Algebra & Number Theory

Volume 7

2013

No. 6

Identifying Frobenius elements in Galois groups

Tim Dokchitser and Vladimir Dokchitser

msp

ALGEBRA AND NUMBER THEORY 7:6 (2013)
dx.doi.org/10.2140/ant.2013.7.1325

Identifying Frobenius elements in Galois groups

Tim Dokchitser and Vladimir Dokchitser

We present a method to determine Frobenius elements in arbitrary Galois extensions of global fields, which may be seen as a generalisation of Euler's criterion. It is a part of the general question how to compare splitting fields and identify conjugacy classes in Galois groups, which we will discuss as well.

1. Introduction

Take a Galois extension $L/\mathbb{Q}$. Associated to each (unramified) prime p is a Frobenius element Frob_p, an element of the Galois group that reduces to $x \mapsto x^p$ modulo a prime above p. In the setting when L is the splitting field of a polynomial f, this element is intimately connected to the factorisation of f mod p: Viewed as a permutation of the roots, Frob_p is a product of disjoint cycles whose lengths are the degrees of the irreducible factors.

In this paper, we address the question of how to determine Frob_p. Generally, we study the problem of how to compare splitting fields and identify conjugacy classes in Galois groups; see Sections 2–4. Our motivation was computing L-series of Artin representations for arbitrary Galois groups, which requires the knowledge of Frobenius elements at all primes; see Remark 5.8 and Example 7.7. Obtaining

MSC2010: primary 11R32; secondary 11R42, 12F10.
Keywords: Frobenius elements, Artin representations, Galois groups.

them directly from the definition is impractical unless L either has small degree or is particularly simple to work with.

Let us briefly illustrate the various standard techniques for computing Frobenius elements. As before, L is the splitting field of a polynomial $f \in \mathbb{Z}[x]$, and we write $G = \mathrm{Gal}(L/\mathbb{Q})$.

Quadratic fields. Suppose $f(x) = x^2 - d$, so $L = \mathbb{Q}(\sqrt{d})$. For a prime $p \nmid 2d$, the Frobenius element is given by the Legendre symbol:

$$\mathrm{Frob}_p = \mathrm{id} \quad \Longleftrightarrow \quad f(x) \bmod p \text{ is reducible} \quad \Longleftrightarrow \quad \left(\frac{d}{p}\right) = 1.$$

There are two essentially different methods to compute it:

(A) Euler's criterion $\left(\frac{d}{p}\right) \equiv d^{(p-1)/2} \bmod p$.

(B) Quadratic reciprocity.

Kummer extensions. Suppose $f(x) = x^3 - 2$, so $L = \mathbb{Q}(\zeta_3, \sqrt[3]{2})$ and $G = \mathrm{S}_3$. For $p \neq 2, 3$ the number of cube roots of 2 mod p determines whether Frob_p is trivial, a 3-cycle or a transposition. It is easy to see that the last case is equivalent to $p \equiv 2 \bmod 3$. There are analogues of both (A) and (B) to distinguish between the first two cases:

(A) Euler's criterion: Since $\mathbb{F}_p^\times$ is cyclic,

$$\begin{aligned} 2 \text{ is a cube mod } p \quad &\Longleftrightarrow \quad 2^{(p-1)/3} \equiv 1 \bmod p. \\ 2 \text{ not a cube mod } p \quad &\Longleftrightarrow \quad 2^{(p-1)/3} \text{ is another third root of unity } z \in \mathbb{F}_p. \end{aligned}$$

To link this criterion to our main theorem below, let us rephrase it: Let

$$M = \left(\begin{smallmatrix} 0 & 0 & 2 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{smallmatrix}\right) \in \mathrm{GL}_3(\mathbb{F}_p), \quad \text{so that } M^3 = \left(\begin{smallmatrix} 2 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{smallmatrix}\right) \quad \text{and} \quad f(M) = 0.$$

Then

$$\begin{aligned} \mathrm{Frob}_p = \mathrm{id} &\iff M^{p-1} = \left(\begin{smallmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{smallmatrix}\right) \iff \tfrac{1}{3}\operatorname{Tr} M^{p-1} = 1, \\ \mathrm{Frob}_p \in [(123)] &\iff M^{p-1} = \left(\begin{smallmatrix} z & 0 & 0 \\ 0 & z & 0 \\ 0 & 0 & z \end{smallmatrix}\right) \iff \tfrac{1}{3}\operatorname{Tr} M^{p-1} \text{ satisfies } t^2+t+1 = 0, \\ \mathrm{Frob}_p \in [(12)] &\iff M^{p-1} = \left(\begin{smallmatrix} 0 & 0 & * \\ * & 0 & 0 \\ 0 & * & 0 \end{smallmatrix}\right) \iff \tfrac{1}{3}\operatorname{Tr} M^{p-1} = 0. \end{aligned}$$

(B) Class field theory over $\mathbb{Q}(\zeta_3)$:

Factorise $p = (a + b\zeta_3)(a + b\bar{\zeta}_3)$. Then 2 is a cube mod p if and only if the ideal $(a + b\zeta_3)$ splits in L, and class field theory says that this is a congruence condition on a and b. In fact, it is easy to verify that

$$2 \text{ is a cube mod } p \quad \Longleftrightarrow \quad a + b\zeta_3 \equiv \pm 1, \pm\zeta_3 \text{ or } \pm\zeta_3^2 \bmod 6.$$

Modular forms. See [Zagier 2008, §4.3]. Suppose $f(x) = x^3 - x - 1$, so $G = \mathrm{S}_3$ and L is the Hilbert class field of $\mathbb{Q}(\sqrt{-23})$. Let ρ be the 2-dimensional irreducible representation of G. It has an associated Artin L-series

$$L(\rho, s) = \sum_{n=1}^{\infty} \frac{a_n}{n^s},$$

whose coefficient a_p for a prime $p \neq 23$ is 2, -1 or 0 depending on whether Frob_p is trivial, a 3-cycle or a transposition. The theory of modular forms tells us that

$$\sum_{n=1}^{\infty} a_n q^n = q \prod_{n=1}^{\infty} (1 - q^n)(1 - q^{23n}),$$

and is a cusp form of weight 1, level 23 and character $\left(\frac{\cdot}{23}\right)$. Moreover, for all integers n not divisible by 23,

$$a_n = \tfrac{1}{2}\big(\#\{x, y \in \mathbb{Z} \mid n = x^2 + xy + 6y^2\} - \#\{x, y \in \mathbb{Z} \mid n = 2x^2 + xy + 3y^2\}\big).$$

Let us remark that in an arbitrary Galois group G, the L-series of the irreducible representations of G also pin down the Frobenius elements. The global Langlands conjecture predicts that, as in this example, all such L-series come from automorphic forms. This is a massive conjectural generalisation of "method (B)". Moreover, like quadratic reciprocity and class field theory, this approach gives expressions for the L-series coefficients a_n that do not depend on n being prime. This is crucial for theoretical applications such as analytic continuation of L-functions. (Note, however, that formulas such as the one above are not practical for numerically computing Frobenius elements.)

