Compositio Math. 144 (2008) 89–106
doi:10.1112/S0010437X0700303X

Linear series and the existence of branched covers

Brian Osserman

Abstract

In this paper, we use the perspective of linear series, and in particular results following from the degeneration tools of limit linear series, to give a number of new results on the existence and non-existence of tamely branched covers of the projective line in positive characteristic. Our results are both in terms of ramification indices and the sharper invariant of monodromy cycles, and the first class of results are obtained by intrinsically algebraic and positive-characteristic arguments.

1. Introduction

Over the complex numbers, the classical theory of branched covers and the Riemann existence theorem give a complete description of branched covers of curves in terms of, in the case of a base of genus zero, the monodromy around branch points. Techniques of lifting to characteristic zero and comparing with the transcendental situation allow one to conclude that when (the order of) monodromy groups are prime to p, the situation in characteristic p remains the same as the classical situation. However, when monodromy groups are divisible by p, even when one studies tame covers the situation becomes far more delicate. The main issue is that, although it remains true that tame covers always lift to characteristic zero, it is no longer the case that a cover in characteristic zero necessarily has good reduction to characteristic p. One subtlety in this context is that the existence and non-existence of tame covers can depend quite strongly on the moduli of the curve; see Tamagawa [Tam04]. We avoid this issue entirely by focusing on the question of which tame covers exist for generic curves, but even in this context very little is known.

Much of the work to date on the existence of tame covers (see, e.g., Raynaud's [Ray99]) focuses on situations where one can still show that one has good reduction from characteristic zero. In this paper we pursue an entirely different tack, using degeneration techniques and the point of view of linear series. We therefore obtain results without lifting to characteristic zero and invoking transcendental techniques, except that our results on monodromy groups are stated in terms of systems of generators which are only known to exist by transcendental methods. Although the idea of using degeneration techniques to obtain this sort of result is not new (see, e.g., Bouw–Wewers [BW05] and Harbater–Stevenson [HS99]), the introduction of certain key new ingredients from the point of view of linear series allows for far stronger results than had previously been known, including in particular sharp non-existence results in the case of genus-zero covers of the projective line with either all ramification indices less than p, or only three ramification points.

Owing to the nature of the arguments, our results are all stated in terms of covers of the projective line with only one ramified point over each of r general branch points. This means that the group-theoretic situation is rather special and, in particular, monodromy groups are always cyclic, alternating, or symmetric groups. However, the restrictions on both genus and branching

Received 24 October 2006, accepted in final form 26 March 2007, published online 23 January 2008.
2000 Mathematics Subject Classification 14H30.
Keywords: tame covers, limit linear series.
The author was supported by fellowships from the Japan Society for the Promotion of Science and the NSF.

type can be relaxed to a substantial degree via a combination of standard techniques; see § 7 below.

Finally, in § 6, we examine two elementary examples which demonstrate two phenomena: first, unlike in the case of characteristic zero, in positive characteristic the particular possibilities for local monodromy cycles of covers can depend in a strong sense on which generators of the fundamental group of the base are used to obtain them; second, degeneration techniques seem to be 'unreasonably effective' in a sense that will be made precise below.

We now fix some terminology so that we can state our main theorem more precisely.

Definition 1.1. Let X be an r-marked curve of genus zero over an algebraically closed field, with marked points $Q_1, \dots, Q_r$. We say that a tuple $(\gamma_1, \dots, \gamma_r) \in (\pi_1^{\text{tame}}(X))^r$ is a *local generating system* for $\pi_1^{\text{tame}}(X)$ if:

(i) the γ_i generate $\pi_1^{\text{tame}}(X)$;

(ii) $\gamma_1 \cdots \gamma_r = 1$;

(iii) each γ_i is a generator of an inertia group at Q_i (i.e. the algebraic equivalent of a small loop around Q_i).

When we write $\pi_1^{\text{tame}}(X)$ for a marked curve X, we always mean the tame fundamental group of the curve with the marked points removed.

We recall [Gro71, Corollaire XIII.2.12] that by specializing from topological generators in characteristic zero, one sees that there always exists a local generating system for $\pi_1^{\text{tame}}(X)$.

Definition 1.2. We say that an r-tuple of cycles $(\sigma_1, \dots, \sigma_r) \in (S_d)^r$ is a *Hurwitz factorization* for $(d, r, \{e_1, \dots, e_r\})$ if:

(i) σ_i has length e_i for all i;

(ii) the product $\sigma_1 \cdots \sigma_r$ is trivial;

(iii) the σ_i generate a transitive subgroup of S_d.

We also sometimes say simply that $(\sigma_1, \dots, \sigma_r)$ is a Hurwitz factorization when no d or e_i are specified, and conditions (ii) and (iii) are satisfied.

If we are given a tame cover f of X, and a local generating system, then we obtain a Hurwitz factorization $(\sigma_1, \dots, \sigma_r)$ by labeling the fiber over the base point, and letting σ_i be the monodromy of f induced by γ_i.

Recall that the pure braid group acts on local generating systems: the ith generator of the braid group B_r acts by replacing (γ_i, γ_{i+1}) by $(\gamma_{i+1}, \gamma_{i+1}^{-1}\gamma_i\gamma_{i+1})$, but does not respect the ordering. The pure braid group is the kernel of the natural map $B_r \to S_r$, so respects the ordering and gives a well-defined action. In the classical setting, any two topological local generating systems are related by a pure braid transformation.

In the case that $r = 3$ or that $e_i < p$ for all i, in § 5 we define a purely group-theoretic condition for a Hurwitz factorization to be p-admissible. We will see in Corollary 5.5 below that p-admissibility is in fact a condition only on the e_i, and can be roughly summarized as the requirement that the e_i are not too close to certain multiples of powers of p in the case that $r = 3$, or that there are enough cycles of moderate length in the case that $e_i < p$ for all i (see also Lemma 2.3 and Example 4.4 below).

For the sake of comparison, we recall the genus-zero case of the following theorem from SGA.

Theorem 1.3 [Gro71, Corollaire XIII.2.12]. *Let X be a curve of genus zero over an algebraically closed field k, with marked points $Q_1, \dots, Q_r$. Then there exists a local generating system $(\gamma_1, \dots, \gamma_r)$ for $\pi_1^{\text{tame}}(X)$ such that if $(\sigma_1, \dots, \sigma_r)$ is any Hurwitz factorization for $(d, r, \{e_1, \dots, e_r\})$ satisfying:*

https://doi.org/10.1112/S0010437X0700303X Published online by Cambridge University Press

(i) $2d-2=\sum_i(e_i-1)$;

(ii) *either* $\operatorname{char} k=0$, *or* $\operatorname{char} k=p>0$, *and the group* $\langle\sigma_1,\dots,\sigma_r\rangle\subseteq S_d$ *has order prime to* p;

there exists a cover $f:\mathbb{P}^1\to X$ *of degree* d, *with* f *branched only over the* Q_i, *and such that for all* i, *the local monodromy around* γ_i *is given by* σ_i.

Our main theorem is then the following.

Theorem 1.4. *Let* $(\sigma_1,\dots,\sigma_r)$ *be a Hurwitz factorization for* $(d,r,\{e_1,\dots,e_r\})$, *with* $2d-2=\sum_i(e_i-1)$, *every* e_i *prime to* p, *and where we suppose in addition that either* $r=3$, *or* $e_i<p$ *for all* i. *Fix also a local generating system* $(\gamma_1,\dots,\gamma_r)$ *for* $\pi_1^{\mathrm{tame}}(X^{\mathrm{gen}})$, *where* X^{gen} *is the geometric generic* r*-marked curve of genus zero, with marked points* $Q_1,\dots,Q_r$.

Then the following are equivalent:

(a) *the tuple* $(\sigma_1,\dots,\sigma_r)$ *is* p*-admissible;*

(b) *there exists a map* $f:\mathbb{P}^1\to\mathbb{P}^1$ *of degree* d, *and distinct* $P_1,\dots,P_r$ *on the source* $\mathbb{P}^1$, *such that* f *is ramified to order* e_i *at each* P_i;

(c) *there exists a cover* $f:\mathbb{P}^1\to X^{\mathrm{gen}}$ *of degree* d, *and a choice of local generating system* $(\gamma_1',\dots,\gamma_r')$, *with* f *branched only over the* Q_i, *and such that for all* i, *the local monodromy around* γ_i' *is given by* σ_i;

(d) *there exists a cover* $f:\mathbb{P}^1\to X^{\mathrm{gen}}$ *of degree* d, *and a pure-braid transformation* $(\gamma_1',\dots,\gamma_r')$ *of* $(\gamma_1,\dots,\gamma_r)$, *with* f *branched only over the* Q_i, *and such that for all* i, *the local monodromy around* γ_i' *is given by* σ_i.

