

The relative class number one problem for function fields, III

Kiran S. Kedlaya

ABSTRACT. We complete the solution of the relative class number one problem for function fields of curves over finite fields. Using work from two earlier papers, this reduces to finding all function fields of genus 6 or 7 over \mathbb{F}_2 with one of 40 prescribed Weil polynomials; one may then verify directly that three of these fields admit an everywhere unramified quadratic extension with trivial relative class group. The search is carried out by carefully enumerating curves based on the Brill–Noether stratification of the moduli spaces of curves in these genera, and particularly Mukai’s descriptions of the open strata.

1. Introduction

This paper continues and concludes the work done in [17, 18] on the *relative class number one problem* for function fields of curves over finite fields (hereafter simply “function fields”), building upon work of Leitzel–Madan [21] and Leitzel–Madan–Queen [22]. That is, we seek to identify finite extensions F'/F of function fields for which the two class numbers are equal.

To state the main result, we recall some context from the introduction of [17]. Given a finite extension F'/F of function fields, we write C, C' for the curves corresponding to F, F' ; $q_F, q_{F'}$ for the orders of the base fields of C, C' ; $g_F, g_{F'}$ for the genera of C, C' ; and $h_F, h_{F'}$ for the class numbers of F, F' . Since the relative class number $h_{F'/F} = h_{F'}/h_F$ is an integer (it is the order of the Prym variety of the covering $C' \rightarrow C$), the relative class number one problem reduces to the cases where $g_{F'} = g_F$ (a *constant extension*) and where $q_F = q_{F'}$ (a *purely geometric extension*). Excluding the trivial cases of a constant extension of genus-0 function field and an extension with $F' \cong F$, one has the following result (see [17] for the tables).

THEOREM 1.1 (Solution of the relative class number one problem). *Let F'/F be an extension of function fields of degree $d > 1$ of relative class number 1.*

- (a) *If F'/F is constant and $g_F > 0$, then q_F, d, g_F , and the isogeny class of $J(C)$ appear in [17, Theorem 1.1]. In particular,*

$$(q_F, d, g_F) \in \{(2, 2, 1), (2, 2, 2), (2, 2, 3), (2, 3, 1), (3, 2, 1), (4, 2, 1)\}.$$

Thanks to Samir Canning, John Cremona, Xander Faber, Joan-Carles Lario, Bjorn Poonen, and David Roe for helpful discussions. The author was supported by NSF (grants DMS-1802161, DMS-2053473) and UC San Diego (Warschawski Professorship).

- (b) If F'/F is purely geometric, $g_F \leq 1$, and $g_{F'} > g_F$, then $q_F, g_F, g_{F'}$, and the isogeny classes of $J(C)$ and $J(C')$ appear in [17, Table 3]. In particular,
- $$(q_F, g_F, g_{F'}) \in \{(2, 0, 1-4), (2, 1, 2-6), (3, 0, 1), (3, 1, 2), (3, 1, 3), (4, 0, 1), (4, 1, 2)\}.$$
- (c) If F'/F is purely geometric, $g_F > 1$, and $q_F > 2$, then $d, g_F, g_{F'}, F$ appear in [17, Table 4]. In particular,
- $$(q_F, d, g_F, g_{F'}) \in \{(3, 2, 2, 3), (3, 2, 2, 4), (3, 2, 3, 5), (3, 3, 2, 4), (4, 2, 2, 3), (4, 3, 2, 4)\}.$$
- (d) If F'/F is purely geometric, $g_F > 1$, $q_F = 2$, and $d > 2$, then $d, g_F, g_{F'}, F$ appear in [17, Table 5]. In particular,
- $$(d, g_F, g_{F'}) \in \{(3, 2, 4), (3, 2, 6), (3, 3, 7), (3, 4, 10), (4, 2, 5), (5, 2, 6), (7, 2, 8)\}.$$
- (e) If F'/F is purely geometric, $g_F > 1$, $q_F = 2$, and $d = 2$, then $g_F, g_{F'}, F$ appear in [17, Table 6]. In particular,
- $$(g_F, g_{F'}) \in \{(2, 3), (2, 4), (2, 5), (3, 5), (3, 6), (4, 7), (4, 8), (5, 9), (6, 11), (7, 13)\}.$$
- (f) If F'/F is neither constant nor purely geometric and $g_{F'} > g_F$, then $q_F = 2, q_{F'} = 4$, and $(g_F, g_{F'}, J(C), J(C'))$ is one of $(0, 1, 0, 1.4.ae)$ or $(1, 2, 1.2.c, 2.4.ae.i)$ (using LMFDB labels to represent isogeny classes of abelian varieties).

This statement is covered by [17, Theorems 1.1, 1.2, 1.3] except for the following points.

- Part (b) requires classifying curves of genus 6 over \mathbb{F}_2 with one particular Weil polynomial. It is shown in [18, Lemma 10.2] that there is a unique such curve.
- Part (d) requires showing that when $q_F = 2$, the extension F'/F is cyclic. This is done in [18, Theorem 1.1] using constraints on the Weil polynomials found in [17].
- Part (e) requires finding all curves of genus 6 and 7 over \mathbb{F}_2 with Weil polynomials in a specific list of 40 entries found in [17] (see Table 2). This brings us to the main result of the present paper, stated as Theorem 1.2 below.

THEOREM 1.2. *The following statements hold.*

- (a) *There are two isomorphism classes of curves C of genus 6 over \mathbb{F}_2 admitting an étale double covering $C' \rightarrow C$ such that $\#J(C')(\mathbb{F}_2) = \#J(C)(\mathbb{F}_2)$. The curves C are Brill–Noether general with automorphism groups C_3 and C_5 .*
- (b) *There is a unique isomorphism class of curves C of genus 7 over \mathbb{F}_2 admitting an étale double covering $C' \rightarrow C$ such that $\#J(C')(\mathbb{F}_2) = \#J(C)(\mathbb{F}_2)$. The curve C is bielliptic with automorphism group D_6 .*

As in [17], given a candidate for C it is straightforward to use MAGMA to generate all of the étale double coverings $C' \rightarrow C$; thus the main computational issue is to “invert the Weil polynomial function” on the output values indicated in Table 2. Unfortunately, the Weil polynomial function is in some sense a “secure hash function”, in that its value generally does not reveal much useful information about the input. Examples where one can invert the function are often of some extremal nature, as in Lauter’s approach to bounding the maximum number of

points on a curve of fixed genus over a fixed finite field: one first enumerates the Weil polynomials consistent with a given point count, then attempts to rule in or out the various candidates. Much work has been done on the second step by Howe; see [13] for a recent survey of this problem.

Unfortunately, the techniques described in [13] do not seem to be applicable to the cases relevant to Theorem 1.2. Fortunately, for genera up to 7 it is feasible to deploy a brute force strategy, i.e., to enumerate a collection of schemes known to include all curves with the given Weil polynomials and then filter through the results. One way to build such a collection is using singular plane curves; see [8] and [9] for recent examples of this approach.

Here we take an alternate approach that accounts for the known geometry of moduli spaces of curves based on Petri’s analysis of linear systems (see §2). This amounts to a natural extension of the computation of the set of isomorphism classes of curves of genus g over \mathbb{F}_2 for $g = 4$ by Xarles [34], based on the fact that a general canonical curve of genus 4 is a complete intersection of type $(2) \cap (3)$ in \mathbf{P}^3 ; and $g = 5$ by Dragutinović [6], based on the fact that a general canonical curve of genus 5 is a complete intersection of type $(2) \cap (2) \cap (2)$ in \mathbf{P}^4 .

While one cannot hope to give similar such descriptions in arbitrary genus (see Remark 3.4), they are available in genus 6 and 7 by work of Mukai [26, 27], although some care is required to use them over a nonclosed base field. As in [8] and [9], we short-circuit the searches using the Weil polynomial constraints, especially the number of \mathbb{F}_2 -rational points. See Lemma 4.1 for more detailed internal references.

One technical innovation introduced along the way (see Appendix A) is a light-weight method for computing orbits of the action of a group G on subsets of a set carrying a G -action; for instance, in the generic genus-7 case we compute orbits of 6-element sets of \mathbb{F}_2 -rational points on the 10-dimensional orthogonal Grassmannian. This construction may be of independent interest for other applications, including extending the tabulation of genus- g curves over \mathbb{F}_2 to a few larger values of g for which the Brill–Noether stratification on moduli can again be made explicit (see Remark 3.4), or finding supersingular genus- g curves over \mathbb{F}_2 for g in a similar range. See §8 for more discussion of the relevant issues.