The purpose of this paper is to extend "method (A)" to arbitrary Galois groups. Here is an illustration for cubic polynomials of the type of criterion that we obtain. Note its similarity to the Kummer case.

General cubic. Suppose $f(x) = x^3 + bx + c$. Pick a prime $p \nmid 3b\Delta$, where $\Delta = -4b^3 - 27c^2$ is the discriminant of f. Let

$$M = \begin{pmatrix} 0 & 0 & -c \\ 1 & 0 & -b \\ 0 & 1 & 0 \end{pmatrix} \in \mathrm{GL}_3(\mathbb{F}_p).$$

Then

$$\begin{aligned} f(x) \text{ has 3 roots mod } p &\iff \operatorname{Tr} M^{p+1} = -2b, \\ f(x) \text{ has 1 root mod } p &\iff \operatorname{Tr} M^{p+1} \text{ satisfies } (t + 2b)(t - b)^2 = -\Delta, \\ f(x) \text{ is irreducible mod } p &\iff \operatorname{Tr} M^{p+1} = b. \end{aligned}$$

This can be easily checked by hand; alternatively, see Theorem 7.2.

Our main result for Frobenius elements is the following generalisation of Euler's criterion. Note that taking the class of x in $\mathbb{F}_q[x]/f(x)$ is the same as taking a matrix M with characteristic polynomial $f(x)$, like in the examples above.

Theorem 1.1. *Let K be a global field and $f(x) \in K[x]$ a separable polynomial with Galois group G. There is a polynomial $h(x) \in K[x]$ and polynomials $\Gamma_C \in K[X]$ indexed by the conjugacy classes C of G such that*

$$\mathrm{Frob}_{\mathfrak{p}} \in C \quad \Longleftrightarrow \quad \Gamma_C\big(\mathrm{Tr}_{\frac{\mathbb{F}_q[x]}{f(x)}/\mathbb{F}_q}(h(x)x^q)\big) = 0 \mod \mathfrak{p}$$

for almost all primes $\mathfrak{p}$ of K; here $\mathbb{F}_q$ is the residue field at $\mathfrak{p}$.

This is proved in Section 5; see Theorem 5.3. Usually one can take $h(x) = x^2$ (see below); in particular $\mathrm{Tr}(x^{q+2})$ then determines the conjugacy class of $\mathrm{Frob}_{\mathfrak{p}}$. In Section 6 we explain how the theorem recovers classical formulas for Frobenius elements in cyclotomic and Kummer extensions. In Section 7 we give explicit examples for nonabelian Galois groups, including general cubics, general quartics and quintics with Galois group D_{10}.

The polynomials Γ_C are explicitly given by

$$\Gamma_C(X) = \prod_{\sigma \in C} \big(X - \sum_{j=1}^{n} h(a_j)\sigma(a_j)\big),$$

where $a_1, \ldots, a_n$ are the roots of f in some splitting field. The "almost all primes" in the theorem are those not dividing the denominators of the coefficients of f, its leading coefficient and the resultants $\mathrm{Res}(\Gamma_C, \Gamma_{C'})$ for $C \neq C'$; the latter simply says that the Γ_C mod $\mathfrak{p}$ are pairwise coprime. (This condition always fails for ramified primes; see Remark 5.6.) Finally, the only constraint on the polynomial h is that the resulting Γ_C are coprime over K. This holds for almost all h, in the sense that the admissible ones of degree at most $n-1$ form a Zariski dense open subset of K^n. Also, a fixed h with $1 < \deg h < n$ (for instance $h(x) = x^2$) will work for almost all f that define the same field; see Section 8.

Remark 1.2. The method of using polynomials in the roots of f to recognise conjugacy classes is also used in "Serre's trick" for alternating groups. For example, $G = A_5$ has 5 conjugacy classes and all but the two classes of 5-cycles have their own cycle type. (Recall that the cycle type of Frobenius can be recovered from the degrees of the factors of the defining quintic f mod p; in practice, these are readily determined by computing $\gcd(x^{p^d} - x, f(x))$ for $d = 1, 2$.) It was pointed out by Serre (see Buhler [1978, p. 53]) that the classes of 5-cycles can be distinguished by evaluating the square root of the discriminant of f modulo p; see Example 3.9. This has been generalised by Roberts [2004] to all alternating groups, and was used for instance by Booker [2005] in his work on L-series for icosahedral representations.

Finally, let us illustrate our approach to Frobenius elements with a simple case:

Example 1.3. The polynomial $f(x) = x^5 + 2x^4 - 3x^3 + 1$ has Galois group $G = \mathrm{D}_{10}$ over $K = \mathbb{Q}$. If we number its complex roots by

$$a_1 \approx -3.01, \quad a_2 \approx -0.35 - 0.53i, \quad a_3 \approx 0.85 - 0.31i, \quad a_4 = \overline{a_3}, \quad a_5 = \overline{a_2},$$

then G is generated by the 5-cycle (12345) and complex conjugation (25)(34). It is easy to see that $f(x)$ is irreducible over $\mathbb{F}_2$, so $\mathrm{Frob}_2 \in G$ is in one of the two conjugacy classes of 5-cycles, either $[(12345)]$ or $[(12345)^2]$. How can we check which one it is?