Furthermore, if $r=3$, *no pure braid transformation is required in part (d).*

The proof of this theorem is completed in § 5; see also Theorems 2.4 and 4.2 below for a purely numerical criterion on the e_i determining whether part (b) holds. The case $r=3$ is sufficiently combinatorial that we can understand the situation completely. However, for $r>3$, we use degeneration arguments, and subtleties arise. In particular, if we allow $e_i>p$, pathologies occur as discussed in Remark 3.2, and not every cover arises via degeneration to the $r=3$ case. The necessity of working with general branch points in order to give any purely group-theoretic criterion for the existence of covers with given monodromy is, as mentioned before, well known, but we give elementary examples in § 6. The same examples also justify the need for a pure-braid operation in part (d) as soon as $r>3$. Specifically, we will see the following.

Proposition 1.5. *Let* Σ *be a set of local generating systems for* $\pi_1^{\mathrm{tame}}(X^{\mathrm{gen}})$ *which is closed under pure-braid operations (e.g. the set of local generating systems arising as specializations of topological systems). Then for* $r=4$ *and* $p=3$ *the total set of Hurwitz factorizations* $(\sigma_1,\dots,\sigma_r)$ *which arise as monodromy of tame covers (with* d *and* e_i *allowed to vary arbitrarily) around* $(\gamma_1,\dots,\gamma_r)\in\Sigma$ *depends on the choice of* $(\gamma_1,\dots,\gamma_r)$.

In particular, there is no group-theoretic criterion for recognizing when a Hurwitz factorization occurs as monodromy of a tame cover around a local generating system which works simultaneously for all systems in Σ.

The general thrust of the argument for our main theorem is to use the point of view of linear series, and the associated degeneration tools of limit linear series, to draw conclusions in the context of branched covers. There are four key ingredients: the exploitation of the existence or non-existence of inseparable maps to draw conclusions about separable maps, as in [Oss06b, Theorem 4.2]; the precise description [Oss06b, Theorem 6.1] of when separable maps can degenerate to inseparable maps; a finiteness result [Oss07, Theorem 5.3], proved by relating the maps in question to certain logarithmic connections with vanishing p-curvature on the projective line, and applying results of Mochizuki; and, finally, a result of Liu and the author [LO] in the classical setting showing that in the situation we study, all Hurwitz factorizations always lie in a single braid orbit. The first is used for the case of three ramification points, while the second and third ingredients are used to conclude sharp non-existence results in the case that all ramification indices are less than p. The last result is used to go from numerical results in terms of ramification indices to statements in terms of monodromy cycles.

2. Numerical results: the case of three points

In this section, we prove our first sharp result, in the case of covers $\mathbb{P}^1 \to \mathbb{P}^1$ with only three ramification points. We use the techniques of [Oss06b, Theorem 4.2] to generalize the result there, giving a complete answer in this case. Our results here, as in § 4, are expressed numerically in terms of ramification indices; we translate into monodromy group statements in § 5.

As in [Oss06b, Theorem 4.2], the main idea is to evaluate the existence of a separable map with the given ramification by examining the possibility of an inseparable linear series with the same ramification. We therefore begin with a brief review of linear series in the case of dimension one, which is far simpler than the general case. For the general definitions, see [Oss06a].

For our purposes, a *linear series* of dimension one and degree d on a curve C consists of a map $f : C \to \mathbb{P}^1$ of degree $d' \leqslant d$, together with an effective divisor of *base points* on C of degree $d - d'$. However, we consider two maps as giving the same linear series if they are related by an automorphism of $\mathbb{P}^1$. Thus, while branched covers are considered up to automorphism of the source, linear series are considered up to automorphism of the target.

We say that a linear series is *separable* or *inseparable* depending on whether the map f is separable or inseparable. We say that a linear series is *ramified* to order (at least) e at a point P if the sum of the ramification index of f at P and the number of base points at P is (at least) e.

A direct computation with the Hurwitz formula implies that if we have $P_1, \ldots, P_r$ on $\mathbb{P}^1$, and $e_1, \ldots, e_r$ such that $2d-2 = \sum_i (e_i - 1)$, then we can only have a separable linear series on $\mathbb{P}^1$ ramified to order e_i at every P_i if the associated map f has degree d, so that we have no base points. Thus, we will typically be working with linear series that have no base points, which are simply maps to $\mathbb{P}^1$ up to automorphism of the image. The exception is when we work with inseparable linear series, as in the proof of Theorem 2.4 below.

We first introduce some useful notation and terminology.

Notation 2.1. Given a positive integer e prime to p, we write $\bar{e}^{[m,u]} := \lceil e/p^m \rceil$ and $\bar{e}^{[m,d]} := \lfloor e/p^m \rfloor$. Also write $e^{[m,u]} := p^m \bar{e}^{[m,u]} - e$ and $e^{[m,d]} := e - p^m \bar{e}^{[m,d]}$.

Definition 2.2. Given positive integers (e_1, e_2, e_3) prime to p, satisfying the triangle inequality, and with $e_1 + e_2 + e_3$ odd, we say that the triple (e_1, e_2, e_3) is *numerically p-admissible* if for any $m > 0$, and any $S \subseteq \{1, 2, 3\}$ such that:

(i) $p^m \leqslant d$;

(ii) $e_i > p^m$ for all $i \in S$;

(iii) $\sum_{i \in S} \bar{e}_i^{[m,d]} + \sum_{i \notin S} \bar{e}_i^{[m,u]}$ is odd;

the following inequality is always satisfied:

$$\sum_{i \in S} e_i^{[m,d]} + \sum_{i \notin S} e_i^{[m,u]} \geqslant p^m. \tag{2.1}$$

Here d is the integer with $2d - 2 = \sum_i (e_i - 1)$.

Note that the triangle inequality condition on (e_1, e_2, e_3), i.e. $e_1 \leqslant e_2 + e_3, e_2 \leqslant e_1 + e_3$ and $e_3 \leqslant e_1 + e_2$, is equivalent to the condition that $e_i \leqslant d$ for all i.

The notion of numerical p-admissibility intuitively corresponds to not having all three ramification indices too close to certain positive multiples of p^m, for any m. The following reformulation is also useful.

LEMMA 2.3. *We may replace (2.1) in the above definition as follows: for a given m, S, if we denote by $d^{[m,S]}$ the integer satisfying*

$$2d^{[m,S]} - 2 = \sum_{i \in S}(\bar{e}_i^{[m,d]} - 1) + \sum_{i \notin S}(\bar{e}_i^{[m,u]} - 1),$$

then (2.1) is equivalent to

$$d < p^m d^{[m,S]} + \sum_{i \in S} e_i^{[m,d]}. \tag{2.2}$$

Proof. The equivalence of the two conditions may be checked directly from the definitions, by verifying the identity

$$d - p^m d^{[m,S]} - \sum_{i \in S} e_i^{[m,d]} = \frac{1}{2}\left(p^m - 1 - \sum_{i \in S} e_i^{[m,d]} - \sum_{i \notin S} e_i^{[m,u]}\right). \qquad \square$$

Our main result gives an explicit and purely numerical criterion for the existence of maps with three ramification points in positive characteristic. The main idea is that an inseparable map will exist (and, hence, a separable map will not exist) with the desired ramification only if the e_i are too close to appropriate positive multiples of p^m, violating numerical p-admissibility.

