

Mochizuki's indigenous bundles and Ehrhart polynomials

Fu Liu · Brian Osserman

Received: 12 April 2004 / Revised: 26 January 2005 / Accepted: 7 July 2005
© Springer Science + Business Media, Inc. 2006

Abstract Mochizuki's work on torally indigenous bundles [1] yields combinatorial identities by degenerating to different curves of the same genus. We rephrase these identities in combinatorial language and strengthen them, giving relations between Ehrhart quasi-polynomials of different polytopes. We then apply the theory of Ehrhart quasi-polynomials to conclude that the number of dormant torally indigenous bundles on a general curve of a given type is expressed as a polynomial in the characteristic of the base field. In particular, we conclude the same for the number vector bundles of rank two and trivial determinant whose Frobenius-pullbacks are maximally unstable, as well as self-maps of the projective line with prescribed ramification.

Keywords Mochizuki · Indigenous bundles · Ehrhart polynomials · Identities

1. Introduction

In this paper, we bring together work of Mochizuki in algebraic geometry and the theory of Ehrhart quasi-polynomials in combinatorics, obtaining results in both fields. There is already a combinatorial result implicit in [1]; we strengthen it and state it explicitly in terms of familiar combinatorial objects: namely, we obtain in Theorem 2.4 below an infinite family of identities of numbers of lattice points in different polytopes. Conversely, we use the theory of Ehrhart quasi-polynomials to conclude in Theorem 2.1 below that the number of dormant torally indigenous bundles of Mochizuki's theory (see the following section for references and discussion, including relationships to certain rational functions and Frobenius-destabilized vector bundles) may be expressed as a polynomial in the characteristic of the base field. Both these phenomena are observed (but not pursued) in [1, p. 46] in a slightly different case, so

The second author was supported by a fellowship from the Japan Society for the Promotion of Science during the preparation of this paper.

F. Liu (✉)
Massachusetts Institute of Technology

B. Osserman
University of California, Berkeley

the present paper may be considered a more complete exploration of these phenomena in the case of dormant torally indigenous bundles.

In fact, the ideas of our paper make no use of the precise definition of dormant torally indigenous bundles, but rather of formal properties which one might expect to find in a number of other settings in the geometry of algebraic curves. The basic idea of Mochizuki's work is that he counts the dormant torally indigenous bundles on general curves of a given type (g, r) (that is, having genus g and r marked points) by degenerating to totally degenerate curves (Definition 3.7); such curves are determined entirely by the combinatorial data of their dual graphs (Definition 3.8). He shows that the number of dormant torally indigenous bundles on a totally degenerate curve can be described combinatorially, essentially as the number of lattice points inside a polytope whose dimensions depend on the characteristic of the base field, and he also shows that the number of dormant torally indigenous bundles on a general curve of type (g, r) is equal to the number on any totally degenerate curve of the same type. In particular, the number agrees for any two totally degenerate curves, which is how we obtain our combinatorial formulas. In addition, the relationship to lattice points of polytopes allows the application of the theory of Ehrhart quasi-polynomials to conclude that these numbers are given by polynomials in the characteristic of the base field.

We remark that for any situation where enumerative invariants are associated to curves in such a way that the invariant for a general curve can be computed at totally degenerate curves, one can expect the computations at totally degenerate curves to be of a combinatorial nature, and then one can hope to obtain non-trivial combinatorial identities by comparing these formulas at different totally degenerate curves, as is done here. Thus, the easiest generalizations of our combinatorial results are likely to arise not from attempting to generalize dormant torally indigenous bundles (although that possibility is discussed briefly at the end of the final section below), but from finding examples of completely unrelated algebro-geometric objects which are associated to curves and satisfy the same formal properties with respect to degeneration.

2. Statements

We first state our theorem in algebraic geometry. Dormant torally indigenous bundles are algebro-geometric objects associated to curves with marked points, and their precise definition, which is rather technical and given in [1, Def. I.1.2, p. 89, Def. I.4.1, p. 113, Def. II.1.1, p. 127], is not relevant to the present paper. However, for the purposes of motivating our Theorem 2.1, we remark on two important cases in which dormant torally indigenous bundles correspond to more concrete and widely-studied objects. In the case of a smooth curve of genus $g \geq 2$ with no marked points, dormant torally indigenous bundles are equivalent up to a factor of 2^{2g} to semistable vector bundles of rank 2 with trivial determinant whose pullbacks under the relative Frobenius morphism are maximally unstable—specifically, contain a line-bundle of degree $g - 1$. In the case that C is smooth with $g = 0$ and r general marked points, dormant torally indigenous bundles are equivalent up to a factor of 2^{r-1} to rational functions on \mathbb{P}^1 ramified to order less than p at the marked points, and unramified elsewhere, up to linear fractional transformation. For both these assertions, see [2]. Theorem 2.1 is new and no easier to prove in both these special cases, and we state the general case partly because it is the most natural level of generality given the arguments, and partly because the phenomenon of invariants in algebraic geometry being expressible by polynomials in the characteristic of the base field is ubiquitous and poorly understood, and rather than

simply providing two apparently unrelated examples of this phenomenon, it seems preferable to have a single example which simultaneously generalizes the special cases.

