)} d_{\mu_c^i(n-e_c)} d_{\mu_d^i(n-e_d)}} \cdot \left[\Upsilon\left(\left\{\sigma_f^\pm, \tau_f^\pm\right\}, v^i(n-v), \left\{\mu_f^i(n-e_f)\right\}, \lambda^i(n-f)\right) \right. \\ &\quad \left. + \Upsilon\left(\left\{\sigma_f^\pm, \tau_f^\pm\right\}, v^i(\check{n}-v), \left\{\mu_f^i(\check{n}-e_f)\right\}, \lambda^i(\check{n}-f)\right) \right] \\ &= 2 \sum_{i=1}^{\ell} \frac{d_{\lambda^i(n-f)} d_{v^i(n-v)}}{d_{\mu_a^i(n-e_a)} d_{\mu_b^i(n-e_b)} d_{\mu_c^i(n-e_c)} d_{\mu_d^i(n-e_d)}} \cdot \Upsilon\left(\left\{\sigma_f^\pm, \tau_f^\pm\right\}, v^i(n-v), \left\{\mu_f^i(n-e_f)\right\}, \lambda^i(n-f)\right), \end{aligned}$$

where the last line used Lemma 5.9. For each i , the ratio of dimensions above agrees with a rational function $Q_i(n) \in \mathbf{Q}(n)$ of n (at least when $n - v \geq 2b$) by Lemma 4.7. Combining this with Lemma 5.12 gives that

$$\Xi_n^{(b)}(Y) = 2 \sum_{i=1}^{\ell} Q_i(n) \Upsilon^*\left(v^i, \left\{\mu_f^i\right\}, \lambda^i, n^{-1}\right)$$

agrees with $F(n^{-1})$, where F is a function of a complex variable z that is meromorphic in an open disc with center 0, and with coefficients of its Laurent series in \mathbf{Q} . Hence, $F(n^{-1})$ itself can be approximated to order $O(n^{-M})$ as $n \rightarrow \infty$ by a Laurent polynomial in n with coefficients in \mathbf{Q} . \square

5.9. Estimating $\Xi_n(Y)$

In this §5.9, we give estimates for $\Xi_n(Y)$ for fixed Y which is boundary reduced or strongly boundary reduced.

Proposition 5.25. *If Y is a boundary reduced tiled surface then as $n \rightarrow \infty$,*

$$\Xi_n(Y) = O_Y(1).$$

Proof. By Proposition 5.23, there is some $b = b(Y)$ such that

$$\Xi_n(Y) = \Xi_n^{(b)}(Y) + O_Y(n^{-1}).$$

As before, let $b_\lambda \stackrel{\text{def}}{=} |\lambda| - \lambda_1$. As in the proof of Proposition 5.24, for $n - \mathbf{v} > 2b$ we can write

$$\begin{aligned} \Xi_n^{(b)}(Y) = 2 \sum_{\substack{\nu \vdash n - \mathbf{v} : b_\nu < b \\ \nu \subset_{\mathbf{v}-\mathbf{e}_f} \mu_f \subset_{\mathbf{e}_f-\mathbf{f}} \lambda}} \frac{d_\lambda d_\nu}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \Upsilon\left(\left\{\sigma_f^\pm, \tau_f^\pm\right\}, \nu, \{\mu_f\}, \lambda\right). \end{aligned} \tag{5.31}$$

Note that in equation (5.31), each of

$$\lambda/\nu, \mu_a/\nu, \mu_b/\nu, \mu_c/\nu, \mu_d/\nu$$

has $\leq \mathbf{v} - \mathbf{f}$ boxes outside their first row, so Lemma 4.3 implies that

$$\frac{d_\lambda d_\nu}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \ll_Y \frac{1}{d_\nu^2} n^{b_\lambda - b_{\mu_a} - b_{\mu_b} - b_{\mu_c} - b_{\mu_d} + 3b_\nu} = \frac{1}{d_\nu^2} n^{b_{\lambda/\nu} - b_{\mu_a/\nu} - b_{\mu_b/\nu} - b_{\mu_c/\nu} - b_{\mu_d/\nu}}, \tag{5.32}$$