THEOREM 2.4. *Suppose that we are given positive integers d, e_1, e_2, e_3 with $2d - 2 = \sum_i (e_i - 1)$, each e_i prime to p, and the $e_i \leqslant d$ for all i, together with distinct points Q_1, Q_2, Q_3 on $\mathbb{P}^1$. Then the triple (e_1, e_2, e_3) is numerically p-admissible if and only if there exists a (necessarily unique up to automorphism) separable cover $f : \mathbb{P}^1 \to \mathbb{P}^1$ of degree d, branched over each Q_i with a single ramification point of index e_i.*

Proof. We first observe that it does not matter whether we fix ramification points P_i or branch points Q_i. Indeed, any three points on either the source or target are automorphism-equivalent, and when there are only three ramification points, we have $e_i + e_j > d$ for any i, j, so no two ramification points can lie above a single branch point.

Since all ramification is specified, the existence of a separable map is equivalent to the existence of a separable linear series with the given degree and ramification. By [Oss06b, Theorem 4.2(i)], such a separable linear series exists if and only if no inseparable linear series exists with (at least) the specified ramification.

Now, suppose that we have such an inseparable linear series, corresponding to a map f of degree d'' with $d - d''$ base points; we can then write $d'' = p^m d'$, where f is the composition of the mth power of (the relative) Frobenius with a separable map of degree d'. Write e'_i for the ramification indices at the P_i of this separable map. Let $S \subseteq \{1, 2, 3\}$ be the subset consisting of i with $p^m e'_i < e_i$. Thus, for each P_i in S, to have the required ramification we need to have at least $e_i - p^m e'_i$ base points in our inseparable linear series. Therefore, we must have $d \geqslant p^m d' + \sum_{i \in S}(e_i - p^m e'_i)$, which we check directly may be rewritten as

$$\sum_{i \in S}(e_i - p^m e'_i) + \sum_{i \notin S}(p^m e'_i - e_i) \leqslant p^m - 1 < p^m.$$

This implies that we must have $e_i - p^m e'_i = e_i^{[m,d]}$ for all $i \in S$ and $p^m e'_i - e_i = e_i^{[m,u]}$ for all $i \notin S$.

We then find that

$$\sum_{i\in S}\bar{e}_i^{[m,d]}+\sum_{i\notin S}\bar{e}_i^{[m,u]}=\sum_i e'_i$$

is odd, and furthermore that (2.1) is violated, completing one direction of the proof.

Conversely, we may suppose that for some m, S, satisfying the conditions of Definition 2.2, we have $d \geqslant p^m d^{[m,S]} + \sum_{i\in S} e_i^{[m,d]}$. It clearly suffices to show that there exists an f' of degree $d^{[m,S]}$ having ramification at least $e'_i := \bar{e}_i^{[m,d]}$ for $i \in S$ and at least $e'_i := \bar{e}_i^{[m,u]}$ for $i \notin S$, since we can then compose with the mth power of Frobenius and add $e_i^{[m,d]}$ base points at P_i for each $i \in S$ to obtain an inseparable map of degree d with ramification at least e_i. In fact, it is enough to have f' any linear series (separable or inseparable, with or without base points), as long as it has the specified degree, and ramification e'_i at each P_i. In this case, we 'compose with the mth power of Frobenius' by composing the associated map to $\mathbb{P}^1$ with the mth power of Frobenius, and multiplying the base point divisor by p^m; this produces a new linear series, multiplying both the degree and all of the ramification indices by p^m.

Thus, as long as the e'_i satisfy the triangle inequality (which we recall is equivalent to the condition that $e'_i \leqslant d^{[m,S]}$ for all i), by [Oss06b, Theorem 4.2(i)] we can find such an f'. Although the triangle inequality for the e_i alone is not enough to show the triangle inequality for the e'_i, one can check using the equivalence of the two inequalities of the lemma that the triangle inequality for the e_i together with the hypothesized inequality $d \geqslant p^m d^{[m,S]} + \sum_{i\in S} e_i^{[m,d]}$ does in fact imply that the e'_i satisfy the triangle inequality, giving us our inseparable map of degree d, and thereby showing that no separable map can exist. □

We also recall the following.

Corollary 2.5. *In the situation of the theorem, if, in addition e_1, e_2 are less than p, numerical p-admissibility is equivalent to the condition that $d < p$.*

Proof. This may be derived directly from the theorem, but also predates it, see [Oss06b, Theorem 4.2(ii)]. □

We now provide some examples to demonstrate the usage of the combinatorial condition of the theorem.

Example 2.6. The indices $(1, d, d)$ for d prime to p provide trivial examples of ramification indices being far enough away from multiples of p^m that an inseparable map never exists, as exhibited by the separable map x^d. Less trivial is the corollary, for instance, that if we take indices $(2, d-1, d)$ for $d \not\equiv 0, 1 \pmod p$ and $p > 2$, we necessarily obtain a separable map.

More substantively, we examine the case that $e_i < 2p$ for all i. The above corollary treats the case that at least two of the e_i are less than p, so we may assume that at most one e_i is less than p. To rule out exceptional cases, we also assume $p > 3$. First, let us suppose that $e_1 < p$, but $e_2, e_3 > p$. In this case, a map exists if and only if $d < 2p$ and $e_2 + e_3 - e_1 \geqslant 2p$. On the other hand, if $e_i > p$ for all i, a map exists if and only if $d \geqslant 2p$ and $e_2 + e_3 - e_1, e_1 + e_3 - e_2, e_1 + e_2 - e_3 \leqslant 2p$.

For both cases, $m = 1$ is the only possibility. In the first case, we see that the possibilities for S are $S = \{2, 3\}$ or $S = \emptyset$, which corresponds to the two inequalities. In the second case, we can have $S = \{1, 2, 3\}, \{1\}, \{2\}$ or $\{3\}$, corresponding to the four given equalities.

We see that for the first case, the requirement is that d is not too large and the e_i are not too symmetric, whereas in the second case, the requirement is that d is sufficiently large and the e_i are sufficiently symmetric. This may seem counterintuitive, but in fact makes sense: if $d \geqslant 2p$, we can compose the Frobenius map with a degree two map to obtain a map ramified to order $2p$

https://doi.org/10.1112/S0010437X0700303X Published online by Cambridge University Press

at two points, and to order p elsewhere; in particular, if $e_1 < p$ and $e_2, e_3 < 2p$, then such a map immediately gives the required ramification. Similarly, the more symmetric the e_i are, the easier it is to obtain the ramification by starting with the Frobenius map and adding base points. In contrast, once we have $e_1, e_2, e_3 > p$, we need them to be large enough that we cannot simply add base points to the Frobenius map, and symmetric enough that the map of degree $2p$ is too asymmetric to yield the required ramification.

3. Branched covers and linear series

Before proceeding to treat the case of covers with more branch points, we need to develop some preliminary results relating the existence questions for linear series and branched covers. That is, we compare the situation for the existence of maps with given ramification points on a fixed source curve, to maps with branch points fixed on the base. The reason for this is two-fold: first, our degeneration techniques for linear series are best-suited to answer existence questions only for general configurations of ramification points, and we would like to be able to strengthen such results; second, the relationship between the two perspectives has certain subtleties, as illustrated by the first remark below.

Our main result is the following.

Theorem 3.1. *Let $d, r, e_1, \ldots, e_r$ be positive integers with $e_i \leqslant d$ and prime to p for all i, and $2d - 2 + 2g = \sum_i (e_i - 1)$. Then the following are equivalent:*

- (a) *there exists a smooth curve C of genus g with distinct $P_i \in C$ and a separable map $f : C \to \mathbb{P}^1$ ramified to order e_i at the P_i;*
- (b) *for general $Q_i \in \mathbb{P}^1$, there exists a smooth curve C of genus g and a separable map $f : C \to \mathbb{P}^1$ of degree d, branched over each Q_i, with a single ramification point of order e_i.*

If, furthermore, $g = 0$ and all of the e_i are less than p, we also have the following equivalent condition:

- (c) *for general $P_i \in \mathbb{P}^1$, there exists a separable map $f : \mathbb{P}^1 \to \mathbb{P}^1$ ramified at each P_i to order e_i.*

Proof. The equivalence of conditions (a) and (b) is well-known via deformation arguments, but also follows immediately from [Oss05, Corollary 3.2]; the idea of the latter is to combine a classical codimension count from the point of view of linear series with easy deformation theory of branched covers to compute the total dimension of the space of maps from curves of genus g to $\mathbb{P}^1$ with the desired ramification.