We now state our first theorem, denoting by F the relative Frobenius morphism:

Theorem 2.1. *Fix $g, r \geq 0$ with $2g - 2 + r > 0$. Then there exists a polynomial $f_{g,r}(n) \in \mathbb{Q}[n]$ such that if k is an algebraically closed field of characteristic $p > 2$, and C a general smooth curve over k of genus g with r general marked points, then the number of dormant totally indigenous bundles on C is given by $f_{g,r}(p)$. Furthermore, $f_{g,r}(n)$ has degree $3g - 3 + 2r$, is even or odd as determined by its degree, and is always strictly positive for $n \geq 2$.*

In particular, if $r = 0$ we have that the number of semistable vector bundles \mathcal{E} of rank two and trivial determinant on C such that $F^\mathcal{E}$ contains a line bundle of degree $g - 1$ is given by $2^{2g} f_{g,0}(p)$. If $g = 0$ we have that the number of maps $f : C = \mathbb{P}^1 \rightarrow \mathbb{P}^1$ ramified to order less than p at the marked points and unramified elsewhere, counted modulo automorphism of the image, is given by $2^{r-1} f_{0,r}(p)$.*

The combinatorial result will require some preliminary definitions. We have:

Definition 2.2. Let V, E be sets, and suppose that we are given φ a map from E to $V \cup \binom{V}{2}$. We then call $G = (V, E, \varphi)$ a **quasi-graph**. The standard notions of **edges**, **vertices**, and edges being **adjacent** to vertices generalize immediately to quasi-graphs. The set of edges E is naturally subdivided into **free** edges, which are $\varphi^{-1}(V)$, and **fixed** edges, given by $\varphi^{-1}(\binom{V}{2})$.

Thus, a quasi-graph may be thought of simply as a graph where some edges—the free edges—are allowed to be attached to only a single vertex. A quasi-graph which consists of only fixed edges is simply a standard graph. Quasi-graphs arise naturally as the dual objects to nodal curves with marked points, where the marked points correspond to free edges of the dual quasi-graph; see Definition 3.8. The usual notions of connectedness, regularity, loops, and simplicity for graphs immediately make sense in the context of quasi-graphs as well. When there is no ambiguity, we will denote by V and E the vertex and edge sets of a quasi-graph G .

We next associate a polytope to certain special quasi-graphs, denoting by $A(v)$ the set of edges adjacent to a vertex v :

Definition 2.3. Let G be a quasi-graph which is regular of degree 3. The convex polytope \mathcal{P}_G associated to G is defined to be the space of real-valued weight functions $w : E \rightarrow \mathbb{R}$ on the edge set of G satisfying the following inequalities:

- (i) for each $e \in E$, $w(e) \geq 0$;
- (ii) for each $v \in V$, $\sum_{e \in A(v)} w(e) \leq 1$;
- (iii) for each $v \in V$ and $e \in A(v)$, $w(e) \leq \sum_{e' \in A(v) \setminus \{e\}} w(e')$.

Note that condition (iii) is just the triangle inequality for the edges adjacent to any given vertex. Note also that (i) and (ii) bound all the $w(e)$ between 0 and 1, so in particular \mathcal{P}_G is in fact a polytope.

We also briefly recall the theory of Ehrhart quasi-polynomials, following [4, p. 235]: a **quasi-polynomial** of degree m and quasi-period d is a function $f : \mathbb{Z} \rightarrow \mathbb{C}$ such that there exist polynomials $f_0, \dots, f_{d-1} \in \mathbb{C}[x]$ of degree m satisfying, for all $x \in \mathbb{Z}$, $f(x) = f_i(x)$, with i determined by $i \equiv x \pmod{d}$. Given a polytope \mathcal{P} of dimension m with rational vertices having common denominator d , if we denote by $n\mathcal{P}$ the n th scaling of \mathcal{P} , one

has the **Ehrhart quasi-polynomial** (see [4, Thm. 4.6.25, Prop. 4.4.1]) of \mathcal{P} , having degree m and quasi-period d , and whose n th value is the number of lattice points in $n\mathcal{P}$. Finally, one has the **Ehrhart reciprocity theorem** (see [4, Thm. 4.6.26]) stating that the number of points in the interior of $n\mathcal{P}$ is given, up to the sign $(-1)^m$, by the $-n$ th value of the Ehrhart quasi-polynomial of \mathcal{P} .

With these definitions, we can now state the combinatorial result:

Theorem 2.4. *Let G, G' be any two quasi-graphs, connected, regular of degree three, and having the same number of vertices and edges. Then the Ehrhart quasi-polynomials for \mathcal{P}_G and $\mathcal{P}_{G'}$ agree at all odd values.*

Ehrhart quasi-polynomials are not in general well-understood: there are formulas for the first two coefficients, but further coefficients, as well as possible criteria for two polytopes to have the same Ehrhart polynomial, remain mysterious. Thus, one might hope that a family of non-trivial identities such as these could help to shed light on the situation.

