PRYM–BRILL–NOETHER LOCI OF SPECIAL CURVES

STEVEN CREECH, YOAV LEN, CAELAN RITTER, AND DEREK WU

ABSTRACT. We use Young tableaux to compute the dimension of V^r, the Prym–Brill–Noether locus of a folded chain of loops of any gonality. This tropical result yields a new upper bound on the dimensions of algebraic Prym–Brill–Noether loci. Moreover, we prove that V^r is pure-dimensional and connected in codimension 1 when $\dim V^r \geq 1$. We then compute the first Betti number of this locus for even gonality when the dimension is exactly 1, and compute the cardinality when the locus is finite and the edge lengths are generic.

CONTENTS

1. INTRODUCTION

Constructing algebraic cycles in abelian varieties can in general be a challenging problem. The problem becomes more tractable if we restrict our attention to varieties that show up in the context of algebraic curves such as Jacobians or Prym varieties. In such cases, algebraic cycles may naturally be constructed by appealing to Brill–Noether theory, and taking advantage of the geometric interpretation of the points of the abelian varieties.

Let $f\colon \widetilde{X} \to X$ be an unramified double cover of either tropical or algebraic curves, and let f_* be the induced map on divisor classes. The corresponding *Prym–Brill–Noether locus* is

$$V^r(X,f) = \{\, [D] \in \mathrm{Jac}(\widetilde{X}) \mid f_*(D) = K_X,\ r(D) \geq r,\ r(D) \equiv r \pmod 2 \,\},$$

where K_X is the canonical divisor of X. It is a variation of the usual Brill–Noether locus $W^d_r(\widetilde{X})$ that also takes symmetries of $\widetilde{X}$ into account. The Prym–Brill–Noether locus naturally lives inside the Prym variety associated with f (see Section 2 for more details). Moreover, since the locus may be described as intersections of translates of the theta class, it is, in particular, tautological [Ara12, Theorem 1.2]. This paper is concerned with the dimension and additional topological properties of $V^r(X,f)$.

Properties of the usual Brill–Noether loci have been studied extensively for curves that are general in moduli in classical algebraic geometry [GH80, Gie82, FL81] and more recently in tropical geometry [CDPR12, JP14, Len14]. When a curve is *not* general in moduli, its Brill–Noether locus is no longer expected to be irreducible or pure-dimensional. Nevertheless, the dimensions of irreducible components of these loci have recently been computed for general k-gonal curves, namely general among curves that admit a k-fold cover of $\mathbb{P}^1$ [CPJ19, JR17, Lar19].

In contrast, much less is known for Prym varieties. Bertram and Welters computed the dimension of the Prym–Brill–Noether locus for curves that are general in moduli [Ber87, Wel85], and Welters has also shown that the locus is generically smooth. The tropical study of Prym varieties was initially introduced in joint work of the second author with Jensen [JL18], and further studied in joint work with Ulirsch [LU19]. As they show, tropical Pryms are abelian of the expected dimension and behave well with respect to tropicalization, leading to a new bound on the dimension of Prym–Brill–Noether loci of general even-gonal algebraic curves.

Our first result is an extension of these techniques to curves of *any* gonality.

Theorem A. *Let $\varphi\colon \widetilde{\Gamma} \to \Gamma$ be a k-gonal uniform folded chain of loops and denote by l the quantity $\lceil \frac{k}{2} \rceil$. Then the codimension of $V^r(\Gamma, \varphi)$ relative to the Prym variety is given by*

$$n(r,k) = \begin{cases} \binom{l+1}{2} + l(r-l) & \text{if } l \le r-1 \\ \binom{r+1}{2} & \text{if } l > r-1 \end{cases}. \tag{1.1}$$

By *uniform* k-gonal we mean that the ratio of the lengths of the upper and lower arcs of each loop is exactly k; see Section 2 for more details. We adopt the convention that a set whose dimension is negative is empty, so $V^r(\Gamma, \varphi)$ is empty if $n(r,k) > g-1$. As it turns out, the odd gonality case is far trickier than the even, necessitating the development of several new combinatorial tools.

As a consequence of the theorem, we obtain an upper bound on the dimensions of Prym–Brill–Noether loci for algebraic curves that are general in the k-gonal locus. In what follows, we work over a non-Archimedean field K with residue field κ whose characteristic is prime to both 2 and k.

Corollary B. *Let $r \ge -1$ and $k \ge 2$. Then there is a nonempty open subset of the k-gonal locus of $\mathcal{R}_g$ such that for every unramified double cover $f\colon \widetilde{C} \to C$ in this open subset we have*

$$\dim V^r(C, f) \le g - 1 - n(r,k). \tag{1.2}$$

We then turn our attention to more subtle *tropological* properties of Prym–Brill–Noether loci of folded chains of loops.[1]

Theorem C. *$V^r(\Gamma, \varphi)$ is pure-dimensional for any gonality k. If $\dim V^r(\Gamma, \varphi) \ge 1$ then it is also connected in codimension 1.*

By "connected in codimension 1," we mean that any two maximal cells are connected by a sequence of cells whose codimension relative to the locus is at most 1. The different properties mentioned in the theorem are proved in Propositions 4.8 and 4.9. The pure-dimensionality of the locus is quite surprising since Brill–Noether loci of general k-gonal curves may very well have maximal components of different dimension (see for instance [JR17, Section 1]). We do not know at this point whether this phenomenon is special to tropical Prym curves or carries on to algebraic ones as well.

If we choose r and k so that $n(r,k) = g-1$, the Prym–Brill–Noether locus is a finite collection of points. If k is also assumed to be even, we may compute the cardinality by constructing a bijection between its points and certain lattice paths (Proposition 5.1). If the dimension is 1, the tropical Prym–Brill–Noether locus is a graph; we compute its first Betti number in the case of generic edge lengths.

Theorem D. *Let $\varphi\colon \widetilde{\Gamma} \to \Gamma$ be a folded chain of loops with generic edge length such that $\dim V^r(\Gamma, \varphi) = 1$. Then the first Betti number of $V^r(\Gamma, \varphi)$ is given by*

$$\frac{r \cdot C(r,0) \cdot \left(\binom{r+1}{2} + 1\right)}{2} + 1. \tag{1.3}$$

Here, $C(r,0)$ is the number of distinct ways to fill a staircase tableau of size r such that each symbol in the set $[\binom{r+1}{2}]$ is used exactly once.[2] Moreover, we calculate the first Betti number in the cases where k is 2 or 4 (Propositions 5.4 and 5.5).

[1] We introduce the descriptor "tropological" to mean "topological" in the context of tropical varieties.

[2] We define $C(r,k)$ more generally in Section 5.1 once we have more tools at our disposal.

Many of our results build on the correspondence between certain Young tableaux and divisors on tropical curves (cf. [CDPR12, Pfl17b]). The key tool that we develop to enumerate such tableaux is the notion of a *non-repeating strip*, a special subset that determines the rest of the tableau (see Section 4.1). We hope that this and other techniques presented in our paper will lead to additional results concerning dimensions and Euler characteristics of tropical and algebraic Brill–Noether loci.

There are numerous interesting avenues for investigation moving forward. The techniques developed in [JR17] for lifting special divisors should be adapted to the current situation to determine the precise dimension of algebraic Prym–Brill–Noether loci (see Conjecture 3.9 for more details). If every maximal cell of $V^r(\Gamma, \varphi)$ can be lifted, Proposition 4.8 would moreover imply that algebraic Prym–Brill–Noether loci are pure-dimensional. It would be intriguing to extend the enumerative results of Section 5 to any gonality, and discover whether an algebraic version holds as well. Note, however, that current degeneration techniques do not immediately imply either an upper or a lower bound on the Betti numbers of algebraic Prym–Brill–Noether curves. Finally, it would be exciting to extend our techniques to ramified double covers and general Galois covers. The latter would be especially challenging since the components of the kernel of such covers do not naturally admit a principal polarization.

Acknowledgements. We thank Dave Jensen for helpful remarks on a previous version of this manuscript. We also thank the nameless referees for their insightful comments and suggestions. This research was conducted at the Georgia Institute of Technology with the support of RTG grant GR10004614 and REU grant GR10004803.

2. Preliminaries

Throughout this paper, we use the terms metric graph and tropical curve interchangeably. We assume that the reader is familiar with the theory of divisors on tropical curves; a beautiful introduction to this topic may be found in [HMY12, Section 2]. Throughout, the *genus* of a graph refers to its first Betti number, which also equals one more than the number of edges minus the number of vertices.

The result of tropicalizing a covering map of algebraic curves is a *harmonic* morphism of metric graphs. Such morphisms induce natural pushforward and pullback maps between divisors that respect the equivalence relation given by chip-firing. A map of graphs is called a *double cover* when it is harmonic of degree 2, and *unramified* if, in addition, it pulls back the canonical divisor of Γ to the canonical divisor of $\widetilde{\Gamma}$. See [LUZ19, Definition 2.7] for precise definitions of harmonic morphisms and their degree.

Fix a divisor class $[D]$ on Γ. The fiber $\varphi_*^{-1}([D])$ consists of either one or two connected components in the Picard group of $\widetilde{\Gamma}$ [JL18, Proposition 6.1]. Each of them is referred to as a *Prym variety*, and their elements are called *Prym divisor classes*. Prym varieties are principally polarized tropical abelian varieties [LU19, Theorem 2.3.7]. We take the divisor D above to be the canonical divisor K_Γ. Fixing an integer r, the *Prym–Brill–Noether locus* $V^r(\Gamma, \varphi)$ consists of the Prym divisors whose rank is at least r and has the same parity as r.

Here we are interested in a particular double cover known as the *folded chain of loops*. In this case, the target Γ of the map φ is the *chain of loops* that recently appeared in various celebrated papers (e.g. [JP16, Pfl17a, JR17]). It consists of g loops, denoted by $\gamma_1, \ldots, \gamma_g$ and connected by bridges. The source graph $\widetilde{\Gamma}$ is a chain of $2g-1$ loops, as exemplified in Fig. 2.1. Each pair of loops $\tilde{\gamma}_a$ and $\tilde{\gamma}_{2g-a}$ (for $a < g$) maps down to γ_a, while each edge of $\tilde{\gamma}_g$ maps isometrically onto the loop γ_g. See [LU19, Section 5.2] for a more detailed explanation.

The *torsion* of a loop γ_a is the least positive integer k such that $\ell_a + m_a$ divides $k \cdot m_a$, where m_a and ℓ_a are the lengths of the lower and upper arcs of γ_a respectively. The chain of loops is *uniform k-gonal* if each loop has torsion k. Note that a uniform k-gonal chain of loops is indeed a k-gonal metric graph

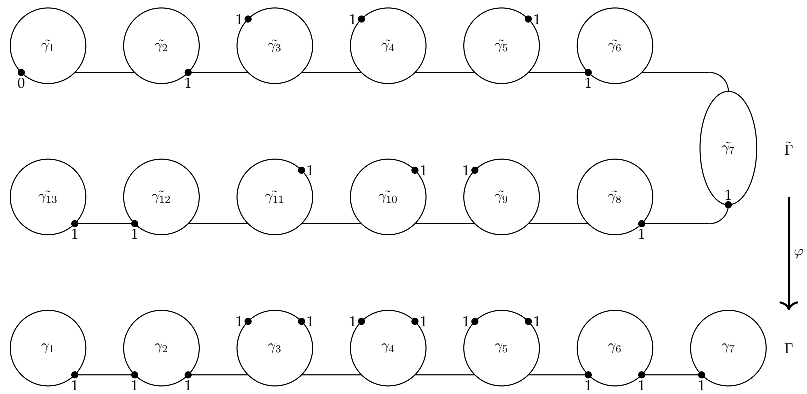

Figure 2.1. A Prym divisor on the 4-gonal folded chain of 13 loops and its image under φ_* on the 4-gonal chain of 7 loops.

in the sense of [ABBR15, Section 1.3.2]. We say that a double cover as above is uniform k-gonal if Γ is, but note that $\widetilde{\Gamma}$ is not in itself uniform k-gonal since the loop $\tilde{\gamma}_g$ has torsion 2.

2.1. **Prym tableaux.** We study divisors only indirectly, making use of a correspondence between sets of divisors on chains of loops and Young tableaux as introduced in [Pfl17b, LU19]; here we shall recall only the essential definitions and introduce some helpful notation.

Let $[n] = \{1, 2, \ldots, n\}$. Given points $(x, y), (x', y') \in \mathbb{N}^2$, which we call *boxes*, we say that (x, y) is *below* (x', y') if $x \le x'$, $y \le y'$, and $(x, y) \ne (x', y')$. For our purposes, a *tableau* on a subset $\lambda \subset \mathbb{N}^2$ is a map $t\colon \lambda \to [n]$ satisfying the *tableau condition*:

for all boxes (x, y) and (x', y') in λ, if (x, y) is below (x', y'), then $t(x, y) < t(x', y')$.

We refer to the elements in the codomain of t as *symbols*. Observe that if λ is a partition of n and t is injective, then t is a standard Young tableau in the usual sense.

A tableau t is a *(k-uniform) displacement tableau* if it also satisfies the *displacement condition*:

whenever $t(x, y) = t(x', y')$, we have that $x - y \equiv x' - y' \pmod{k}$.[3]

This condition partitions λ into k regions. To be precise, we define the *i-th diagonal modulo k*, denoted by $D_{i,k}$, to be the set of boxes $\{(x, y) \in \lambda \mid x - y \equiv i \pmod{k}\}$; then λ is the disjoint union of $D_{i,k}$ for $i \in \mathbb{Z}/k\mathbb{Z}$, and the fiber of each symbol of t is contained in some $D_{i,k}$. Co-opting earlier terminology, we also call k the *torsion* of t.