Now, we want to show that if $g = 0$ and all e_i are less than p, we also have that condition (a) implies condition (c). By [Oss06b, Appendix], we have a moduli scheme $MR := MR(\mathbb{P}^1, \mathbb{P}^1, (e_1, \ldots, e_r))$ parametrizing tuples $(f, (P_1, \ldots, P_r))$, where $f : \mathbb{P}^1 \to \mathbb{P}^1$ is a separable morphism of degree d, the P_i are distinct k-valued points of the domain, and f is ramified to order at least e_i at P_i. This comes with natural forgetful morphisms $\text{ram} : MR \to (\mathbb{P}^1)^r$ and $\text{branch} : MR \to (\mathbb{P}^1)^r$ giving the ramification points and branch points, which is to say the P_i and $f(P_i)$, respectively. From the argument of [Oss05, Corollary 3.2], we know that MR has dimension exactly $r + 3$ if it is non-empty.

The main tool for the argument is a finiteness result, which is currently only available in the situation that all of the e_i are odd, so we will induct on the number of even e_i, noting that it must always be an even number for the degree to be integral. In the base case that all of the e_i are odd, by [Oss07, Theorem 5.3] there can be only finitely many maps for any given choice of the P_i, up to automorphism of the image space. Thus, all of the fibers of the ram morphism are at most three-dimensional, and since MR has dimension $r + 3$ if it is non-empty, we find that the ram morphism must dominate the $(\mathbb{P}^1)^r$ parametrizing ramification point configurations, as desired. For the induction step, without loss of generality we can suppose that e_1, e_2 are even. Suppose that MR is non-empty and fails to dominate $(\mathbb{P}^1)^r$ under the ram morphism. MR is necessarily dominant under the branch morphism, so let $(f, (P_1, \dots, P_r))$ be a point in it such the all of the $f(P_i)$ are distinct. However, then, by [Oss06b, Lemma 5.2] we see that if we replace e_1, e_2 by $p-e_1, p-e_2$, and denote the new moduli scheme by MR', then $\operatorname{ram}(MR')$ still fails to dominate $(\mathbb{P}^1)^r$, and at our specific P_i, we obtain a new f' having ramification $p-e_1, p-e_2, e_3, \dots, e_r$ at the P_i, contradicting the induction hypothesis. □

Remark 3.2. Note that we cannot hope to drop the hypothesis that all e_i are less than p in order to obtain the equivalence of condition (c). Indeed, for $p > 2$ if we consider the family of functions $x^{p+2} + tx^p - x$ as t is allowed to vary arbitrarily, we see that we obtain an infinite (tamely ramified) family of maps for which the ramification points remain fixed, while the branch points (necessarily) move. While such maps do exist, they occur only for special configurations of ramification points; see [Oss06b, Proposition 5.4] for details, and Example 6.2 below for a more detailed examination of a similar example. Note, however, that (still in the case $g = 0$) our argument shows that the equivalence of condition (c) will hold whenever we know that there are only finitely many linear series with the given ramification for arbitrary distinct configurations of the P_i.

Remark 3.3. The finiteness result of [Oss07] used to show the equivalence of condition (c) is atypical in that it is proved via a relationship to certain connections with vanishing p-curvature, Mochizuki's dormant torally indigenous bundles. The fundamental obstruction to obtaining a direct proof is that in the setting of connections, one first enlarges to the category of connections with nilpotent p-curvature before proving finiteness, and it is not clear what the analogous construction would be in the context of rational functions on $\mathbb{P}^1$.

4. Numerical results: the case that $e_i < p$ for all i

We now put together the results of § 3 with the results on linear series of [Oss06b] to obtain a sharp statement on the existence and non-existence of covers $\mathbb{P}^1 \to \mathbb{P}^1$ with given ramification indices, in the case that all indices are less than p. Our methods are algebraic and intrinsically positive characteristic.

We first state the numerical condition which will be equivalent to existence of tame covers in the case that $e_i < p$.

Definition 4.1. Fix $e_1, \dots, e_r$ positive integers, with $\sum_i (e_i - 1)$ even, and further assume that $e_i < p$ for all i.

We say that $(e_1, \dots, e_r)$ is *numerically p-admissible* if there exist $r-3$ positive integers $e'_2, \dots, e'_{r-2}$, prime to p, such that for any m with $1 \leqslant m < r-1$:

(i) the triple $(e'_m, e_{m+1}, e'_{m+1})$ satisfies the triangle inequality;

(ii) $e'_m + e_{m+1} + e'_{m+1}$ is odd and less than $2p$;

where we use the convention $e'_1 := e_1$ and $e'_{r-1} := e_r$.

Note that for $r = 3$, it follows from Corollary 2.5 that the definition of numerical p-admissibility agrees with the definition already given in § 2.

In fact, numerical p-admissibility is equivalent to the existence of a cover coming from smoothing a totally degenerate cover, i.e. one constructed from a collection of maps $\mathbb{P}^1 \to \mathbb{P}^1$, with three ramification points each. The e_i in this situation are the ramification points above marked points, while the e'_i are the ramification points above nodes; see Figure 1.

https://doi.org/10.1112/S0010437X0700303X Published online by Cambridge University Press

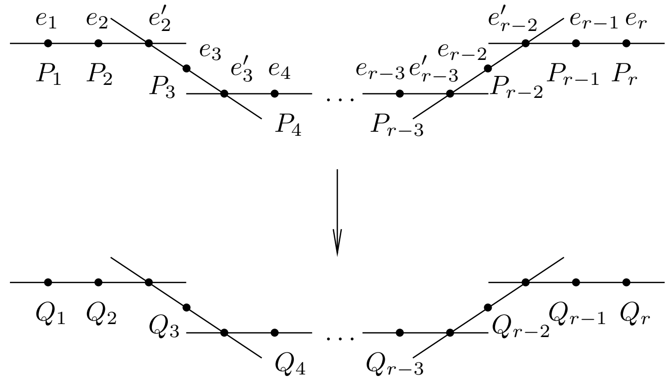

FIGURE 1. The geometry behind numerical p-admissibility.

We further remark that there is a rather concrete interpretation of numerical p-admissibility when $e_i < p$ for all i, which says roughly that there have to be enough e_i which are not too large or small. See Example 4.4 below for details.

Our theorem, which is a direct application of § 3 and the results of [Oss06b], is the following.

THEOREM 4.2. *Fix d, r and $e_1, \ldots, e_r$ positive integers, with $e_i < p$ for all i and satisfying $2d - 2 = \sum_i (e_i - 1)$. Then the following are equivalent.*

(a) *The tuple $(e_1, \ldots, e_r)$ is numerically p-admissible.*
(b) *For general $P_i \in \mathbb{P}^1$ there exists a separable map $f : \mathbb{P}^1 \to \mathbb{P}^1$ of degree d, ramified to order e_i at each P_i.*
(c) *For general $Q_i \in \mathbb{P}^1$ there exists a separable map $f : \mathbb{P}^1 \to \mathbb{P}^1$ of degree d, branched over each Q_i with a single ramification point of index e_i.*
(d) *There exists a separable map $f : \mathbb{P}^1 \to \mathbb{P}^1$ of degree d, ramified to order e_i at some distinct P_i.*

Proof. The $g = 0$ case of Theorem 3.1 tells us that parts (c) and (d) are equivalent, and that because all the e_i are less than p, we also have that parts (b) and (c) are equivalent. The equivalence of parts (a) and (b) is due to prior work using limit linear series and controlling degeneration from separable to inseparable maps; see [Oss06b, Theorem 1.4]. □

In order to illustrate how the (in principle, rather complicated) combinatorial conditions of the theorem may be applied, we examine some examples of existence and non-existence results.

Example 4.3. Existence examples are easy to construct: one simply needs to construct a chain

$$e_1, e_2, e_2', \ldots, e_{r-2}, e_{r-2}', e_{r-1}, e_r$$

such that each triple $(e_1, e_2, e_2'), (e_2', e_3, e_3'), \ldots, (e_{r-2}', e_{r-1}, e_r)$ satisfies numerical p-admissibility for three points. For instance, since the triple $(p-2, (p-1)/2, (p-1)/2)$ is always numerically p-admissible for $p > 2$, if r is a multiple of three we can use the sequence

$$\frac{p-1}{2}, p-2, \frac{p-1}{2}, \frac{p-1}{2}, p-2, \frac{p-1}{2}, \ldots, \frac{p-1}{2}, p-2, \frac{p-1}{2}$$

to obtain a map with $e_i = p - 2$ when $i \equiv 2 \pmod 3$ and $e_i = (p-1)/2$ otherwise (and we can make similar constructions in the cases that $r \equiv 1, 2 \pmod 3$).