Consider the expressions

$$n_1 = a_1a_2 + a_2a_3 + a_3a_4 + a_4a_5 + a_5a_1,$$
$$n_2 = a_1a_3 + a_2a_4 + a_3a_5 + a_4a_1 + a_5a_2.$$

If we think of G as the group of symmetries of a pentagon, the sums are taken over all edges and over all diagonals, respectively. Therefore they are clearly G-invariant, and hence rational numbers. Also, as a_i are algebraic integers, n_1 and n_2 are in fact integers, readily recognised from their complex approximations as being 2 and -5.

Now suppose b_1 is a root of $f(x)$ in $\mathbb{F}_{2^5}$, and $b_i = b_{i-1}^2$ for $i = 2, 3, 4, 5$ are its other roots ordered by the action of the Frobenius automorphism. Then

$$N = b_1b_2 + b_2b_3 + b_3b_4 + b_4b_5 + b_5b_1$$

is in $\mathbb{F}_2$. By considering the reduction modulo a prime $\mathfrak{q}$ above 2 in the splitting field, we see that if $\mathrm{Frob}_\mathfrak{q}$ is (12345) or $(12345)^{-1}$, then $n_1 \equiv N \bmod 2$. Similarly, if $\mathrm{Frob}_\mathfrak{q}$ is $(12345)^2$ or $(12345)^3$, then $n_2 \equiv N \bmod 2$. Computing in $\mathbb{F}_2^5$ (or noting that $N = \mathrm{Tr}_{\mathbb{F}_2[x]/f(x)}(x^3)$) we find that $N = 0$, so Frob_2 must be in $[(12345)]$.

In the language of Theorem 1.1, we took $h(x) = x$ and proved that

$$\Gamma_{[(12345)]} = (X-2)^2 \quad \text{and} \quad \Gamma_{[(12345)^2]} = (X+5)^2$$

distinguish between the two conjugacy classes of 5-cycles: If $f(x)$ is irreducible mod p (and $p \neq 7$, so that $2 \not\equiv -5$), then

$$\mathrm{Frob}_p \in C \quad \Longleftrightarrow \quad \Gamma_C(\mathrm{Tr}_{\mathbb{F}_p[x]/f(x)}(x^{p+1})) = 0 \bmod p.$$

This choice of $h(x)$ was in some sense deceptively simple, because the roots n_i of the Γ_C were integers. (We used that the conjugacy classes of 5-cycles are self-inverse in D_{10}.) Generally, these roots would be algebraic integers of degree $|C|$. For example, $h(x) = x^2$ leads to

$$\Gamma_{[(12345)]} = X^2 + 5X + 18 \quad \text{and} \quad \Gamma_{[(12345)^2]} = X^2 - 11X + 42,$$

and $\mathrm{Tr}(x^{p+2})$ is a root of one of them whenever $f(x) \bmod p$ is irreducible.

Notation. Throughout the paper we use the following notation:

K	ground field
$f(x)$	separable polynomial in $K[x]$ of degree n
L	some extension of K where f splits completely
$\boldsymbol{a} = [a_1, \dots, a_n]$	ordered roots of f in L
$K(\boldsymbol{a})$	field generated by the a_i over K (a splitting field of f)
$G_{\boldsymbol{a}}$	Galois group of f, considered as a subgroup of S_n via its permutation action on $[a_1, \dots, a_n]$.
$[\Psi]$	conjugacy class of $\Psi \in G_{\boldsymbol{a}}$.
$\mathfrak{p}$	prime of K, when K is a global field
$\mathbb{F}_q$	residue field at $\mathfrak{p}$
$\mathrm{Frob}_{\mathfrak{p}}$	any (arithmetic) Frobenius element at $\mathfrak{p}$ in $G_{\boldsymbol{a}}$
$e^F_{\boldsymbol{a}}$, Γ, $M^F_{\boldsymbol{a},\Psi}$	see Definitions 2.2, 2.7, 3.4 and 4.3.

Recall that a global field is a finite extension of either $\mathbb{Q}$ or $\mathbb{F}_p(T)$. The Frobenius element in $\mathrm{Gal}(L/K)$ at $\mathfrak{p}$ is characterised by $\mathrm{Frob}_{\mathfrak{p}}(x) \equiv x^q \bmod \mathfrak{q}$ for all $x \in L$ that are integral at some fixed prime $\mathfrak{q}$ of L above $\mathfrak{p}$. The element $\mathrm{Frob}_{\mathfrak{p}}$ is well-defined modulo inertia and up to conjugation. In particular, its conjugacy class is well-defined if $\mathfrak{p}$ is unramified in L/K.

The symmetric group S_n acts on n-tuples by $[c_1, \dots, c_n]^\sigma = [c_{\sigma^{-1}(1)}, \dots, c_{\sigma^{-1}(n)}]$. It acts on the ring of polynomials in n variables $K[x_1, \dots, x_n]$ by $\sigma(x_i) = x_{\sigma(i)}$; thus, for a polynomial $F \in K[x_1, \dots, x_n]$,

$$F^\sigma([c_1, \dots, c_n]) = F([c_1, \dots, c_n]^{\sigma^{-1}}),$$

where $F([\,\cdot\,])$ is the evaluation of F on the n-tuple. Note that all our actions are left actions.

2. Isomorphisms of splitting fields

In this section we introduce our main tools. The reader who is only interested in applications to Frobenius elements may skip to Section 5 and prove Theorem 5.3 directly (at the expense of not seeing the origins of Γ_C).

As a motivation, consider the following general question:

Problem 2.1. Suppose a given separable polynomial $f(x) \in K[x]$ of degree n splits completely in $L \supset K$ and $L' \supset K$. Given the roots $a_1, \dots, a_n$ and $b_1, \dots, b_n$ of f in L and L', find a bijection between them that comes from an isomorphism of splitting fields of f inside L and L'.

We assume that we know the Galois group of f over K as a permutation group on the roots in L, but we do not want to construct the splitting fields explicitly. Instead,

we will evaluate polynomials in $K[x_1, \ldots, x_n]$ on the roots in L and L' taken in various orders and try to extract information out of the values (as in Example 1.3).

Definition 2.2. For $F \in K[x_1, \ldots, x_n]$ define the *evaluation map* $S_n \to K(\boldsymbol{a})$ by

$$e^F_{\boldsymbol{a}}(\sigma) = F([a_1, \ldots, a_n]^\sigma).$$

Definition 2.3. Let T be a subgroup of S_n. A *T-invariant* F is an element of $K[x_1, \ldots, x_n]$ whose stabiliser is precisely T.