As in [17] and [18], the arguments depend on a number of computations in SAGEMATH [30] and MAGMA [30]; the computations take about 8 hours on a single CPU (Intel i5-1135G7@2.40GHz) and can be reproduced using some Jupyter notebooks found in the repository [19]. (Some functionality used in SAGEMATH is derived from GAP [11] and SINGULAR [5].)

2. The structure of canonical curves

Let C be a curve of genus g over a *finite* field k . Let \bar{k} be an algebraic closure of C . We collect here a few facts about the geometry of C that will be used frequently, and often without comment, in what follows. See [31] for a characteristic-free treatment of much of this material (and [12, §4.3] for some additional details).

A g_d^r on C is a line bundle of degree d whose space of global sections has dimension $r + 1$; if such a bundle is basepoint-free, then it defines a degree- d map $C \rightarrow \mathbf{P}^r$. (If the bundle is not basepoint-free, then the global sections generate a basepoint-free subbundle of degree strictly less than d .) Since k is finite, every Galois-invariant divisor class on C contains a k -rational divisor (see for example [2,

Remark 2.4]). Consequently, if $C_{\bar{k}}$ admits a *unique* g_d^r for some r, d , then so does C .

The Castelnuovo–Severi inequality (see for example [32, Theorem 3.11.3]) asserts that if there exist curves C_1, C_2 of genera g_1, g_2 and morphisms $f_1: C \rightarrow C_1, f_2: C \rightarrow C_2$ of degrees d_1, d_2 such that $k(C)$ is the compositum of $k(C_1)$ and $k(C_2)$ over k , then

$$g \leq d_1 g_1 + d_2 g_2 + (d_1 - 1)(d_2 - 1).$$

We will use this bound to ensure that certain low-degree maps out of C occur in isolation.

We say that C is *hyperelliptic* if C admits a g_2^1 (which is automatically basepoint-free if $g > 0$). By Castelnuovo–Severi, if $g > 1$ then $C_{\bar{k}}$ can admit only one g_2^1 ; consequently, C is hyperelliptic if and only if $C_{\bar{k}}$ is hyperelliptic. Let $\iota: C \rightarrow \mathbf{P}_k^{g-1}$ be the *canonical morphism*, defined by the canonical linear system; then ι is a degree-2 map onto a rational normal curve if C is hyperelliptic and an embedding otherwise. By abuse of language, ι is commonly called the *canonical embedding* even when C is allowed to be hyperelliptic.

For $g > 4$, we say that C is *trigonal* if it admits a g_3^1 but not a g_2^1 (so the former is necessarily basepoint-free). By Castelnuovo–Severi again, $C_{\bar{k}}$ can admit only one g_3^1 ; consequently, C is trigonal if and only if $C_{\bar{k}}$ is trigonal. By Petri’s theorem (a/k/a the Max Noether–Enriques–Petri theorem), if C is not trigonal or a smooth plane quintic (when $g = 6$), then $\iota(C)$ is cut out by quadrics.

By contrast, for C trigonal, the linear system of quadrics containing $\iota(C)$ cuts out a rational normal scroll; the latter is isomorphic to the Hirzebruch surface

$$\mathbf{F}_n := \text{Proj}_{\mathbf{P}_k^1}(\mathcal{O}_{\mathbf{P}_k^1} \oplus \mathcal{O}(n)_{\mathbf{P}_k^1})$$

for a certain integer $n \geq 0$ called the *Maroni invariant*¹ of C . The structure map $\mathbf{F}_n \rightarrow \mathbf{P}_k^1$, whose fibers form a ruling of \mathbf{F}_n , restricts to the trigonal projection $\pi: C \rightarrow \mathbf{P}_k^1$.

- For $n > 0$, \mathbf{F}_n is isomorphic to an $(n, 1)$ -hypersurface in $\mathbf{P}_k^1 \times_k \mathbf{P}_k^2$. Let b be the unique irreducible curve in \mathbf{F}_n with negative self-intersection (the *directrix*) and let f be a fiber of the ruling; then

$$(2.1) \quad b^2 = -n, \quad b \cdot f = 1, \quad f^2 = 0,$$

and blowing down \mathbf{F}_n along b yields the weighted projective space $\mathbf{P}(1 : 1 : n)_k$. Of the linear systems

$$(2.2) \quad |3b + \frac{g+3n+2}{2}f|, \quad |b + \frac{g+n-2}{2}f|, \quad |-2b + (-n-2)f|,$$

the first contains C , the second defines the embedding $\mathbf{F}_n \rightarrow \mathbf{P}_k^{g-1}$, and the third is the canonical linear system.

- For $n = 0$, we have $\mathbf{F}_{n, \bar{k}} \cong \mathbf{P}_{\bar{k}}^1 \times_{\bar{k}} \mathbf{P}_{\bar{k}}^1$. Let b and f be fibers of the two different rulings; then (2.1) and the interpretation of (2.2) remain valid. Since $3 \neq \frac{g+2}{2}$, the symmetry of the two rulings is broken by C and so everything descends from \bar{k} to k .

Since C and b are effective, $0 \leq b \cdot C = -3n + \frac{g+3n+2}{2}$, so

$$0 \leq n \leq \lfloor \frac{g+2}{3} \rfloor, \quad n \equiv g \pmod{2}.$$

¹Here we follow the terminology of [31]. The original definition of Maroni [25] follows a different numbering convention which is also commonly used.

In case $n = \frac{g+2}{3}$, we have $b \cdot C = 0$ and so C also embeds into $\mathbf{P}(1 : 1 : n)$.

We say that C is *bielliptic* if it admits a degree-2 map to a genus-1 curve over k . By Castelnuovo–Severi once more, if $g > 5$ then $C_{\bar{k}}$ can admit at most one such map; consequently, C is bielliptic if and only if $C_{\bar{k}}$ is bielliptic.

3. Brill–Noether stratifications

We now specialize the previous discussion to the genera of direct concern here, following Mukai. We use the conventions that the Grassmannian $\text{Gr}(r, V)$ parametrizes subspaces of dimension r of a specified vector space V and that the Plücker embedding maps into $\mathbf{P}(\wedge^r V^\vee)$.

THEOREM 3.1. *Let C be a curve of genus 6 over a finite field k . Then one of the following holds.*

- (1) *The curve C is hyperelliptic.*
- (2) *The curve C is trigonal of Maroni invariant 2. In this case, C occurs as a complete intersection of type $(2, 1) \cap (1, 3)$ in $\mathbf{P}_k^1 \times_k \mathbf{P}_k^2$, where the $(2, 1)$ -hypersurface is isomorphic to \mathbf{F}_2 .*
- (3) *The curve C is trigonal of Maroni invariant 0. In this case, C occurs as a curve of bidegree $(3, 4)$ in $\mathbf{P}_k^1 \times_k \mathbf{P}_k^1$.*
- (4) *The curve C is bielliptic.*
- (5) *The curve C occurs as a smooth quintic curve in \mathbf{P}_k^2 .*
- (6) *The curve C occurs as a transverse intersection of four hyperplanes, a quadric hypersurface, and the 6-dimensional Grassmannian $\text{Gr}(2, 5)$ in \mathbf{P}_k^9 .*

PROOF. This again follows from Petri’s theorem except for the last case, in which the description can be found in [26, Theorem 5.2]. We recall that argument both to fill in some details that are left to the reader in [26] (by comparison with a similar argument in genus 8), and to see that it is characteristic-free and applies over a finite base field.

By the Brill–Noether theorem on the existence of special divisors (see [20] for a characteristic-free treatment), $C_{\bar{k}}$ admits a g_4^1 , which we call ξ ; let $\eta := \omega_C \xi^{-1}$ be its Serre adjoint. By Riemann–Roch we have

$$h^0(\xi) = 2, \quad h^0(\eta) = h^0(\xi) + g - 1 - \deg(\xi) = 3;$$

that is, η is a g_6^2 . Since C is not trigonal or a plane quintic, the linear system $|\eta|$ is basepoint-free; we thus have a map $\Phi_{|\eta|} : C_{\bar{k}} \rightarrow \mathbf{P}_{\bar{k}}^2$ induced by η . The image of $\Phi_{|\eta|}$ cannot be a singular cubic or a smooth cubic because we are assuming C is neither hyperelliptic nor bielliptic, so it must be a sextic curve \bar{C} . Other than ξ , every g_4^1 of $C_{\bar{k}}$ arises by projection from a double point of \bar{C} ; it follows that the space $W_4^1(C_{\bar{k}})$ parametrizing the g_4^1 ’s of $C_{\bar{k}}$ is finite (recovering [26, Proposition 5.3]).