When presented with cross-disciplinary results such as these, one naturally wonders whether they can be obtained more directly by further exploration of the situation. In algebraic geometry, it is a general phenomenon that answers to enumerative questions are given as polynomials in the characteristic, but as often as not, as in the case here, the only way to show this is to compute the answer and show *a posteriori* that it is a polynomial. We are thus motivated to ask:

Question 2.5. Can one show *a priori* by methods of algebraic geometry that the number of dormant indigenous bundles must be given by a polynomial in p ? Can such an argument be given covering a wider range of enumerative problems?

Correspondingly, we wonder:

Question 2.6. Can one demonstrate directly a combinatorial relationship between the polytopes \mathcal{P}_G for G as in Theorem 2.4 which implies agreement of their Ehrhart quasi-polynomials?

We discuss additional, more concrete combinatorial questions in the final section.

3. Proofs

We start by associating a second polytope to any quasi-graph which is regular of degree 3, which is affinely isomorphic to \mathcal{P}_G , but imbedded in a larger-dimensional space:

Definition 3.1. Let G be a quasi-graph which is regular of degree 3. We describe a second polytope \mathcal{P}'_G associated to G , defined to be the space of real-valued weight functions $w : E \cup V \rightarrow \mathbb{R}$ on the edges and vertices of G satisfying the following inequalities:

- (i) for each $e \in E$, $w(e) \geq 0$;
- (ii) for each $v \in V$, $\sum_{e \in A(v)} w(e) = 2w(v)$;
- (iii) for each $v \in V$ and $e \in A(v)$, $w(e) \leq w(v)$;
- (iv) for each $v \in V$, $w(v) \leq 1$.

Indeed, one checks that points of \mathcal{P}'_G correspond to points of $2\mathcal{P}_G$, by leaving the $w(e)$ unchanged and setting $w(v)$ as determined by (iii) above. The $w(v)$ act as ‘dummy variables’ to insure that lattice points of $n\mathcal{P}'_G$ are merely lattice points of $2n\mathcal{P}_G$ with $w(e)$ ’s having even sum at any v .

We also specify:

Definition 3.2. Let G be a quasi-graph. A **sub-quasi-graph** H of G is a quasi-graph obtained by restricting the adjacency function φ for G to subsets of the vertex and edge sets on which φ remains well-defined.

In particular, a sub-quasi-graph may not change a fixed edge to a free edge.

Lemma 3.3. *Let G be a quasi-graph which is connected and regular of degree 3. Then the odd values of the Ehrhart quasi-polynomials for \mathcal{P}_G and $\mathcal{P}'_{G'}$ differ by an integer multiple determined by G .*

Proof: Let e_1, \dots, e_d be the edges of G , and v_1, \dots, v_m be the vertices. Then the n th value of the Ehrhart quasi-polynomial of \mathcal{P}_G (respectively, $\mathcal{P}'_{G'}$) is by definition the number of possible integer values for the $w(e_i)$ (respectively, the $w(e_i)$ and $w(v_j)$) lying inside the (closed) polytope $n\mathcal{P}_G$ (respectively, $n\mathcal{P}'_G$), which is obtained by replacing the 1 in the definition of \mathcal{P}_G (respectively, $\mathcal{P}'_{G'}$) by n . We claim that for n odd, both of these are equivalent (up to constant integer multiple) to:

$$\# \left\{ (\lambda_1, \dots, \lambda_d, d_1, \dots, d_m) \in \mathbb{Z}^{d+m} : \forall i, 0 < \lambda_i < n + 2; \quad \forall j, d_j < n + 2; \right. \\ \left. \forall i, j \text{ such that } e_i \in A(v_j), \lambda_i \leq d_j; \quad \forall j, 2d_j + 1 = \sum_{i: e_i \in A(v_j)} \lambda_i \right\} \quad (3.1)$$

First, note that the conditions $\lambda_i < n + 2$ are superfluous. If we set $\lambda_i = w(e_i) + 1$ for each i and $d_j = w(v_j) + 1$ for each j , we recover the description of the n th value of the Ehrhart quasi-polynomial of $\mathcal{P}'_{G'}$. Next, it is easily checked that if we set $\lambda_i = 2w(e_i) + 1$ (at which point the d_j are all uniquely determined) as the $w(e_i)$ range over all possibilities for the n th value of the Ehrhart quasi-polynomial of \mathcal{P}_G , we recover all possible assignments of the λ_i, d_j for which λ_i are all odd. The key observation is that for n odd, the number of possibilities with all λ_i odd is $\frac{1}{N_G}$ times the total number of possibilities, where N_G is the number of (not necessarily spanning) sub-quasi-graphs of G which are regular of degree 2. Indeed, if one starts with an arbitrary assignment of λ_i, d_j , the edges e_i for which λ_i are even gives such a sub-quasi-graph, and if all λ_i which are even are replaced by $n + 2 - \lambda_i$ (and the d_j adjusted accordingly), one can check that all conditions are preserved, and one obtains an assignment with all λ_i odd. This sets up a natural, visibly invertible correspondence between arbitrary assignments of λ_i and d_j satisfying the required inequalities, and pairs of assignments with all λ_i odd together with an arbitrary sub-quasi-graph of G which is regular of degree 2. This completes the proof of the claim, and the lemma. □