For non-existence, the basic idea is that having e_i very large or very small rigidifies the possibilities for the e_i' (the former because of the condition $e_m' + e_{m+1} + e_{m+1}' < 2p$, and the latter because of the triangle inequalities), and can be used to ensure that no sequence of e'_i can satisfy all conditions at once. We will see that in essence, a map will exist unless all of our e_i are sufficiently close to p or 1, where 'sufficiently close' depends on p and on r.

Example 4.4. For our first example, if r is odd, and we set $e_i = p-1$ for $i < r-1$, and $e_{r-1}+e_r > p$ and odd, we see that the summation condition on (e_1, e_2, e'_2) determines $e'_2 = 1$, and the triangle inequalities on (e'_2, e_3, e'_3) determine $e'_3 = p-1$, and so forth, alternating until e'_{r-2} is determined as $p-1$. However, then (e'_{r-2}, e_{r-1}, e_r) violates the condition that their sum is less than $2p$.

Similarly, if r is even, with $e_i = p-1$ for $i < r-1$, and $e_{r-1} \neq e_r$ with $e_{r-1}+e_r$ even, we find that e'_{r-2} is determined as 1, and then (e'_{r-2}, e_{r-1}, e_r) violates the triangle inequality.

One can construct many more examples in this way: if the e_i are less than $p-1$ but still close to p, there is more flexibility, so one has to put restrictions on r. Similarly, one can use several small e_i, although as long as $e_i \geqslant 2$ each one will introduce some flexibility.

Finally, we mention one other example of our results: the case of four points.

Example 4.5. Suppose that $r = 4$, with $e_i < p$ for all i. We claim that we have the explicit formula that a cover exists if and only if $e_i > d+1-p$ for all i. Indeed, applying Theorem 3.1 and [Oss06b, Corollary 8.1] (which is just an explicit computation with the combinatorial condition of the corollary), we see that a map exists if and only if

$$\min\{e_i, d+1-e_i, p-e_i, p-d-1+e_i\}_i > 0.$$

We have $e_i, p-e_i, d+1-e_i > 0$ by hypothesis, so the only possibility for non-existence is that $p-d-1+e_i \leqslant 0$ for some i, which gives the desired statement.

Remark 4.6. Although the existence of maps with given ramification indices clearly does not depend on the order of the indices, our criterion is asymmetric. This reflects the fact that one can obtain the same result by degenerating to different totally degenerate curves, and obtain non-trivial combinatorial relations as a result. For details on some of these relations in a slightly different setting, see [LO06].

5. Group-theoretic results

In this section, we reformulate our previous numerical results on branched covers in terms of the sharper invariant of monodromy groups, ultimately proving Theorem 1.4. Unlike the prior results, our results here will be dependent on transcendental techniques, as even our statements work systematically with local generating systems of tame fundamental groups. In § 7, we discuss how the various restrictions of our main theorem can be relaxed in various ways to obtain large families of existence and non-existence results for additional covers.

The translation from numerical results to group-theoretic results is much simpler in the three-point case, thanks to the following lemma.

Lemma 5.1. *Given $(d, 3, \{e_1, e_2, e_3\})$, there exists a unique Hurwitz factorization $(\sigma_1, \sigma_2, \sigma_3)$, up to simultaneous relabelling.*

Proof. One can argue the uniqueness using the discussion of § 2, but in fact the unique Hurwitz factorization in this case may also be described explicitly; see [LO, Lemma 2.1]. □

We next state the promised definition of p-admissibility. First, note that the pure braid group acts on Hurwitz factorizations in exactly the same manner as on local generating systems.

https://doi.org/10.1112/S0010437X0700303X Published online by Cambridge University Press

DEFINITION 5.2. Let $(\sigma_1, \dots, \sigma_r)$ be a Hurwitz factorization for $(d, r, \{e_1, \dots, e_r\})$, where all e_i are prime to p, and $2d - 2 = \sum_i (e_i - 1)$.

If $r = 3$, we say that $(\sigma_1, \sigma_2, \sigma_3)$ is *p-admissible* if (e_1, e_2, e_3) is numerically p-admissible (Definition 2.2).

If $e_i < p$ for all i, we say that $(\sigma_1, \dots, \sigma_r)$ is *p-admissible* if there exists a pure-braid transformation replacing $(\sigma_1, \dots, \sigma_r)$ by $(\sigma'_1, \dots, \sigma'_r)$ and such that:

(i) for any m with $1 \leqslant m \leqslant r - 1$, the partial product $\sigma''_m := \sigma'_1 \cdots \sigma'_m$ is a cycle;

(ii) for any m with $1 \leqslant m < r - 1$, the sum of the lengths of $\sigma''_m, \sigma'_{m+1}, \sigma''_{m+1}$ is less than $2p$.

The condition for $r > 3$ has a geometric interpretation in terms of totally degenerate covers; see the discussion following Definition 4.1 as well as Figure 1 above.

As with numerical p-admissibility, if $r = 3$ and $e_i < p$ for all i, Corollary 2.5 implies that the two above definitions are equivalent (and, indeed, that no pure braid transformation is necessary).

We also mention that if $(\sigma_1, \dots, \sigma_r)$ is p-admissible, then one checks easily that the tuple of lengths $(e_1, \dots, e_r)$ of the cycles is numerically p-admissible, by letting the e'_i be the lengths of the σ''_i.

The basic lemma required to go from our numerical results to the final group-theoretic statements is the following. Although the proof of the lemma is by explicit construction, the intuition comes from considering monodromy groups of certain admissible covers. See also Figure 1 above and Remark 7.5 below.

LEMMA 5.3. *Suppose that $e_i < p$ for all i, and $(e_1, \dots, e_r)$ is numerically p-admissible (Definition 4.1). Let d be determined by $2d - 2 = \sum_i (e_i - 1)$. Then there exists a Hurwitz factorization for $(d, r, \{e_1, \dots, e_r\})$ which is p-admissible.*

More precisely, given any $e'_2, \dots, e'_{r-2}$ verifying numerical p-admissibility, there is a Hurwitz factorization $(\sigma_1, \dots, \sigma_r)$ such that the partial products σ''_i are cycles of length e'_i.

Proof. In fact, we produce a Hurwitz factorization which satisfies the conditions for p-admissibility without any braid transformation. The proof is inductive, with a base case of $r = 3$. Indeed, the $r = 3$ case is immediate from the definition, once we know that a Hurwitz factorization exists, which follows from Lemma 5.1.

For the induction step, suppose that our assertion holds for $r - 1$. Suppose also we are given $(e_1, \dots, e_r)$, together with $e'_2, \dots, e'_{r-2}$ satisfying the conditions for the r-tuple to be numerically p-admissible. We then note that $(e_1, \dots, e_{r-2}, e'_{r-2})$ is also numerically p-admissible, with degree $d' = d - (e_{r-1} + e_r - e'_{r-2} - 1)/2$. By the induction hypothesis, we can find some Hurwitz factorization $(\sigma_1, \dots, \sigma_{r-2}, \sigma'_{r-2})$ for $(d', r - 1, \{e_1, \dots, e_{r-2}, e'_{r-2}\})$ which satisfies the conditions for p-admissibility without any braid transformation. We then note that (e'_{r-2}, e_{r-1}, e_r) is also numerically p-admissible, with degree $d'' = (e'_{r-2} + e_{r-1} + e_r - 1)/2$. By the base case, we can find a corresponding Hurwitz factorization $(\sigma''_{r-2}, \sigma_{r-1}, \sigma_r)$. In particular, the existence of the two Hurwitz factorizations implies that we have $e'_{r-2} \leqslant d', d''$. Now, we observe that $d' + d'' = d + e'_{r-2}$, so we see that $d', d'' \leqslant d$, and we can map $\{1, \dots, d'\}$ and $\{1, \dots, d''\}$ into $\{1, \dots, d\}$ such that:

(i) $\{1, \dots, d'\}$ maps into $\{1, \dots, d'\}$ and $\{1, \dots, d''\}$ maps into $\{d' - e'_{r-2} + 1, \dots, d\}$;

(ii) σ'_{r-2} maps to the inverse of σ''_{r-2}.