We now emulate [26, Lemma 3.6]. The extensions $0 \rightarrow \xi \rightarrow E \rightarrow \eta \rightarrow 0$ are parametrized by $\text{Ext}(\eta, \xi) \cong H^1(\eta^{-1}\xi)$, which is Serre dual to $H^0(\eta^2)$. For an extension e , let $\delta_e : H^0(\eta) \rightarrow H^1(\xi)$ be the corresponding connecting homomorphism; then the linear map

$$\Delta : \text{Ext}(\eta, \xi) \rightarrow H_0(\eta)^\vee \otimes H^1(\xi), \quad e \mapsto \delta_e$$

is dual to the multiplication map

$$\mu : H^0(\eta) \otimes H^0(\eta) \rightarrow H^0(\eta^2).$$

By Riemann–Roch again, $h^0(\eta^2) = \deg(\eta^2) - g + 1 = 7$. Since the image of $\Phi_{|\eta|}$ cannot be contained in a conic, the linear map $\text{Sym}^2 H^0(\eta) \rightarrow H^0(\eta^2)$ is injective, so its cokernel has codimension 1. We conclude that $\ker(\Delta)$ is one-dimensional, and so there is a unique nontrivial extension of η by ξ which is a stable bundle with five linearly independent global sections. Because k is perfect, this uniqueness property ensures that the resulting vector bundle of rank 2 on $C_{\bar{k}}$ descends to a unique vector bundle E on C . We have now recovered [26, Theorem 5.1(1)]. We deduce from this an analogue of [26, Lemma 3.10]: if ξ' is any g_4^1 on $C_{\bar{k}}$, then $\dim \text{Hom}(\xi', E) \leq 1$.

Since $\delta_e = 0$ and ξ and η are generated by global sections, so is E . Hence for each point in C , the fiber of E at this point is a 2-dimensional quotient of the 5-dimensional space $H^0(E)$; this defines a map $\Phi_{|E|}: C \rightarrow \text{Gr}(2, H^0(E)^\vee)$. Let $\lambda: \wedge^2 H^0(E) \rightarrow H^0(\wedge^2 E) = H^0(\omega_C)$ be the natural map. We then have a commutative diagram

$$\begin{array}{ccc} C_{\bar{k}} & \xrightarrow{\Phi_{|E|}} & \text{Gr}(2, H^0(E)^\vee) \\ \downarrow & & \downarrow \\ \mathbf{P}(H^0(\omega_C)) & \xrightarrow{\mathbf{P}(\lambda)} & \mathbf{P}(\wedge^2 H^0(E)) \end{array}$$

where the left vertical arrow is the canonical embedding and the right vertical arrow is the Plücker embedding. The hyperplanes of $\mathbf{P}(\wedge^2 H^0(E))$ are parametrized by $\mathbf{P}((\wedge^2 H^0(E))^\vee)$; the hyperplanes among these which containing the image of C are parametrized by $\mathbf{P}((\ker \lambda)^\vee)$.

We now emulate [26, Theorem B]. Suppose that $U \subset H^0(E)$ is a 2-dimensional subspace such that $\lambda(\wedge^2 U) = 0$. Then the evaluation map $U \otimes \mathcal{O}_C \rightarrow E$ is not generically surjective; its image is a line subbundle L of E satisfying $h^0(L) \geq 2$. The stability of E forces $\deg(L) < 5$, and $h^0(L) = 2$ since C has no g_4^1 by the adjunction formula. Since C is not hyperelliptic or trigonal, L must be a g_4^1 . That is, this construction defines a map from $\mathbf{P}((\ker \lambda)^\vee) \cap \text{Gr}(2, H^0(E)^\vee)$ to $W_4^1(C_{\bar{k}})$; this map is injective by our analogue of [26, Lemma 3.10] and surjective by [26, Proposition 3.1]. We have now recovered [26, Theorem 5.1(2)].

We now follow the proof of [26, Theorem 5.2] as written. To wit, since $\mathbf{P}((\ker \lambda)^\vee) \cap \text{Gr}(2, H^0(E)^\vee) \cong W_4^1(C_{\bar{k}})$ is finite and $\text{Gr}(2, H^0(E)^\vee)$ has codimension 3 in $\mathbf{P}(\wedge^2 H^0(E))$, $\dim(\ker \lambda) \leq 4$; hence λ is surjective and so $\Phi_{|E|}$ is an embedding. By Petri's theorem (§2), the image of $\Phi_{|E|}$ is cut out by the hyperplanes in $\mathbf{P}((\ker \lambda)^\vee)$ plus a single quadric. \square

THEOREM 3.2. *Let C be a curve of genus 7 over a finite field k . Then one of the following holds.*

- (1) *The curve C is hyperelliptic.*
- (2) *The curve C is trigonal of Maroni invariant 3. In this case, C occurs as a hypersurface of degree 9 in $\mathbf{P}(1 : 1 : 3)_k$.*
- (3) *The curve C is trigonal of Maroni invariant 1. In this case, C occurs as a complete intersection of type $(1, 1) \cap (3, 3)$ in $\mathbf{P}_k^1 \times_k \mathbf{P}_k^2$.*
- (4) *The curve C is bielliptic.*
- (5) *The curve C is not bielliptic but admits a g_6^2 which is self-adjoint (squares to the canonical class). In this case, C is a complete intersection of type $(3) \cap (4)$ in $\mathbf{P}(1 : 1 : 1 : 2)_k$, where the degree 3 hypersurface can be taken to be defined by $x_0 x_3 + P_3(x_1, x_2) = 0$ for some separable cubic P_3 .*

- (6) *The curve C admits a pair of distinct g_6^2 's. In this case, C occurs as a complete intersection of type $(1, 1) \cap (1, 1) \cap (2, 2)$ in $\mathbf{P}_k^2 \times_k \mathbf{P}_k^2$.*
- (7) *The curve C does not admit a g_6^2 but $C_{\bar{k}}$ does. In this case, C occurs as a complete intersection of type $(1, 1) \cap (1, 1) \cap (2, 2)$ in the quadratic twist of $\mathbf{P}_k^2 \times_k \mathbf{P}_k^2$.*
- (8) *The curve C admits a g_4^1 but $C_{\bar{k}}$ does not admit a g_6^2 . In this case, C occurs as complete intersection of type $(1, 1) \cap (1, 2) \cap (1, 2)$ in $\mathbf{P}_k^1 \times_k \mathbf{P}_k^3$ in which the $(1, 1)$ -hypersurface is a \mathbf{P}^2 -bundle over \mathbf{P}^1 . (It is also true that all of the $(1, 2)$ -hypersurfaces vanishing on C are geometrically irreducible, but we won't use this here.)*
- (9) *The curve C does not admit a g_4^1 . In this case, C occurs as a transverse intersection of 9 hyperplanes and the orthogonal Grassmannian $\text{OG}^+(5, 10)$ in \mathbf{P}_k^{15} .*

PROOF. Petri's theorem covers cases (1)–(3). We treat cases (4)–(8) as summarized in [27, Table 1], postponing case (9) until §7 where we introduce the relevant notation.

Suppose that $C_{\bar{k}}$ is not hyperelliptic or trigonal but admits a g_4^1 ; let ξ be one such and let $\eta := \omega_C \xi^{-1}$ be its Serre adjoint, which by Riemann–Roch is a g_8^2 . Since C cannot admit a g_5^2 , $|\eta|$ is basepoint-free. Let $\pi: C_{\bar{k}} \rightarrow \mathbf{P}_{\bar{k}}^1$ and $\tau: C_{\bar{k}} \rightarrow \mathbf{P}_{\bar{k}}^3$ be the maps defined by $|\xi|$ and $|\eta|$.

If $C_{\bar{k}}$ has no g_6^2 , then τ is an embedding; its image cannot lie in a quadric by the adjunction formula, so η does not factor as a product of two g_4^1 's. By [26, Corollary 3.2], any g_4^1 other than ξ would have to occur as a subbundle of η , so in fact ξ is the unique g_4^1 on $C_{\bar{k}}$. This means that both ξ and η descend from $C_{\bar{k}}$ to C . We can now follow the proofs of [27, Lemma 6.1, Proposition 6.3] to the desired conclusion.

Suppose instead that $C_{\bar{k}}$ has a g_6^2 ; let α be one such and let β be its Serre adjoint, which is also a g_6^2 . Since $C_{\bar{k}}$ is not hyperelliptic or trigonal, the map $f: C_{\bar{k}} \rightarrow \mathbf{P}_{\bar{k}}^2$ defined by $|\alpha|$ is either birational onto a sextic or a double cover of a smooth cubic. In the latter case $C_{\bar{k}}$ is evidently bielliptic, as then is C . In the former case, from the proof of [26, Proposition 3.1] we see that there are no g_6^2 's on $C_{\bar{k}}$ other than α and β . Namely, if ζ is a third g_6^2 , then for $E = \xi \oplus \eta$ we have $h^0(\zeta^{-1}E) = 0$ and so

$$h^0(\zeta\xi) + h^0(\zeta\eta) = h^0(\zeta E) = h^0(\omega_C \zeta E^\vee) = h^0(\zeta^{-1}E) + 2 \deg(\zeta) = 12;$$

this is only possible if one of $\zeta\xi$ or $\zeta\eta$ is special, which is impossible because they are both of degree 12 and not canonical.