We now proceed to describe the vertices of the polytopes we have constructed. We start with:

Lemma 3.4. *Let G be a regular quasi-graph of degree 3. Then any vertex of \mathcal{P}_G whose coordinates are all non-zero has coordinates equal to $\frac{1}{4}$ or $\frac{1}{2}$. More precisely, for any vertex of \mathcal{P}_G , the weights associated to the three edges adjacent to any vertex of G are $\{\frac{1}{4}, \frac{1}{4}, \frac{1}{2}\}$.*

Proof: Denote by E_2 and E_1 the sets of fixed and free edges of G respectively. Then one has $3(\#V) = 2(\#E_2) + \#E_1$. A vertex of \mathcal{P}_G must satisfy all of the inequalities listed in Definition 2.3; moreover, by replacing the inequalities with equalities, one obtains a collection of linear constraints, and among these, the vertex must satisfy some $\#E$ independent constraints. By hypothesis, none of these constraints are of the form $w(e) = 0$, so they are chosen from the $4(\#V)$ constraints of the form $w(e) + w(e') + w(e'') = 1$ or $w(e) = w(e') + w(e'')$ where e, e', e'' are the three edges (possibly with multiplicity) adjacent to some vertex of G . We note that for any given vertex of G , we cannot have two constraints of the second form, since that would force one of the weights to be 0. Therefore, each vertex of G can supply at most two constraints, and in the case of two, one of the two is necessarily $w(e) + w(e') + w(e'') = 1$. Moreover, since we need $\#E = \#E_1 + \#E_2$ constraints, we must have at least $\frac{1}{3}(\#E_2) + \frac{2}{3}(\#E_1)$ vertices of G with two constraints; denote these vertices by V^2 . For any given such vertex in V^2 , we may write the two constraints as $w(e) + w(e') + w(e'') = 1$ and $w(e) = w(e') + w(e'')$, which yields $w(e) = \frac{1}{2}$. Let E^2 denote the set of edges e forced to have weight $\frac{1}{2}$ in this manner.

Now, by the hypothesis that all weights are non-zero, for any given vertex of G we cannot have two edges with weight $\frac{1}{2}$, so every vertex in V (and in particular in V^2) is adjacent to a unique edge of E^2 . The next observation is that conversely, every edge $e \in E^2$ is adjacent to a unique vertex of V^2 . Indeed, if both vertices adjacent to e were in V^2 , the two constraints at each would both necessarily force e to have weight $\frac{1}{2}$, which would imply that they were linearly dependent. We claim that in fact every vertex of G is adjacent to an edge in E^2 ; that is, E^2 gives a perfect matching of G . We subdivide V^2 into V^{21} and V^{22} , according to whether the corresponding edge of E^2 is free or fixed, respectively. We therefore want to show that $\#V^{21} + 2(\#V^{22}) \geq \#V$. But $2(\#V^{21}) + 2(\#V^{22}) = 2(\#V^2) \geq \frac{2}{3}(\#E_2) + \frac{4}{3}(\#E_1)$, and by definition $\#V^{21} \leq \#E_1$, so subtracting we find that $\#V^{21} + 2(\#V^{22}) \geq \frac{2}{3}(\#E_2) + \frac{1}{3}(\#E_1) = \#V$, as claimed. Finally, we can conclude that $\#V^{21} + 2(\#V^{22}) = \#V$, and it follows that we must have had $\#V^{21} = \#E_1$, and no vertices of G without any associated constraints.

The final step is to show that given our description so far, if one assigns $\frac{1}{4}$ to all edges not in E^2 , then this assignment satisfies every constraint which is permissible based on the hypotheses that all weights are non-zero and that the weights of E^2 are predetermined as $\frac{1}{2}$. Indeed, let v be any vertex of G , and e, e', e'' its adjacent edges. Suppose without loss of generality that $e \in E^2$. Then under the assignment $w(e') = w(e'') = \frac{1}{4}$, both constraints $w(e) + w(e') + w(e'') = 1$ and $w(e) = w(e') + w(e'')$ will be satisfied. All that remains is to note that with $w(e) = \frac{1}{2}$, no valid assignment of $w(e')$ and $w(e'')$ can achieve $w(e') = w(e) + w(e'')$ or $w(e'') = w(e) + w(e')$, since with all three weights positive, their sum would have to be greater than 1. Thus, our assignment satisfies any possible choice of constraints associated to the vertices of G not in V^2 , and we conclude that our chosen vertex of \mathcal{P}_G has coordinates of the desired form. □

We can now conclude:

Proposition 3.5. *Let G be a regular quasi-graph of degree 3. Then any vertex of \mathcal{P}_G has each of its coordinates equal to 0, $\frac{1}{4}$ or $\frac{1}{2}$, with the only possible weights associated to the*

edges adjacent to a given vertex of G being $\{0, 0, 0\}$, $\{0, \frac{1}{2}, \frac{1}{2}\}$, $\{0, \frac{1}{4}, \frac{1}{4}\}$, and $\{\frac{1}{4}, \frac{1}{4}, \frac{1}{2}\}$. Any vertex of \mathcal{P}'_G has each of its coordinates equal to $0, \frac{1}{2}$ or 1 .