(recall the notation \ll from §§1.7). Since all λ, ν in the sum (5.31) have a bounded number of boxes outside their first row, depending on Y , for sufficiently large n , the condition $\lambda_1 + \nu_1 > n - \mathbf{f} + (\mathbf{v} - \mathbf{f})^2$ of Proposition 5.21 is met for large n . Hence, by Proposition 5.21 Part 1,

$$\begin{aligned} \left| \Xi_n^{(b)}(Y) \right| &\ll_Y \sum_{\nu \vdash n - \mathbf{v} : b_\nu < b} \frac{1}{d_\nu^2} \left[\sum_{\nu \subset_{\mathbf{v}-\mathbf{e}_f} \mu_f \subset_{\mathbf{e}_f-\mathbf{f}} \lambda} \left(n \cdot \frac{(\mathbf{v} - \mathbf{f})^2}{\lambda_1 + \nu_1 - (n - \mathbf{f})} \right)^{b_{\lambda/\nu} - b_{\mu_a/\nu} - b_{\mu_b/\nu} - b_{\mu_c/\nu} - b_{\mu_d/\nu}} \right] \\ &\ll_Y \sum_{\nu \vdash n - \mathbf{v} : b_\nu < b} \frac{1}{d_\nu^2} \left[\sum_{\nu \subset_{\mathbf{v}-\mathbf{e}_f} \mu_f \subset_{\mathbf{e}_f-\mathbf{f}} \lambda} 1 \right] \ll_Y \sum_{\nu \vdash n - \mathbf{v} : b_\nu < b} \frac{1}{d_\nu^2} \leq \zeta^{S_{n-\mathbf{v}}}(2) = 2 + O\left(\frac{1}{n^2}\right), \end{aligned}$$

where the second asymptotic inequality follows as $n \cdot \frac{(\mathbf{v}-\mathbf{f})^2}{\lambda_1 + \nu_1 - (n-\mathbf{f})} \ll 1$ uniformly over all ν, λ in play and $b_{\lambda/\nu} - b_{\mu_a/\nu} - b_{\mu_b/\nu} - b_{\mu_c/\nu} - b_{\mu_d/\nu}$ is bounded from above by a constant, the third asymptotic inequality follows as the number of YDs that extend a given ν by at most \mathbf{v} boxes is bounded independently of n , and the last asymptotic inequality follows by Proposition 4.5 with $b = 1$. \square

If Y is strongly boundary reduced, then we get a finer estimate.

Proposition 5.26. *If Y is a strongly boundary reduced tiled surface, then as $n \rightarrow \infty$,*

$$\Xi_n(Y) = 2 + O_Y(n^{-1}). \tag{5.33}$$

Proof. We begin as in the proof of Proposition 5.25 by choosing $b(Y)$ such that $\Xi_n(Y) = \Xi_n^{(b)}(Y) + O_Y(n^{-1})$ and equation (5.31) holds. It now suffices to prove the proposition with $\Xi_n(Y)$ replaced with $\Xi_n^{(b)}(Y)$.

There are summands of equation (5.31) corresponding to λ/ν having all boxes in the first row. By Proposition 5.21 Part 3, each of these summands contributes $2 \cdot \frac{d_\lambda d_\nu}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}}$ to $\Xi_n^{(b)}$, but in this case ν, μ_f, λ all belong to the same family of YDs, so this contribution is $\frac{2}{d_\nu^2} (1 + O(n^{-1}))$. As there is one of these summands for each $\nu \vdash n - \mathbf{v}$ with $b_\nu < b$, together these contribute

$$2 \left(1 + O \left(\frac{1}{n} \right) \right) \sum_{v+n-v : b_v < b} \frac{1}{d_v^2} = 2 + O \left(\frac{1}{n} \right) \tag{5.34}$$

by Lemma 4.7. The constant term 2 appearing in equation (5.34) is the main term of equation (5.33).

For any other summand of equation (5.31) $b_{\lambda/v} > 0$ and so by Proposition 5.21 Part 2 combined with equation (5.32),

$$\frac{d_\lambda d_v}{d_{\mu_a} d_{\mu_b} d_{\mu_c} d_{\mu_d}} \Upsilon \left(\left\{ \sigma_f^\pm, \tau_f^\pm \right\}, \nu, \{ \mu_f \}, \lambda \right) \ll_Y \frac{1}{d_v^2} \cdot \frac{1}{n},$$

so as argued in the proof of Proposition 5.25, the total contribution of these summands is $O_Y(n^{-1})$. Hence