If α and β are isomorphic, then they both descend to C ; otherwise, they descend either to C or to its quadratic base extension. We can now follow the proof of [27, Proposition 6.5] to conclude. \square

REMARK 3.3. In [27, Proposition 6.4], it is also shown that bielliptic curves occur as complete intersections of type $(3) \cap (4)$ in $\mathbf{P}(1 : 1 : 1 : 2)_k$. We will not use this in our computations.

REMARK 3.4. While we will not need to do so here, it is possible to push this treatment through to a few higher genera. For example, Mukai showed that (over an algebraically closed field) a genus-8 curve with no g_7^2 is a linear section of the 8-dimensional Grassmannian $\text{Gr}(2, 6) \subset \mathbf{P}^{14}$ [27]; building on this, the complete

Brill–Noether stratification in genus 8 has been described by Ide–Mukai [15]. Similarly, a genus-9 curve with no g_5^1 is a linear section of the 6-dimensional symplectic Grassmannian $\mathrm{SpG}(3, 6) \subset \mathbf{P}^{13}$ [28], and an analogous assertion holds in genus 10 [29]. Pushing this even further would amount to establishing unirationality of the moduli space of genus- g curves, which is known to hold for $g \leq 14$ [33] and to fail for $g \geq 22$ [7], [10].

4. Overview of the proof

We now give an overview of the proof of Theorem 1.2.

LEMMA 4.1. *For the various strata in moduli described above, the number of isomorphism classes of curves C over \mathbb{F}_2 in each stratum admitting étale double coverings $C' \rightarrow C$ such that $\#J(C')(\mathbb{F}_2) = \#J(C)(\mathbb{F}_2)$ is given in Table 1. In particular, Theorem 1.2 holds.*

Type of C	$g = 6$					$g = 7$				
	Dim	See	$\#C$	$\#C'$	Time	Dim	See	$\#C$	$\#C'$	Time
hyperelliptic	11	§5	0	0	—	13	§5	0	0	—
trigonal, Maroni ≥ 2	12	§6	9	0	10m	13	§5	0	0	—
trigonal, Maroni ≤ 1	13	§6	9	0	2m	15	§6	0	0	5m
bielliptic	10	§5	0	0	—	12	§5	2	1	5m
plane quintic	12	§6	1	0	1m	—	—	—	—	—
self-adjoint g_6^2	—	—	—	—	—	15	§6	0	0	5m
rational g_6^2 pair	—	—	—	—	—	16	§6	0	0	30m
irrational g_6^2 pair	—	—	—	—	—	16	§6	0	0	45m
tetragonal, no g_6^2	—	—	—	—	—	17	§6	1	0	2h
generic	15	§6	38	2	4h	18	§7	1	0	1h

TABLE 1. Outline of the use of the Brill–Noether stratification in the proof of Lemma 4.1. Of the columns, “Dim” records the dimension of the stratum in moduli, “See” locates the description of this case in the text, “ $\#C$ ” counts curves whose point counts appear in Table 2, and “ $\#C'$ ” counts double covers with relative class number 1.

PROOF. To begin with, we recall from [17, Theorem 1.3(b)] that the Weil polynomials of C and C' are restricted to an explicit finite list. In Table 2, we list the possible values of the tuple $(\#C(\mathbb{F}_{2^i}))_{i=1}^g$.

For each stratum, we exhibit a set T of schemes of finite type over \mathbb{F}_2 of size at most 10^6 , such that every curve C over \mathbb{F}_2 belonging to the specified stratum whose point counts are consistent with Table 2 is isomorphic to some scheme in T . In most cases, all of the schemes in T will be presented as subschemes of a single ambient scheme X ; see Table 1 for internal cross-references.

Given a set T as indicated, we conclude as follows (iterating over all $C \in T$). All computations are done in SAGEMATH except as indicated.

- Optionally, for one or more $i \geq 1$, compute $\#C(\mathbb{F}_{2^i})$ using a lookup table of $X(\mathbb{F}_{2^i})$, retaining cases consistent with Table 2. We typically do this when we have at least 10^5 cases to deal with.

4, 14, 16, 18, 14, 92	5, 11, 11, 31, 40, 53	6, 10, 9, 38, 11, 79	
4, 14, 16, 18, 24, 68	5, 11, 11, 31, 40, 65	6, 10, 9, 38, 21, 67	
4, 14, 16, 26, 14, 68	5, 11, 11, 39, 20, 53	6, 10, 9, 38, 31, 55	6, 18, 12, 18, 6, 60, 174
4, 16, 16, 20, 9, 64	5, 11, 11, 39, 20, 65	6, 14, 6, 26, 26, 68	6, 18, 12, 18, 6, 72, 132
5, 11, 11, 31, 20, 65	5, 13, 14, 25, 15, 70	6, 14, 6, 26, 26, 80	6, 18, 12, 18, 6, 84, 90
5, 11, 11, 31, 20, 77	5, 13, 14, 25, 15, 82	6, 14, 6, 26, 36, 56	7, 15, 7, 31, 12, 69, 126
5, 11, 11, 31, 20, 89	5, 13, 14, 25, 15, 94	6, 14, 6, 34, 16, 56	7, 15, 7, 31, 22, 45, 112
5, 11, 11, 31, 30, 53	5, 13, 14, 25, 25, 46	6, 14, 6, 34, 26, 44	7, 15, 7, 31, 22, 57, 70
5, 11, 11, 31, 30, 65	5, 13, 14, 25, 25, 58	6, 14, 12, 26, 6, 44	7, 15, 7, 31, 22, 57, 84
5, 11, 11, 31, 30, 77	5, 13, 14, 25, 25, 70	6, 14, 12, 26, 6, 56	
5, 11, 11, 31, 30, 89	5, 15, 5, 35, 20, 45	6, 14, 12, 26, 6, 66	

TABLE 2. Tuples $(\#C(\mathbb{F}_{2^i}))_{i=1}^g$ allowed by [17, Theorem 1.3(b)] for $g = 6, 7$.

- Optionally, for one or more $i \geq 1$, compute $\#C(\mathbb{F}_{2^i})$ by computing the length of the intersection in $C \times_{\mathbb{F}_2} C$ of the diagonal with the graph of the i -th power of the Frobenius morphism, retaining cases consistent with Table 2. We typically do this when we have between 10^4 and 10^5 cases to deal with.
- Use MAGMA to check whether C is one-dimensional and integral, and if so whether its normalization has genus g . If so, compute $\#C(\mathbb{F}_{2^i})$ for $i = 1, \dots, g$ by enumerating places of the function field of C , retaining cases consistent with Table 2.
- Use MAGMA to compute isomorphism class representatives among the remaining curves. The count of these is reported in Table 1.
- Use MAGMA to identify quadratic extensions of the remaining function fields with relative class number 1. The count of these is reported in Table 1; this yields the claimed results. \square

Table 1 also includes in each case a rough timing of the computation. The timings should not be taken too seriously; they reflect some combination of the dimensions of the strata in moduli (included in Table 1), the special nature of the Weil polynomials in question (which we exploit especially heavily for generic curves of genus 7), the highly nonuniform extent to which we attempted to optimize the calculation in the various cases, the imbalance between genus 6 and 7 in Table 2, and variable load on the machine in question.

5. Point counts

In a few cases of Lemma 4.1, we can confirm that the options listed in Table 2 imply a nontrivial lower bound on the gonality of C . This amounts to settling some cases of Lemma 4.1 with $T = \emptyset$.

- If $g = 6$, then C cannot be hyperelliptic: we have $\#C(\mathbb{F}_4) > 10 = 2\#\mathbf{P}^1(\mathbb{F}_4)$ except in three cases where $\#C(\mathbb{F}_{16}) = 38 > 2\#\mathbf{P}^1(\mathbb{F}_{16})$.
- If $g = 7$, then C cannot be hyperelliptic: we have $\#C(\mathbb{F}_4) \geq 15 > 2\#\mathbf{P}^1(\mathbb{F}_4)$.
- If $g = 7$ and $\#C(\mathbb{F}_2) = 6$, then C cannot be trigonal: we have $\#C(\mathbb{F}_4) = 18 > 15 = 3\#\mathbf{P}^1(\mathbb{F}_4)$.