The n-th *anti-diagonal* A_n is the set of all boxes (x, y) such that $x + y = n + 1$. Define the *lower triangle of size n* to be $T_n := \bigcup_{i=1}^{n} A_i$. For example, Fig. 2.2 shows a lower-triangular displacement tableau of size 6 and torsion 3.[4] $D_{1,3}$ is blue, A_6 is red, and their intersection is purple. Every box (x, y) here not colored red or purple is below (some box of) A_6.

As explained in [Pfl17b, Section 3], k-uniform displacement tableaux on the rectangle $[g-d+r] \times [r+1]$ with codomain $[g]$ give rise to divisors of degree d and rank at least r on the uniform k-gonal chain of g loops, as we now recall. The location (x, y) of the symbol $a \in [g]$ in the tableau indicates where to place a chip on the a-th loop. Whenever a does not appear in the tableau, the chip may be placed arbitrarily on that loop, thereby allowing the locus a single degree of freedom. Otherwise, the a-th loop will have a chip at distance $m_a \cdot (x - y)$ counter-clockwise from its rightmost vertex (where the loops are arranged from left to right, as in the bottom of Fig. 2.1). Finally place $d - g$ chips at the rightmost

[3]Or equivalently, if (x, y) and (x', y') contain the same symbol, then (x, y) and (x', y') must be separated by a lattice distance that is a multiple of k.

[4]We adopt the French notation, where the bottom-left box is $(1, 1)$, the first coordinate increases to the right, and the second coordinate increases upwards.

11					
9	10				
7	8	9			
5	6	7	8		
3	4	5	6	7	
1	2	3	4	5	6

Figure 2.2. A typical example of a lower-triangular displacement tableau of size 6 with torsion 3.

vertex of the g-th loop. The displacement condition guarantees that this is well-defined when a symbol appears in the tableau more than once.

The tableau–divisor correspondence naturally extends to the folded chain of loops, although the chips on the lower loops of $\widetilde{\Gamma}$ (as depicted in Fig. 2.1) are measured *clockwise* from the *leftmost* vertex, and the stack of $d-g$ chips are placed at the *leftmost* vertex of the $(2g-1)$-th loop. As the genus of the folded chain $\widetilde{\Gamma}$ is $2g-1$, the symbols should be taken from $[2g-1]$ and the domain should have shape $[2g-1-d+r]\times[r+1]$. Since the parity of the g-th loop is different than the rest, the fiber of g must be contained in $D_{i,2}$ for some i. By a slight abuse of terminology, we shall still refer to such tableaux as "k-uniform."

We wish to produce Prym divisors; these map down to K_Γ and so must have degree $2g-2$. Hence, any tableau that yields Prym divisors under the correspondence must be defined on the square domain $[r+1]\times[r+1]$. Moreover, the counter-clockwise distance of the chip on the a-th loop (for $a\in[g-1]$) must equal the clockwise distance of the chip on the $(2g-a)$-th loop. This motivates the following *Prym condition*:

$$t(x,y)=2g-t(x',y') \text{ only if } (x,y) \text{ and } (x',y') \text{ both lie in the same diagonal modulo } k.$$

Definition 2.1. A tableau t is *Prym of type* (g,r,k) if it has shape $[r+1]\times[r+1]$ and codomain $[2g-1]$, it is k-uniform (see above), and it satisfies the Prym condition.

The two tropical Prym varieties arising from a folded chain of loops are distinguished by the parity of the rank of the divisors that they classify [LU19, Theorem 5.3.8]. The parity, in turn, is determined by the placement of the chip on the g-th loop, or equivalently, the position of the symbol g in the tableau. We denote by $P(t)$ the set of Prym divisors obtained from t via the tableau–divisor correspondence whose rank coincides with r modulo 2. Explicitly, if $t^{-1}(g)$ is contained in $D_{r,2}$, then $P(t)$ coincides with the set of divisors obtained from the correspondence. If $t^{-1}(g)$ is contained in $D_{r+1,2}$, then $P(t)$ is empty.[5] Finally, if $t^{-1}(g)$ is empty, then $P(t)$ is a proper subset of the divisors obtained from the correspondence. Either way, $P(t)$ is a cell in the Prym variety.

Remark 2.2. By [LU19, Corollary 5.3.10], for fixed gonality g and torsion k, the Prym–Brill–Noether locus $V^r(\Gamma,\varphi)$ is the union of the subspaces $P(t)$, where t ranges over the Prym tableaux of type (g,r,k). Moreover, it suffices to consider the tableaux for which the symbol g is in the "correct" diagonal modulo 2, namely, $D_{r,2}$.

3. Dimensions of Prym–Brill–Noether loci

Our primary focus in this section is to prove Theorem A by constructing Prym tableaux that—in a sense we shall make precise—minimize the number of symbols used. In Section 3.1, we describe a restricted class of Prym tableaux, called reflective, that are easier to work with and determine sets of

[5] The special role that the symbol g plays in determining $P(t)$ will cause minor headaches in Section 3.1, but thereafter, we avoid the issue entirely by working (almost) exclusively with a different sort of tableau whose symbols only go up to $g-1$.

divisors that are maximal with respect to containment. In Section 3.2, we compute the largest dimension of any cell determined by a reflective tableau (of fixed type) and thereby compute the dimension of the Prym–Brill–Noether locus.

3.1. **Reflective tableaux.** Fix a Prym tableau t of type (g,r,k). We define the *codimension* of t to be number of integers $a \in [g-1]$ for which either of the symbols a or $2g-a$ appears in t. By the tableau–divisor correspondence, the codimension of t coincides with the codimension of the cell $P(t)$ relative to the Prym variety (provided that $t^{-1}(g) \subset D_{r,2}$); indeed, there are at most $g-1$ degrees of freedom—one for each loop $\tilde{\gamma}_a$ with $a \in [g-1]$—and the chip on the a-th loop is free just in case neither a nor $2g-a$ appears in the tableau. Then the path to proving Theorem A is clear:

> To compute the codimension of $V^r(\Gamma,\varphi)$, it suffices to compute the minimal codimension of any Prym tableau of type (g,r,k).

To that end, it is beneficial to consider tableaux with a stronger symmetry than Prym tableaux. Given $\lambda \subset [r+1]\times[r+1]$, consider the map $\rho\colon \lambda \to [r+1]\times[r+1]$ defined by $\rho(x,y) = (r+2-y, r+2-x)$; in other words, ρ picks out the box that is the reflection of (x,y) across the *main anti-diagonal*, A_{r+1}. Fixing a map $t\colon \lambda \to [2g-1]$, we say that a box $(x,y)\in\lambda$ is *reflective* (in t) provided that $\rho(x,y)\in\lambda$ and $t(x,y) = 2g - t(\rho(x,y))$, i.e., the symbol in the box is the dual of the symbol in its reflection.

Definition 3.1. A displacement tableau t is said to be *reflective* if every box of t is reflective.

Note that reflective tableaux defined on $[r+1]\times[r+1]$ are Prym. Moreover, if t is such a tableau, then each box along the main anti-diagonal of t must contain the symbol g, and g appears nowhere else. In particular, $t^{-1}(g)\subset D_{r,2}$, so $P(t)$ is nonempty.

Our goal in the remainder of this section is to prove that, in our search for Prym tableaux of minimal codimension, it suffices to restrict our attention to the class of reflective tableaux. Proposition 3.3 makes this precise, although we first need the notion of tableaux containment that the next definition provides.

Definition 3.2. Given Prym tableaux t and s of type (g,r,k), we say that t *dominates* s if $g\in t(D_{i,2})$ implies that $g \in s(D_{i,2})$ and if, for any $a\neq g$ and $i\in\mathbb{Z}/k\mathbb{Z}$, $a\in t(D_{i,k})$ implies that either $a\in s(D_{i,k})$ or $2g-a\in s(D_{i,k})$. If t and s each dominate the other, then we call them *equivalent*.

It follows from the tableau–divisor correspondence that t dominates s only if $P(t)\supset P(s)$. Indeed, this containment holds whenever each chip that is fixed in $P(t)$ is also fixed in $P(s)$ at the same coordinate.[6] It follows that t and s are equivalent only if $P(t)=P(s)$. If s dominates t, then $\operatorname{codim}(s)\leq\operatorname{codim}(t)$. Therefore, for the purpose of computing the dimension of $V^r(\Gamma,\varphi)$, we may restrict our attention to tableaux that are maximal with respect to the partial order given by dominance. The main result of this section is the following.

Proposition 3.3. *Let t be a Prym tableau such that $t^{-1}(g)\subset D_{r,2}$. Then there exists a reflective tableau s that dominates t.*

The following definition from [Pfl17b] will be used repeatedly during the proof. Given a partition λ and a subset $S\subset\mathbb{Z}/k\mathbb{Z}$, the *upward displacement of λ by S*, denoted $\operatorname{disp}^+(\lambda,S)$, is equal to $\lambda\cup L$, where L consists precisely of those boxes $(x,y)\notin\lambda$ such that all of the following conditions hold:

- $(x-1,y)\in\lambda$ or $x=1$,
- $(x,y-1)\in\lambda$ or $y=1$, and
- $(x,y)\in D_{i,k}$ for some $i\in S$.

The boxes in L are known as the *loose boxes of λ with respect to S*. When $S=\mathbb{Z}/k\mathbb{Z}$, we use the shorthand $\operatorname{disp}^+(\lambda)$ and note the following: if λ is a partition, then so is $\operatorname{disp}^+(\lambda)$; L is nonempty; and every box

[6] The condition on the symbol g in Definition 3.2 ensures that if $P(t)$ is empty, then so is $P(s)$.

in $\mathbb{N}^2 \smallsetminus \mathrm{disp}^+(\lambda)$ is above some box in L. The usefulness of this operation on partitions is made evident in the following example, which outlines the subsequent proof of Proposition 3.3.

Example 3.4. Consider the first Prym tableau of type $(g, r, k) = (11, 4, 3)$ in the sequence illustrated in Fig. 3.1. This tableau is far from being reflective, but at each step we make small changes so that the resulting tableau is closer to being reflective and dominates the preceding one.

At each step, the boxes previously dealt with are colored blue. We look at the symbols in the loose boxes with respect to the lower-left blue partition and choose the minimum a; we look at the symbols contained in the reflection of the loose boxes and choose the maximum b; then denote by c the minimum of a and $2g - b$. Now, wherever c or $2g - c$ appears, color the corresponding box and its reflection red. To produce the next tableau in the sequence, replace each symbol in the red-colored boxes with c or $2g - c$ as appropriate. The final tableau is reflective and dominates the initial tableau.

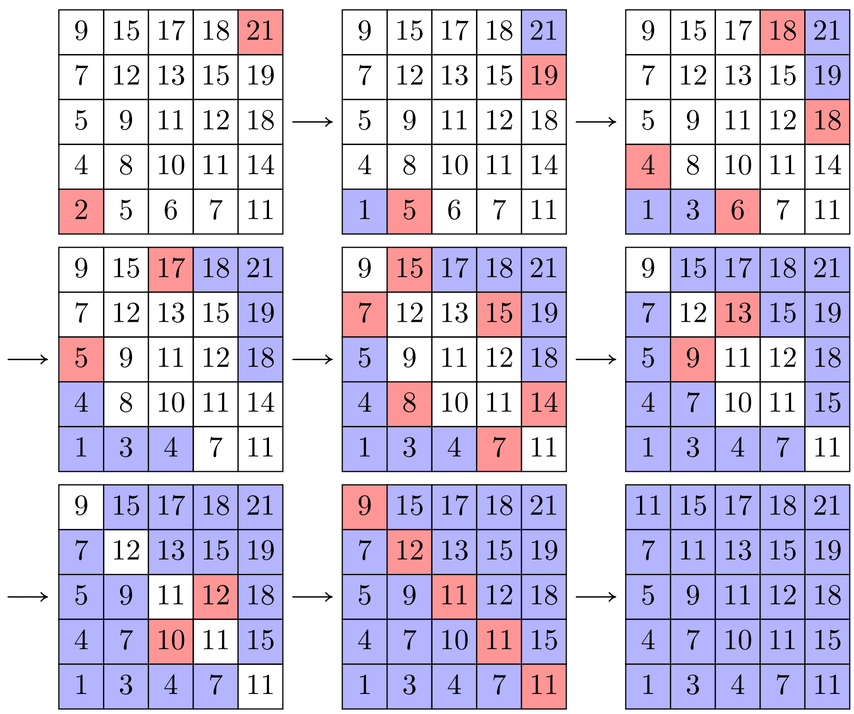

Figure 3.1. Replacing a non-reflective tableau with a dominant reflective one.

The basic operation of the algorithm is to repeatedly *reflect* symbols, i.e., given a box (x, y), to insert the dual symbol, $2g - t(x, y)$, into the reflection, $\rho(x, y)$. The following lemma ensures that the result is still a Prym tableau, granted that the tableau condition holds; then the proof of Proposition 3.3 will make the rest of the algorithm precise.

Lemma 3.5. *Given a Prym tableau t, fix a box (x, y). Then the map s obtained by defining*

$$s(\omega) = \begin{cases} 2g - t(x, y) & \text{for } \omega = \rho(x, y) \\ t(\omega) & \text{otherwise} \end{cases}$$

satisfies the displacement and Prym conditions.

Proof. The only box at which either of the conditions might fail is at $\rho(x, y)$. However, taking the difference of the coordinates of $\rho(x, y) = (r + 2 - y, r + 2 - x)$, we find that $\rho(x, y) \in D_{x-y,k}$. The Prym condition is immediately satisfied, and it is not hard to see that, since any other box containing the symbol $2g - t(x, y)$ would need to be in $D_{x-y,k}$, the displacement condition is also satisfied. □

Proof of Proposition 3.3. Let $\lambda := [r+1] \times [r+1]$ be the domain of t, and let $s_0 := t$. We describe an algorithm which at each step, given a Prym tableau s_i, produces a Prym tableau s_{i+1} that dominates s_i and is reflective on a larger subset of λ. After a finite number of steps, we obtain a Prym tableau s_m that is reflective away from the main anti-diagonal and that dominates t by transitivity. In the final step, the symbols along the main anti-diagonal of s_m are replaced with g to obtain a reflective tableau s.