Remark 2.4. Any $F \in K[x_1, \ldots, x_n]$ is evidently T-invariant if we take for T its stabiliser in S_n. Also, any subgroup $T < S_n$ has a T-invariant, for example,

$$F = \sum_{t \in T} m^t, \quad \text{where } m = x_1^{n-1} x_2^{n-2} \cdots x_{n-1},$$

since clearly the stabiliser of m in S_n is $\{1\}$.

Lemma 2.5. *Let F be a T-invariant and $\sigma, \tau \in S_n$.*

(1) $e^F_{\boldsymbol{a}^\tau}(\sigma) = e^F_{\boldsymbol{a}}(\sigma\tau)$.

(2) $g(e^F_{\boldsymbol{a}}(\sigma)) = e^F_{\boldsymbol{a}}(\sigma g^{-1})$ *for* $g \in G_{\boldsymbol{a}}$.

(3) *The map* $e^F_{\boldsymbol{a}} : S_n \to K(\boldsymbol{a})$ *is constant on the right cosets* $T\sigma$.

Proof. (1) $e^F_{\boldsymbol{a}^\tau}(\sigma) = F((\boldsymbol{a}^\tau)^\sigma) = F(\boldsymbol{a}^{\sigma\tau}) = e^F_{\boldsymbol{a}}(\sigma\tau)$.

(2) For $g \in G_{\boldsymbol{a}}$,

$$\begin{aligned} g(e^F_{\boldsymbol{a}}(\sigma)) &= g(F([a_1, \ldots, a_n]^\sigma)) = F([g(a_1), \ldots, g(a_n)]^\sigma) \\ &= F(([a_1, \ldots, a_n]^{g^{-1}})^\sigma) = F([a_1, \ldots, a_n]^{\sigma g^{-1}}) \quad = e^F_{\boldsymbol{a}}(\sigma g^{-1}). \end{aligned}$$

(3) For $\tau \in T$,

$$\begin{aligned} e^F_{\boldsymbol{a}}(\tau\sigma) &= F([a_1, \ldots, a_n]^{\tau\sigma}) = F(([a_1, \ldots, a_n]^\sigma)^\tau) \\ &= F^{\tau^{-1}}([a_1, \ldots, a_n]^\sigma) = F([a_1, \ldots, a_n]^\sigma) \quad = e^F_{\boldsymbol{a}}(\sigma). \end{aligned}$$

□

Remark 2.6. Part (3) of the lemma says that the values of F on the various permutations $\boldsymbol{a}^\sigma$ of the roots are essentially the right cosets of T in S_n. It may accidentally happen that the same value occurs on two right cosets, but it is always possible to adjust the original polynomial f to prevent this (see Lemma 8.1c). Part (2) of Lemma 2.5 says that the action of the Galois group $\mathrm{Gal}(K(\boldsymbol{a})/K)$ on these values translates into right multiplication by $G_{\boldsymbol{a}}$. This motivates the following:

Definition 2.7. For a double coset $D = T\sigma_0 G_{\boldsymbol{a}}$ in S_n, define the corresponding "minimal polynomial"

$$\Gamma^F_{\boldsymbol{a},\sigma_0} = \Gamma^F_{\boldsymbol{a},D}(X) = \prod_{\sigma \in T \backslash D} (X - e^F_{\boldsymbol{a}}(\sigma)) \in K[X].$$

By Lemma 2.5(3), this is well-defined.

Remark 2.8. Note that by Lemma 2.5(2), $G_{\boldsymbol{a}}$ permutes the linear factors of $\Gamma^F_{\boldsymbol{a},D}$ transitively, so it is a power of an irreducible polynomial in $K[X]$. If $e^F_{\boldsymbol{a}} : T\backslash \mathrm{S}_n \to K(\boldsymbol{a})$ is injective, then $\Gamma^F_{\boldsymbol{a},D}(X)$ is irreducible, and hence the minimal polynomial of $e^F_{\boldsymbol{a}}(\sigma_0)$.

Remark 2.9. The point is that the $\Gamma^F_{\boldsymbol{a},D}(X)$ are K-rational objects, and they can be used to compare different splitting fields:

Proposition 2.10. *Let $\boldsymbol{a}, \boldsymbol{b}$ be orderings of roots of f in two splitting fields of f, and let $\phi : K(\boldsymbol{a}) \to K(\boldsymbol{b})$ be an isomorphism. If $e^F_{\boldsymbol{a}} : T \setminus \mathrm{S}_n \to K(\boldsymbol{a})$ is injective, then for every double coset $D \in T\backslash \mathrm{S}_n / G_{\boldsymbol{a}}$,*

$$\Gamma^F_{\boldsymbol{a},D}(F(\boldsymbol{b})) = 0 \quad \Longleftrightarrow \quad \boldsymbol{b} = [\phi(a_1), \ldots, \phi(a_n)]^\sigma \quad \text{for some } \sigma \in D.$$

Proof. We have that $\Gamma^F_{\boldsymbol{a},D}(F(\boldsymbol{b})) = 0$ if and only if $F(\boldsymbol{b}) = \phi(x)$ for some root x of $\Gamma^F_{\boldsymbol{a},D}$ in $K(\boldsymbol{a})$. Such roots are $e^F_{\boldsymbol{a}}(\sigma)$ for some $\sigma \in D$, so

$$\begin{aligned}
\Gamma^F_{\boldsymbol{a},D}(F(\boldsymbol{b})) = 0 \quad &\Longleftrightarrow \quad F(\boldsymbol{b}) = \phi(e^F_{\boldsymbol{a}}(\sigma)) && \text{for some } \sigma \in D \\
&\Longleftrightarrow \quad F(\phi^{-1}(\boldsymbol{b})) = e^F_{\boldsymbol{a}}(\sigma) = F(\boldsymbol{a}^\sigma) && \\
&\Longleftrightarrow \quad \phi^{-1}(\boldsymbol{b}) = (\boldsymbol{a}^\sigma)^\tau = \boldsymbol{a}^{\tau\sigma} && \text{for some } \tau \in T \\
&\Longleftrightarrow \quad \boldsymbol{b} = \phi(\boldsymbol{a}^{\sigma'}) = \phi(\boldsymbol{a})^{\sigma'} && \text{for some } \sigma' \in D. \quad \square
\end{aligned}$$

