The relative class number one problem for function fields, I

Kiran S. Kedlaya

Department of Mathematics, University of California San Diego
kedlaya@ucsd.edu

These slides can be downloaded from https://kskedlaya.org/slides/.
Jupyter notebooks available from https://github.com/kedlaya/same-class-number.

Algorithmic Number Theory Symposium (ANTS-XV)
University of Bristol
August 9, 2022

Supported by NSF (grants DMS-2053473 and prior) and UC San Diego (Warschawski Professorship).

The UC San Diego campus sits on unceded ancestral land of the Kumeyaay Nation.
The University of Bristol was chartered using funds predominantly derived from the transatlantic slave trade.

The problem

Let F'/F be a finite extension of function fields of curves over finite fields. Let $g_F, g_{F'}$ be the genera of F and F'. Let $q_F, q_{F'}$ be the cardinalities of the base fields* of F, F'.

Let $h_F, h_{F'}$ be the class numbers of F and F'. The ratio $h_{F'/F} := h_{F'}/h_F$ is always an integer (more on this shortly). Following Leitzel–Madan (1976), we ask: in what cases does $h_{F'/F} = 1$?

To make this a potentially finite problem, we only specify the isomorphism classes of F and F', not the inclusion (this only makes a difference when $g_F \leq 1$). We also ignore the trivial cases:

- $F' \cong F$;
- $g_F = g_{F'} = 0$.

*By "base field" I mean the integral closure of the prime subfield.

Contrast with the number field case

In the number field setting, class number 1 is much more common, because class groups are always "incomplete". The product

$$\text{class number} \times \text{unit regulator}$$

behaves much more predictably, and can be interpreted as the volume of a natural compact topological group (the **Arakelov class group**).

For relative class number 1, one can only hope for a finiteness result for (nontrivial) extensions which preserve the unit rank, i.e., CM fields.[†] For **normal** CM fields, finiteness was proved by Odlyzko and the full classification (under GRH) by Hoffman–Sircana.

By contrast, the full Picard group of a function field looks like $\mathbb{Z} \times$ (finite) and removing one point always takes out $\mathbb{Z}$.

[†] A **CM field** is a totally imaginary quadratic extension of a totally real field.

Constant vs. geometric extensions

We say that:

- F'/F is **constant** if $F' = F \cdot \mathbb{F}_{q_{F'}}$;
- F'/F is **purely geometric** (hereafter **geometric**) if $q_F = q_{F'}$.

Let E be the compositum $F \cdot \mathbb{F}_{q_{F'}}$; then E/F is constant and F'/E is geometric. Since the relative class number is always an integer, $h_{F'/F} = 1$ if and only if $h_{E/F} = h_{F'/E} = 1$.

The relative class number one problem thus reduces to the constant and geometric cases. The constant case is relatively easy, so in this talk I will focus on the geometric case. Hereafter, unless specified assume F'/F is geometric and write

$$q := q_F = q_{F'}, \qquad g := g_F, \qquad g' := g_{F'}.$$

The Prym variety

Let C, C' be the curves with function fields F, F'. We have an isogeny of abelian varieties

$$J(C') \cong J(C) \times A$$

for some abelian variety A over $\mathbb{F}_q$, called the **Prym variety**. We have[‡]

$$h_{F'/F} = \#A(\mathbb{F}_q) \in \mathbb{Z}.$$

In particular, if $\#A(\mathbb{F}_q) = 1$ and $F' \neq F$, then:

- we have $q \leq 4$ by the Weil bounds;
- for $q = 3, 4$, A is isogenous to a product of the unique elliptic curve E over $\mathbb{F}_q$ with $\#E(\mathbb{F}_q) = 1$;
- for $q = 2$, A is isogenous to a product of simple factors classified by Madan–Pal–Robinson in 1977.

[‡]This holds even if F'/F is not geometric, and explains why $h_{F'/F} \in \mathbb{Z}$ as promised.