If we consider $(\sigma_1, \dots, \sigma_r)$ as lying in S_d via these maps, we then check easily that they give a Hurwitz factorization for $(d, r, \{e_1, \dots, e_r\})$, and furthermore satisfy the conditions for p-admissibility without any braid transformation. □

We also recall the main theorem of [LO].

THEOREM 5.4 (Liu–Osserman). *Given* $(d, r, \{e_1, \ldots, e_r\})$ *with* $2d-2 = \sum_i (e_i - 1)$, *any two Hurwitz factorizations are related by a pure braid transformation.*

Putting together the theorem and the lemma (together with the earlier observation that p-admissibility implies numerical p-admissibility), we conclude the following.

COROLLARY 5.5. *A Hurwitz factorization* $(\sigma_1, \ldots, \sigma_r)$ *for* $(d, r, \{e_1, \ldots, e_r\})$ *is* p*-admissible if and only if* $(e_1, \ldots, e_r)$ *is numerically* p*-admissible.*

We are now ready to prove Theorem 1.4.

Proof of Theorem 1.4. If $(\sigma_1, \ldots, \sigma_r)$ is p-admissible, then $(e_1, \ldots, e_r)$ is numerically p-admissible, so we see by Theorems 2.4 and 4.2 that there exists a cover f with ramification indices $(e_1, \ldots, e_r)$, and we conclude that part (a) implies (b).

Now, a cover as in part (b) has monodromy $(\sigma'_1, \ldots, \sigma'_r)$ around $(\gamma_1, \ldots, \gamma_r)$ for some Hurwitz factorization $(\sigma'_1, \ldots, \sigma'_r)$. We then have by Theorem 5.4 that $(\sigma_1, \ldots, \sigma_r)$ is related to $(\sigma'_1, \ldots, \sigma'_r)$ by some pure braid transformation, so it follows that if we replace $(\gamma_1, \ldots, \gamma_r)$ by $(\gamma'_1, \ldots, \gamma'_r)$ under the same transformation, the local monodromy of f around γ'_i is given by σ_i, so we have that part (b) implies part (d).

On the other hand, it is clear that part (d) implies part (c) which implies part (b), so it only remains to check that part (b) implies part (a) under our hypotheses. Accordingly, suppose that we have a cover f with ramification indices $(e_1, \ldots, e_r)$ at points $P_1, \ldots, P_r$, and suppose that either $e_i < p$ for all i, or that $r = 3$. We claim that $(e_1, \ldots, e_r)$ is numerically p-admissible: indeed, in the first case, this follows by Theorem 4.2, while in the second, it follows from Theorem 2.4. Thus, by Corollary 5.5, we see that $(\sigma_1, \ldots, \sigma_r)$ is p-admissible, so in either of these cases, we have that part (b) implies part (a), as desired.

The last assertion is that for $r = 3$, no braid operations are necessary: this follows from the uniqueness of the Hurwitz factorization (Lemma 5.1) in that case. □

We give a simple example of the theorem.

Example 5.6. From the existence portion of our main theorem, we recover examples along the lines of covers constructed by Bouw and Wewers [BW05], and Harbater and Stevenson [HS99].

Suppose $p > 3$ and $r \geqslant 3$, and fix general points $Q_1, \ldots, Q_r$ on $\mathbb{P}^1$. Then there exists a tame cover $f : \mathbb{P}^1 \to \mathbb{P}^1$ of degree $d = r + 1$, with monodromy group A_d, and with a single ramified point over each branch point, of order three.

By [LO, Theorem 5.3], any such cover has monodromy group A_d, so it suffices to note that the tuple $(\underbrace{3, \ldots, 3}_{r \text{ times}})$ is always numerically p-admissible for $p > 3$ and $r \geqslant 3$. Indeed, if we set $e'_i = 3$ for all i, we see that the conditions will be satisfied.

For some non-existence examples following from the theorem, see § 7.

Remark 5.7. Guralnick conjectures (see [GS, p. 2] for details) that other than finitely many possibilities, all primitive genus-zero groups should be cyclic, alternating, symmetric or one of a few other possible families. Thus, even though all of the monodromy groups in Theorem 1.4 are one of the three types mentioned [LO, Theorem 5.3], since they are all primitive (except in the cyclic case with composite degree), the situation is actually reasonably general. That said, our non-existence results can also be immediately applied to draw conclusions on a far wider class of monodromy groups arising as imprimitive subgroups of S_d; see Theorem 7.1 below.

6. Two examples

In this section we compute two elementary examples which shed light on the subtleties of the situation in characteristic p, particularly regarding: the difficulty of using Riemann existence-style theorems to compute Hurwitz numbers; the peculiarly good behavior of degenerations; and (for the sake of completeness, although this is well-known) the difficulties of dropping the generality of branch points in existence statements. In particular, we prove Proposition 1.5.

Example 6.1. We explicitly compute the situation for covers of $\mathbb{P}^1$ of degree three, with four simple branch points. We may assume that 0, 1, and ∞ are three of the four ramification points, and each is mapped to itself. We are thus considering functions of the form $f(x) = (ax^3 + bx^2)/(cx + d)$ where a, b, c, d are all non-zero, and we can therefore choose to set $a = 1$. Furthermore, since $f(1) = 1$, we have $d = b+1-c$. We also suppose that f is simply ramified at some λ, and that $f(\lambda) = \mu$; we then want to compute f in terms of μ to find all possibilities with particular prescribed branch points. Differentiating f, and imposing the desired conditions, we find that $c = (2b^2 + 5b + 3)/(b + 1) = 2b + 3$, and that $\lambda = (-b^2 - 2b)/(2b + 3)$. Finally, the condition for $f(\lambda) = \mu$ can be simplified to

$$b^4 + (2 + 8\mu)b^3 + 36\mu b^2 + 54\mu b + 27\mu = 0.$$

Over $\mathbb{C}$, this quartic corresponds to the four solutions one obtains by classical cycle decompositions. Specifically, if we fix a local generating system $(\gamma_1, \dots, \gamma_4)$ for $\pi_1^{\mathrm{top}}(\mathbb{P}^1 \smallsetminus \{P_i\}_i)$, by considering monodromy around the γ_i, we have that covers are in one-to-one correspondence with Hurwitz factorizations for $(d = 3, r = 4, \{2, 2, 2, 2\})$, which up to equivalence are

$$(12)(12)(23)(23)$$
$$(12)(23)(23)(12)$$
$$(12)(23)(31)(23)$$
$$(12)(23)(12)(31)$$

We now turn to the case of characteristic 3. Here, the formula for b in terms of μ reduces to

$$b^4 + (-\mu - 1)b^3 = 0,$$

and since b must be non-zero, we obtain the unique solution $b = 1 + \mu$, which corresponds to the function $f = (x^2(x + 1 + \mu))/((-\mu - 1)x - \mu)$, and which gives $\lambda = \mu$.

Example 6.2. We next consider a similar example, with $d = 4$, $e_1 = 4$ and $e_2 = e_3 = e_4 = 2$. Normalizing as before, we can write $f(x) = ax^4 + bx^3 + cx^2$ with a, c non-zero and $a + b + c = 1$. Differentiating to impose ramification, we find $b = 4 - 2c$, and $\lambda = 2c/4(c - 3)$. Finally, imposing $f(\lambda) = \mu$ gives us that c is a root of

$$c^4 - (4 + 16\mu)c^3 + 144\mu c^2 - 432\mu c + 432\mu = 0.$$

The quartic polynomial corresponds to the Hurwitz factorizations

$$(1234)(12)(43)(31)$$
$$(1234)(12)(14)(43)$$
$$(1234)(12)(31)(14)$$
$$(1234)(13)(14)(23)$$

As before, in characteristic zero (or with $p > 3$) we have four covers. However, once again, in characteristic 3 we see that we get a single cover, with $c = \mu + 1$. Note here the oddity that in characteristic 3, we have $\lambda = -1$ is always fixed, so that although a cover exists for general μ, in fact a map does not exist for general λ, and there are infinitely many for $\lambda = -1$. In particular,

https://doi.org/10.1112/S0010437X0700303X Published online by Cambridge University Press

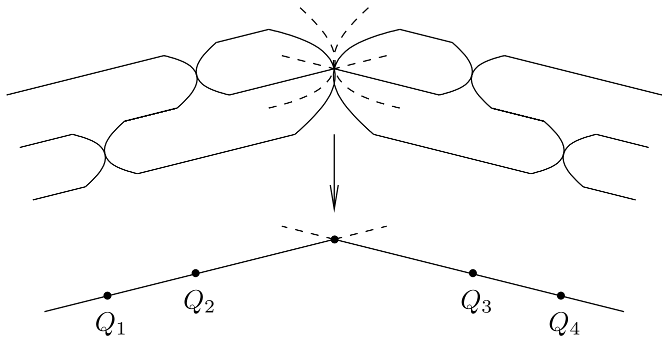

Figure 2. The phantom admissible cover.

this map does not come from a totally degenerate map as in Figure 1. This is part of a much more general phenomenon; see [Oss06b, Proposition 5.4].