Induction hypotheses. Suppose that after the i-th step we have a Prym tableau s_i that dominates s_{i-1}. Furthermore, suppose that we have an integer $0 \leq n_i \leq g-1$ and a subset $\kappa_i \subset \lambda$ (where, by convention, $n_0 = 0$ and $\kappa_0 = \varnothing$) such that

- $\omega \in \kappa_i$ just if $\omega \in T_r$ and $s_i(\omega) \leq n_i$,
- $\omega \in \rho(\kappa_i)$ just if $\omega \in \rho(T_r)$ and $s_i(\omega) \geq 2g - n_i$,
- $s_i|_{\kappa_i \cup \rho(\kappa_i)}$ is reflective.[7]

If $\kappa_i = T_r$, then $i = m$ and we are ready to perform the final step. Otherwise, note that κ_i must be a partition by the tableau condition. Therefore, let L_i be the set of loose boxes of κ_i that lie below the main anti-diagonal, and observe that L_i is nonempty.

Definition of s_{i+1}. Consider the minimal positive integer n_{i+1} among the set of symbols $s_i(L_i) \cup (2g - s_i(\rho(L_i)))$. We claim that n_{i+1} exists and is at most $g-1$. Indeed, given any $\omega \in L_i$, if $s_i(\omega) \leq g-1$, then we are done. Otherwise, $s_i(\omega) \geq g$; since ω lies below its reflection $\rho(\omega)$, it follows that $s_i(\rho(\omega)) \geq g+1$; moreover, s_i is Prym, so it must be the case that $s_i(\rho(\omega)) \leq 2g-1$; from both of these inequalities, we find that $1 \leq 2g - s_i(\rho(\omega)) \leq g-1$. In either case, the claim holds. We also know from the assumptions on κ_i and its reflection that $n_{i+1} \geq n_i + 1$.

We now define

$$N_{i+1} := \{\, \omega \in L_i \mid s_i(\omega) = n_i \text{ or } s_i(\rho(\omega)) = 2g - n_i \,\}.$$

In going from s_i to s_{i+1}, we modify only the boxes in $N_{i+1} \cup \rho(N_{i+1})$.[8] In particular, the symbol n_{i+1} is placed in each box of N_{i+1} (that did not already contain it) while $2g - n_{i+1}$ is placed in $\rho(N_{i+1})$. Thus, we define s_{i+1} by

$$s_{i+1}(\omega) = \begin{cases} n_{i+1} & \text{for } \omega \in N_{i+1} \\ 2g - n_{i+1} & \text{for } \omega \in \rho(N_{i+1}) \\ s_i(\omega) & \text{otherwise} \end{cases}.$$

Proof that s_{i+1} is a Prym tableau. By applying Lemma 3.5 every time a symbol is replaced, we know that s_{i+1} satisfies the displacement and Prym conditions. It remains to show that it satisfies the tableau condition at the modified boxes. We shall consider only the case where $\omega \in N_{i+1}$; the case where $\omega \in \rho(N_{i+1})$ follows in a similar way. Observe first that every box below ω contains a symbol that is smaller than n_{i+1}. Indeed, since $\omega \in L_i$, it follows that every box below ω lies in κ_i; since the maximum value of a symbol in κ_i is n_i and we know that $n_i < n_{i+1}$, every box below ω contains a symbol that is strictly smaller than $s_{i+1}(\omega)$.

We also claim that the boxes immediately above ω contain symbols that are greater than n_{i+1}. Writing $\omega = (x, y)$, observe that $(x+1, y) \notin L_i \cup \rho(L_i)$, so $s_{i+1}(x+1, y) = s_i(x+1, y)$. By the tableau condition on s_i, we have that $s_i(x+1, y) > s_i(x, y)$. Finally, (x, y) is in L_i and n_{i+1} was chosen to be minimal among the symbols of L_i (in particular), so we get that $s_i(x, y) \geq n_{i+1} = s_{i+1}(x, y)$. Chaining these inequalities together yields $s_{i+1}(x+1, y) > s_{i+1}(x, y)$; the same argument works for $(x, y+1)$. Hence, s_{i+1} satisfies the tableau condition and so is a Prym tableau.

Proof that s_{i+1} dominates s_i. Let $\omega \in \lambda$. If $s_{i+1}(\omega) \notin \{\, n_{i+1}, 2g - n_{i+1} \,\}$, then $s_{i+1}(\omega) = s_i(\omega)$. This implies that for all symbols besides n_{i+1} and $2g - n_{i+1}$ (including g), the conditions of Definition 3.2 are satisfied.

[7]In Example 3.4, $\kappa_i \cup \rho(\kappa_i)$ is represented by the blue boxes.

[8]In Example 3.4, $N_{i+1} \cup \rho(N_{i+1})$ is represented by the red boxes.

If $s_{i+1}(\omega) \in \{ n_{i+1}, 2g - n_{i+1} \}$, then either $s_{i+1}(\omega) = s_i(\omega)$ or $s_{i+1}(\omega) = 2g - s_i(\rho(\omega))$. In the first case, we are done as above; in the second case, the desired condition still holds on account of the fact that the dual of the symbol $s_{i+1}(\omega)$ appears in s_i and, in particular, is contained in $\rho(\omega)$, which occupies the same diagonal modulo k as ω. Hence, s_{i+1} dominates s_i.

Proof that s_{i+1} satisfies the induction hypotheses. Define κ_{i+1} to be $\kappa_i \cup N_{i+1}$. Then every box in T_r that contains a symbol at most n_i is in κ_i, while any box containing n_{i+1} is in N_{i+1}. Using the definition of loose boxes and the fact that n_{i+1} minimizes the symbols in L_i, we find that no symbol strictly between n_i and n_{i+1} appears in s_{i+1}. Moreover, any box in T_r is above some box of L_i, so the tableau condition precludes n_{i+1} from appearing in $T_r \smallsetminus L_i$; if n_{i+1} appears in L_i, then it appears in N_{i+1} by definition. From these observations, we find that κ_{i+1} contains precisely those boxes of T_r with symbols at most n_i. A similar argument shows that $\rho(\kappa_{i+1})$ contains precisely those boxes of $\rho(T_r)$ with symbols at least $2g - n_i$.

Finally, the restriction of s_{i+1} to $\kappa_{i+1} \cup \rho(\kappa_{i+1})$ is reflective. Indeed, it is reflective on $\kappa_i \cup \rho(\kappa_i)$ because s_i is, and s_{i+1} and s_i agree on that subset. Moreover, s_{i+1} is reflective on $N_{i+1} \cup \rho(N_{i+1})$ by construction: every symbol in this subset is the dual of the symbol in its reflection. Therefore, all the inductive hypotheses are satisfied.

Final step. Since κ_{i+1} strictly contains κ_i, after a finite number of steps m, we have that $\kappa_m = T_r$. In other words, s_m is reflective everywhere but (possibly) the main anti-diagonal, A_{r+1}. Then we replace the symbols in all of the boxes in A_{r+1} with the symbol g; the resulting tableau s is reflective and dominates s_m, completing the proof.[9] □

A reflective tableau is uniquely determined by its restriction to T_r, so we may as well only consider this subset.

Definition 3.6. A *staircase Prym tableau of type* (g, r, k) is a k-uniform displacement tableau $t\colon T_r \to [g-1]$.

We extend all definitions regarding Prym tableaux to staircase Prym tableaux in the natural way; for instance, denoting by $\hat{t}$ the reflective tableau which extends a given staircase Prym tableau t, we define $P(t)$ to be just $P(\hat{t})$. Certain other definitions become more intuitive: the codimension of t is just the number of distinct symbols appearing in t, and t dominates another staircase Prym tableau s just in case $t(D_{i,k}) \subset s(D_{i,k})$.

3.2. **Proof of Theorem A.** Throughout this section, $\varphi\colon \widetilde{\Gamma} \to \Gamma$ will represent a folded chain of loops of genus g, where the edge lengths of Γ are either generic or the torsion of each loop is k. For the sake of brevity, we will refer to the folded chain of loops and its corresponding Prym tableaux in the former case as *generic* and in the latter as *k-gonal*.

The dimension of $V^r(\Gamma, \varphi)$ is known in the generic case and when k is even; see [LU19, Theorem 6.1.4, Corollary 6.2.2]. When k is odd, [LU19, Remark 6.2.3] provides an upper and a lower bound for the dimension. In this section we show that the dimension of $V^r(\Gamma, \varphi)$ in fact coincides with the lower bound. We restate the precise result here.

Theorem A. *Let $\varphi\colon \widetilde{\Gamma} \to \Gamma$ be a k-gonal uniform folded chain of loops and denote by l the quantity $\lceil \frac{k}{2} \rceil$. Then the codimension of $V^r(\Gamma, \varphi)$ relative to the Prym variety is given by*

$$n(r,k) = \begin{cases} \binom{l+1}{2} + l(r-l) & \text{if } l \le r-1 \\ \binom{r+1}{2} & \text{if } l > r-1 \end{cases}. \tag{1.1}$$

[9] It may not be obvious at first glance why the statement of Proposition 3.3 requires that $t^{-1}(g) \subset D_{r,2}$. If $g \in t(D_{r+1,2})$, then the algorithm described in this proof still produces a reflective tableau s. However, the final step forces $g \in s(D_{r,2})$, so s would fail to dominate t in this case.

To prove the theorem, we need to compute the minimal codimension of $P(t)$ over all Prym tableaux t of type (g,r,k) such that $t^{-1}(g) \subset D_{r,2}$ (see Remark 2.2). By Proposition 3.3, it suffices to consider staircase Prym tableaux: given any Prym tableau (with the correct g-fiber), we apply the reflection algorithm to obtain a dominating reflective Prym tableau. Per the discussion at the end of Section 3.1, it then suffices to consider the staircase Prym tableau that constitutes its restriction to T_r.

The expression $\binom{r+1}{2}$ in the second case in Eq. (1.1) counts the number of boxes in T_r. In this subset, the lattice distance between any two boxes is at most $2r-2 \le 2l-2 < k$, so each must contain a unique symbol; it follows that the number of symbols in any such tableau is precisely $\binom{r+1}{2}$. The same reasoning explains the presence of the $\binom{l+1}{2}$ term in the first case: it counts the number of symbols in T_l, which are all necessarily unique. Any repeats occur above T_l. In fact, we claim that a tableau of minimal codimension contains precisely l new symbols on each subsequent anti-diagonal, of which there are $r-l$; this accounts for the $l(r-l)$ term. Precisely, we say that a set of symbols $S \subset t(A_n)$ is *new* if $S \cap t(T_{n-1})$ is empty. If our claim is true, then the tableau depicted in Fig. 2.2, which is a staircase Prym tableau of type $(12,6,3)$,[10] has minimal codimension.

Proposition 3.7. *Given a staircase Prym tableau t of type (g,r,k), there exist at least l new symbols in A_n for each $n \ge l+1$.*

The following lemma establishes a restriction on symbols which will go most of the way toward proving Proposition 3.7, from which the proof of Theorem A quickly follows.

Lemma 3.8. *Let t be a staircase Prym tableau of type (g,r,k), and fix $n \le r$. For any boxes $(x,y) \in D_{i,k}$ and $(x',y') \in D_{i+1,k}$ that lie below A_n, there exists a box $\omega \in A_n \cap (D_{i,k} \cup D_{i+1,k})$ such that $t(\omega)$ is greater than both $t(x,y)$ and $t(x',y')$.*

Proof. Let $a = t(x,y)$ and $b = t(x',y')$. Since a and b lie in different diagonals modulo k, we know that $a \ne b$. We will assume that $a < b$; the proof follows in the same way when the converse inequality holds. We want to show that there is a box $\omega := (\omega_1, \omega_2)$ in $A_n \cap (D_{i,k} \cup D_{i+1,k})$ that lies above (x',y'), since this would force $t(\omega) > b$.

Indeed, define $\delta = n+1-x'-y'$. We know that $x'+y' \le n$ because (x',y') sits below A_n, so $\delta \ge 1$. If δ is even, then we define

$$\omega := \left(x' + \frac{\delta}{2}, y' + \frac{\delta}{2}\right).$$

Note that ω_1 and ω_2 are both positive integers, $\omega_1 + \omega_2 = n+1$, and $\omega_1 - \omega_2 = x' - y' \equiv i+1 \pmod{k}$; moreover, ω sits above (x',y'), as desired.

Suppose instead that δ is odd; then define

$$\omega := \left(x' + \frac{\delta-1}{2}, y' + \frac{\delta+1}{2}\right).$$

The desired properties once again hold (although in this case, $\omega \in D_{i,k}$). □

Proof of Proposition 3.7. Given n such that $l+1 \le n \le r$, we note first that $T_{n-1} \cap D_{i,k}$ is nonempty. Indeed, we may write $i \in \{-l+1, \ldots, l-1\}$. If $i \ge 0$, we have that $(1+i,1) \in T_{n-1} \cap D_{i,k}$; if $i < 0$, then $(1,1-i) \in T_{n-1} \cap D_{i,k}$.

For each i, choose $\omega_i \in T_{n-1} \cap D_{i,k}$ such that $t(\omega_i)$ is maximal among $t(T_{n-1} \cap D_{i,k})$. Then apply Lemma 3.8 to each pair $\{\omega_i, \omega_{i+1}\}$ to obtain a box $\eta_i \in A_n \cap (D_{i,k} \cup D_{i+1,k})$ such that $t(\eta_i) > t(\omega_i)$ and $t(\eta_i) > t(\omega_{i+1})$. Hence, $t(\eta_i) > t(\omega)$ for every box $\omega \in T_{n-1} \cap (D_{i,k} \cup D_{i+1,k})$ and so is new in A_n.