$$\Xi_n^{(b)}(Y) = 2 + O_Y \left(\frac{1}{n} \right). \quad \square$$

6. Proofs of main theorems

6.1. Proof of Theorem 1.1 and its extension to finitely generated subgroups

We give the proof when $g = 2$; the extension to other $g \geq 2$ is clear. We are given a finitely generated subgroup $J \leq \Gamma = \Gamma_2$ and $M \in \mathbf{N}$ and we wish to show that $\mathbb{E}_n[\text{fix}_J] \stackrel{\text{def}}{=} \mathbb{E}_{2,n}[\text{fix}_J]$ can be approximated to order $O(n^{-M})$ by a Laurent polynomial of n with rational coefficients. Given this, the trivial bound $\mathbb{E}_n[\text{fix}_J] \leq n$ implies that the Laurent polynomial takes the form (1.2). The fact that the $a_i(J)$ do not depend on M is clear. By setting $J = \langle \gamma \rangle$, we obtain Theorem 1.1 from this.

By Lemma 2.7, we have

$$\mathbb{E}_n[\text{fix}_J] = \mathbb{E}_n(\text{Core}(J)). \tag{6.1}$$

Now, let \mathcal{R} be any finite resolution of $\text{Core}(J)$; by Theorem 2.14 at least one exists.

Each element of this resolution is a morphism $h : \text{Core}(J) \rightarrow W_h$ of tiled surfaces. By Lemma 2.9,

$$\mathbb{E}_n(\text{Core}(J)) = \sum_{h \in \mathcal{R}} \mathbb{E}_n^{\text{emb}}(W_h).$$

Now, using Theorem 5.10 for each of the terms $\mathbb{E}_n^{\text{emb}}(W_h)$ gives

$$\mathbb{E}_n[\text{fix}_J] = \frac{(n!)^3}{|\mathbb{X}_n|} \sum_{h \in \mathcal{R}} \frac{\binom{n}{\mathbf{v}(W_h)} \binom{n}{\mathbf{f}(W_h)}}{\prod_{f \in \{a,b,c,d\}} \binom{n}{e_f(W_h)}} \mathbb{E}_n(W_h),$$

where $\mathbf{v}(W_h)$, $e_f(W_h)$ and $\mathbf{f}(W_h)$ are the number of vertices, f -labeled edges ($f \in \{a, b, c, d\}$) and octagons, respectively, of W_h . Also, recall the definition of \mathbb{E}_n from Theorem 5.10. By equation (1.3), $|\mathbb{X}_n| = (n!)^3 \cdot \zeta^{S_n}(2)$, and so

$$\mathbb{E}_n[\text{fix}_J] = \frac{1}{\zeta^{S_n}(2)} \sum_{h \in \mathcal{R}} \frac{\binom{n}{\mathbf{v}(W_h)} \binom{n}{\mathbf{f}(W_h)}}{\prod_{f \in \{a,b,c,d\}} \binom{n}{e_f(W_h)}} \mathbb{E}_n(W_h). \tag{6.2}$$

Now, we note:

- By Corollary 4.9, there is a polynomial $Q_{2,M} \in \mathbf{Z}[t]$ with $\frac{1}{\zeta^{S_n}(2)} = \frac{1}{2} Q_{2,M}(n^{-1}) + O(n^{-M})$.
- For any fixed $\ell \geq 0$, both $\binom{n}{\ell}$ and $\binom{n}{\ell}^{-1}$ agree with Laurent polynomials of n with \mathbf{Q} -coefficients up to order $O(n^{-M})$ as $n \rightarrow \infty$.
- By Proposition 5.24, there is a Laurent polynomial $\Xi_M^*(W_h) \in \mathbf{Q}[t, t^{-1}]$ such that $\mathbb{E}_n(W_h) = \Xi_M^*(W_h)[n] + O(n^{-M})$ as $n \rightarrow \infty$.

Hence, all terms in equation (6.2) can be approximated by Laurent polynomials of n with rational coefficients to order $O(n^{-M})$ as $n \rightarrow \infty$. This proves Theorem 1.1. \square

6.2. Proof of Theorems 1.2 and 1.3

Again, we give the proofs of Theorems 1.2 and 1.3 when $g = 2$. Given a finitely generated subgroup $J \leq \Gamma = \Gamma_2$, let $\chi_{\max}(J)$ be as in equation (1.4). We can assume J is nontrivial since Theorem 1.3 is obvious in this case (as is Theorem 1.2 when γ is the identity).