Theorem 2.11. *Let F be a $G_{\boldsymbol{a}}$-invariant with $e^F_{\boldsymbol{a}} : G_{\boldsymbol{a}}\backslash \mathrm{S}_n \to K(\boldsymbol{a})$ injective. If $F(\boldsymbol{b}) = F(\boldsymbol{a}) \in K$, then $a_i \mapsto b_i$ defines an isomorphism $K(\boldsymbol{a}) \to K(\boldsymbol{b})$.*

Proof. Take $T = G_{\boldsymbol{a}}$ and D the principal double coset $G_{\boldsymbol{a}} 1 G_{\boldsymbol{a}}$, and apply the proposition. Since $\Gamma^F_{\boldsymbol{a},D}(X) = X - F(\boldsymbol{a})$, we have $\Gamma^F_{\boldsymbol{a},D}(F(\boldsymbol{b})) = 0$, so $\boldsymbol{b} = \phi(\boldsymbol{a})^\sigma$ for some $\sigma \in G_{\boldsymbol{a}}$ and some isomorphism $\phi : K(\boldsymbol{a}) \to K(\boldsymbol{b})$. Then $\phi \circ \sigma$ is the required isomorphism. $\square$

Remark 2.12. This gives a solution to Problem 2.1:

Pick a $G_{\boldsymbol{a}}$-invariant F, for instance using Remark 2.4. Adjusting f if necessary, we may assume that $e^F_{\boldsymbol{a}} : T \setminus \mathrm{S}_n \to K(\boldsymbol{a})$ is injective (Lemma 8.1c). In L', keep permuting the roots of f until $F(\boldsymbol{b})$ becomes $F(\boldsymbol{a}) \in K$. When this happens, $a_i \mapsto b_i$ defines an isomorphism of the two splitting fields.

Note however, that in the worst case we are evaluating a polynomial with $|G|$ terms on $|G\backslash \mathrm{S}_n / G|$ permutations. So the complexity is about $n!$ operations, which is impractical for large n.

Example 2.13 (D_{10}-extensions). Suppose $f(x) \in K[x]$ has degree 5, and $G_{\boldsymbol{a}} = \mathrm{Gal}(f/K)$ is the dihedral group D_{10}, generated by (12345) and (25)(34). Take

$$F(x_1, \ldots, x_5) = x_1x_2 + x_2x_3 + x_3x_4 + x_4x_5 + x_5x_1.$$

This is a T-invariant with $T = G_{\boldsymbol{a}}$: It is clearly invariant under D_{10}, and on the other hand a permutation preserving F is determined by $x_1 \mapsto x_i,\ x_2 \mapsto x_{i\pm1}$, so there are at most 10 choices. In particular, $F(a_1, \ldots, a_5)$ is invariant under the Galois group, and so lies in K. Substituting the a_i into F in all possible orders gives the values

$$e_{\boldsymbol{a}}^F(\sigma^{-1}) = a_{\sigma(1)}a_{\sigma(2)} + a_{\sigma(2)}a_{\sigma(3)} + a_{\sigma(3)}a_{\sigma(4)} + a_{\sigma(4)}a_{\sigma(5)} + a_{\sigma(5)}a_{\sigma(1)}.$$

Clearly each one occurs at least 10 times for varying $\sigma \in \mathrm{S}_5$, corresponding to the fact that $e_{\boldsymbol{a}}^F$ factors through $\mathrm{D}_{10}\backslash\mathrm{S}_5$. The assumption that the map $e_{\boldsymbol{a}}^F : T \setminus \mathrm{S}_n \to K(\boldsymbol{a})$ is injective simply says that there are no more repetitions, and there are $120/10 = 12$ distinct values.

Suppose that this is indeed the case, and let $b_1, \ldots, b_5$ be the roots of f in some other splitting field. If we substitute the b_i in F in all possible orders $\boldsymbol{b}^\sigma$, we get again 12 values, one of which is $F(a_1, \ldots, a_5) \in K$. There are 10 isomorphisms $K(\boldsymbol{a}) \to K(\boldsymbol{b})$ obtained from one another by composing with Galois. They are determined by $\boldsymbol{a} \mapsto \boldsymbol{b}^\sigma$ for 10 permutations $\sigma \in \mathrm{S}_n$. Clearly, for each of these σ, we have $F(\boldsymbol{b}^\sigma) = F(\boldsymbol{a})$. But, since every value is taken exactly 10 times, we have the converse as well: if $F(\boldsymbol{b}^\sigma) = F(\boldsymbol{a})$ for some $\sigma \in \mathrm{S}_n$, then $\boldsymbol{a} \mapsto \boldsymbol{b}^\sigma$ must define an isomorphism of the splitting fields. So to find an isomorphism, we only need to locate $F(\boldsymbol{a})$ among the 12 values $F(\boldsymbol{b}^\sigma)$.

Note that the other values $F(\boldsymbol{b}^\sigma)$ are not in general K-rational, so we cannot compare them with the values on $\boldsymbol{a}$. Their minimal polynomials are the $\Gamma_{\boldsymbol{a},D}^F(X)$ for the 4 double cosets $\mathrm{D}_{10}\backslash\mathrm{S}_5/\mathrm{D}_{10}$.

3. Recognising conjugacy in Galois groups

In questions such as computing Frobenius elements in Galois groups it is not necessary to compare the roots in two splitting fields. It suffices to identify the conjugacy class of a specific Galois automorphism:

Problem 3.1. Let $f(x) \in K[x]$ be a separable polynomial that splits completely in $L \supset K$, and suppose we know $G = \mathrm{Gal}(f/K)$ as a permutation group on the roots in L. If L' is another field where f splits completely and we are given a permutation of the roots of f in L' that comes from some Galois automorphism, find the conjugacy class of this automorphism in G.