A lower bound on point counts

Let T_{A,q^n} be the trace of the q^n-power Frobenius on A; then

$$\#C(\mathbb{F}_{q^n}) = \#C'(\mathbb{F}_{q^n}) + T_{A,q^n} \geq T_{A,q^n}.$$

For $q = 3, 4$, we have $1 = \#E(\mathbb{F}_q) = q + 1 - T_{E,q}$ and so[§]

$$\#C(\mathbb{F}_q) \geq T_{A,q} = q \dim(A) = q(g' - g) \geq q(g - 1).$$

For $q = 2$, we can have $T_{A,q} = 0$, so there is no useful bound on $\#C(\mathbb{F}_2)$. But using the Madan–Pal–Robinson classification, data from LMFDB for $\dim(A) \leq 6$, and a bit of linear programming, we get

$$\begin{aligned} &1.3366\,T_{A,2} + 0.3366\,T_{A,4} + 0.1137(T_{A,8} - T_{A,2}) \\ &\quad + 0.0537(T_{A,16} - T_{A,4}) \geq 1.5612\dim(A) \quad \Longrightarrow \\ &1.3366\#C(\mathbb{F}_2) + 0.3366\#C(\mathbb{F}_4) + 0.1137(\#C(\mathbb{F}_8) - \#C(\mathbb{F}_2)) \\ &\quad + 0.0537(\#C(\mathbb{F}_{16}) - \#C(\mathbb{F}_4)) \geq 1.5612(g' - g) \geq 1.5612(g - 1). \end{aligned}$$

[§] The estimate $g' - g \geq g - 1$ follows from Riemann–Hurwitz.

Comparison with upper bounds on point counts

We now compare with effective "linear programming" upper bounds on $\#C(\mathbb{F}_{q^n})$ (Ihara, Drinfeld–Vlăduț, Oesterlé, Serre).

$$q = 4: \qquad \#C(\mathbb{F}_q) \leq 1.435g + 21.75$$
$$q = 3: \qquad \#C(\mathbb{F}_q) \leq 1.153g + 11.67.$$

For $q = 2$, let a_i be the number of degree-i closed points on C; then

$$a_1 + 0.3366(2a_2) + 0.1382(3a_3) + 0.0537(4a_4) \leq 0.8042g + 5.619.$$

For each q, combining this slide with the previous one limits (g, g') to an explicit finite list.

We have now reduced the relative class number one problem to a finite computation! However, some care is required to make this tractable; the computation is **mostly** finished in this paper, up to some loose ends.

Outline of the finite computation for $g \leq 1$

Reminder: for $g \leq 1$, we are only trying to identify the isomorphism classes of C and C', not the map.

- For each possible pair (g, g'), enumerate candidate Weil polynomials for C and C' in SageMath.[¶]
- For each pair of Weil polynomials, if possible, use LMFDB to identify all C and C' with those Weil polynomials. LMFDB contains data about abelian varieties over finite fields (Dupuy–K–Roe–Vincent) and Jacobians (Howe, Xarles, Dragutinović).

This only fails in two cases with $q = 2, g = 1, g' = 6$. In one of these, C' is ruled out by an argument of Grantham–Howe–Faber (based on Serre's resultant criterion). In the other, there exists a suitable C' which is a cyclic 5-fold étale cover of a certain genus-2 curve. **Loose end:** uniqueness.

[¶]This uses C code of mine dating back to 2008.

Outline of the finite computation for $g > 1$

- For each pair (g, g'), use Riemann–Hurwitz to compute options for $d = [F' : F]$.
- Use further constraints based on d to eliminate some triples (d, g, g').
- For each remaining triple (d, g, g'):
  - Enumerate Weil polynomials for C and C' using SageMath. (The rate-limiting cases are $(d, g, g') = (2, 8, 15), (2, 9, 17)$.)
  - Use LMFDB to identify all C with a suitable Weil polynomial. **Loose end**: do this for $q = 2$, $g = 6, 7$.
  - For each C, use class field theory in Magma to find all cyclic extensions F'/F of the right degree and genus, then check the relative class number.
  - If $d > 2$, use the Weil polynomial constraints to rule out all noncyclic extensions. For $q > 2$, we only need to handle $d = 3$. **Loose end**: do this for $q = 2$.