Proof: We begin with the assertion for \mathcal{P}_G , working by induction on the number of coordinates of a given vertex which are equal to 0. The base case is that all coordinates are non-zero, which we handled in the previous lemma. Now, suppose we have an edge e of G whose weight is zero for our chosen vertex of \mathcal{P}_G . Suppose e is a fixed edge, and let v_1, v_2 be the two adjacent vertices; note that it suffices to consider the case that v_1 and v_2 are distinct, since if e is a loop, the triangle inequalities in the definition of \mathcal{P}_G will force the other edge adjacent to $v_1 = v_2$ to have weight 0 as well, and we could instead choose this edge to be e . Next, note that if e_1, e'_1 and e_2, e'_2 are the adjacent edges other than e at v_1, v_2 respectively, then the triangle inequalities at v_1 and v_2 force $w(e_1) = w(e'_1)$ and $w(e_2) = w(e'_2)$. We define a new graph G' obtained by removing $e, v_1,$ and $v_2,$ and replacing each pair e_1, e'_1 and e_2, e'_2 by single edges e''_1 and e''_2 in the obvious way: that is, if e_i, e'_i are each adjacent to vertices other than v_1 or $v_2,$ replace them with an edge adjacent to those two vertices; if only one is adjacent to a vertex other than v_1 or $v_2,$ replace them by a single free edge; if neither is adjacent to a vertex other than v_1 or $v_2,$ remove them entirely.

Now, we claim that we obtain a vertex of $\mathcal{P}_{G'}$ by assigning weights to the edges of G' which are the same as G where the graphs are the same, and which assign the common weight of e_i, e'_i to the new edges e''_i for $i = 1, 2$. It suffices to show that we can provide constraints from the definition of $\mathcal{P}_{G'}$ to replace any constraints that were lost when v_1 and v_2 were removed. The constraints coming from triangle inequalities at v_i are easily replaced: they can either require $w(e_i) = w(e'_i),$ or $w(e_i) = w(e_i) = 0$. The first condition is superfluous, while the second can be replaced by the constraint $w(e''_i) = 0$. So we need only show that we can effectively replace the condition that the sum of the weights at a v_i be equal to 1, which gives $w(e_i) = w(e'_i) = \frac{1}{2},$ so is equivalent to requiring $w(e''_i) = \frac{1}{2}$. Choose v_3 to be a vertex adjacent to $e''_i;$ without loss of generality, suppose this was the vertex adjacent to e_i in G . Since we had $w(e_i) = \frac{1}{2}$ in $G,$ the inequality requiring the sum of the three weights at v_3 to be at most 1, together with the triangle inequality for $w(e_i),$ implies that in fact both of these inequalities are sharp, giving corresponding constraints satisfied in $G,$ so we can then require them also in G' in order to force the weight of e''_i to be $\frac{1}{2},$ as desired.

By the induction hypothesis, we can assume that the vertex of $\mathcal{P}_{G'}$ we have constructed has weights only equal to $0, \frac{1}{4}$ or $\frac{1}{2},$ with the weights of the asserted form for edges adjacent to a particular vertex. One checks easily that our description of weights of the edges adjacent to a given vertex is preserved by the construction, as long as we verify that if in constructing G' we removed an edge of G other than $e,$ its weight must also have been one of $0, \frac{1}{4}, \frac{1}{2}$. Such a removal only occurred if both e_i and e'_i were both adjacent only to v_1, v_2 for $i = 1$ or 2 . There are only two possibilities: either $e_i = e'_i$ is a loop, or e_i and e'_i are both adjacent to both v_1 and $v_2,$ in which case these three edges and two vertices are necessarily all of G . In the first case, one checks that in order for the weight of the loop to be uniquely determined, given that $w(e) = 0,$ its weight is necessarily either 0 or $\frac{1}{2}$. In the second case, one can check the assertion of the Proposition directly for G (see also Example 4.1). Finally, note that although we carried out this process in the case that e was fixed, the argument works equally well (and is in fact simpler) in the case that e is free. This completes the induction argument for \mathcal{P}_G .