Therefore, for each pair $\{i, i+1\} \subset \mathbb{Z}/k\mathbb{Z}$, the set $A_n \cap (D_{i,k} \cup D_{i+1,k})$ contains at least one new symbol, which we shall denote by b_i. Note that if $\{i,i+1\}$ and $\{j,j+1\}$ are disjoint, then their respective symbols b_i and b_j must lie in different diagonals modulo k, and so must be distinct. Thus,

[10] In fact, it is staircase Prym of type $(g,6,3)$ for any $g \ge 12$.

the minimum number of new symbols in A_n coincides with the minimum number of elements we can choose from $\mathbb{Z}/k\mathbb{Z}$ such that we have at least one element in each pair $\{i, i+1\}$. Suppose for the sake of contradiction that we could achieve this with $l-1$ elements. Each is a member of two pairs, so we cover at most $2(l-1) < k$ pairs. This is insufficient, as there are k pairs, so the minimum size of such a set is l. ☐

Proof of Theorem A. We have already proved the case where $l > r$, so assume otherwise. From Proposition 3.7 and our earlier remarks, we get that T_r contains at least $\binom{l+1}{2} + l(r-l)$ distinct symbols. Hence, $\operatorname{codim} V^r(\Gamma, \varphi)$ is bounded below by this quantity. Meanwhile, [LU19, Corollary 6.2.2, Remark 6.2.3] implies that it is also an upper bound, so we are done. ☐

3.3. Relation to algebraic geometry.

We are now in a position to prove Corollary B, restated below.

Corollary B. *Let $r \geq -1$ and $k \geq 2$. Then there is a nonempty open subset of the k-gonal locus of $\mathcal{R}_g$ such that for every unramified double cover $f\colon \widetilde{C} \to C$ in this open subset we have*

$$\dim V^r(C, f) \leq g - 1 - n(r,k). \tag{1.2}$$

Proof. Having established Theorem A, the proof of the Corollary is almost identical to the proof of [LU19, Theorem B] and similar to analogous results from [CDPR12, JR17, Pfl17b]. We illustrate the general idea, and leave the details to the reader. First, due to our assumption that the characteristic of the residue field is prime to both 2 and k, we may lift the folded chain of loops $\varphi : \widetilde{\Gamma} \to \Gamma$ to a k-gonal unramified double cover $f\colon \widetilde{X} \to X$ [LU19, Lemma 7.0.1]. By Baker's specialization lemma [Bak08, Corollary 2.11], the tropicalization of $V^r(X, f)$ (if non-empty) lies within $V^r(\Gamma, \varphi)$. By Gubler's Bieri–Groves Theorem [Gub07, Theorem 6.9], dimensions are preserved under tropicalization, so the codimension of $V^r(X, f)$ inside the Prym variety is bounded from below by $n(r, k)$. A standard upper semicontinuity argument shows that $n(r, k)$ is, in fact, an upper bound on the codimension for a non-empty open set in the k-gonal locus of $\mathcal{R}_g$, as claimed. ☐

Note that this bound is not necessarily strict. For instance, if $g \leq 2k-2$, then the general curve is k-gonal. In this case, the codimension of the Prym–Brill–Noether locus of a general curve is $\binom{r+1}{2}$ [Wel85], which is stronger than the bound provided in Corollary B. However, we believe that our bound is strict when g is sufficiently high.

Conjecture 3.9. *Suppose that $g \gg n(r,k)$, and let $f\colon \widetilde{C} \to C$ be a generic Prym curve. Then*

$$\dim V^r(C, f) = g - 1 - n(r,k).$$

4. Tropological properties

As before, fix a folded chain of loops $\varphi\colon \widetilde{\Gamma} \to \Gamma$ of genus g and gonality k. In this section, we prove that the Prym–Brill–Noether locus $V^r(\Gamma, \varphi)$ is pure-dimensional (Proposition 4.8) and connected in codimension 1 when the dimension is greater than zero (Proposition 4.9). In Section 4.1, we develop the notions of strips and non-repeating tableaux, which will also be necessary for computing the Betti number of $V^r(\Gamma, \varphi)$ in Section 5.[11] The proof of pure-dimensionality then comes as an easy corollary of Proposition 4.7. We tackle connectedness in Section 4.2.

[11] We use the term *Betti number* for the genus of the Prym Brill–Noether locus to distinguish it from the genus of our underlying graphs.

4.1. **Strips and non-repeating tableaux.** We focus our attention on Prym tableaux of minimal codimension. Since Proposition 3.3 implies that any such tableau is equivalent to a reflective tableau and hence a staircase Prym tableau, it suffices to consider this restricted type. To simplify our terminology, we shall say that a tableau is *minimal* if it is staircase Prym of minimal codimension.

In the generic case (which, by a slight abuse of terminology, we take to include both the case of generic edge lengths and the non-generic case with $l \geq r$), minimal tableaux are relatively easy to classify, since they are precisely the standard Young tableaux on T_r. The cases of even and odd torsion elude such a concise description; nonetheless, as we will presently make precise, there are subsets of T_r that we call strips on which minimal tableaux are determined up to equivalence.

Definition 4.1. A subset $\mu \subset T_r$ is a *strip* if $T_l \subset \mu$ and there exists a unique box in $\mu \cap A_n$ for each $n \in \{ l, l+1, \ldots, r \}$ called the *n-th leftmost box* that satisfies the following properties:

- $(1, l)$ is the l-th leftmost box,
- if (x, y) is the n-th leftmost box, then the $(n+1)$-th leftmost box is $(x, y+1)$ or $(x+1, y)$, and
- if (x, y) is the n-th leftmost box, then the boxes of $\mu \cap A_n$ are precisely those of the form $(x+i, y-i)$ for each $i \in \{ 0, 1, \ldots, l-1 \}$.

If (x, y) is the n-th leftmost box, then we call $(x+l-1, y-l+1)$ the *n-th rightmost box*. We call r and l the *length* and *width* of μ, respectively.

Note that $\mu \cap A_n$ contains precisely $\min\{ n, l \}$ boxes, any two of which are separated by lattice distance at most $2l-2$. This implies that any k-uniform tableau defined on T_r must be injective on each $\mu \cap A_n$. Moreover, since we designate $(1, l)$ as the l-th leftmost box and choose each subsequent leftmost box out of two possibilities, it follows that μ may take on any of 2^{r-l} distinct shapes.

$T_r \smallsetminus \mu$ consists of two (possibly empty) contiguous components, which we shall call the *left* and *right*, respectively. In particular, the left component of $T_r \smallsetminus \mu$ (if it exists) is the one that contains the box $(1, r)$. We refer to the strip whose right component is empty as the *horizontal strip* and denote it by μ_0.

We now introduce a subclass of maps $T_r \to [g-1]$ that will play a key role for the rest of the paper. The even and odd cases differ; in what follows, let ϵ be 0 if k is even and 1 if k is odd.

Definition 4.2. Given a strip μ and a map $t\colon T_r \to [g-1]$ such that $t|_\mu$ satisfies the tableau and displacement conditions, we say that t is *non-repeating in μ* if

(a) $t(x, y) = t(x+l-\epsilon, y-l)$ for each (x, y) in the left component of $T_r \smallsetminus \mu$,
(b) $t(x, y) = t(x-l, y+l-\epsilon)$ for each (x, y) in the right component of $T_r \smallsetminus \mu$, and
(c) writing the n-th leftmost box as (x, y), if $(x+1, y)$ is the $(n+1)$-leftmost box, then $t(x, y) < t(x+l-\epsilon, y-l+1)$; otherwise, $t(x+l-1, y-l+1) < t(x, y+1-\epsilon)$.

We refer to (a) and (b) as the *left* and *right repeating conditions* respectively and to (c) as the *gluing condition*.

See Fig. 4.1 for an example. It is straightforward to check that these conditions are symmetrical with respect to transposing the first and second coordinates.[12] This fact will simplify the proofs of several properties of non-repeating maps.

As a small convenience, we shall use the cardinal directions to refer to boxes relative to a given box, with east and north corresponding to increasing first and second coordinates, respectively. So for example, the north neighbor of (x, y) is $(x, y+1)$, and the box three steps west of (x, y) is $(x-3, y)$. We shall also call the north and east neighbors the *upper neighbors* and the south and west neighbors the *lower neighbors*.

Proposition 4.3. *Given a strip μ, if t is non-repeating in μ, then t is a minimal tableau.*

[12] To be precise, we transpose by switching the coordinates of each box, replacing x with y and vice versa, and replacing "left" with "right" and vice versa.

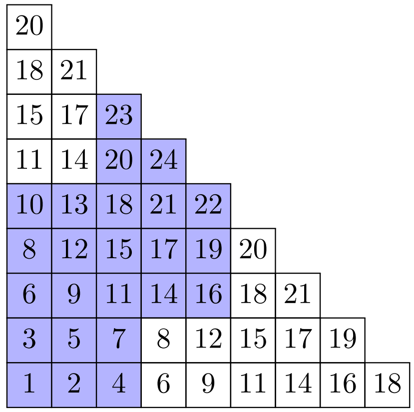

Figure 4.1. A minimal tableau of size 9 and torsion 5 that is non-repeating on the strip depicted in blue.

Proof. We first show that t is staircase Prym. The displacement condition holds in μ by definition and in $T_r \smallsetminus \mu$ by the repeating conditions, since each symbol in $T_r \smallsetminus \mu$ is copied from a box that is distance k away.

Recall from the definition that $t|_\mu$ satisfies the tableau condition. We need to check that t satisfies the tableau condition. We observe first that $t(1,1)$ is smaller than both $t(2,1)$ and $t(1,2)$. This follows in the case that $k = 2$ because both $(2,1)$ and $(1,2)$ contain the same symbol and one of the two is in μ along with $(1,1)$; in the case that $k > 2$, all three boxes are in μ. Suppose for the sake of induction that the tableau condition holds for all boxes in T_{n-1} (and in particular, every box in A_{n-1} contains a symbol smaller than the symbols of its upper neighbors). Let (x,y) be a box in A_n. If $n = r$, we are done; otherwise, it suffices to show that $t(x,y) < t(x+1,y)$ and $t(x,y) < t(x,y+1)$.

By transposing the coordinates if necessary, we may assume that the $(n+1)$-th leftmost box is east of the n-th leftmost box. Suppose first that (x,y) is in μ. If it is not the n-th leftmost box, both of the desired inequalities follow from the fact that $(x+1,y)$ and $(x,y+1)$ are both also in μ. Otherwise, its east neighbor is in μ while its north neighbor is in the left component of $T_r \smallsetminus \mu$. We use the left repeating condition followed by the gluing condition to obtain the desired inequality:

$$t(x,y+1) = t(x+l-\epsilon, y-l+1) > t(x,y).$$

Now suppose that (x,y) is in the left component. If k is odd, then the symbols in (x,y) and its upper neighbors are copied from the respective symbols in $(x+l-1, y-l)$ and *its* upper neighbors. Since $(x+l-1,y-l)$ is in A_{n-1}, it satisfies the tableau condition by the induction hypothesis. If k is even, we observe via repeated application of the left repeating condition that there is some box (x',y') in $\mu \cap A_n$ such that the symbols in (x,y) and its upper neighbors are copied from the respective symbols in (x',y') and *its* upper neighbors. We checked that the desired inequalities hold for every box in $\mu \cap A_n$, so they hold at (x,y) as well. Analogous arguments hold in both the odd and even cases when (x,y) is in the right component.

By induction, t satisfies the tableau condition. Thus, t is staircase Prym. It remains to show that t has minimal codimension. Indeed, observe that μ consists of $n(r,k) = \binom{l+1}{2} + l(r-l)$ boxes, so $t|_\mu$ contains at most that many distinct symbols. Every symbol of t in $T_r \smallsetminus \mu$ is repeated from within μ, so t as a whole contains at most $n(r,k)$ symbols. By Theorem A, t is minimal. □

Corollary 4.4. *If t is non-repeating in μ, then $t|_\mu$ is injective.*

Proof. By applying Theorem A, we know that t cannot contain fewer than $n(r,k)$ symbols. Every symbol of t appears in μ, and μ consists of precisely $n(r,k)$ boxes. Hence, each of those boxes must contain a distinct symbol. □

Lemma 4.5. *Fix a strip μ and $i \in \mathbb{Z}/k\mathbb{Z}$. Then for any map t non-repeating in μ and any boxes $\omega \in \mu \cap A_m \cap D_{i,k}$ and $\omega' \in \mu \cap A_n \cap D_{i,k}$ with $m < n$, it must be the case that $t(\omega) < t(\omega')$.*

Proof. The statement is true for $n \le l$ by the tableau condition on $t|_\mu$ since $D_{i,k} \cap T_l$ is contained in a single diagonal. We proceed by induction on n. Suppose that the statement is true in T_n, and let ω' be a box in $\mu \cap A_{n+1} \cap D_{i,k}$. Since $\mu \cap A_m \cap D_{i,k}$ contains at most one box, it suffices to show that $t(\omega) < t(\omega')$ for $\omega \in \mu \cap A_m \cap D_{i,k}$ where $m \le n$ is the maximum index such that $\mu \cap A_m \cap D_{i,k}$ is nonempty.

Suppose that k is odd. Let (x, y) be the n-th leftmost box, and assume without loss of generality that $(x+1, y)$ is the $(n+1)$-th leftmost box. If ω' is the $(n+1)$-th rightmost box, $(x+l, y-l+1)$, then $m = n$ and $\omega = (x, y)$. Then $t(\omega) < t(\omega')$ by the gluing condition. If ω' is any other box (x', y') in $\mu \cap A_{n+1} \cap D_{i,k}$, it is not hard to see that $\mu \cap A_n \cap D_{i,k}$ is empty and $\mu \cap A_{n-1} \cap D_{i,k}$ contains precisely one box; namely, $(x'-1, y'-1)$. Then $\omega = (x'-1, y'-1)$, and the tableau condition on $t|_\mu$ implies that $t(\omega) < t(\omega')$.