Let $\mathcal{R} = \mathcal{R}(\text{Core}(J), \chi_{\max}(J))$ be the resolution of $\text{Core}(J)$ defined in Definition 2.13, certified to be a resolution by Theorem 2.14. Let $\mathcal{R}_{\max}(\gamma)$ (resp. $\mathcal{R}_{<\max}(\gamma)$) be the morphisms $h : \text{Core}(J) \rightarrow W_h$ of \mathcal{R} with $\chi(W_h) = \chi_{\max}(J)$ (resp. $\chi(W_h) < \chi_{\max}(J)$), so

$$\mathcal{R} = \mathcal{R}_{\max} \sqcup \mathcal{R}_{<\max}.$$

By Theorem 2.14, all elements of \mathcal{R}_{\max} are strongly boundary reduced and all elements of \mathcal{R} are boundary reduced. Repeating the argument of §§6.1 we obtain equation (6.2) again.

Now, we observe:

- By Proposition 4.5 with $b = 1$, $\zeta^{S_n}(2) = 2 + O(n^{-2})$ as $n \rightarrow \infty$.
- For each $h : \text{Core}(J) \rightarrow W_h \in \mathcal{R}$, the ratio of Pochhammer symbols satisfies as $n \rightarrow \infty$

$$\frac{(n)_{\mathfrak{v}(W_h)}(n)_{\mathfrak{f}(W_h)}}{\prod_{f \in \{a,b,c,d\}}(n)_{\mathfrak{e}_f(W_h)}} = n^{\chi(W_h)} + O\left(n^{\chi(W_h)-1}\right). \tag{6.3}$$

- By Proposition 5.26, for each $h : \text{Core}(J) \rightarrow W_h \in \mathcal{R}_{\max}$ we have $\Xi_n(W_h) = 2 + O_{W_h}(n^{-1})$.
- By Proposition 5.25, for each $h : \text{Core}(J) \rightarrow W_h \in \mathcal{R}_{<\max}$ we have $\Xi_n(W_h) = O_{W_h}(1)$.

Hence, from equation (6.2),

$$\begin{aligned} \mathbb{E}_n[\text{fix}_J] &= \frac{1}{\zeta^{S_n}(2)} \left[\sum_{h \in \mathcal{R}_{\max}} n^{\chi_{\max}(J)} \left(1 + O(n^{-1})\right) \cdot \left(2 + O(n^{-1})\right) + \sum_{h \in \mathcal{R}_{<\max}} O\left(n^{\chi_{\max}(J)-1}\right) \cdot O(1) \right] \\ &= \frac{1}{2 + O(n^{-2})} \left[2 \cdot |\mathcal{R}_{\max}| \cdot n^{\chi_{\max}(J)} + O\left(n^{\chi_{\max}(J)-1}\right) \right] \\ &= |\mathcal{R}_{\max}| \cdot n^{\chi_{\max}(J)} + O\left(n^{\chi_{\max}(J)-1}\right), \end{aligned}$$

where all implied constants depend on J . By Proposition 2.15, $|\mathcal{R}_{\max}| = |\mathfrak{M}\mathfrak{D}\mathfrak{G}(J)|$ which proves Theorem 1.3.

Finally, if $J = \langle \gamma \rangle$, and q is maximal such that $\gamma = \gamma_0^q$ for some $\gamma_0 \in \Gamma$, then Corollary 2.17 tells us that $|\mathcal{R}_{\max}| = d(q)$. This proves Theorem 1.2. \square

Acknowledgements. We thank Nir Avni, Frédéric Naud, Mark Powell and Henry Wilton for helpful discussions related to this work. The research was supported by the Israel Science Foundation: ISF grant 1071/16 of D.P. This project 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).

Competing interest. The authors have no competing interest to declare.

References

[AB83] M. F. Atiyah and R. Bott, ‘The Yang–Mills equations over Riemann surfaces’, *Philos. Trans. Roy. Soc. London Ser. A* **308**(1505) (1983), 523–615.
 [ABB+11] M. Abert, N. Bergeron, I. Biringer, T. Gelander, N. Nikolov, J. Raimbault and I. Samet, ‘On the growth of Betti numbers of locally symmetric spaces’, *C. R. Math. Acad. Sci. Paris* **349**(15–16) (2011), 831–835.