We may now conclude the desired statement for $\mathcal{P}'_G:$ for finding vertices of our polytopes, we work over $\mathbb{R},$ and in this setting, as mentioned above, the introduction of the $w(v)$ coordinates are irrelevant, and if we ignore these coordinates, the polytope \mathcal{P}'_G is the same as $2\mathcal{P}_G$. In particular, with the possible exception of the $w(v),$ all coordinates of vertices

of \mathcal{P}'_G are equal to 0, $\frac{1}{2}$ or 1. But the same follows for $w(v)$ from our sharp description of the possible weights associated to edges adjacent to a vertex of G in the case of a vertex of \mathcal{P}_G . □

This gives us our key result:

Corollary 3.6. *The odd values of the Ehrhart quasi-polynomial of \mathcal{P}'_G , and hence \mathcal{P}_G , are given by a single polynomial, of degree equal to $\#E$.*

Proof: We note that the dimension of \mathcal{P}_G (and hence \mathcal{P}'_G) is equal to $\#E$: indeed, it is easily verified that the $(\#E)$ -cube with all weights between $\frac{1}{6}$ and $\frac{1}{3}$ lies inside \mathcal{P}_G . The assertion for \mathcal{P}'_G is then immediate from Proposition 3.5 and the existence theorem for Ehrhart-quasi-polynomials. The assertion for \mathcal{P}_G then follows by Lemma 3.3. □

Recall the following:

Definition 3.7. A **nodal curve** C is a curve obtained from a (not necessarily connected) smooth curve by gluing together pairs of points transversely, creating **nodes** at these points. The smooth curve from which C is obtained is unique, and called the **normalization** of C , and will be denoted \tilde{C} . When considering curves with marked points, we set the marked points of \tilde{C} to be the points lying above marked points or nodes of C . Finally, a **totally degenerate curve** is a nodal curve such that \tilde{C} consists of disjoint copies of \mathbb{P}^1 with three marked points each.

Because any three points on \mathbb{P}^1 are equivalent up to automorphism, a totally degenerate curve is determined by combinatorial data, and specifically by the dual quasi-graph:


Definition 3.8. Let C be a nodal curve with marked points. Then the **dual quasi-graph** associated to C is defined to be the quasi-graph whose vertices are the components of C , whose fixed edges correspond to nodes of C and are adjacent to the components intersecting at a given node, and whose free edges correspond to marked points of C , and are adjacent to the component on which the marked point lies.

One checks directly that a C such that \tilde{C} is a disjoint union of \mathbb{P}^1 's is totally degenerate if and only if its dual quasi-graph is regular of degree 3, and that conversely given a quasi-graph which is regular of degree 3, there is a unique totally degenerate curve having the chosen dual quasi-graph. Finally, one checks that the type (g, r) of the curve is related to the number of vertices and edges of the dual quasi-graph by the formulas $\#V = 2g - 2 + r$, $\#E = 3g - 3 + 2r$.

The theorem we will use which is implicit in Mochizuki's work may be stated as:

Theorem 3.9. (Mochizuki) *Fix $g, r \geq 0$ with $2g - 2 + r > 0$, and p an odd prime. Then the number of dormant torally indigenous bundles on a general curve of type (g, r) over an algebraically closed field of characteristic p is given as the $(p - 2)$ nd value of the Ehrhart quasi-polynomial of \mathcal{P}_G , where G is any connected regular quasi-graph of degree 3 satisfying $\#V = 2g - 2 + r$, $\#E = 3g - 3 + 2r$; in particular, these values depend only on $\#V$ and $\#E$.*

Proof: The first relevant statement is that the number of dormant torally indigenous bundles on a general curve of type (g, r) over an algebraically closed field may be computed at any




totally degenerate curve of type (g, r) . This follows immediately from the assertions of [1, Thm 2.8, p. 153] in the $n = 0$ case that the stack of dormant torally indigenous bundles is finite and flat over $\overline{\mathcal{M}}_{g,r}$ and is étale over points corresponding to totally degenerate curves. Next, one needs to know that a dormant torally indigenous bundle on a totally degenerate curve C is equivalent to dormant torally indigenous bundles on each component of \tilde{C} having radii which agree at any two points which are glued together; this is immediate from [1, §I.4.4, p. 118] when one takes into account that the dormancy condition is simply a condition of vanishing p -curvature, and will not be affected by gluing.

The final ingredient is the description of dormant torally indigenous bundles on \mathbb{P}^1 with three marked points, given as the $n = 0$ case of [1, Thm. IV.2.3, p. 211]. If we are given λ_i as in this last theorem (these always exist, since the radii are only defined up to ± 1 , so we could choose all the λ_i to be odd), we have to check that the existence of a finite separable morphism from \mathbb{P}^1 to itself ramified to orders λ_i at $0, 1, \infty$ and unramified elsewhere is equivalent to the conditions that the degree d , which by the Riemann-Hurwitz formula is determined by $2d + 1 = \sum \lambda_i$, must also satisfy $d < p$ and $d \geq \lambda_i$ for all i . This is shown in [3], but could also be deduced directly from [1, (2), p. 232]. Given all of this, one can verify directly that the dormant torally indigenous bundles on a totally degenerate curve C are nearly counted by setting $n = p - 2$ in Eq. (3.1) as applied to the dual quasi-graph G of C . The only discrepancy is a factor of the N_G of the proof of Lemma 3.3, since radii are only defined up to ± 1 and therefore assignments of λ_i which differ by $\pm 1 \pmod p$ give the same dormant torally indigenous bundle. Since we saw in the proof of Lemma 3.3 that $\frac{1}{N_G}$ times the value of Eq. (3.1) computed the odd values of the Ehrhart quasi-polynomial of \mathcal{P}_G , we thus conclude the desired result. \square