We wish to emphasize two phenomena. The first is the following: although it is a feature of the classical situation that, thanks to the Riemann existence theorem, which Hurwitz factorizations are realized as the monodromy of a cover around a local generating system is independent of the choice of system, the same statement fails in positive characteristic, even for tame covers. Indeed, both these examples demonstrate the phenomenon of Proposition 1.5 (and, in particular, either example proves the proposition). In both cases, all four Hurwitz factorizations lie in a single braid orbit, as can be checked directly or follows from [LO]. Since there is only a single cover in characteristic three, if we start with any local generating system, we see that which Hurwitz factorization we get varies between all four possibilities as we change the generating system by pure braid operations.

This phenomenon also makes it quite difficult to compute Hurwitz numbers using a Riemann-existence-style theorem, as we see that although we can identify which Hurwitz factorizations arise as the monodromy of covers around some local generating system, there is no group-theoretic criterion which works independently of the choice of local generating system, so it seems quite subtle to use such criteria to actually count covers.

The second phenomenon that we wish to discuss is that degenerations in these examples behave 'unreasonably well', in a different sense for each example. In the case of the first example, we had originally hoped to give an example where covers always degenerate from separable to inseparable. Roughly, we start with a cover having monodromy cycles $(1,2),(2,3),(2,3),(1,2)$ at $P_1,\dots,P_4$ for some local generating system (which we know exists). We then degenerate the base to a curve with two components, with P_1,P_2 on one component and P_3,P_4 on the other. If the cover remained separable under this degeneration, it seems we would have to have a cover as in Figure 2, with monodromy at the node given by $(3,2,1)$ and $(1,2,3)$ on the two components, and this would not be possible in characteristic 3. However, we see that we in fact cannot have bad degeneration: there is one cover in the smooth case and also one cover (corresponding to monodromy $(1,2),(1,2),(2,3),(2,3)$, unramified over the node) after degeneration. What is going on is that for our analysis to work, the local generating system we choose on the generic fiber has to specialize to a 'geometric' system in the sense of Remark 7.5 below, so the only possible explanation is that while a cover with the specified monodromy exists for certain local generating systems, none of these systems specialize to geometric ones under the given degeneration.

On the other hand, in the second example, if we consider the same degeneration of the base, we see again that covers degenerate well, because again there is a single cover of both the smooth and degenerate curves. In this case, the surprising aspect is that degeneration is well-behaved from the point of view of covers, despite the fact that from the linear series point of view, the map cannot degenerate to a separable map (in fact, we saw that λ cannot move at all). We thus see that in two different ways, situations where we might expect to find a separable cover degenerating to an inseparable one do not in fact give examples of bad degeneration.

We do however mention in this context that Bouw has an example, obtained from [Bou01, Proposition 7.8], of a higher-genus cover with four branch points such that degenerations are always inseparable. However, the argument in that case involves p-rank considerations, which one might hope to understand completely. If there are no other obstructions to having good behavior under (suitably generic) degeneration, it should be possible to use degeneration techniques to obtain a range of results beyond those presented here.

Finally, we mention that when $\mu = -1$ in either example, the unique map specializes to an inseparable map. We thus find that even though a separable cover exists for μ general, we have a configuration of four distinct points, namely $-1, 0, 1, \infty$ over which no separable cover exists. Again, although this is already well-known, we see that we have very simple examples showing the necessity of considering general branch points in order to obtain any group-theoretic condition for a Hurwitz factorization to be realized as monodromy of a cover.

7. Further results and discussion

We begin with a discussion of the variety of ways in which the restrictions of Theorem 1.4 can be relaxed. We state one generalized non-existence result below, Theorem 7.1, and give examples. This is followed by a purely group-theoretic result which comes out of our arguments, and a discussion of how one might use geometric arguments to prove a result such as Theorem 1.4 without the benefit of the main result of [LO].

We first remark that standard techniques can be used to vastly generalize our existence results. We do not state any general results here, because they tend to be complicated and we do not have any evidence that they are sharp. However, we do point out that by taking the Galois closures of the covers constructed in Theorem 1.4, we obtain large families of existence results on Galois covers, many of which will have higher genus. Furthermore, by starting from the genus-zero covers we have already constructed, gluing them together to obtain an admissible cover with higher genus, and then deforming the admissible cover to obtain a smooth cover, we obtain many more higher-genus covers.

Along these lines, Fujiwara [Fuj90] has used similar techniques to show that the prime-to-p part of the fundamental group of any curve can be computed starting from the fundamental group of $\mathbb{P}^1$ with three marked points. One could broaden the array of examples even further by using the array of other covers known to existence, for instance from the papers [Ray99], [BW05] and [HS99] cited earlier, as well as additional work of Stevenson [Ste96], Saidi [Sai07] and Raynaud [Ray00].

Next, we observe that Theorem 1.4(b) makes no hypotheses on the ramification points mapping to distinct branch points. Thus, by letting branch points coincide, and composing covers, we can considerably generalize our non-existence results as follows.

Theorem 7.1. *Let $\{\sigma_1, \dots, \sigma_r\}$ be a Hurwitz factorization of degree d and genus g. That is, the σ_i have trivial product and generate a transitive subgroup of S_d, and if $(e_1, \dots, e_m)$ are the lengths (with multiplicity) of all of the cycles in the disjoint cycle representations of the σ_i, we have $2d-2+2g = \sum_i (e_i - 1)$.*

Suppose that the σ_i act on blocks of size m, so that they define permutations $\sigma'_1, \dots, \sigma'_r$ in $S_{d/m}$, and suppose further that if $e_1, \dots, e_n$ are the lengths of all of the cycles in the disjoint cycle

https://doi.org/10.1112/S0010437X0700303X Published online by Cambridge University Press

representations of the σ'_i, that:

(i) *either $e_i < p$ for all i, or $n = 3$;*
(ii) $2(d/m) - 2 = \sum_i (e_i - 1)$;
(iii) *the tuple $(e_1, \ldots, e_n)$ is not numerically p-admissible.*

Then in characteristic p, there is no cover of degree d and genus g and with monodromy $\{\sigma_1, \ldots, \sigma_r\}$ (around any choice of local generating system).

Note that as a special case of this we obtain non-existence results for Galois covers as well, by considering the Galois closures of the covers considered in Theorem 1.4. Combining this with the above remark on existence, we have sharp existence results on those Galois covers which are Galois closures of the covers treated by Theorem 1.4.

Proof. Suppose that $\bar{f}$ is a cover with monodromy $\{\sigma_1, \ldots, \sigma_r\}$ in characteristic p. Then we can lift it to a cover f in characteristic zero with the same monodromy. Under the stated hypotheses, we have that f factors through a cover $g : \mathbb{P}^1 \to \mathbb{P}^1$ with monodromy $\{\sigma'_1, \ldots, \sigma'_r\}$, and since f specializes to a separable cover in characteristic p, it follows that g also specializes to a separable cover $\bar{g}$, having monodromy $\{\sigma'_1, \ldots, \sigma'_r\}$. However, the latter cover cannot exist by the implication that part (b) implies part (a) of Theorem 1.4. □

We give two examples of the theorem, the first a genus-zero primitive cover, and the second a higher-genus imprimitive cover.