$\#E(\mathbb{F}_{2^i})_{i=1}^4$	$\#C(\mathbb{F}_{2^i})_{i=1}^4$	Disposition
(1, 5, 13, 25)	(6, 10, 9, 38)	$\#C(\mathbb{F}_2) > 2\#E(\mathbb{F}_2)$
(3, 9, 9, 9)	(5, 13, 41, 25)	$\#C(\mathbb{F}_{16}) > 2\#E(\mathbb{F}_{16})$
(3, 9, 9, 9)	(6, 10, 9, 38)	$\#C(\mathbb{F}_{16}) > 2\#E(\mathbb{F}_{16})$
(5, 5, 5, 25)	(5, 13, 14, 25)	$\#C(\mathbb{F}_4) > 2\#E(\mathbb{F}_4)$
(5, 5, 5, 25)	(6, 10, 9, 38)	$\#C(\mathbb{F}_4) = 2\#E(\mathbb{F}_4), \#C(\mathbb{F}_2) \not\equiv 0 \pmod{2}$

TABLE 3. Possible point counts for C bielliptic of genus 6 covering the genus-1 curve E .

- If $g = 7$ and $\#C(\mathbb{F}_2) = 7$, then C cannot be trigonal of Maroni invariant 3: we have $\#C(\mathbb{F}_2) = 7$ which exceeds the number of *smooth* points of $\mathbf{P}(1 : 1 : 3)(\mathbb{F}_2)$.

We can use similar logic in the case where C is bielliptic. Suppose that $C \rightarrow E$ is a double covering of an elliptic curve. Then the Weil polynomial of E must divide that of C , and moreover must satisfy the resultant criterion [17, Corollary 9.4]. For $g = 6$, the possibilities are listed in Table 3; in most cases, we find that $\#C(\mathbb{F}_{2^i}) > 2\#E(\mathbb{F}_{2^i})$ for some i , an impossibility. In one case, $\#C(\mathbb{F}_4) = 2\#E(\mathbb{F}_4)$, which ensures that $C \rightarrow E$ does not ramify over any degree-1 places, but this is inconsistent with the fact that $\#C(\mathbb{F}_2) \not\equiv 0 \pmod{2}$. (Alternatively, the unique degree-3 place of E must map to a degree-1 place of C , which again contradicts $\#C(\mathbb{F}_4) = 2\#E(\mathbb{F}_4)$.) We thus again settle this case of Lemma 4.1 with $T = \emptyset$.

For $g = 7$, we may make a similar application of the resultant criterion to see that $\#C(\mathbb{F}_2) = 6$ and $\#E(\mathbb{F}_2) \in \{3, 5\}$. We can rule out $\#E(\mathbb{F}_2) = 5$ by noting that $\#C(\mathbb{F}_4) = 18 > 10 = 2\#E(\mathbb{F}_4)$; we must thus have $\#E(\mathbb{F}_2) = 3$. Now note that E has p -rank 0 and C has p -rank 5, so by the Deuring–Shafarevich formula [17, (7.2)] the map $E \rightarrow C$ must ramify over six distinct geometric points. Since $\#C(\mathbb{F}_4) = 18 = 2\#E(\mathbb{F}_4)$, the map $C \rightarrow E$ cannot ramify over any degree-1 or degree-2 points of C ; the ramification is thus either over a single degree-6 place or over the two distinct degree-3 places of C . We may thus settle this case of Lemma 4.1 by computing the set T of double covers of E with the indicated ramification divisors using MAGMA.

REMARK 5.1. Although we did not exploit this systematically in our calculations, we point out that for every entry of Table 2 with $g = 7$, [13, Theorem 4.15] implies the existence of a map from C to a particular elliptic curve of degree at most 5. For example, when $\#C(\mathbb{F}_{2^i})_{i=1}^7 = (6, 18, 12, 18, 6, 72, 132)$, C must admit a degree-2 map to the elliptic curve E with $\#E(\mathbb{F}_2) = 3$; consequently, this option can be ignored in all but the bielliptic case.

6. The use of orbit lookup trees

In most of the remaining cases, we use a uniform paradigm to make an exhaustive calculation over the relevant term of the Brill–Noether stratification. Again, all computations are done in SAGEMATH except as indicated.

- Let X be the ambient variety indicated in Table 4. Compute the set $S := X(\mathbb{F}_2)$ and the group $G := \text{Aut}(X)(\mathbb{F}_2)$.

- Use the method of orbit lookup trees (Appendix A) to compute orbit representatives for the action of G on subsets of S of size up to g . In some cases, we can impose extra conditions on the set S .
 - For $g = 7$ with a rational g_6^2 , no three points of S have the same projection onto either \mathbf{P}_k^2 .
 - For $g = 7$ tetragonal, no five points of S have the same projection in \mathbf{P}_k^1 .
- For each orbit representative for subsets of size in $\{4, 5, 6\}$ (if $g = 6$) or $\{6, 7\}$ (if $g = 7$), use linear algebra to find all tuples of hypersurfaces X_1, \dots, X_{m-1} of the indicated degrees passing through these \mathbb{F}_2 -points. In the case of $g = 7$ trigonal of Maroni invariant 1, we require X_1 to be smooth.
- For each choice, impose linear conditions on X_m to ensure that $X_1 \cap \dots \cap X_m$ has *exactly* the specified set of \mathbb{F}_2 -rational points. (This crucially exploits the fact that the base field is \mathbb{F}_2 ; a similar strategy is used in [8, §6].) Take T to be the resulting set of schemes $X_1 \cap \dots \cap X_m$.

See Table 4 for how the notation maps to the various cases. Some additional clarifications:

- In the case of $g = 6$ trigonal of Maroni invariant 2, we take $X = X_{2,1}$ to be defined by $x_0^2 y_0 + x_0 x_1 y_1 + x_1^2 y_2$.
- In the case of $g = 6$ generic, we find candidates for the intersection of type $(1)^4$ by computing orbits for the action on sets of 4 k -points of the dual of \mathbf{P}_k^9 . We then apply generators of $\mathrm{GL}(4, \mathbb{F}_2)$ to these subsets to identify cases where the linear spans are G -equivalent (compare Remark A.8); this yields 55 candidates for $X_1 \cap \dots \cap X_{m-1}$. We finally enumerate subsets of $X \cap X_1 \cap \dots \cap X_{m-1}$ of size in $\{4, 5, 6\}$, without further use of the group action.
- In the case of $g = 7$ with a self-adjoint g_6^2 , we take $X = X_3$ to be defined by a polynomial of the form $x_0 x_3 + P(x_1, x_2)$ with

$$P \in \{x_1 x_2 (x_1 + x_2), x_1 (x_1^2 + x_1 x_2 + x_2^2), x_1^3 + x_1 x_2^2 + x_2^3\}$$
 and ignore the group action.
- In the case of $g = 7$ tetragonal, we take $X = X_{1,1}$ to be defined by $x_0 y_0 + x_1 y_1$. We then break symmetry when choosing the defining polynomials P_1, P_2 of X_1, X_2 by fixing a total ordering on the quotient of the space of $(1, 2)$ -polynomials by the multiples of $x_0 y_0 + x_1 y_1$ and then forcing an ordering on the classes of P_1, P_2 .

The generic case in genus 7 is handled slightly differently to avoid the computational bottleneck of enumerating orbits of 7-element subsets of X ; see §7.

7. The generic case in genus 7

We now describe a variant of the paradigm from §6 to handle generic (non-tetragonal) curves of genus 7. In the process, we summarize the proof of [27, Theorem 0.4] and so confirm case (9) of Theorem 3.2.

Let k be a finite field (of any characteristic). Let V be the vector space k^{10} equipped with the quadratic form $\sum_{i=1}^5 x_i x_{5+i}$. We write $\mathrm{SO}(V)$ for the unique index-2 subgroup of the orthogonal group of V ; it admits a characteristic-free characterization as the kernel of the *Dickson invariant*.

g	Case	X	X_1, \dots, X_{m-1}	X_m
6	trigonal, Maroni 2	$X_{2,1} \subset \mathbf{P}^1 \times \mathbf{P}^2$	\emptyset	(1, 3)
6	trigonal, Maroni 0	$\mathbf{P}^1 \times \mathbf{P}^1$	\emptyset	(3, 4)
6	plane quintic	\mathbf{P}^2	\emptyset	(5)
6	generic	$\text{Gr}(2, 5) \subset \mathbf{P}^9$	$(1)^4$	(2)
7	trigonal, Maroni 1	$\mathbf{P}^1 \times \mathbf{P}^2$	(1, 1)	(3, 3)
7	self-adjoint g_6^2	$X_3 \subset \mathbf{P}(1 : 1 : 1 : 2)$	\emptyset	(4)
7	rational g_6^2	$\mathbf{P}^2 \times \mathbf{P}^2$	$(1, 1)^2$	(2, 2)
7	irrational g_6^2	twist of $\mathbf{P}^2 \times \mathbf{P}^2$	$(1, 1)^2$	(2, 2)
7	tetragonal	$X_{1,1} \subset \mathbf{P}^1 \times \mathbf{P}^3$	(1, 2)	(1, 2)
7	generic	$\text{OG}^+ \subset \mathbf{P}^{15}$	$(1)^8$	(1)

TABLE 4. Group actions associated to Brill–Noether strata.