We can now easily give the proofs of our main theorems:

Proof of Theorem 2.1: We note that given any specified (g, r) with $2g - 2 + r > 0$, we can find a connected regular quasi-graph of degree 3 having the number of vertices and edges required by Theorem 3.9. This is equivalent to the standard algebraic geometry statement that there exist totally degenerate curves of any hyperbolic type, but one can easily verify it directly. Indeed, the figure demonstrates how to increase either g or r by 1 while keeping the other fixed, and given this it suffices to check the base cases of $(g, r) = (0, 3), (1, 1), (2, 0)$, which is easily accomplished. Putting together Corollary 3.6 with Theorem 3.9 then gives the existence and degree of the desired polynomial. The positivity follows from the fact that for any $n \geq 0$, \mathcal{P}_G contains the lattice point with all weights equal to 0. Lastly, to see that the polynomial is always odd or even, we note that simply by translating all coordinates by 1 one sees that the number of lattice



The case $(g, r) = (2, 0)$

points in the interior of $n\mathcal{P}_G$ is equal to the number in the the closed polytope $(n - 4)\mathcal{P}_G$. Applying the reciprocity theorem for Ehrhart polynomials then easily gives the desired result.

Proof of Theorem 2.4: This is almost the same as Theorem 3.9, except that it asserts agreement of the n th value of the Ehrhart quasi-polynomial for all odd n rather than those for which $n + 2$ is a prime. The stronger statement then follows from Corollary 3.6, although in fact for this application one need not consider \mathcal{P}'_G at all: it is enough to use the existence of Ehrhart quasi-polynomials once one knows that the Ehrhart quasi-polynomial of \mathcal{P}_G has quasi-period 4, since Mochizuki’s values then give infinitely many values for n congruent to either 1 or 3 mod 4. Thus, to prove Theorem 2.4 it suffices to know the statement of Proposition 3.5 for \mathcal{P}_G only.

Remark 3.10. Translating from self-maps of \mathbb{P}^1 to indigenous bundles and back in order to obtain the statement for $g = 0$ of Theorem 2.1 may seem superfluous, and indeed one could argue directly using the results of [3] that the number of such maps is counted by the Ehrhart polynomial of \mathcal{P}_G . However, there is something to be said for concluding the statement as a special case of a more general result.

Remark 3.11. In fact, an induction argument similar to that carried out for Proposition 3.5 can be used to show that the number N_G of Lemma 3.3 also depends only on the number of vertices and edges of G , and is in fact given explicitly by 2^g for $r = 0$ and 2^{g+r-1} for $r > 0$. The basic idea is to observe that the operation of exclusive or puts the structure of an abelian group with all elements having order 2 on the set of sub-quasi-graphs which are regular of degree 2, and then to make use of the fact that for any given edge, a generating set for this group may always be chosen to have at most one element which contains the chosen edge. It then follows that one also obtains identities for the odd values of the Ehrhart quasi-polynomials of \mathcal{P}'_G . We have chosen to phrase our main result in terms of \mathcal{P}_G partly for the sake of simplicity, and partly because it seems like the more natural object, in that its Ehrhart quasi-polynomial computes the number of dormant torally indigenous bundles directly, and it is imbedded in a space of its own dimension.

4. Further remarks and questions

In this section, we discuss some explicit examples and possible directions of further investigation, with a particular focus towards the combinatorial side. We begin by describing the simplest example of our results, seeing that the combinatorial identities obtaining in Theorem 2.4 do in fact appear to be non-trivial.

Example 4.1. Consider the case of $(g, r) = (2, 0)$, or equivalently graphs with three edges and two vertices. One checks that we get only two graphs: the G_1 and G_2 of the figure. The corresponding polytopes \mathcal{P}_{G_1} and \mathcal{P}_{G_2} are, respectively: a regular tetrahedron with vertices at $(0, 0, 0)$, $(\frac{1}{2}, \frac{1}{2}, 0)$, $(0, \frac{1}{2}, \frac{1}{2})$ and $(\frac{1}{2}, 0, \frac{1}{2})$; and a square pyramid with vertices $(0, 0, 0)$, $(\frac{1}{2}, 0, 0)$, $(0, \frac{1}{2}, 0)$, $(\frac{1}{2}, \frac{1}{2}, 0)$ and $(\frac{1}{4}, \frac{1}{4}, \frac{1}{2})$. One finds that in fact not only the odd values, but the entire Ehrhart quasi-polynomials of \mathcal{P}_{G_1} and \mathcal{P}_{G_2} agree, and are given by $\frac{1}{24}(n^3 + 6n^2 + 20n + 24)$ for even n and $\frac{1}{24}(n^3 + 6n^2 + 11n + 6)$ for odd n . The number of dormant torally indigenous bundles in this case is thus given by $\frac{1}{24}(p^3 - p)$.