The case where k is even follows in a similar way; we omit the details here. ☐

Beginning with the following proposition, we start to see that the odd and even cases are fundamentally different. In particular, even tableaux that are non-repeating on some strip are in fact non-repeating on every strip; shortly, we will restrict our attention to the horizontal strip whenever we talk about the even case.

Proposition 4.6. *Let t and s be tableaux that are non-repeating in μ and ν respectively. For k odd, t and s are equivalent if and only if $\mu = \nu$ and $t|_\mu = s|_\nu$. For k even, if $\mu = \nu$, then t and s are equivalent if and only if $t|_\mu = s|_\nu$.*

Proof. The converse of each statement trivially follows from Definition 4.2. For the forward direction, suppose that t and s are equivalent, and assume for the time being that $\mu = \nu$. Because $t(D_{i,k}) = s(D_{i,k})$ for each i, the total ordering on the boxes of $\mu \cap D_{i,k}$ given by Lemma 4.5 forces $t|_\mu = s|_\nu$. This proves the statement in the even case. For the odd case, suppose for the sake of contradiction that $\mu \ne \nu$, and let n be the smallest index such that $\mu \cap A_{n+1} \ne \nu \cap A_{n+1}$ (and note in particular that $n \ge l$). Applying the argument above to the restricted domain T_n, we have that $t|_{\mu \cap T_n} = s|_{\nu \cap T_n}$. Then the gluing condition on μ and ν forces the $(n+1)$-th leftmost box of each to be the same, so $\mu \cap A_{n+1} = \nu \cap A_{n+1}$, a contradiction. ☐

Proposition 4.7. *Given a staircase Prym tableau t, there exists a strip μ and a tableau s that is non-repeating in μ such that s dominates t. Moreover, in the even case, this strip may be chosen to be horizontal.*

Proof. First suppose that k is odd. We begin by defining a tableau $s_l = t|_{T_l}$ and a strip $\mu_l = T_l$, and proceed by induction: suppose that we have defined a tableau s_n on T_n that is non-repeating on a strip $\mu_n \subset T_n$, and suppose that $s_n|_{\mu_n} = t|_{\mu_n}$. Let (x, y) be the n-th leftmost box; then $(x+l-1, y-l+1)$ is the n-th rightmost box. These two boxes, separated by distance less than k, cannot contain the same symbol in t. Extend μ_n to a strip $\mu_{n+1} \subset T_{n+1}$ by defining the $(n+1)$-leftmost box to be $(x+1, y)$ if $t(x, y) < t(x+l-1, y-l+1)$, and $(x, y+1)$ otherwise. (This ensures that s_{n+1}, which we shall presently define, satisfies the gluing condition.) Then extend s_n to the map s_{n+1} that agrees with t on $\mu_{n+1} \cap A_{n+1}$ and is defined elsewhere according to the repeating conditions. It is not hard to see that s_{n+1} is non-repeating in μ_{n+1} and dominates $t|_{T_{n+1}}$. Take $s = s_r$ for the desired result.

We now consider the case that k is even. Given each $\omega \in A_n \cap D_{i,k}$, define $s(\omega) = \max t(A_n \cap D_{i,k})$. It is clear that s dominates t. We claim that s is non-repeating on the horizontal strip μ_0. Note first that the displacement and tableau conditions are satisfied everywhere. The first follows immediately from the definition of s; to prove the second, it suffices to show that $s(x, y) < s(x+1, y)$ for each box (x, y) since the other inequality, $s(x, y) < s(x, y+1)$, follows by transposing coordinates. Indeed, suppose that $(x, y) \in A_n \cap D_{i,k}$. Choose a box $(x', y') \in A_n \cap D_{i,k}$ that satisfies $t(x', y') = \max t(A_n \cap D_{i,k})$; then

$$s(x, y) = t(x', y') < t(x'+1, y') \le \max t(A_{n+1} \cap D_{i+1,k}) = s(x+1, y).$$

The equalities follow from the definition of s, the strict inequality follows from the tableau condition on t, and the weak inequality follows because $(x'+1, y') \in A_{n+1} \cap D_{i+1,k}$.

The left repeating condition holds because, given a box (x, y) in the left component, $(x, y) \in A_n \cap D_{i,k}$ implies that $(x+l, y-l) \in A_n \cap D_{i,k}$, so $s(x, y) = s(x+l, y-l)$. The right repeating condition is vacuously satisfied because the right component is empty. The gluing condition follows from the tableau and left repeating conditions: $s(x, l) < s(x, l+1) = s(x+l, 1)$. ☐

The fact that the Prym–Brill–Noether locus is pure-dimensional readily follows from the results of this section.

Proposition 4.8. *$V^r(\Gamma, \varphi)$ is pure-dimensional for any gonality k.*

Proof. Given a Prym tableau t, we want to find a Prym tableau that dominates t and has codimension $n(r, k)$. Indeed, apply the reflection algorithm of Proposition 3.3 to t; the resulting tableau u dominates t. In the generic case, $u|_{T_r}$ is injective, so we are done. Otherwise, apply Proposition 4.7 to $u|_{T_r}$ to obtain a map s defined on T_r that dominates $t|_{T_r}$ by transitivity. This map is a minimal tableau by Proposition 4.3. Extend s uniquely to a reflective tableau. This is the desired tableau. ☐

The ultimate motivation for defining non-repeating tableaux comes from Propositions 4.6 and 4.7, which for fixed parameters (g, r, k) yield the following powerful correspondence:

$$\{\text{maximal cells of } V^r(\Gamma, \varphi)\} \leftrightarrow \{\text{non-repeating tableaux of type } (g, r, k)\}. \tag{4.1}$$

This result will be invaluable in the remainder of this section and in Section 5.

For $k \le 2r$, define a *strip tableau of type* (g, r, k) to be an injective tableau t defined on a strip μ of length r and width l such that t satisfies the gluing condition on μ and takes values in $[g-1]$.[13] In the odd case, μ may take any of 2^{r-l} possible shapes; in the even case, we require that $\mu = \mu_0$. (We adopt the convention that $\mu = \mu_0$ whenever we refer to non-repeating tableaux in the even case.)

Clearly, any strip tableau extends uniquely to a non-repeating tableau by applying the repeating conditions. Conversely, any non-repeating tableau determines a unique strip tableau. Hence, the two classes of tableaux are equivalent, and we may use either one depending on the circumstances. Keeping in mind the nuances that distinguish the odd and even cases, Eq. (4.1) yields the correspondence

$$\{\text{maximal cells of } V^r(\Gamma, \varphi)\} \leftrightarrow \{\text{strip tableaux of type } (g, r, k)\} \tag{4.2}$$

as expected. Just as we extended the definitions pertaining to Prym tableaux to staircase Prym tableaux (see the discussion at the end of Section 3.1), so too may we extend these definitions even further to strip tableaux.

4.2. **Connectedness.** In this section, we shall occupy ourselves with the following result, which we prove in three seperate cases—generic, even, and odd.

Proposition 4.9. *If $\dim V^r(\Gamma, \varphi) \ge 1$, then $V^r(\Gamma, \varphi)$ is connected in codimension 1.*

As we explain in Remark 4.11, the following definition captures the analogous notion of connectedness on the level of tableaux.

Definition 4.10. Suppose that t and s are k-uniform displacement tableaux having the same shape, the same number of distinct symbols, and image contained in $[g-1]$ for some g. We say that t and s are *adjacent* if there exist two symbols $a, b \in [g-1]$ and two indices $i, j \in \mathbb{Z}/k\mathbb{Z}$ such that

- $s(D_{i,k}) = t(D_{i,k}) \cup \{a\}$,

[13] We require that $k \le 2r$ because otherwise μ, which contains T_l by definition, is larger than T_r—this does not make sense! In the case that $k > 2r$, the "strip tableaux" are really just the standard Young tableaux defined on T_r. Observe that Eq. (4.2) still holds in this case.

- $t(D_{j,k}) = s(D_{j,k}) \cup \{ b \}$,
- $t(D_{h,k}) = s(D_{h,k})$ for all $h \in \mathbb{Z}/k\mathbb{Z} \smallsetminus \{ i, j \}$, and
- $a = b$ only if $i = j$.

We say that t and s are *connected* if there exists a sequence $(t_i)_{i=0}^{n}$ of tableaux such that $t_0 = t$, $t_n = s$, and t_i is adjacent to t_{i+1} for each $i \in \{ 0, \dots, n-1 \}$.

Remark 4.11. Using the tableau–divisor correspondence (Section 2.1), it is straightforward to show that, in the particular case where t and s are minimal tableaux of type (g, r, k), t and s are adjacent just in case either $P(t) \cap P(s)$ is a torus of codimension 1 or $P(t) = P(s)$. The former corresponds to the case $a \neq b$, the latter to the case $a = b$.[14] Likewise, the tableaux t and s are connected just if $P(t)$ and $P(s)$ are connected in codimension 1 by a sequence of cells of $V^r(\Gamma, \varphi)$. It then follows from the correspondence in Eq. (4.1) that in order to prove Proposition 4.9, it suffices to show that any two non-repeating tableaux (of the same type) are connected.

Definition 4.10 is unintuitive and unwieldy, but in fact, there are relatively straightforward methods by which we may produce connected tableaux. As we shall see, to prove that any two non-repeating tableaux are connected, we only need the following three operations and combinations thereof. Fix a k-uniform displacement tableau t and a symbol $a \in [g-1]$ that does not appear in t.

(a) Choose a box ω and define $s(\omega) = a$ and $s(\omega') = t(\omega')$ for all $\omega' \neq \omega$. We call this procedure *swapping* a *into* ω. In general, s will not satisfy the tableau condition unless a is greater than the symbols in the lower neighbors of ω and smaller than those in the upper neighbors. Moreover, s has the same number of symbols as t just in case the original symbol, $t(\omega)$, does not appear elsewhere in t. Given that s satisfies these conditions, s is adjacent to t.

(b) To ensure that s does not have more symbols than t, we introduce the related notion of *swapping* a *in for* b, where b is any symbol. If b does not appear in the tableau, define $s = t$ (i.e., do nothing). Otherwise, for each box ω containing b, define $s(\omega) = a$, and for all other boxes ω', define $s(\omega') = t(\omega')$. The tableau condition must again be checked, this time at each box ω. Supposing that it holds, s is adjacent to t.

(c) It is straightforward to check that we may always swap a in for $a+1$ and a in for $a-1$. Hence, if there is a symbol $b > a$ that we want to pull out of the tableau, we iterate the following procedure: after the i-th step, $a+i$ is not in the tableau, so swap $a+i$ in for $a+i+1$. After $b-a$ steps, each symbol $a+i$ in t for $i \in [b-a]$ has been decremented by 1. In particular, b no longer appears in the tableau. An analogous procedure may be used in the case that $b < a$; in either case, we call this *cycling out* b *using* a. The resulting tableau s is connected to t.

The following lemma outlines the first step in proving Proposition 4.9.

Lemma 4.12. *Any two injective tableaux of the same shape containing at most* $g-2$ *symbols are connected.*

To prove it, we introduce one more tool. Given any shape $\lambda \subset \mathbb{N}^2$, we establish a total order on its boxes as follows: for any $(x, y) \in A_m$ and $(x', y') \in A_n$, say that $(x, y) < (x', y')$ if $m < n$, or if both $m = n$ and $x < x'$.[15] Let $Q_\lambda(\omega)$ be the place of the box ω in the order, i.e., the number of boxes $\omega' \in \lambda$ for which $\omega' \leq \omega$. Then we define an $\mathbb{N}$-valued function R_λ on injective tableaux of shape λ so that

$$R_\lambda(t) := |\lambda| - \max \{ Q_\lambda(\omega) \mid t(\omega') = Q_\lambda(\omega') \text{ for all } \omega' \leq \omega \}.$$

There is a unique tableau $\bar{t}$ for which $R_\lambda(t) = 0$; call it the *standard increasing tableau*. For example, if $\lambda = T_4$, then $\bar{t}$ is the final tableau in Fig. 4.2, while the value of R_{T_4} at the first tableau in the sequence is

[14] If we did not require that $i = j$ whenever $a = b$, then we could obtain $P(t) \cap P(s) = \varnothing$, which is undesirable.

[15] When $(x, y) < (x', y')$, we say that (x, y) is "smaller" than (x', y') in order to avoid confusion with the previous terminology "below" (see Section 2.1); the latter implies the former, but the converse does not hold in general.

8, since the first and second boxes contain the correct symbols but the third box does not—it contains a 5 rather than a 3. Intuitively, R_λ measures how far a given tableau is from being identical to $\bar{t}$.

Proof of Lemma 4.12. We will show by induction on the values of R_λ that any injective tableau t of shape λ is connected to the standard increasing tableau $\bar{t}$.

If $R_\lambda(t) = 0$, then the statement is trivially true since it must be the case that $t = \bar{t}$. Otherwise, suppose that any tableau s with $R_\lambda(s) < R_\lambda(t)$ is connected to $\bar{t}$. Then it suffices to show that t is connected to some such s.