The *orthogonal Grassmannian* of V , denoted OG , parametrizes Lagrangian (isotropic 5-dimensional) subspaces of V . Let L_0 be the subspace spanned by the first 5 coordinate vectors e_1, \dots, e_5 , which by construction is isotropic. Then OG splits into two connected components, each of which parametrizes Lagrangian subspaces of V whose intersection with L_0 has a specified parity. Let OG^+ be the component containing L_0 ; it carries an action of $\text{SO}(V)$.

The space OG^+ admits an analogue of the Plücker embedding called the *spinor embedding*. The target of the spinor embedding can be described as the projectivization of the even orthogonal algebra $\wedge^{\text{ev}} L_0$. The spinor embedding can be computed easily using the following (characteristic-free) recipe described in [27, §1]. Let L_∞ be the subspace spanned by e_6, \dots, e_{10} ; this is an isotropic subspace lying on the other component of OG . Split the orthogonal algebra $S = \wedge^\bullet L_\infty$ into even and odd components $S^{\text{ev}}, S^{\text{odd}}$. The *Clifford map* $V \rightarrow \text{End } S$ then carries each element of L_∞ to a creation operator (taking the wedge product with that element) and each element of $L_0 \cong L_\infty^\vee$ to an annihilation operator (contraction). For each Lagrangian subspace L , the Clifford map carries the elements of L to endomorphisms of S whose joint kernel is one-dimensional, and is contained in either S^{ev} or S^{odd} according to whether or not $L \in \text{OG}^+$; this yields the spinor embedding.

The following combines [27, Proposition 1.16, Proposition 2.2].

LEMMA 7.1 (Mukai). *Let $P \subset \mathbf{P}_k^{15}$ be a 6-dimensional linear subspace which passes through L_0 and meets OG^+ transversely.*

- (a) *No 4 points of $C = \text{OG}^+ \cap P$ lie in a 2-plane.*
- (b) *The scheme C is a canonical curve of genus 7 with one marked point with no g_4^1 .*

Although for our purposes we only need the opposite implication to Lemma 7.1, this direction is actually crucial to the argument.

Now let C be a curve of genus 7 with one marked point admitting no g_4^1 . We set notation as in [27, §3]. Let $C \subset \mathbf{P}_k^6$ be the canonical embedding and set $W := H^0(\mathbf{P}_k^6, I_C(2))$; by Petri's theorem (§2), $\dim W = (g-2)(g-3)/2 = 10$. Set $E := N_{C/\mathbf{P}_k^6}^\vee \otimes \omega_C^2$; this is a bundle of rank $g-2 = 5$. From the exact sequence

$$0 \rightarrow N_{C/\mathbf{P}_k^6}^\vee \rightarrow \Omega_{\mathbf{P}_k^6}|_C \rightarrow \omega_C \rightarrow 0$$

we see that $\det E \cong \omega_C^2$. Since $N_C^\vee/\mathbb{P}_k^6 \cong I_C/I_C^2$ and $\omega_C \cong \mathcal{O}_C(1)$, we obtain a linear map $W \rightarrow H^0(C, E)$. Since C is not trigonal, by Petri's theorem again, for every closed point $p \in C$ the kernel W_p of the induced map from W to the fiber E_p is 5-dimensional; this defines a map $C \rightarrow \mathrm{Gr}(5, W)$.

The following is [27, Proposition 3.3].

LEMMA 7.2. *With notation as above (i.e., C is a curve of genus 7 with no g_4^1), for any two distinct closed points $p, q \in C$, the intersection $W_p \cap W_q$ is 1-dimensional (not just odd-dimensional).*

The following is [27, Theorem 4.2].

LEMMA 7.3 (Mukai). *With notation as above (i.e., C is a curve of genus 7 with no g_4^1), let $f: \mathrm{Sym}^2 W \rightarrow H^0(C, \mathrm{Sym}^2 E)$ be the natural map. Then every nonzero element of $\ker f$ is nondegenerate. Consequently, $\dim \ker f \leq 1$.*

When C arises as in Lemma 7.1, then we have a natural identification $V \cong W$ (see [27, Corollary 2.5]) and so the quadratic form on V defines a nonzero element of $\ker f$. The upshot of this is that the embedding $C \rightarrow \mathrm{Gr}(5, W)$ factors through OG^+ . As in [27, (5.1)], this defines an injective map from the space of linear sections of OG^+ passing through L_0 to the moduli space of genus 7 curves with one marked point. As these spaces are both 11-dimensional and the latter is irreducible [4], the map is dominant; that is, the *generic* curve of genus 7 with one marked point occurs as a linear section of OG^+ passing through L_0 . As in [27, Corollary 5.3], it then follows that $\dim \ker f = 1$ in all cases (where we still assume that C has no g_4^1). Thus we end up with a map $C \rightarrow \mathrm{Gr}(5, W)$, and it is straightforward to check that it is an embedding [27, Theorem 0.4]. This establishes case (9) of Theorem 3.2 modulo the following remark.

REMARK 7.4. We comment briefly on what happens if we consider a curve of genus 7 *without* a marked point. In this case, $H^0(E)$ still carries a distinguished nondegenerate quadratic form, but it is not guaranteed to be isomorphic *over* k to the form we are using.

When k is finite, however, this issue does not arise for the following reason. There are only two isomorphism classes of such forms, distinguished by the discriminant in odd characteristic and the Arf invariant in even characteristic; moreover, these become isomorphic over the quadratic extension of k . By passing to a finite extension of k of suitably large *odd* degree over which C acquires a rational point, we can see that $H^0(E)$ must carry the quadratic form which admits Lagrangian subspaces over k .

We now specialize the previous discussion to the case $k = \mathbb{F}_2$ and describe the computation that proves Lemma 4.1 for curves of genus 7 with no g_4^1 . We first build an orbit lookup tree (Appendix A) to depth 6 for the action of $G := \mathrm{SO}(V)$ on $S := \mathrm{OG}^+(\mathbb{F}_2)$ with the following tuples forbidden.

- Any pair of points corresponding to Lagrangian subspaces whose intersection has dimension greater than 1 (ruled out by Lemma 7.2).
- Any triple of collinear points or 4-tuple of coplanar points (ruled out by Lemma 7.1).
- Any tuple whose linear span has positive-dimensional intersection with OG^+ (ruled out by the fact that the canonical embedding does not factor through a hyperplane).

- Any tuple whose linear span meets OG^+ in 8 or more \mathbb{F}_2 -rational points (ruled out because Table 2 requires $\#C(\mathbb{F}_2) \leq 7$).

This yields 494 orbit representatives. Inspecting the results, we find that each orbit representative spans either a 4-plane or a 5-plane in \mathbf{P}_k^{15} .

Let V be an orbit representative. We now separate into cases depending on whether $\#C(\mathbb{F}_2) = 6$ or $\#C(\mathbb{F}_2) = 7$, whether or not V spans a 4-plane or a 5-plane, and whether or not this span contains any more points of OG^+ .

- If $\#C(\mathbb{F}_2) = 7$ and V spans a 4-plane, we verify that the span of V does not contain a seventh \mathbb{F}_2 -point of OG^+ . This means that without loss of generality, we may ignore this case.
- If $\#C(\mathbb{F}_2) = 6$, V spans a 5-plane, and this 5-plane does not contain a 7th point of OG^+ , we hash the remaining \mathbb{F}_2 -points of OG^+ according to their joint linear span with V , retaining 6-planes that do not appear at all.
- If $\#C(\mathbb{F}_2) = 7$, V spans a 5-plane, and this 5-plane does not contain a 7th point of OG^+ , we build a similar hash table, retaining 6-planes that occur exactly once.
- If $\#C(\mathbb{F}_2) = 6$ and V spans a 4-plane, we build a similar hash table, retaining 5-planes that do not appear at all. We then hash pairs of 5-planes in this list according to their span, retaining 6-planes that appear 3 times.
- If $\#C(\mathbb{F}_2) = 7$, V spans a 5-plane, and this 5-plane contains a 7th point of OG^+ , we may assume without loss of generality that the maximum intersection multiplicity of the linear span with OG^+ among the \mathbb{F}_2 -rational intersection points occurs for the 7th point. We then build a similar hash table, retaining 6-planes that do not appear at all.