While we have not explicitly presented \mathcal{P}_G for different G in further cases, one can compute that in the case of $(g, r) = (3, 0)$, there are five different graphs, for which the corresponding \mathcal{P}_G have 8, 10 or 14 vertices depending on G . Thus, in this situation there are at least three polytopes which are not combinatorially equivalent for which we obtain relations. It seems reasonable to expect that the number of non-trivial identities obtained will grow with (g, r) .

In computing further examples, there are two phenomena which stand out. The first is that in all examples computed so far, for any two G, G' as in Theorem 2.4 we find that not only the odd values of the Ehrhart quasi-polynomial agree, but the even values agree as well. This holds for examples with (g, r) up to $(4, 0)$, as well $(0, 6)$ and $(1, 2)$. The data for $(5, 0)$ is also consistent with this conclusion, although computation of the entire Ehrhart polynomial for even a single graph in this case appears unfeasible. We therefore conjecture:

Conjecture 4.2. In Theorem 2.4, the restriction to odd values of the Ehrhart quasi-polynomials is unnecessary.

We also remark that the same seems to hold for \mathcal{P}'_G in the few examples we have computed thus far. This is interesting in its own right, as there is no obvious relation between the even values of the Ehrhart quasi-polynomials of \mathcal{P}_G and \mathcal{P}'_G .

We also make some observations on the period of the Ehrhart quasi-polynomial of \mathcal{P}_G . First, in certain cases with $r > 0$ (for instance, when $(g, r) = (1, 1), (1, 2)$, or $(3, 1)$) one can compute that the Ehrhart quasi-polynomial in fact has quasi-period 4. This is as expected based on the fact that for any graph G of the corresponding type, some vertices of \mathcal{P}_G have quarter-integer coordinates, but it is interesting in that it means that Theorem 2.4 is producing potentially infinitely many examples of rational polytopes for which the Ehrhart quasi-polynomial has different quasi-periods when restricting attention to even or odd values. In contrast, we have observed that when $r = 0$ the Ehrhart quasi-polynomials of \mathcal{P}_G and \mathcal{P}'_G always appear to have quasi-period 2 and 1 respectively, rather than the *a priori* expected quasi-periods of 4 and 2. This would follow from the above conjecture for $g \leq 4$, since in this case there is always a G with all vertices of \mathcal{P}_G lying on half-integers (and correspondingly, all vertices of \mathcal{P}'_G lying on integers). However, for $g = 5$ and above it is possible to show by strengthening the induction argument of Proposition 3.5 that every \mathcal{P}_G must have some vertices with quarter-integer coordinates. Yet, the data we have for $g = 5$ is consistent with the Ehrhart quasi-polynomial of \mathcal{P}_G having quasi-period 2. Thus, we seem to have a separate pattern not explained by our previous conjecture, and we ask:

Question 4.3. Is it true that for $r = 0$, the quasi-period of the Ehrhart quasi-polynomial of \mathcal{P}_G is always 2? Are there other cases where the quasi-period is smaller than expected based

on the denominators of the vertices of \mathcal{P}_G as G ranges over all quasi-graphs corresponding to (g, r) ? Is the quasi-period of the Ehrhart quasi-polynomial for \mathcal{P}'_G always half the quasi-period for \mathcal{P}_G ?

Finally, we would like to add that Mochizuki had already remarked on the existence of apparently non-trivial combinatorial identities implicit in his work (see [1, pp. 238–239]), but in a more general setting than we have treated here. Additional families of combinatorial identities are obtainable by considering more general nilpotent torally indigenous bundles than the dormant ones we have examined. For instance, Mochizuki treats a few cases of these identities for $g = 0, 1$ in the situation of ordinary torally indigenous bundles; see [1, p. 24] and [1, Cor. V.1.3, p. 237]. He also develops the basic combinatorial algorithms necessary to describe the intermediate cases, called spiked torally indigenous bundles; see [1, Def. II.3.1, p. 160], [1, Thm. V.1.1, p. 236], as well as the diagram [1, p. 270]. Translating these identities into combinatorial language is likely to be more complicated than for the dormant case, but may yield correspondingly more interesting identities.

Acknowledgment We would like to thank Shinichi Mochizuki for bringing the relevant results of [1] to our attention.

References

1. S. Mochizuki, *Foundations of p -adic Teichmüller theory*, American Mathematical Society, 1999.
2. B. Osserman, *Mochizuki's crys-stable bundles: A lexicon and applications*, to appear in Publications of the Research Institute for Mathematical Sciences.
3. B. Osserman, *Rational functions with given ramification in characteristic p* , to appear in *Compositio Mathematica*.
4. R. P. Stanley, *Enumerative Combinatorics: Vol. I*, Cambridge Studies in Advanced Mathematics, no. 49, Cambridge University Press, 1997.