Remark 3.2. An isomorphism ϕ of the two splitting fields of f induces an isomorphism of Galois groups G and G'. We would like to identify an element $\mathfrak{B} \in G'$ as an element $\mathfrak{A} \in G$. Note, however, that $\mathfrak{A}$ depends on the choice of ϕ. As any two isomorphisms differ by a Galois automorphism, the conjugacy class $[\mathfrak{A}]$ is well-defined and this is what we are after.

It is easy to see that a solution to Problem 2.1 answers Problem 3.1 as well, so this is a weaker question. However, we aim for a more practical solution (see Remark 2.12). We may clearly restrict our attention to one cycle type in S_n. For convenience, throughout the section we also fix a representative:

Notation 3.3. Fix an element $\xi \in \mathrm{S}_n$ and write $Z_\xi < \mathrm{S}_n$ for its centraliser.

Definition 3.4. Suppose $\Psi \in \mathrm{S}_n$ is conjugate to ξ, in other words they have the same cycle type, say $\xi = \sigma_0 \Psi \sigma_0^{-1}$. For a T-invariant F and an ordering $\boldsymbol{a}$ of the roots of f, define the polynomial

$$M^F_{\boldsymbol{a},\Psi}(X) = \prod_{\sigma \in (Z_\xi \cap T)\backslash Z_\xi \sigma_0} \Gamma^F_{\boldsymbol{a},\sigma}(X).$$

It is well-defined by Lemma 2.5(3). Note that $Z_\xi \sigma_0$ is the set of all permutations that conjugate Ψ to ξ; in particular it is independent of the choice of σ_0.

Remark 3.5. The situation we have in mind is that we have two sets of roots $\boldsymbol{a}$ and $\boldsymbol{b}$ of f in different splitting fields. So there is an isomorphism $\phi : K(\boldsymbol{a}) \to K(\boldsymbol{b})$, but we do not have it explicitly. However, suppose we know that an automorphism $\mathscr{A} \in \mathrm{Gal}(K(\boldsymbol{a})/K)$ corresponds to $\mathscr{B} \in \mathrm{Gal}(K(\boldsymbol{b})/K)$ under ϕ, and that they permute the roots by

$$\mathscr{A}(\boldsymbol{a}) = \boldsymbol{a}^\Psi, \quad \mathscr{B}(\boldsymbol{b}) = \boldsymbol{b}^\xi, \quad \Psi, \xi \in \mathrm{S}_n.$$

Then $\{\boldsymbol{a}^\sigma\}_{\sigma \in Z_\xi \sigma_0}$ is the set of all reorderings of $\boldsymbol{a}$ on which $\mathscr{A}$ acts as ξ, and $M^F_{\boldsymbol{a},\Psi}(X)$ is the smallest K-rational polynomial that has $F(\boldsymbol{a}^\sigma)$ as roots for all such σ. But $\phi^{-1}(\boldsymbol{b})$ must be one of these reorderings because $\mathscr{B}$ acts on $\boldsymbol{b}$ as ξ. The upshot is that $M^F_{\boldsymbol{a},\Psi}(X)$ has $F(\boldsymbol{b})$ as a root, and its construction does not require the knowledge of ϕ. In other words, if $M^F_{\boldsymbol{a},\Psi}(F(\boldsymbol{b})) \neq 0$, then we know that $\mathscr{A}$ does not correspond to $\mathscr{B}$ under any isomorphism. (In Section 4 we will take $T = Z_\xi$ and turn this into an if and only if statement.)

Lemma 3.6. *Let $\phi : K(\boldsymbol{a}) \to K(\boldsymbol{b})$ be an isomorphism of two splitting fields of f, and define $\rho \in \mathrm{S}_n$ by $\boldsymbol{b} = \phi(\boldsymbol{a}^\rho)$. Then $M^F_{\boldsymbol{a},\rho^{-1}\Phi\rho} = M^F_{\boldsymbol{b},\Phi}$.*

Proof. Write $\Psi = \rho^{-1}\Phi\rho$. Pick σ_Φ with $\xi = \sigma_\Phi \Phi \sigma_\Phi^{-1}$, and let $\sigma_\Psi = \sigma_\Phi \rho$, so that

$$\sigma_\Psi \Psi \sigma_\Psi^{-1} = \sigma_\Phi \rho \Psi \rho^{-1} \sigma_\Phi^{-1} = \sigma_\Phi \Phi \sigma_\Phi^{-1} = \xi.$$

By definition,

$$M^F_{\boldsymbol{b},\Phi} = \prod_{\sigma \in (Z_\xi \cap T)\backslash Z_\xi \sigma_\Phi} \Gamma^F_{\boldsymbol{b},\sigma}, \qquad M^F_{\boldsymbol{a},\Psi} = \prod_{\sigma \in (Z_\xi \cap T)\backslash Z_\xi \sigma_\Psi} \Gamma^F_{\boldsymbol{a},\sigma}.$$

We claim that $\Gamma^F_{\boldsymbol{a},s\sigma_\Psi} = \Gamma^F_{\boldsymbol{b},s\sigma_\Phi}$ for $s \in Z_\xi$. First we show that they have the same degree. Because $G_{\boldsymbol{b}} = \rho G_{\boldsymbol{a}} \rho^{-1}$ by the definition of ρ,

$$\begin{aligned}\deg \Gamma^F_{\boldsymbol{a},s\sigma_\Psi} &= |T\backslash T s\sigma_\Psi G_{\boldsymbol{a}}| = |T\backslash T s\sigma_\Psi G_{\boldsymbol{a}} \rho^{-1}| \\ &= |T\backslash T s\sigma_\Phi \rho G_{\boldsymbol{a}} \rho^{-1}| = |T\backslash T s\sigma_\Phi G_{\boldsymbol{b}}| = \deg \Gamma^F_{\boldsymbol{b},s\sigma_\Phi}.\end{aligned}$$

Since both polynomials are powers of irreducible ones, it now suffices to identify one of the roots:

$$\begin{aligned}e^F_{\boldsymbol{a}}(s\sigma_\Psi) &= e^F_{\boldsymbol{a}}(s\sigma_\Phi \rho) = F(\boldsymbol{a}^{s\sigma_\Phi \rho})) = F(\phi^{-1}(\boldsymbol{b})^{s\sigma_\Phi})) \\ &= F(\phi^{-1}(\boldsymbol{b}^{s\sigma_\Phi})) = \phi^{-1}(F(\boldsymbol{b}^{s\sigma_\Phi})) = \phi^{-1}(e^F_{\boldsymbol{b}}(s\sigma_\Phi)).\end{aligned}$$ □