We then take T to be the set of intersections of the resulting 6-planes with OG^+ .

8. Towards a full census in genera 6 and 7

It would be extremely desirable to refine the methods used here to complete a full census of curves of genus 6 and 7 over \mathbb{F}_2 , both to provide a consistency check on our own work and to make the tables available for other purposes. One important aspect of such a census is the process of making it simultaneously reliable and rigorous. In other words, given a putative list of isomorphism classes of curves of genus g over \mathbb{F}_2 , how can one verify that this list is accurate?

In one direction, it is easy from the data first to compute individual properties of the curves in question, such as their automorphism groups, and to test pairs of curves to confirm that they are not isomorphic; the latter can be accelerated by first hashing curves according to their zeta function (and any other data that may have been computed, such as the order and structure of the automorphism group).

In the other direction, one may check completeness of the list by computing the count of \mathbb{F}_q -points on the moduli space \mathcal{M}_g using the Lefschetz trace formula for the moduli space \mathcal{M}_g of stable curves of genus g and the combinatorics of the boundary strata. (Note that the point count is “stacky” in that each equivalence class of \mathbb{F}_q -points is weighted inversely by the order of its automorphism group.) For $g = 6$, it is known that $\#\mathcal{M}_g(\mathbb{F}_q)$ equals a fixed polynomial in q [1, Corollary 1.6], which

in principle can be computed using the SAGEMATH package described in [3]. For $g = 7$, while $\#\overline{\mathcal{M}}_g(\mathbb{F}_q)$ is known to equal a fixed polynomial in q [1, Corollary 1.5], the same does not immediately follow for $\#\mathcal{M}_g(\mathbb{F}_q)$ due to a possible contribution from the level-1 modular form Δ in the boundary of $\overline{\mathcal{M}}_7$.

As noted in the introduction, this discussion in principle also applies in some higher genera, using the parametrizations indicated in Remark 3.4. However, it is not clear to us to what extent the resulting computations are feasible.

Appendix A. Orbit lookup trees

Throughout this appendix, fix a finite group G and a finite set S equipped with a left G -action. We exhibit a combinatorial structure that allows us to efficiently compute orbit representatives for the action of G on k -element subsets of S for various small values of k . Such computations are already implemented in software (notably in MAGMA); however, the approach we take here seems to be well-adapted to our present work, as it avoids instantiating in memory the entire set of k -element subsets of S .

DEFINITION A.1. Let Γ be a directed graph with loops with vertex set $S \sqcup \{\bullet\}$, in which each edge is either of the form $\bullet \rightarrow v$ for some $v \in S$, or of the form $v_1 \xrightarrow{g} v_2$ for some $v_1, v_2 \in S$ with a label $g \in G$ satisfying $g(v_1) = v_2$. Defining connected components of Γ in terms of the underlying undirected graph, we say that a component is *eligible* if it does not contain \bullet and a vertex is *eligible* if it lies in an eligible component.

A *group retract* of Γ consists of a subset V of Γ consisting of one vertex in each eligible connected component, together with a function h from the union of the eligible components to G with the property that for each $v' \in \Gamma$ in the connected component of $v \in V$, we have $h(v')(v) = v'$.

One way to compute a group retract, given a choice of V , is to fix a spanning tree in each eligible component, then compute the unique function satisfying:

- for all $v \in V$, $h(v) = 1_G$;
- for each chosen spanning tree, for every edge $v_1 \xrightarrow{g} v_2$ in the tree, $h(v_2) = gh(v_1)$.

Note that this function can be computed in linear time in the input length.

DEFINITION A.2. Let F be a subset of the power set of S (the *forbidden subsets*). We say that a subset of S is *eligible* if it contains no forbidden subset.

For any positive integer n , an *orbit lookup tree* of depth n (for G, S, F) is a rooted tree T_n of depth n with the following properties (and additional data as indicated).

- Each node at depth k is labeled by a k -element subset U of S , colored either red or green. In what follows, we freely conflate nodes with their labels.
- The parent of every node U is a green node which is a subset of U . In particular, there is a unique ordering of the elements x_1, \dots, x_k of U such that each initial segment of this sequence is also a node; we write $U = [x_1, \dots, x_k]$ instead of $U = \{x_1, \dots, x_k\}$ when we need to indicate this choice of ordering.
- For $k = 0, \dots, n$, the green nodes at depth k form a set of G -orbit representatives for the eligible k -element subsets of S .

- For each eligible node U , we further record an element $g_U \in G$ such that $g_U^{-1}(U)$ is a green node.
- For each green node U , we further record the stabilizer G_U .
- Every green node U at depth $k < n$ has children which form a set of G_U -orbit representatives of the $(k+1)$ -element subsets of S containing U . We further record a function $h_U: S \setminus U \rightarrow G_U$ such that for each $y \in S \setminus U$, the element $x_U(y) := h_U^{-1}(y)$ has the property that $U \cup \{x_U(y)\}$ is a node.

As the name suggests, the structure of an orbit lookup tree makes it easy to find the chosen G -orbit representative of a subset of S .

ALGORITHM A.3. *Given an orbit lookup tree T_n of depth n , for any $k \in \{0, \dots, n\}$ and any sequence x_1, \dots, x_k of distinct elements of S , the following recursive algorithm determines whether $\{x_1, \dots, x_k\}$ is eligible, and if so produces a green node U of T_n and an element $g \in G$ such that $g(U) = \{x_1, \dots, x_k\}$.*

- (1) *If $k = 0$, return $U := \emptyset$, $g := 1_G$ and stop.*
- (2) *If $\{x_1, \dots, x_{k-1}\}$ is a node of T , let U' be this node and set $g_0 := 1_G$. Otherwise, apply the algorithm to x_1, \dots, x_{k-1} to obtain a green node U' of T and an element $g_0 \in G$ for which $g_0(U') = \{x_1, \dots, x_{k-1}\}$. If instead we find that $\{x_1, \dots, x_{k-1}\}$ is not eligible, report that U is not eligible and stop.*
- (3) *Set $y := g_0^{-1}(x_k)$, $U_1 := U' \cup \{x_{U'}(y)\}$, $g_1 := g_0 h_{U'}(y)$.*
- (4) *Set $g_2 := g_{U_1}$ and return $U := g_2^{-1}(U_1)$, $g := g_1 g_2$. If instead we find that g_{U_1} is undefined, then report that U is not eligible.*

REMARK A.4. In Algorithm A.5, we will use Algorithm A.3 in a situation where the elements g_{U_1} are not yet computed at depth k . By omitting step (4) and returning U_1, g_1 , we still obtain a node U of T_n and an element $g \in G$ such that $g(U) = \{x_1, \dots, x_k\}$, but without the guarantee that U is green. However, at deeper steps of the recursion we must execute Algorithm A.3 in full, including step (4).

The key point is that we can use group retracts to build an orbit lookup tree in an efficient fashion.

ALGORITHM A.5. *Given an orbit lookup tree T_n of depth n , the following algorithm extends T_n to an orbit lookup tree T_{n+1} of depth $n+1$.*

- (1) *For each green node U at depth n :*
 - (a) *Choose a sequence h_1, \dots, h_m of generators of G_U by picking random elements.*
 - (b) *Construct the Cayley graph Γ_U on $S \setminus U$ for the sequence h_1, \dots, h_m .*
 - (c) *Compute a group retract (V, g) of $\Gamma_U \cup \{\bullet\}$ (with no edges incident to \bullet) and set $h_U := g$.*
 - (d) *For each $y \in V$, add to T_{n+1} an uncolored node $U \cup \{y\}$ with parent U .*
- (2) *Construct a directed graph with loops Γ on the nodes of depth $n+1$ plus the dummy vertex \bullet as follows. For each node $U = [x_1, \dots, x_{n+1}] \in \Gamma$ (optionally in parallel):*
 - (a) *If U is forbidden, then add an edge $\bullet \rightarrow U$.*

- (b) If U is not forbidden, then for $j = 1, \dots, n$, apply Algorithm A.3, as modified in Remark A.4, to the sequence

$$x_1, \dots, x_{j-1}, x_{n+1}, x_{j+1}, \dots, x_n, x_j$$

to find a node U_1 and an element $g_1 \in G$ such that $g_1(U_1) = U$, and add the edge $U_1 \xrightarrow{g_1} U$ to Γ . If instead we find that U is not eligible, then add an edge $\bullet \rightarrow U$.