Let ω be the smallest box (relative to the total order) such that $t(\omega) \neq Q_\lambda(\omega)$. Denote by a the value $Q_\lambda(\omega)$ and by S the set of boxes strictly smaller than ω. Our goal is to find a sequence of swapping operations that leaves the boxes of S (which already contain the correct symbols) untouched while inserting the symbol a into ω. Choose the smallest symbol b not in t. Observe that $a \leq b \leq g-1$; the lower bound holds because every symbol less than a appears in (the correct box of) the tableau, while the upper bound follows from the assumption that t uses at most $g-2$ symbols. First, cycle out a using b and call the resulting tableau u. This has the effect of incrementing any symbol $\{a, a+1, \ldots, b-1\}$ appearing in t, so symbols in S are unaffected. Furthermore, u is injective and connected to t. Second, swap a into ω and call the resulting tableau s. This operation trivially leaves S unaffected. Moreover, s satisfies the tableau condition: the fact that $a < u(\omega)$ covers the upper neighbors, while the lower neighbors are both in S (provided that they are in λ at all) and so contain symbols that are smaller than a. Since u is injective, $u(\omega)$ appears only once, so s contains the same number of symbols as u. It follows that s is connected to u and hence to t as well by transitivity. Since $R_\lambda(s) \leq R_\lambda(t) - 1$, this completes the proof. □

Proof of Proposition 4.9 in the generic case. In the generic case, the non-repeating tableaux are precisely the injective tableaux on T_r. Each such tableau contains $g-1-\dim V^r(\Gamma,\varphi)$ symbols; since $\dim V^r(\Gamma,\varphi) \geq 1$, we apply Lemma 4.12 to find that any two are connected. By Remark 4.11, we are done. □

Example 4.13. In Fig. 4.2, we outline the proof of Lemma 4.12 in the case that $\lambda = T_4$ and $g \geq 12$. We begin with the tableau on the top right and terminate at the standard increasing tableau, each step either a cycle or a swap. The set of boxes S at each step is colored blue. The first step cycles out 3 using 11; notice that each symbol greater than or equal to 3 is incremented. The second step swaps 3 into $(2,1)$, thereby removing 6 from the tableau. We continue cycling and swapping as appropriate until every symbol is in the correct position according to the order.

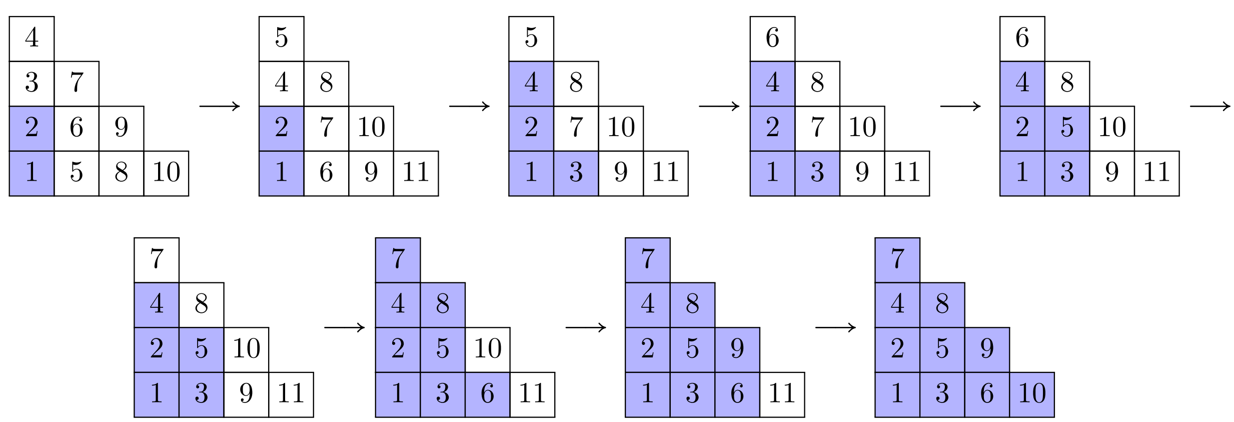

Figure 4.2. Example application of the algorithm from the proof of Lemma 4.12.

Lemma 4.14. *If $\dim V^r(\Gamma,\varphi) \geq 1$, then for any gonality k, any two tableaux of type (g,r,k) non-repeating in the horizontal strip μ_0 are connected.*

The proof is very similar to that of Lemma 4.12, so we highlight only the major differences. The key idea is that any tableau non-repeating in μ_0 is uniquely determined by its restriction to μ_0. Hence, we may naturally define R_{μ_0} on such tableaux, and we take as our base point $\bar{t}$ the unique tableau which extends the standard increasing tableau on μ_0. We also make use of the following essential observations.

Remark 4.15. Suppose that t is non-repeating in a strip μ and that the map s is produced from t by swapping a in for b. Then s trivially satisfies the repeating conditions. Meanwhile, by Corollary 4.4, we know that b appears exactly once in $t|_\mu$; provided that s satisfies the tableau and gluing conditions at the box containing b, it is non-repeating in μ. Now, the gluing condition is satisfied whenever the repeating and tableau conditions are (everywhere) satisfied. Hence, if s is produced from t by cycling out b using a, then s is non-repeating in μ.

Proof of Lemma 4.14. As before, we induct on the values of R_{μ_0}, noting that $\bar{t}$ is the unique tableau satisfying $R_{\mu_0}(\bar{t}) = 0$. Let t be a tableau non-repeating in μ_0 and $\omega \in \mu_0$ the smallest box such that $t(\omega) \neq Q_{\mu_0}(\omega)$; much as before, denote $Q_{\mu_0}(\omega)$ by a, and let S be the set of boxes in μ_0 smaller than ω. First, cycle out a using b, where $b \in \{a, a+1, \ldots, g-1\}$ does not appear in t (and exists by the assumption that $\dim V^r(\Gamma,\varphi) \geq 1$). The resulting tableau u is connected to t and, by Remark 4.15, is non-repeating in μ_0.

Here we diverge from the proof of Lemma 4.12: instead of merely swapping a into the box ω, we swap it in for the symbol $u(\omega)$; call the result s. By Remark 4.15, to show that s is non-repeating in μ_0, it suffices to check the tableau and gluing conditions at ω. The former follows as in the proof of Lemma 4.12, since $a < u(\omega)$ and the lower neighbors (if they exist) are in S. The latter is trickier. Because μ_0 is horizontal, every leftmost box is of the form (x,l) for some x. The gluing condition on u says that $u(x,l) < u(x+l-\epsilon,1)$. If $\omega = (x,l)$, then the gluing condition is satisfied for s since $a < u(x,l)$. If $\omega = (x+l-\epsilon,1)$, then $(x,l) < \omega$. Therefore, $(x,l) \in S$ and contains a symbol smaller than a. In any other case, the gluing condition is trivially satisfied.

Hence, s is a tableau non-repeating in μ_0 that is connected to t and satisfies $R_{\mu_0}(s) \leq R_{\mu_0}(t) - 1$. ☐

Proof of Proposition 4.9 in the even case. The non-repeating tableaux for even values of k are each non-repeating in μ_0 in particular (by the discussion following Eq. (4.1)). By Lemma 4.14 and Remark 4.11, we are done. ☐

The odd case is more difficult than the even case because we cannot just consider the horizontal strip: by the correspondence in Eq. (4.1), each of the 2^{r-l} strips determines a distinct set of maximal cells of $V^r(\Gamma,\varphi)$. Therefore, we introduce a height function H and—as we did for R_λ—show that any tableau is connected to another with a lower H value. Given a tableau t of odd torsion which is non-repeating in μ, define $H(t)$ to be the second coordinate of the r-th leftmost box of μ. Note that H is well-defined by Proposition 4.6. Moreover, $H(t) = l$ if and only if $\mu = \mu_0$.

We again take $\bar{t}$ to be unique non-repeating tableau that extends the standard increasing tableau on μ_0. To simplify our notation, we introduce the unit vectors $\hat{x}$ and $\hat{y}$ to describe boxes relative to other boxes. For example, if $\omega = (x,y)$, then $\omega + \hat{x} = (x+1,y)$ and $\omega - 2\hat{y} = (x, y-2)$.

Proof of Proposition 4.9 in the odd case. We shall prove it by induction on the values of H. First, suppose that $H(t) = l$. Then t is non-repeating in μ_0, hence connected to $\bar{t}$ by Lemma 4.14. For the induction step, suppose that t is non-repeating in a strip μ and that every tableau s satisfying $H(s) < H(t)$ is connected to $\bar{t}$. Denote by (x,y) the unique box in μ for which $y = H(t)$ and $(x-1,y) \notin \mu$. Denote its anti-diagonal by A_q, and let $n = r - q$. Define $\psi_i := (x+i, y)$ for each $i \in \{0, 1, \ldots, n\}$. Since $H(t) = y$, ψ_i is the $(q+i)$-th leftmost box for all i, and in particular, ψ_n is the r-th leftmost box.

Our goal is to show that t is connected to a tableau s that is non-repeating in ν, where ν is the strip that agrees with μ up to A_{q-1} but has every subsequent leftmost box east of the previous one. (In particular, the q-th leftmost box of ν is $(x+1, y-1)$ rather than (x, y).) Then we will be done, since $H(s) = H(t) - 1$.

Preliminary observations. Notice that each ψ_i is in the left component of $T_r \smallsetminus \nu$. Hence, in order for s to satisfy the left repeating condition, we need $s(\psi_i) = s(\omega_{i,0})$, where $\omega_{i,0} := (x+i+l-1, y-l)$. Note that $\omega_{0,0}$ is the $(q-1)$-th rightmost box of μ, while for each $i \geq 1$, $\omega_{i,0}$ is in the right component of $T_r \smallsetminus \mu$. Hence, for each $i \geq 1$, the right repeating condition yields $t(\omega_{i,0}) = t(\psi_i - \hat{x} - \hat{y})$. More generally, for each $j \geq 0$ we define $\omega_{i,j} = (x+i+l-1+jl, y-l-j(l-1))$. Then for all i and $j \geq 1$, $\omega_{i,j}$ is in the right component of both $T_r \smallsetminus \mu$ and $T_r \smallsetminus \nu$, so we have $t(\omega_{i,j}) = t(\omega_{i,0})$ and $s(\omega_{i,j}) = s(\omega_{i,0})$ by the right repeating condition. See Fig. 4.3 for a schematic diagram of our notations.

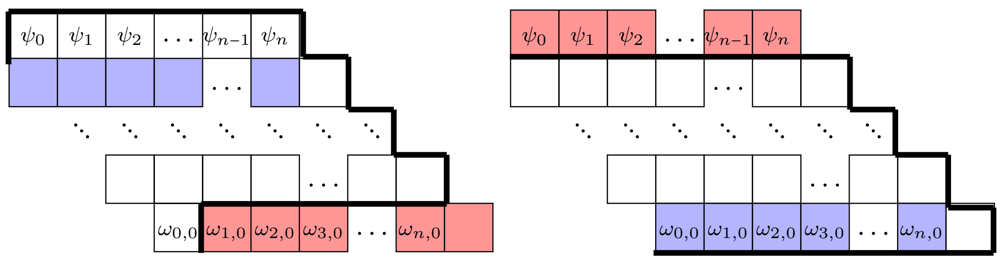

Figure 4.3. The left and right tableaux are restrictions of t and s respectively to the same subset of T_r. The rightmost anti-diagonal in the subset is A_r. The bold lines outline the strips μ and ν respectively. The labels are box names, not symbols; only the ψ_i and $\omega_{i,0}$ are shown. Symbols in the red boxes outside of each strip are copied from the respective blue boxes within the strip.

First attempt, using t. To go from t to s, we could try to replace the symbol in $\omega_{i,j}$ with the symbol in ψ_i for each i and j, and leave all other symbols unchanged. Unfortunately, the tableau condition would necessarily fail at $\omega_{n,0}$, which lies on A_{r-1}. Indeed, we have $t(\psi_n) > t(\psi_n - \hat{y}) = t(\omega_{n,0} + \hat{x})$. It is also possible that $t(\psi_n) > t(\omega_{n,0} + \hat{y})$. We shall modify t so that these issues are avoided; in particular, we will put the two largest symbols, $g-2$ and $g-1$, into $\omega_{n,0} + \hat{y}$ and $\omega_{n,0} + \hat{x}$, respectively.

Construction of u. Cycle out $g-2$ using any symbol that does not appear in t. By Remark 4.15, the resulting tableau is still non-repeating in μ. The only symbol greater than $g-2$ is $g-1$, and if that appears in the tableau, then it appears on A_r (since, being the largest symbol, it cannot be in a box with upper neighbors). Hence, when we swap $g-2$ into $\omega_{n,0} + \hat{y}$, we know that the tableau condition is satisfied since it is necessarily larger than both of its lower neighbors' symbols. Moreover, the symbol it replaces is unique in the tableau; indeed, any symbol in $A_r \cap \mu$ in a tableau non-repeating in μ is unique. Thus, the resulting tableau is still non-repeating in μ. Next, cycle out $g-1$ (using whatever symbol is free) and call the resulting tableau v. Again, v is non-repeating in μ and connected to t by the previous operations.

Now we swap $g-1$ into $\omega_{n,0} + \hat{x}$ to produce a k-uniform displacement tableau u. Observe that u is not non-repeating in μ; indeed, $\omega_{n,0} + \hat{x}$ is in the right component of $T_r \smallsetminus \mu$, and the box it should be repeated from, $\psi_n - \hat{y}$, contains a symbol other than $g-1$ (which we had cycled out). It is important to note that, as a result, the codimension of $P(u)$ relative to $V^r(\Gamma, \varphi)$ is 1. Thus, we need to show that u, which is dominated by v, is also dominated by some tableau s non-repeating in ν; this will imply that t and s are adjacent.