Example 7.2. We see that the genus-zero cover of degree nine, branched over three points with monodromy given by

$$\sigma_1 = (1, 2, 3, 4)(5, 6, 7, 8), \quad \sigma_2 = (8, 9, 2, 1)(4, 3, 6, 5), \quad \sigma_3 = (1, 5)(9, 8, 7, 3),$$

does not exist in characteristic $p = 5$, as the tuple $(4, 4, 4, 4, 4, 2)$ is not numerically p-admissible.

Example 7.3. We consider the Hurwitz factorization with $d = 10$ and $g = 1$ and three branch points, with monodromy

$$\sigma_1 = (1, 3, 5, 8, 2, 4, 6, 7), \quad \sigma_2 = (10, 8, 6, 4, 9, 7, 5, 3), \quad \sigma_3 = (10, 3, 1, 9, 4, 2)(7, 8).$$

We see that this is imprimitive, acting on the blocks $[1, 2], [3, 4], [5, 6], [7, 8], [9, 10]$ as cycles of length 4, 4 and 3, so this cover in characteristic zero factors through the genus zero cover of degree five corresponding to the Hurwitz factorization

$$\sigma'_1 = (1, 2, 3, 4), \quad \sigma'_2 = (5, 4, 3, 2), \quad \sigma'_3 = (5, 2, 1).$$

In characteristic five, we have that the latter cover does not exist, so we conclude that the genus-one cover with the given monodromy also does not exist in characteristic 5.

We next make the following simple group-theoretic observation which we do not believe is obvious without some type of geometric argument.

Proposition 7.4. *Given a Hurwitz factorization $(\sigma_1, \ldots, \sigma_r)$ for $(d, r, \{e_1, \ldots, e_r\})$, where $2d - 2 = \sum_i (e_i - 1)$, then there exists a pure braid transformation $(\sigma'_1, \ldots, \sigma'_r)$ of $(\sigma_1, \ldots, \sigma_r)$ such that each partial product $\prod_{i=1}^m \sigma'_i$ is a cycle, for $1 \leqslant m \leqslant r$.*

Proof. Indeed, if we start with such a Hurwitz factorization, if we consider p sufficiently large, the tuple is automatically p-admissible, so applying Lemma 5.3 we know there exists some $(\sigma'_1, \ldots, \sigma'_r)$ of the desired form. Then, by Theorem 5.4, the two Hurwitz factorizations are related by a pure braid operation. □

In fact, the argument for Lemma 5.3 shows that any $p > d$ is sufficiently large for the above argument. Of course, considering large p above is equivalent to simply thinking about the situation in characteristic zero. Geometrically, it might be simpler to argue that we can switch to the linear series point of view, degenerate to obtain a totally degenerate limit linear series, and then switch back to the point of view of admissible covers. This ensures that there is only one ramification point over each node, and hence that the partial products are cycles. This avoids any use of Theorem 5.4, but requires a good understanding of local generating systems for admissible fundamental groups, discussed in more detail in the following remark.

Remark 7.5. Owing to Theorem 5.4, we have been able to considerably simplify the transition from numerical to group-theoretic results, compared with the arguments originally envisioned. However, the original arguments should hold more generally than Theorem 5.4, so we briefly sketch here how we would expect them to go. We suppose we have a given branched cover, with given local generating system on the base. We wish to be able to say something about the monodromy of the cover, perhaps after pure braid transformation; for these purposes, it is probably best to assume (as is done in [Gro71]) that the local generating system arises as the specialization of a set of topological generators from characteristic zero. The program is then as follows.

For the most general step, we suppose that we have a family of semistable curves and a local generating system on the (smooth) geometric generic fiber, arising from specialization from topological generators. We say that a local generating system for the admissible fundamental group of the special fiber is 'geometric' if it arises by gluing local generating systems on each component of the normalization; in this case, one can express monodromy of admissible covers in terms of the monodromy of components of the normalization. The basic assertion is that after a pure braid transformation, the given local generating system on the geometric generic fiber specializes to a geometric local generating system on the special fiber. Given appropriate definitions, the statement is clear in the topological setting, since any two choices of local generating systems on the smooth fiber are related by pure braid transformations. One then has to check that the definitions behave well with respect to algebraization and specialization. Finally, one would lift the original family of curves to characteristic zero, and compare with the topological setting to obtain to desired result.

We return to considering a given branched cover of $\mathbb{P}^1$. We first think of the cover as a one-dimensional linear series, and show that it can be degenerated to a limit linear series on some degeneration of the original curve. We then show that a family of limit linear series can be realized geometrically as a family of admissible covers. Applying the previous discussion, after pure braid transformation the local generating system specializes to a geometric local generating system, so we can use the geometry of the admissible cover to describe the monodromy of the smooth cover.

Acknowledgements

I would like to thank Akio Tamagawa for encouraging me to study the question addressed in this paper, urging me to pursue the sharper results in terms of monodromy groups, and describing existing results; I am also grateful for his and Shinichi Mochizuki's patience in explaining over many conversations the background theory on monodromy groups and admissible covers. Finally, I would like to thank David Harbater, Jakob Stix, Robert Guralnick, Michel Raynaud and Robin Hartshorne for their helpful conversations.

References

Bou01 I. Bouw, *The p-rank of ramified covers of curves*, Compositio Math. **126** (2001), 295–322.

BW05 I. Bouw and S. Wewers, *Alternating groups as monodromy groups in positive characteristic*, Pacific J. Math. **222** (2005), 185–200.

Fuj90 K. Fujiwara, *Etale topology and the philosophy of log*, in *Algebraic Geometry Symposium* (Kinosaki, 1990), 116–123 (Japanese).

Gro71 A. Grothendieck, *Revêtements étales et groupe fondamental*, in *Séminaire de Géométrie Algébrique du Bois Marie 1960–1961 (SGA 1)*. With two papers by M. Raynaud, Lecture Notes in Mathematics, vol. 224 (Springer, Berlin, 1971).

GS R. Guralnick and J. Shareshian, *Symmetric and alternating groups as monodromy groups of Riemann surfaces I: generic covers and covers with many branch points*, Mem. Amer. Math. Soc., to appear.

HS99 D. Harbater and K. F. Stevenson, *Patching and thickening problems*, J. Algebra **212** (1999), 272–304.

LO F. Liu and B. Osserman, *The irreducibility of certain pure-cycle Hurwitz spaces*, Amer. J. Math., to appear.

LO06 F. Liu and B. Osserman, *Mochizuki's indigenous bundles and Ehrhart polynomials*, J. Algebraic Combin. **23** (2006), 125–136.

Oss05 B. Osserman, *Deformations of covers, Brill–Noether theory, and wild ramification*, Math. Res. Lett. **12** (2005), 483–491.

Oss06a B. Osserman, *A limit linear series moduli scheme*, Ann. Inst. Fourier (Grenoble) **56** (2006), 1165–1205.

Oss06b B. Osserman, *Rational functions with given ramification in characteristic p*, Compositio Math. **142** (2006), 433–450.

Oss07 B. Osserman, *Mochizuki's crys-stable bundles: a lexicon and applications*, Publ. Res. Inst. Math. Sci. **43** (2007), 95–119.

Ray99 M. Raynaud, *Spécialisations des revêtements en caractéristique* $p > 0$, Ann. Sci. École Norm. Sup. (4) **32** (1999), 87–126.

Ray00 M. Raynaud, *Revêtements des courbes en caractéristique* $p > 0$ *et ordinarité*, Compositio Math. **123** (2000), 73–88.

Sai07 M. Saidi, *Revêtements modérés et groupes fondamental de graphe*, Compositio Math. **107** (1007), 319–338.

Ste96 K. F. Stevenson, *Galois groups of unramified covers of projective curves in characteristic p*, J. Algebra **182** (1996), 770–804.

Tam04 A. Tamagawa, *Finiteness of isomorphism classes of curves in positive characteristic with prescribed fundamental groups*, J. Algebraic Geom. **13** (2004), 675–724.

Brian Osserman osserman@math.ucdavis.edu
Department of Mathematics, University of California at Davis, One Shields Ave., Davis, CA 95616, USA

https://doi.org/10.1112/S0010437X0700303X Published online by Cambridge University Press