- (3) Compute a group retract (V, h) of Γ . Color each vertex in V green and color each remaining vertex (other than \bullet) red. For each vertex U in an eligible component, set $g_U := h(U)$.
- (4) For each green node $U = [x_1, \dots, x_{n+1}]$ at depth $n+1$, let G_U be the group generated by:
- the stabilizer of x_{n+1} in $G_{\{x_1, \dots, x_n\}}$;
 - for each edge $U_1 \xrightarrow{g} U_2$ in Γ , the element $g_{U_2}^{-1} g g_{U_1}$.

PROOF. The key point here is to confirm that the group computed in step (4), which is evidently contained in the stabilizer of the green node $U = [x_1, \dots, x_{n+1}]$, is actually equal to it. Let H_U be the group computed in (4) and let G_U be the full stabilizer; then the inclusions

$$\begin{aligned} G_{\{x_1, \dots, x_n\}} \cap G_{x_{n+1}} &\subseteq H_U \cap G_{\{x_1, \dots, x_n\}} \subseteq G_U \cap G_{\{x_1, \dots, x_n\}} = G_{\{x_1, \dots, x_n\}} \cap G_{x_{n+1}} \\ G_{\{x_1, \dots, x_n\}} \cap G_{x_{n+1}} &\subseteq H_U \cap G_{x_{n+1}} \subseteq G_U \cap G_{x_{n+1}} = G_{\{x_1, \dots, x_n\}} \cap G_{x_{n+1}} \end{aligned}$$

show that all of these groups coincide. That is, x_{n+1} has the same stabilizers in H_U and G_U , and so the orbit-stabilizer formula implies that the index $[G_U : H_U]$ equals the size of the G_U -orbit of x_{n+1} divided by the size of the H_U -orbit. Consequently, it suffices to check that the orbits coincide.

We again identify orbits with the connected components of the Cayley graph. If the G_U -orbit of x_{n+1} consists of x_{n+1} itself, there is nothing to check. Otherwise, the orbit also contains x_j for some $j \in \{1, \dots, n\}$, and the edges arising from the index j in step (2b) guarantee that x_j and x_{n+1} are joined in the Cayley graph. \square

REMARK A.6. Algorithm A.3 provides a consistency check for the computation of an orbit lookup tree, as one can spot-verify that a random k -element subset is indeed G -equivalent to some green node. For our purposes this is sufficient, as we only need a set that covers all orbits, not necessarily a set of orbit representatives.

If one really wants to verify that no two distinct green nodes are G -equivalent, it may be easiest to do this using some *ad hoc* computable invariants of the G -action. Alternatively, if no subsets are forbidden, we may verify the orbit-stabilizer formula: the sum of $[G : G_U]$ over green nodes U at depth n should equal $\binom{|S|}{n}$.

REMARK A.7. We have made no systematic effort to optimize Algorithm A.5 or even to give a careful costing. In step (2b) of Algorithm A.5, a further optimization is possible: instead of taking all $j \in \{1, \dots, n\}$, it suffices to take a set of orbit representatives for the action of $G_{\{x_1, \dots, x_n\}}$ on $\{1, \dots, n\}$. Our initial experiments were inconclusive as to whether this yielded a meaningful speedup in practice, so we did not pursue it.

REMARK A.8. In some applications, the set S carries the structure of a k -vector space for some finite field k , G acts k -linearly on S , and one is interested in the action of G on subspaces rather than subsets. One can treat this situation

by considering orbits of linearly independent subsets, but in practice it would be more efficient to adapt the algorithms to the linear setting. As we will not need this here, we omit the details.

References

- [1] S. Canning and H. Larson, On the Chow and cohomology rings and moduli spaces of stable curves, arXiv:2208.02357v1 (2022).
- [2] W. Castryck and J. Tuitman, Point counting on curves using a gonality preserving lift, *Quart. J. Math.* **69** (2018), 33–74.
- [3] V. Delecroix, J. Schmitt, and J. van Zelm, admcycles—a Sage package for calculations in the tautological ring of the moduli space of stable curves, *J. Software Alg. Geom.* **11** (2021), 89–112.
- [4] P. Deligne and D. Mumford, The irreducibility of the moduli space of curves of given genus, *Publ. Math. IHÉS* **36** (1969), 75–109.
- [5] W. Decker, G.-M. Greuel, G. Pfister, and H. Schönemann, Singular – A computer algebra system for polynomial computations, version 4.3.1p1, 2022, <http://www.singular.uni-kl.de>.
- [6] D. Dragutinović, Computing binary curves of genus five, arXiv:2202.07809v1 (2022); associated repository https://github.com/DusanDragutinovic/MT_Curves.
- [7] D. Eisenbud and J. Harris, The Kodaira dimension of the moduli space of curves of genus ≥ 23 , *Invent. Math.* **90** (1987), 359–387.
- [8] X. Faber and J. Grantham, Binary curves of small fixed genus and gonality with many rational points, *J. Algebra* **597** (2022), 24–46.
- [9] X. Faber, J. Grantham, and E.W. Howe, On the maximum gonality of a curve over a finite field, arXiv:2207.14307v1 (2022).
- [10] G. Farkas, Birational aspects of the geometry of $\overline{\mathcal{M}}_g$, *Surveys Diff. Geom.* **14** (2009), 57–110.
- [11] The GAP Group, GAP – Groups, Algorithms, and Programming, version 4.11.1, 2021, <https://www.gap-system.org>.
- [12] P. Griffiths and J. Harris, *Principles of Algebraic Geometry*, Wiley Interscience, 1978.
- [13] E.W. Howe, Deducing information about curves over finite fields from their Weil polynomials, arXiv:2110.04221v3 (2022).
- [14] E.W. Howe and K.E. Lauter, New methods for bounding the number of points on curves over finite fields, in *Geometry and arithmetic*, Eur. Math. Soc., Zürich, 2012, 173–212.
- [15] M. Ide and S. Mukai, Canonical curves of genus eight, *Proc. Japan Acad. Ser. A* **79** (2003), 59–64.
- [16] K.S. Kedlaya, Search techniques for root-unitary polynomials, in *Computational Arithmetic Geometry*, Contemporary Math. 463, Amer. Math. Soc., 2008, 71–82.
- [17] K.S. Kedlaya, The relative class number one problem for function fields, I, *Res. Num. Theory* **8** (2022), proceedings of Algorithmic Number Theory Symposium (ANTS-XV), article 79.
- [18] K.S. Kedlaya, The relative class number one problem for function fields, II, arXiv:2206.02084v1 (2022).
- [19] K.S. Kedlaya, GitHub repository <https://github.com/kedlaya/same-class-number>.
- [20] S.L. Kleiman and D. Laksov, On the existence of special divisors, *Amer. J. Math.* **94** (1972), 431–436.
- [21] J.R.C. Leitzel and M.L. Madan, Algebraic function fields with equal class number, *Acta Arith.* **30** (1976), 169–177.
- [22] J.R.C. Leitzel, M.L. Madan, and C.S. Queen, Algebraic function fields with small class number, *J. Number Theory* **7** (1975), 11–27.
- [23] The LMFDB Collaboration, L-Functions and Modular Forms Database, <https://lmfdb.org>.
- [24] The Magma Group, University of Sydney, MAGMA version 2.27-1, 2022, <http://magma.maths.usyd.edu.au>.
- [25] A. Maroni, Le serie lineari speciali sulle curve trigonali, *Ann. Mat. Pura Appl.* **25** (1946), 341–343.
- [26] S. Mukai, Curves and Grassmannians, in *Algebraic Geometry and Related Topics*, International Press, Cambridge, MA, 1993, 19–40.
- [27] S. Mukai, Curves and symmetric spaces, I, *Amer. J. Math.* **117** (1995), 1627–1644.
- [28] S. Mukai, Curves and symmetric spaces, II, *Annals of Math.* **172** (2010), 1539–1558.

- [29] S. Mukai, Curves, K3 surfaces, and Fano 3-folds of genus ≤ 10 , in *Algebraic Geometry and Commutative Algebra*, Kinokuniya, Tokyo, 1988, 357–377.
- [30] The Sage Developers, SageMath version 9.7, 2022, <https://www.sagemath.org>.
- [31] B. Saint-Donat, On Petri’s analysis of the linear system of quadrics through a canonical curve, *Math. Ann.* **206** (1973), 157–175.
- [32] H. Stichtenoth, *Algebraic Function Fields and Codes*, second edition, Graduate Texts in Math. 254, Springer–Verlag, Berlin, 2009.
- [33] A. Verra, The unirationality of the moduli spaces of curves of genus 14 or lower, *Compos. Math.* **141** (2005), 1425–1444.
- [34] X. Xarles, A census of all genus 4 curves over the field with 2 elements, arXiv:2007.07822v1 (2020).