Second attempt, using u. We now construct s from u in the same way that we attempted to construct s from t. Precisely, for each i and j, we define $s(\omega_{i,j}) = u(\psi_i)$ and let s coincide with u everywhere else. It is not hard to see that $s|_\nu$ is injective and so trivially satisfies the displacement condition. Moreover, the left and right repeating conditions are satisfied since $s(\psi_i) = s(\omega_{i,0})$ and $s(\omega_{i,j}) = s(\omega_{i,j-1})$ for each $j \geq 1$. Moreover, the leftmost boxes of ν that are not leftmost boxes of μ are of the form $\psi_i - \hat{y}$ for each i. The gluing condition on s then follows by the tableau condition on u, since the $\omega_{i,0}$ are the corresponding rightmost boxes of ν; explicitly,

$$s(\psi_i - \hat{y}) = u(\psi_i - \hat{y}) < u(\psi_i) = s(\psi_i) = s(\omega_{i,0}).$$

It remains to show that $s|_\nu$ satisfies the tableau condition. This amounts to checking it at each $\omega_{i,0}$. The south neighbor of $\omega_{i,0}$ is not in ν, so we may safely ignore it. Moreover, since the entire block of symbols in the set $\{\omega_{i,0}\}$ is copied from $\{\psi_i\}$, we know that the condition is satisfied between each pair $\{\omega_{i,0}, \omega_{i+1,0}\}$. For $i \neq n$, we need to check that the north neighboring symbol is larger:

$$s(\omega_{i,0}) = u(\psi_i) < u(\psi_i + \hat{y}) = u(\omega_{i,0} + \hat{y}) = s(\omega_{i,0} + \hat{y}).$$

For $i = 0$, the west must be smaller:

$$s(\omega_{0,0}) = u(\psi_0) > u(\psi_0 - \hat{y}) = u(\omega_{1,0}) > u(\omega_{0,0} - \hat{x}) = s(\omega_{0,0} - \hat{x}).$$

Finally, for $i = n$, the upper neighbors' symbols are $g-1$ and $g-2$, the two largest symbols in the tableau. Hence, s is non-repeating in ν. Moreover, it clearly dominates u, which completes the proof. ☐

The next example demonstrates the algorithm that lowers the height of the strip by one.

Example 4.16. Consider the first tableau t in Fig. 4.4, where $g = 23$, $r = 8$, and $k = 5$. We color the strip μ blue. In our example, we note that $H(\mu) = 5$, so $\psi_0 = (3, 5)$. The first step of the algorithm is to cycle out $g - 2 = 21$ using 22 and then swap it into $\omega_{n,0} + \hat{y} = (6, 3)$. (These two operations do not change the tableau, since 21 was already in the correct box.) The next step is to cycle out $g - 1 = 22$ (this does nothing) and swap it into $\omega_{n,0} + \hat{x} = (7, 2)$, thereby producing the second tableau u. Since u is not minimal, we do not color any strip. For the last step, we copy the symbols from the boxes ψ_0 and ψ_1 (17 and 19 respectively) into the boxes $\omega_{0,0}$ and $\omega_{1,0}$. This produces the third tableau s. We color the strip ν; observe that the height has decreased by one. Another iteration of this process would yield a tableau that is non-repeating in the horizontal strip.

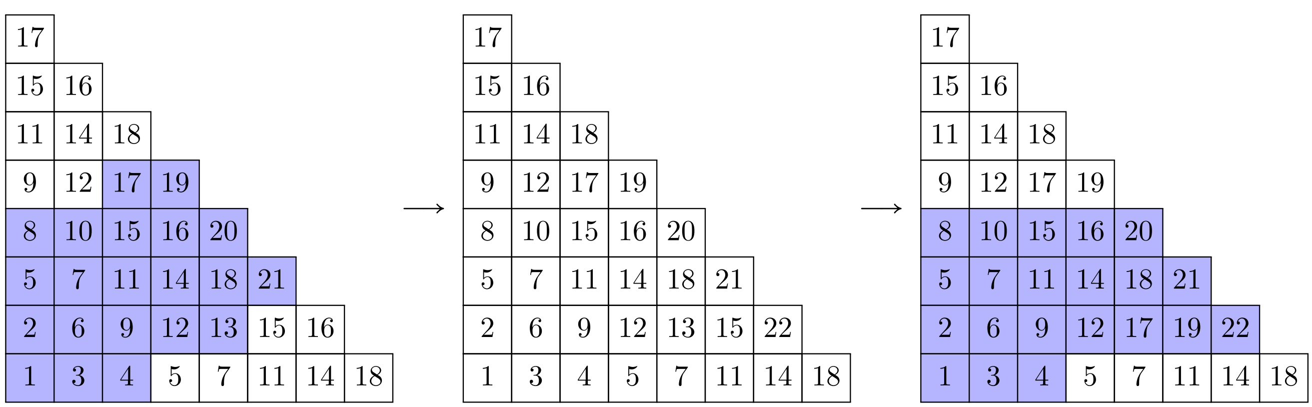

Figure 4.4

5. Enumerative properties

Now that we have established a few general facts about Prym–Brill–Noether loci, we begin to look at some of their enumerative properties. We start by counting the number of divisors in 0-dimensional loci before examining the 1-dimensional case.

5.1. **Cardinality of finite Prym–Brill–Noether loci.** In this section, we fix the parameters g, r, and k so that $g-1 = n(r,k)$. By Theorem A, this condition ensures that $\dim V^r(\Gamma,\varphi) = 0$, so every point of $V^r(\Gamma,\varphi)$ is in itself a maximal cell.

Recall that, by Eq. (4.2), the maximal cells of the Prym–Brill–Noether locus are in bijection with strip tableaux of the corresponding type. It is clear that there are finitely many strips and finitely many ways to fill each one, so the cardinality of $V^r(\Gamma,\varphi)$, which we denote by $C(r,k)$, is finite. This number has been computed [LU19, Corollary 6.1.5] for generic edge lengths (where $k = 0$ by convention) and equivalently for $k > 2r-2$ using the hook-length formula. We now compute it in the case that k is even and at most $2r-2$.

Proposition 5.1. *For even $k \le 2r-2$, the number of divisor classes in the 0-dimensional locus is*

$$C(r,k) = n! \sum \begin{vmatrix} \frac{1}{(r+\alpha_1 k)!} & \frac{1}{(r-2+\alpha_2 k)!} & \cdots & \frac{1}{(r-k+2+\alpha_l k)!} \\ \frac{1}{(r+1+\alpha_1 k)!} & \frac{1}{(r-1+\alpha_2 k)!} & \cdots & \frac{1}{(r-k+3+\alpha_l k)!} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{1}{(r+l-1+\alpha_1 k)!} & \frac{1}{(r+l-3+\alpha_2 k)!} & \cdots & \frac{1}{(r-l+1+\alpha_l k)!} \end{vmatrix} \tag{5.1}$$

where $n = n(r,k) = g-1$ and the sum is taken over all l-tuples $(\alpha_i)_{i=1}^l$ for which $\alpha_i \in \mathbb{Z}$ and $\sum_{i=1}^l \alpha_i = 0$.

Proof. Using the correspondence in Eq. (4.2) and the definition of a strip tableau for k even, we find that $C(r,k)$ equals the number of ways to fill out the horizontal strip μ_0 of length r and width l using each symbol in $[g-1]$ exactly once while adhering to the tableau and gluing conditions. We aim to construct a bijection between these strip tableaux and lattice paths in $\mathbb{Z}^l$ joining $(l, l-1, \dots, 1)$ to $(r+l, r+l-2, \dots, r-l+2)$ such that each step is in a positive unit direction and every point $(z_1, z_2, \dots, z_l)$ satisfies the constraints $z_1 > z_2 > \dots > z_l > z_1 - k$. Once we have this, we are done: by [Bón15, Theorem 10.18.6], the number of such lattice paths is exactly given by Eq. (5.1).

Given a strip tableau t, we obtain a lattice path in the following way. The path begins at $(l, l-1, \dots, 1)$. Suppose that the first $a-1$ steps in the path have been defined and satisfy the conditions above. Identify the unique box $(x,y) \in \mu_0$ containing the symbol a. Then define the a-th step of the path to be a positive unit step in the y-th coordinate.

By the tableau condition, there are precisely $x-1$ values of $i \in \{\dots, -1, 0, 1, \dots, x-1\}$ for which $t(x-i, y) < a$ (namely, the positive ones); this implies that the a-th step of the lattice path is the x-th step in the y-th coordinate. Applying similar reasoning to the boxes $(x, y-1)$ and $(x, y+1)$, which contain symbols less than a and greater than a respectively, we may conclude that, by the a-th step, at least x steps have been taken in the $(y-1)$-th coordinate and at most $x-1$ steps have been taken in the $(y+1)$-th coordinate. Since the initial point of the path satisfies $z_{y-1} > z_y > z_{y+1}$, it follows that the a-th point does as well. By induction, every point in the path satisfies the constraints $z_1 > z_2 > \dots > z_l$.

Next, the gluing condition forces $t(x+l, 1) > t(x, l)$ for each x. On the lattice path, this means that the $(x+l)$-th step in the first coordinate must come after the x-th step in the l-th coordinate. At the starting point, the first coordinate is already greater by $l-1$ compared to the l-th coordinate, and the gluing condition allows this gap to grow to at most $k-1$, giving us the final inequality $z_l > z_1 - k$.

Counting the number of boxes in each row of the strip demonstrates that the endpoint is $(r+l, r+l-2, \dots, r-l+2)$, as expected. Hence, the procedure defined above, in fact, yields a lattice path of the desired form. Conversely, given a lattice path, we may reverse the construction to get a strip tableau: if the a-th step in the lattice path is the x-th step in the y-th coordinate, then the symbol a goes into box (x,y). The first $l-1$ inequalities on the coordinates verify the tableau condition and the final inequality verifies the gluing condition. ☐

For convenience, we include $C(r,k)$ for small values of r and k in Fig. 5.1.

r	$C(r,0)$	$C(r,2)$	$C(r,4)$	$C(r,6)$	$C(r,8)$
1	1	1	1	1	1
2	2	1	2	2	2
3	16	1	4	16	16
4	768	1	8	128	768
5	292864	1	16	1024	35480
6	1100742656	1	32	8178	1671168

Figure 5.1. $C(r,k)$ for several values of r and k.

Example 5.2. For low values of k, we may exhibit all the horizontal strips directly. For instance, we claim that $C(r,2) = 1$ for every r. Indeed, the Prym tableaux with minimal codimension are uniquely determined by the bottom row, and the only way to fill out the row is by using the symbols 1 through $g-1$ in increasing order.

To see that $C(r,4) = 2^{r-1}$, we note first that the tableau condition forces the symbol 1 to be placed into the box $(1,1)$. The tableau and gluing conditions together force each subsequent pair of symbols $\{2n-2, 2n-1\}$ for $n \in \{2,3,\dots,r\}$ to be placed into $A_n \cap \mu_0$, which contains two boxes. Thus, each value of n yields 2 possibilities for symbol placement. This choice is independent of previous choices, so the total number of possibilities is 2^{r-1}.

When $g-1 > n(r,k)$, the Prym–Brill–Noether locus has positive dimension. Its maximal cells still correspond to strip tableaux, but in this case, each tableau uses only $n(r,k)$ of the $g-1$ available symbols. Keeping in mind that $C(r,k)$ counts the number of strip tableaux when the set of symbols is fixed and every symbol must be used, we easily obtain the following result.

Proposition 5.3. *The number of maximal cells of $V^r(\Gamma,\varphi)$ equals*

$$C(r,k) \cdot \binom{g-1}{n(r,k)}.$$

$C(r,k)$ remains unknown in the case that k is odd and at most $2r-3$. The difficulty lies in counting the number of ways to fill strips that are not horizontal, which is not in general equal to $C(r,k)$.

5.2. **First Betti number of 1-dimensional loci.** We now choose g, r, and k so that $g-1 = n(r,k)+1$. Then $\dim V^r(\Gamma,\varphi) = 1$. In particular, the Prym–Brill–Noether locus is a metric graph that consists of finitely many circles. Each circle corresponds to a strip tableau that uses all but a single symbol a, which we call the *free* symbol. The circle consists of divisors with a fixed chip on each loop, except for $\tilde{\gamma}_a$ and $\tilde{\gamma}_{2g-a}$. We refer to these loops as free as well.

The only way that two different circles intersect is if they have different free loops and agree on the fixed location of the chips on the other loops. We also see that if two circles intersect, then they do so at exactly one point. It follows that $V^r(\Gamma,\varphi)$ has a 4-regular model. Since the graph is 4-valent, the number of edges e equals twice the number of vertices v. The Betti number is therefore

$$e - v + 1 = 2v - v + 1 = v + 1.$$

In terms of strip tableaux, t and t' with free symbols a and a' respectively give rise to non-trivially intersecting circles precisely when $a \neq a'$ and t' is obtained from t by swapping a in for a'.

The rest of the section is devoted to calculating the Betti number of this graph in the generic case and when k is 2 or 4. We begin with the generic case.

Theorem D. *Let* $\varphi\colon \widetilde{\Gamma} \to \Gamma$ *be a folded chain of loops with generic edge length such that* $\dim V^r(\Gamma,\varphi) = 1$. *Then the first Betti number of* $V^r(\Gamma,\varphi)$ *is given by*

$$\frac{r \cdot C(r,0) \cdot \left(\binom{r+1}{2} + 1\right)}{2} + 1. \tag{1.3}$$

Proof. Since the edge lengths are generic, we have a correspondence between maximal cells of $V^r(\Gamma,\varphi)$ and injective tableaux defined on T_r. We know from Theorem A that $n(r,k) = \binom{r+1}{2}$. Since $g-1 = n(r,k)+1$, it follows from Proposition 5.3 that the number of maximal cells is $C(r,0)\cdot\left(\binom{r+1}{2}+1\right)$.

The vertices of the 4-regular model of $V^r(\Gamma,\varphi)$ are precisely the intersection points between circles. Let E^{T_r} denote the average number of intersection points on each circle. Then the total number of vertices is given by $\frac{1}{2}E^{T_r}\cdot C(r,0)\cdot\left(\binom{r+1}{2}+1\right)$. Notice the similarity to the first term in Eq. (1.3); since the Betti number is given by $v+1$ (where v is the number of vertices), it suffices to show that $E^{T_r} = r$.