- [ABB+17] M. Abert, N. Bergeron, I. Biringer, T. Gelander, N. Nikolov, J. Raimbault and I. Samet, ‘On the growth of L^2 -invariants for sequences of lattices in Lie groups’, *Ann. of Math. (2)* **185**(3) (2017), 711–790.
- [Bau62] G. Baumslag, ‘On generalised free products’, *Math. Zeit.* **78**(1) (1962), 423–438.
- [BP20] E. Baker and B. Petri, ‘Statistics of finite degree covers of torus knot complements’, Preprint, 2020, [arXiv:2005.11956](https://arxiv.org/abs/2005.11956).
- [BR10] K. Bou-Rabee, ‘Quantifying residual finiteness’, *J. Algebra* **323**(3) (2020), 729–737.
- [BS87] 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)*, Ann. of Math. Stud., vol. 111 (Princeton Univ. Press, Princeton, NJ, 1987), 451–478.
- [BS01] I. Benjamini and O. Schramm, ‘Recurrence of distributional limits of finite planar graphs’, *Electron. J. Probab.* **6**(23) (2001), 13.
- [Bus10] P. Buser, *Geometry and Spectra of Compact Riemann Surfaces*, Modern Birkhäuser Classics, (Birkhäuser Boston, Ltd., Boston, MA, 2010). Reprint of the 1992 edition.
- [CS06] B. Collins and P. Śniady, ‘Integration with respect to the Haar measure on unitary, orthogonal and symplectic group’, *Comm. Math. Phys.* **264**(3) (2006), 773–795.
- [CSST10] T. Ceccherini-Silberstein, F. Scarabotti and F. Tolli, *Representation Theory of the Symmetric Troups: The Okounkov–Vershik Approach, Character Formulas, and Partition Algebras* (Cambridge University Press, Cambridge, 2010).
- [Deh12] M. Dehn, ‘Transformation der Kurven auf zweiseitigen Flächen’, *Math. Ann.* **72**(3) (1912), 413–421.
- [Dix69] J. D. Dixon, ‘The probability of generating the symmetric group’, *Math. Z.* **110**(3) (1969), 199–205.
- [EWPS21] D. Ernst-West, D. Puder and M. Seidel, ‘Word measures on $GL_n(q)$ and free group algebras’, Preprint, 2021, [arXiv:2110.11099](https://arxiv.org/abs/2110.11099).
- [FRT54] J. S. Frame, G. de B. Robinson and R. M. Thrall, ‘The hook graphs of the symmetric groups’, *Canad. J. Math.* **6** (1954), 316–324.
- [Gam06] A. Gamburd, ‘Poisson–Dirichlet distribution for random Belyi surfaces’, *Ann. Probab.* **34**(5) (2006), 1827–1848.
- [Gol84] W. M. Goldman, ‘The symplectic nature of fundamental groups of surfaces’, *Adv. in Math.* **54**(2) (1984), 200–225.
- [Hat05] A. Hatcher, *Algebraic Topology* (Cambridge University Press, Cambridge, 2005).
- [Hem72] J. Hempel, ‘Residual finiteness of surface groups’, *Proc. Amer. Math. Soc.* **32** (1972), 323.
- [HP22] L. Hanany and D. Puder, ‘Word measures on symmetric groups’, *Int. Math. Res. Not. IMRN* (2022). doi: [10.1093/imrn/rnac084](https://doi.org/10.1093/imrn/rnac084).
- [Hur02] A. Hurwitz, ‘Ueber die anzahl der Riemann’schen flächen mit gegebenen verzweigungspunkten’, *Math. Ann.* **55**(1) (1902), 53–66.
- [Lab13] F. Labourie, *Lectures on Representations of Surface Groups*, Zurich Lectures in Advanced Mathematics (Eur. Math. Soc. (EMS), Zürich, 2013).
- [LLM23] N. Lazarovich, A. Levit and Y. Minsky, ‘Surface groups are flexibly stable’, *J. Eur. Math. Soc. (JEMS)* 2023. To appear. [arXiv:1901.07182](https://arxiv.org/abs/1901.07182).
- [LS04] M. W. Liebeck and A. Shalev, ‘Fuchsian groups, coverings of Riemann surfaces, subgroup growth, random quotients and random walks’, *J. Algebra* **276**(2) (2004), 552–601.
- [Lul96] N. Lulov, ‘Random walks on symmetric groups generated by conjugacy classes’, PhD thesis, Harvard University, 1996.
- [Mag21] M. Magee, ‘Random unitary representations of surface groups II: The large n limit’, *Geom. Topol.* To appear, 2023, [arXiv:2101.03224](https://arxiv.org/abs/2101.03224).
- [Mag22] M. Magee, ‘Random unitary representations of surface groups I: Asymptotic expansions’, *Comm. Math. Phys.* **391**(1) (2022), 119–171.
- [Med78] A. D. Mednykh, ‘Determination of the number of nonequivalent coverings over a compact Riemann surface’, *Dokl. Akad. Nauk SSSR* **239**(2) (1978), 269–271.
- [Mir07] M. Mirzakhani, ‘Simple geodesics and Weil–Petersson volumes of moduli spaces of bordered Riemann surfaces’, *Invent. Math.* **167**(1) (2007), 179–222.
- [MNP22] M. Magee, F. Naud and D. Puder, ‘A random cover of a compact hyperbolic surface has relative spectral gap $\frac{3}{16} - \varepsilon$ ’, *Geom. Funct. Anal. (GAFA)* **32**(3) (2022), 595–661.
- [MP02] T. W. Müller and J.C. Puchta, ‘Character theory of symmetric groups and subgroup growth of surface groups’, *J. Lond. Math. Soc.* **66**(3) (2002), 623–640.
- [MP19] M. Magee and D. Puder, ‘Matrix group integrals, surfaces, and mapping class groups I: $\mathcal{U}(n)$ ’, *Invent. Math.* **218**(2) (2019), 341–411.
- [MP21] M. Magee and D. Puder, ‘Surface words are determined by word measures on groups’, *Israel Journal of Mathematics* **241** (2021), 749–774.
- [MP22a] M. Magee and D. Puder, ‘Core surfaces’, *Geom. Dedicata* **216**(46) (2022), 1–25.
- [MP22b] M. Magee and D. Puder, ‘Matrix group integrals, surfaces, and mapping class groups II: $O(n)$ and $Sp(n)$ ’, *Math. Ann.* (2022). doi: [10.1007/s00208-022-02542-1](https://doi.org/10.1007/s00208-022-02542-1).
- [Nic94] A. Nica, ‘On the number of cycles of given length of a free word in several random permutations’, *Random Structures Algorithms* **5**(5) (1994), 703–730.
- [PP15] D. Puder and O. Parzanchevski, ‘Measure preserving words are primitive’, *J. Amer. Math. Soc.* **28**(1) (2015), 63–97.