Corollary 3.7. *The map $\Psi \mapsto M^F_{\boldsymbol{a},\Psi}$ is constant on every conjugacy class of $G_{\boldsymbol{a}}$ with cycle type ξ.*

Proof. By the lemma above, $M^F_{\boldsymbol{a},\Psi} = M^F_{\boldsymbol{a},g\Psi g^{-1}}$ for $g \in G_{\boldsymbol{a}}$. □

We now have an approach to Problem 3.1:

Proposition 3.8. *Let $\boldsymbol{a}, \boldsymbol{b}$ be orderings of the roots of f in two different splitting fields, and suppose $\Psi \in G_{\boldsymbol{a}}$ and $\Phi \in G_{\boldsymbol{b}}$ have cycle type ξ. If the polynomials $M^F_{\boldsymbol{a},\psi}$ are distinct for ψ in different conjugacy classes of $G_{\boldsymbol{a}}$ of cycle type ξ, then*

$$\begin{array}{l}\text{there is an isomorphism } K(\boldsymbol{a}) \to K(\boldsymbol{b}) \\ \text{under which } \Psi \text{ corresponds to } \Phi\end{array} \quad \Longleftrightarrow \quad M^F_{\boldsymbol{a},\Psi} = M^F_{\boldsymbol{b},\Phi}.$$

If, moreover, the $M^F_{\boldsymbol{a},\psi}$ are pairwise coprime, then this occurs precisely when $M^F_{\boldsymbol{a},\Psi}(F(\boldsymbol{b}^\sigma)) = 0$ for some (any) $\sigma \in \mathrm{S}_n$ with $\xi = \sigma\Phi\sigma^{-1}$.

Proof. "$\Longrightarrow$" is Lemma 3.6. For "$\Longleftarrow$", pick any isomorphism $\phi : K(\boldsymbol{a}) \to K(\boldsymbol{b})$. The polynomial $M^F_{\boldsymbol{b},\Phi}$ agrees with some $M^F_{\boldsymbol{a},\psi}$ by the lemma, and Ψ lies in the conjugacy class of ψ by assumption. Composing ϕ with an automorphism of $K(\boldsymbol{a})/K$ (which corresponds to conjugating ψ) we obtain the required isomorphism. □

Example 3.9 (Serre's trick [Buhler 1978; Roberts 2004]). Suppose $\mathrm{char}\, K \neq 2$, $f \in K[x]$ has degree n, and $G_{\boldsymbol{a}} = \mathrm{Gal}(f/K)$ is the alternating group A_n. There is a particularly nice T-invariant with $T = \mathrm{A}_n$, a "square root of the discriminant"

$$F(x_1, \ldots, x_n) = \prod_{i<j}(x_i - x_j).$$

The only double cosets $T\sigma G_{\boldsymbol{a}}$ in S_n are $D = \mathrm{A}_n$ and its complement D' in S_n. Clearly $\Gamma^F_{\boldsymbol{a},D}(X) = X - F(\boldsymbol{a})$ and $\Gamma^F_{\boldsymbol{a},D'}(X) = X + F(\boldsymbol{a})$, and $F(\boldsymbol{a})^2 = \Delta_f$ is the discriminant of f. So if $\boldsymbol{b}$ is the list of roots of f in some other splitting field, we find that

$$\begin{array}{l}a_i \mapsto b_i \text{ defines an isomorphism} \\ K(\boldsymbol{a}) \to K(\boldsymbol{b})\end{array} \quad \Longleftrightarrow \quad \prod_{i<j}(a_i - a_j) = \prod_{i<j}(b_i - b_j).$$

This illustrates Theorem 2.11 in the case of A_n. To explain Proposition 3.8 in this setting, suppose $\xi \in \mathrm{S}_n$ is a product of cycles of distinct odd degrees, so that there are two conjugacy classes $[\Psi_1]$, $[\Psi_2]$ in $G_{\boldsymbol{a}} = \mathrm{A}_n$ of cycle type ξ (for example 5-cycles in A_5). Say $\sigma_1\Psi_1\sigma_1^{-1} = \xi = \sigma_2\Psi_2\sigma_2^{-1}$ with $\sigma_1 \in \mathrm{A}_n$ and $\sigma_2 \notin \mathrm{A}_n$. In this case $Z_\xi \subset \mathrm{A}_n = T$, so

$$M^F_{\boldsymbol{a},\Psi_1}(X) = \Gamma^F_{\boldsymbol{a},\sigma_1}(X) = \Gamma^F_{\boldsymbol{a},D}(X) = X - F(\boldsymbol{a}),$$
$$M^F_{\boldsymbol{a},\Psi_2}(X) = \Gamma^F_{\boldsymbol{a},\sigma_2}(X) = \Gamma^F_{\boldsymbol{a},D'}(X) = X + F(\boldsymbol{a}).$$

Suppose again that $\boldsymbol{b}$ is the list of roots of f in some other splitting field, and $\mathcal{B} \in \mathrm{Gal}(K(\boldsymbol{b})/K)$ is an automorphism of cycle type ξ. Rearranging the b_i if necessary, assume that $\mathcal{B}$ acts on the b_i as ξ, that is, $\mathcal{B}(\boldsymbol{b}) = \boldsymbol{b}^\xi$. The statement of the proposition is that

$$\mathcal{B} \text{ comes from } [\Psi_1] \text{ under an isomorphism } K(\boldsymbol{a}) \to K(\boldsymbol{b}) \quad \Longleftrightarrow \quad \prod_{i<j}(a_i - a_j) = \prod_{i<j}(b_i - b_j),$$

which is precisely Serre's trick. The same invariant F may sometimes be used in other subgroups of S_n to distinguish between the conjugacy classes of such cycle types. (It determines whether the two classes are conjugate in A_n or not.)

4. The directed edges invariant

As before, suppose $f(x) \in K[x]$ is separable and $\boldsymbol{a} = [a_1, \dots, a_n]$ are its (ordered) roots in a splitting field. We apply the results of Section 3 when $T = Z_\xi$, the centraliser of ξ. This is particularly nice for two reasons: First, the polynomials $M^F_{\boldsymbol{a},\psi}$ of Proposition 3.8 are irreducible and distinct, and second, it is easy to write down a T-invariant with just n terms and of degree 3 (compare the polynomials in Remark 2.4 and Example 4.2).