For any skew shape λ, denote by f^λ the number of distinct injective tableaux defined on λ that take values in $[n]$, where n is the number of boxes in λ. From [CLMPTiB18, Theorem 2.9], it follows that the average number of intersection points per circle is

$$E^\lambda := 2\left(r + \sum_{i=1}^{r} \frac{r-i}{n+1}\cdot\frac{f^{{}^i\lambda}}{f^\lambda} - \sum_{i=1}^{r}\frac{r+1-i}{n+1}\cdot\frac{f^{\lambda^i}}{f^\lambda}\right), \tag{5.2}$$

where the terms ${}^i\lambda$ and λ^i describe the tableaux obtained by adding a box to the left or the right respectively in the i-th row. Taking $\lambda = T_r$, f^λ reduces to $C(r,0)$.

Consider the first summation in Eq. (5.2). When $i = r$, the term is clearly 0. Provided that $i \neq r$, the resulting shape of ${}^i\lambda$ is not a skew tableau, so $f^{{}^i\lambda} = 0$. Thus, this summation vanishes.

1				
3	1			
5	3	1		
7	5	3	1	
9	7	5	3	1

1				
3	1			
6	4	2	1	
7	5	3	2	
9	7	5	4	1

Figure 5.2. Hook lengths of each box in T_5 and in $(T_5)^3$.

Next, we look at the second summation. We need to enumerate the tableaux obtained by adding a box to the end of each row of T_r. Each f^{λ^i} can be computed using the hook length formula. We note that in T_r, fixing $q \in [r]$, the boxes in A_q each have hook length $2(r-q)+1$. When a box is added, the hook length of every box in its row and column increases by 1, while all other hook lengths remain the same. (See Fig. 5.2 for an example.) Thus, the fraction $f^{\lambda^i}/(n+1)f^\lambda$ simplifies down to the ratio of the differing hook lengths:

$$\frac{f^{\lambda^i}}{(n+1)f^\lambda} = \frac{(n+1)!\prod h_\lambda(i,j)}{(n+1)n!\prod h_{\lambda^i}(i,j)} = \frac{(2(r-i)+1)!!(2i-3)!!}{(2(r-i+1))!!(2i-2)!!},$$

where $h_\lambda(i,j)$ is the hook length of box (i,j) in shape λ and $(-1)!!$ is defined as 1. We observe that

$$\frac{(2i-3)!!}{(2i-2)!!} = \frac{(2i-3)!!}{2^{i-1}(i-1)!} = \frac{(2i-2)!}{2^{i-1}(i-1)!(2i-2)!!} = \frac{(2i-2)!}{2^{i-1}(i-1)!2^{i-1}(i-1)!} = \frac{\binom{2i-2}{i-1}}{2^{2(i-1)}}.$$

A similar calculation gives us

$$\frac{(2(r-i)+1)!!}{(2(r-i+1))!!} = \frac{\binom{2(r-i+1)}{r-i+1}}{2^{2(r-i+1)}},$$

so

$$\frac{f^{\lambda^i}}{(n+1)f^\lambda} = \frac{\binom{2i-2}{i-1}\binom{2(r-i+1)}{r-i+1}}{2^{2r}}.$$

Setting $j = r - i + 1$, the sum becomes

$$\sum_{i=1}^{r}(r-i+1)\cdot\frac{\binom{2i-2}{i-1}\binom{2(r-i+1)}{r-i+1}}{2^{2r}} = \sum_{j=1}^{r} j\cdot\frac{\binom{2j}{j}\binom{2(r-j)}{r-j}}{2^{2r}}.$$

For each j we have

$$j\cdot\frac{\binom{2j}{j}\binom{2(r-j)}{r-j}}{2^{2r}} + (r-j)\cdot\frac{\binom{2(r-j)}{(r-j)}\binom{2(r-(r-j))}{r-(r-j)}}{2^{2r}} = r\cdot\frac{\binom{2j}{j}\binom{2(r-j)}{r-j}}{2^{2r}}.$$

Thus, grouping j and $r-j$ together, and adding $j = 0$ to match $j = r$, we get

$$\sum_{j=0}^{r} j\cdot\frac{\binom{2j}{j}\binom{2(r-j)}{r-j}}{2^{2r}} = \frac{r}{2}\sum_{j=0}^{r}\frac{\binom{2j}{j}\binom{2(r-j)}{r-j}}{2^{2r}}.$$

Finally, by [Sve84], the sum is equal to 1, so the entire term is equal to $\frac{r}{2}$. Plugging this value back into Eq. (5.2), we conclude that the average number of vertices on each circle is r, as desired. □

We conclude the paper by computing the first Betti number of the Prym–Brill–Noether curve for low even gonality.

Proposition 5.4. *Suppose that $k = 2$ and that the Prym–Brill–Noether locus is 1-dimensional. Then it contains $r + 1 = g - 1$ circles, and has first Betti number $r + 1$.*

Proof. In this case, each tableau contains $g - 2$ symbols and is determined by the bottom $1 \times r$ rectangle; the positions of the symbols in the strip are determined after choosing a symbol to leave out. When 1 or $g - 1$ is the free symbol, it may only swap into the first or last box in the strip, respectively, so the corresponding circle only has a single vertex. If any other symbol m is left out, it can swap with either the symbol $m - 1$ or $m + 1$, so the corresponding circle has two vertices. Thus, the locus is a chain of $r + 1 = g - 1$ circles wedged together, which has Betti number of $r + 1$. □

The last case that we deal with is $k = 4$.

Proposition 5.5. *Suppose that $k = 4$ and that the Prym–Brill–Noether locus is 1-dimensional. Then it has the following structure.*

(i) The circles corresponding to the free symbol 1 have a single vertex.
(ii) The circles corresponding to any other odd free symbol have two vertices.
(iii) The circles corresponding to the free symbol 2 have three vertices.
(iv) The circles corresponding to the free symbol $2r$ have two vertices.
(v) The circles corresponding to any other even free symbol have four vertices.

The graph has $2^{r-1}\cdot 2r$ circles and first Betti number $2^{r-1}(3r-2)+1$.

Proof. Since $k = 4$, the genus and rank are related by $g = 2r + 1$. From the gluing condition, it follows that the pair of symbols in each of the boxes $(m+1, 1)$ and $(m, 2)$ is strictly bigger than the pair of symbols in $(m, 1)$ and $(m-1, 2)$ (see Fig. 5.3). In total, for any missing symbol there are 2^{r-1} tableaux (see Example 5.2), giving rise to $2^{r-1}\cdot(2r)$ circles.

Next, we calculate the number of vertices in the graph, by finding the number of ways of swapping in a free symbol. If the free symbol is 1, it may only be swapped with 2, which must be in the bottom left corner. Therefore, any circle corresponding to a tableau with missing symbol 1 has exactly one vertex. Similarly, a missing 2 may only be swapped for the first three boxes, and a missing $2r$ may only be swapped for the two rightmost boxes.

Suppose that the strip is missing an even symbol $2 < 2m < 2r$. Then the symbols in the boxes $(m+1, 1)$ and $(m, 2)$ are $2m - 2$ and $2m - 1$, and the symbols in the boxes to to right are $2m + 1$ and $2m + 2$. The symbol $2m$ may be swapped in for any of them. If, on the other hand, the strip is missing the odd symbol $2 < 2m + 1 < 2r$, then the boxes $(m + 1, 1)$ and $(m, 2)$ are $2m$ and $2m + 2$, and the symbols to the right are $2m + 3$ and $2m + 4$. Our symbol $2m + 1$ may only be swapped in for of $2m$ or $2m + 2$ without violating either the tableau or gluing condition.

Altogether, we see that there are

$$\frac{(4(r-2) + 2(r-1) + 1 + 3 + 2) \cdot 2^{r-1}}{2} = 2^{r-1}(3r - 2)$$

vertices, so the Betti number is $2^r \cdot (3r - 2) + 1$. □

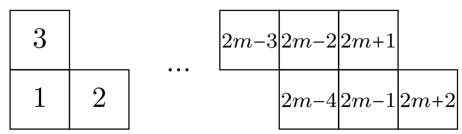

Figure 5.3. The bottom 2 rows of a tableau. The symbol $2m$ is missing, and may be swapped in four different boxes.

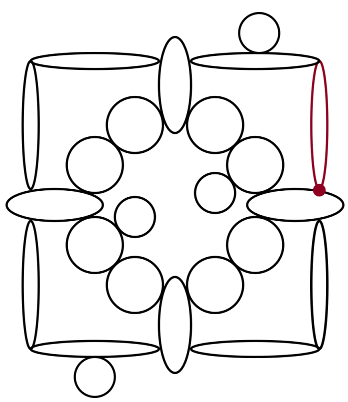

(a) V^r for $(g, k, r) = (7, 4, 3)$.

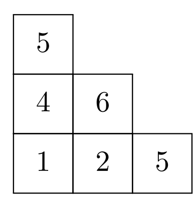

(b) The tableau corresponding to the highlighted circle in the locus.

Figure 5.4

Example 5.6. Let $g = 7$, $k = 4$, and $r = 3$. The Prym–Brill–Noether locus is depicted in Fig. 5.4a. In this case, $n(r, k) = 5$, and $C(r, k) = 2^{3-1} = 4$. Proposition 5.3 shows that the locus consists of $4 \cdot \binom{6}{5} = 24$ circles, and Proposition 5.5 implies that the Betti number is $4(3(3) - 2) + 1 = 29$.

Each of the four circles with 4 vertices corresponds to a tableau with free symbol 4. The four circles with only a single vertex correspond to the free symbol 1, and the circles they intersect with correspond to the free symbol 2. The highlighted circle in red is the circle corresponding to the tableau on the right, which has free symbol 3. The highlighted point of intersection corresponds to swapping the symbols 4 and 3.

References

[ABBR15] O. Amini, M. Baker, E. Brugallé, and J. Rabinoff. Lifting harmonic morphisms I: metrized complexes and Berkovich skeleta. *Res. Math. Sci.*, 2:Art. 7, 67, 2015.

[Ara12] M. Arap. Algebraic cycles on prym varieties. *Math. Ann.*, 353(3):707–726, 2012.

[Bak08] M. Baker. Specialization of linear systems from curves to graphs. *Algebra Number Theory*, 2(6):613–653, 2008.

[Ber87] A. Bertram. An existence theorem for Prym special divisors. *Invent. Math.*, 90(3):669–671, 1987.

[Bón15] M. Bóna. *Handbook of enumerative combinatorics*, volume 87. CRC Press, 2015.

[CDPR12] F. Cools, J. Draisma, S. Payne, and E. Robeva. A tropical proof of the Brill–Noether theorem. *Adv. Math.*, 230(2):759–776, 2012.

[CLMPTiB18] M. Chan, A. López Martín, N. Pflueger, and M. Teixidor i Bigas. Genera of Brill–Noether curves and staircase paths in young tableaux. *Trans. Amer. Math. Soc.*, 370(5):3405–3439, 2018.

[CPJ19] K. Cook-Powell and D. Jensen. Components of Brill–Noether loci for curves with fixed gonality. *arXiv preprint:1907.08366*, 2019.

[FL81] W. Fulton and R. Lazarsfeld. On the connectedness of degeneracy loci and special divisors. *Acta Math.*, 146(3-4):271, 1981.

[GH80] P. Griffiths and J. Harris. On the variety of special linear systems on a general algebraic curve. *Duke Math. J.*, 47(1):233–272, 1980.

[Gie82] D. Gieseker. Stable curves and special divisors: Petri's conjecture. *Invent. Math.*, 66(2):251, 1982.

[Gub07] W. Gubler. Tropical varieties for non-Archimedean analytic spaces. *Invent. Math.*, 169(2):321–376, 2007.

[HMY12] C. Haase, G. Musiker, and J. Yu. Linear systems on tropical curves. *Math. Z.*, 270(3-4):1111–1140, 2012.

[JL18] D. Jensen and Y. Len. Tropicalization of theta characteristics, double covers, and Prym varieties. *Selecta Math.*, 24(2):1391–1410, 2018.

[JP14] D. Jensen and S. Payne. Tropical independence I: Shapes of divisors and a proof of the Gieseker-Petri theorem. *Algebra Number Theory*, 8(9):2043–2066, 2014.

[JP16] D. Jensen and S. Payne. Tropical independence ii: The maximal rank conjecture for quadrics. *Algebra Number Theory*, 10(8):1601–1640, 2016.

[JR17] D. Jensen and D. Ranganathan. Brill–noether theory for curves of a fixed gonality. *Preprint arXiv:1701.06579*, 2017.

[Lar19] H. Larson. A refined Brill–Noether theory over hurwitz spaces. *arXiv preprint:1907.08597*, 2019.

[Len14] Y. Len. The Brill–Noether rank of a tropical curve. *J. Algebraic Combin.*, 40(3):841–860, 2014.

[LU19] Y. Len and M. Ulirsch. Skeletons of Prym varieties and Brill–Noether theory. *arXiv preprint:1902.09410*, 2019.

[LUZ19] Y. Len, M. Ulirsch, and D. Zakharov. Abelian tropical covers. *arXiv:1906.04215*, 2019.

[Pfl17a] N. Pflueger. Brill–Noether varieties of k-gonal curves. *Adv. Math.*, 312:46–63, 2017.

[Pfl17b] N. Pflueger. Special divisors on marked chains of cycles. *J. Combin. Theory Ser. A*, 150:182–207, 2017.

[Sve84] M. Sved. Counting and recounting: The aftermath. *Math. Intelligencer*, 6:44–46, 1984.

[Wel85] G.E. Welters. A theorem of Gieseker–Petri type for Prym varieties. *Ann. Sci. École Norm. Sup.*, 18(4):671–683, 1985.

Department of Mathematics, Brown University, Providence, RI 02912, USA
Email address: steven_creech@brown.edu

School of Mathematics and Statistics, University of St Andrews, St Andrews, KY16 9SS, UK
Email address: yoav.len@st-andrews.ac.uk

Department of Mathematics, University of Washington, Seattle, WA 98195, USA
Email address: critter1@uw.edu

School of Mathematics, Georgia Institue of Technology, Atlanta, GA 30332-0160, USA
Email address: dwu96@gatech.edu