- [PZ22] D. Puder and T. Zimhoni, 'Local statistics of random permutations from free products', Preprint 2022, [arXiv:2203.12250](https://arxiv.org/abs/2203.12250).
- [Sco78] P. Scott, 'Subgroups of surface groups are almost geometric', *J. Lond. Math. Soc.* **2**(3) (1978), 555–565.
- [Sta83] J. R. Stallings, 'Topology of finite graphs', *Invent. Math.* **71**(3) (1983), 551–565.
- [Tom67] M. Tomita, 'On canonical forms of von Neumann algebras', in *Fifth Functional Analysis Sympos. (Tôhoku Univ., Sendai, 1967)* (Math. Inst., Tôhoku Univ., Sendai, Japan, 1967), 101–102. Japanese.
- [VDN92] D. V. Voiculescu, K. J. Dykema and A. Nica, *Free Random Variables*, CRM Monograph Series, vol. 1 (American Mathematical Society, Providence, RI, 1992).
- [VO96] A. Vershik and A. Okounkov, 'A new approach to representation theory of symmetric groups', *Selecta Math. (N.S.)* **2**(4) (1996), 581–606.
- [Wit91] E. Witten, 'On quantum gauge theories in two dimensions', *Comm. Math. Phys.* **141**(1) (1991), 153–209.
- [Wri20] A. Wright, 'A tour through Mirzakhani's work on moduli spaces of Riemann surfaces', *Bull. Amer. Math. Soc. (N.S.)* **57**(3) (2020), 359–408.
- [Zag94] D. Zagier, 'Values of zeta functions and their applications', in *First European Congress of Mathematics, Vol. II (Paris, 1992)*, Progr. Math., vol. 120 (Birkhäuser, Basel, 1994), 497–512.