Loose ends

We have completed the finite computation for $q = 3, 4$. For $q = 2$, there are three remaining steps.

- For $g = 1, g' = 6$, we must check that there is only one candidate for C'. This uses a technique of Howe which uses the particular shape of the zeta function to force C' to admit an order-5 automorphism.
- For $g > 1$, we have $d \leq 7$ and this is sharp (!). Ruling out noncyclic extensions requires studying the zeta functions of other quotients of the Galois closure; similar ideas were used by Rigato to sharpen upper bounds on the number of $\mathbb{F}_q$-points on a genus-g curve.
- For $d = 2$, we have $g \leq 7$ and this is sharp (!!). For $g = 6, 7$ we do not (yet!) have a table of isomorphism classes of genus-g curves over $\mathbb{F}_2$, so we make a targeted enumeration over M_g to find these curves.

These three steps are elaborated in two subsequent papers "The relative... II, III" (currently available as preprints).

Summary of the results, part 1

Theorem

Assume F'/F is constant and $g_F > 0$. Then (q_F, d, g_F) is one of

$$(2,2,1),(2,2,2),(2,2,2),(2,2,3),(2,3,1),(2,3,1),(3,2,1),(4,2,1)$$

and all options for F are known.

Theorem

Assume F'/F is geometric, $g_F \leq 1$, and $g_{F'} > g_F$. Then

$$(q_F, g_F, g_{F'}) \in \{(2,0,1),(2,0,2),(2,0,3),(2,0,4),(2,1,2),(2,1,3),(2,1,4),(2,1,5),(2,1,6),(3,0,1),(3,1,2),(3,1,3),(4,0,1),(4,1,2)\}$$

and all options for (F, F') are known except when $g_{F'} = 6$.

Summary of the results, part 2

Theorem

Assume F'/F is geometric, $g_{F'} > g_F > 1$, and $q_F > 2$. Then

$$(q_F, d, g_F, g_{F'}) \in \{(3,2,2,3),(3,2,2,4),\\(3,2,3,5),(3,3,2,4),(4,2,2,3),(4,3,2,4)\}$$

and all options for F'/F are known and cyclic.

Theorem

Assume F'/F is geometric, $g_{F'} > g_F > 1$, $q_F = 2$, and $d > 2$. Then

$$(d, g_F, g_{F'}) \in \{(3,2,4),(3,2,6),(3,3,7),(3,4,10),\\(4,2,5),(4,2,6)\star,(4,3,9)\star,(5,2,6),(6,2,7)\star,(7,2,8)\}$$

and all **cyclic** *options are known (covering all cases not marked $\star$).*

Summary of the results, part 3

Theorem

Assume F'/F is geometric, $g_{F'} > g_F > 1$, $q_F = 2$, and $d = 2$. Then

$$(g_F, g_{F'}) \in \{(2,3),(2,4),(2,5),\\ (3,5),(3,6),(4,7),(4,8),(5,9),(6,11),(7,13)\}$$

and all options with $g_F \leq 5$ are known. There are at least two examples with $g_F = 6$ and at least one with $g_F = 7$.

What about larger relative class numbers?

In principle, one can use similar techniques to solve the relative class number m problem[‖] for any fixed $m > 1$, with two caveats.

- It is probably hopeless to classify abelian varieties A over $\mathbb{F}_2$ with $\#A(\mathbb{F}_2) = m$. However, it should be possible to make a direct linear programming argument to establish a useful lower bound on some linear combination of traces of A.
- We cannot hope to exclude noncyclic extensions. One alternative might be a good method to enumerate degree-d extensions of a fixed function field; for $d = 3, 4, 5$ this should be doable** using Bhargava's parametrizations.

[‖] Again, when the base field has genus 0 or 1, one can only hope to describe the isomorphism classes of the two fields and not the morphism.

** In the number field setting, this was done by Belabas for $d = 3$.