Proposition 4.1. *Let $\xi \in \mathrm{S}_n$ with centraliser Z_ξ. Suppose that F is a Z_ξ-invariant such that $e^F_{\boldsymbol{a}} : Z_\xi \setminus \mathrm{S}_n \to K(\boldsymbol{a})$ is injective. Let Ψ, $\Psi' \in G_{\boldsymbol{a}}$ be two elements of cycle type ξ. Then*

(1) *$M^F_{\boldsymbol{a},\Psi}$ is irreducible, and equals $\Gamma^F_{\boldsymbol{a},\sigma}$ for any $\sigma \in \mathrm{S}_n$ with $\xi = \sigma\Psi\sigma^{-1}$.*

(2) *$M^F_{\boldsymbol{a},\Psi}$ has degree $|[\Psi]|$.*

(3) *$M^F_{\boldsymbol{a},\Psi} = M^F_{\boldsymbol{a},\Psi'}$ if and only if Ψ and Ψ' are conjugate in $G_{\boldsymbol{a}}$.*

Proof. For brevity, write $Z = Z_\xi$. Pick $\sigma, \sigma' \in \mathrm{S}_n$ with $\sigma\Psi\sigma^{-1} = \xi = \sigma'\Psi(\sigma')^{-1}$.

(1) By definition,

$$M^F_{\boldsymbol{a},\Psi} = \prod_{\tau \in (Z\cap Z)\setminus Z\sigma} \Gamma^F_{\boldsymbol{a},\tau} = \Gamma^F_{\boldsymbol{a},\sigma}.$$

It is irreducible by the assumed injectivity of $e^F_{\boldsymbol{a}}$; see Remark 2.8.

(2) By definition,

$$\deg \Gamma^F_{\boldsymbol{a},\sigma} = |Z\backslash Z\sigma G_{\boldsymbol{a}}| = \frac{|Z\sigma G_{\boldsymbol{a}}|}{|Z|} = \frac{|\sigma^{-1}Z\sigma G_{\boldsymbol{a}}|}{|Z|}$$
$$= \frac{|G_{\boldsymbol{a}}|}{|G_{\boldsymbol{a}} \cap \sigma^{-1}Z\sigma|} = \frac{|G_{\boldsymbol{a}}|}{|\operatorname{Cent}_{G_{\boldsymbol{a}}}(\Psi)|} = |[\Psi]|.$$

(3) If Ψ and Ψ' are conjugate, then $M^F_{\boldsymbol{a},\Psi} = M^F_{\boldsymbol{a},\Psi'}$ by Corollary 3.7. Conversely, suppose that $M^F_{\boldsymbol{a},\Psi} = M^F_{\boldsymbol{a},\Psi'}$. Since $e^F_{\boldsymbol{a}}$ is injective, $Z\sigma G_{\boldsymbol{a}} = Z\sigma' G_{\boldsymbol{a}}$, so $\sigma' = s\sigma g$ for some $s \in Z$ and $g \in G_{\boldsymbol{a}}$. Then

$$\Psi' = (\sigma')^{-1}\xi\sigma' = g^{-1}\sigma^{-1}s^{-1}\xi s\sigma g = g^{-1}\sigma^{-1}\xi\sigma g = g^{-1}\Psi g,$$

so $[\Psi'] = [\Psi]$. □

Example 4.2 (the directed edges invariant). Let $\xi \in \mathrm{S}_n$ and fix a polynomial $h \in K[x]$ of degree at least 2. Define

$$F(x_1, \ldots, x_n) = \sum_{j=1}^{n} h(x_j)\, x_{\xi(j)}.$$

It can be visualised as the directed edges in a graph that define the action by ξ. For instance, for $\xi = (1234)(56) \in \mathrm{S}_6$ and $h(x) = x^2$,

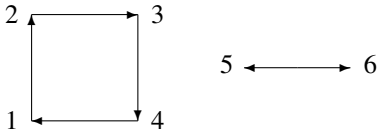

$$F = x_1^2x_2 + x_2^2x_3 + x_3^2x_4 + x_4^2x_1 + x_5^2x_6 + x_6^2x_5$$

It is clearly a Z_ξ-invariant.

Definition 4.3. Fix $h(x) \in K[x]$. For each conjugacy class C in $G_{\boldsymbol{a}}$ define

$$\Gamma_C(X) = \prod_{\sigma \in C} \Big(X - \sum_{j=1}^{n} h(a_j)\sigma(a_j) \Big).$$

Lemma 4.4. *Let F be as in Example 4.2. Then for every $\Psi \in G_{\boldsymbol{a}}$,*

$$M^F_{\boldsymbol{a},\Psi}(X) = \Gamma_{[\Psi]}(X).$$

Proof. Pick $\sigma \in \mathrm{S}_n$ with $\sigma\Psi\sigma^{-1} = \xi$. First, suppose $\tau \in [\Psi]$ and $u_\tau \in \mathrm{S}_n$ satisfies $u_\tau^{-1}\xi u_\tau = \tau$. Then

$$e^F_{\boldsymbol{a}}(u_\tau) = F(\boldsymbol{a}^{u_\tau}) = \sum_i h(a_{u_\tau^{-1}(i)})\, a_{u_\tau^{-1}(\xi(i))} = \sum_j h(a_j)\, a_{u_\tau^{-1}\xi u_\tau(j)} = \sum_j h(a_j)\tau(a_j).$$

On the other hand, note that for $t \in Z_\xi$ and $g \in G_{\boldsymbol{a}}$,

$$(t\sigma g)^{-1}\xi(t\sigma g) = g^{-1}\sigma^{-1}t^{-1}\xi t\sigma g = g^{-1}\sigma^{-1}\xi\sigma g = g^{-1}\Psi g.$$

So for $\tau = g^{-1}\Psi g \in [\Psi]$,

$$\{u_\tau \in S_n | u_\tau^{-1}\xi u_\tau = \tau\} = Z_\xi \sigma g,$$

because the left-hand side is clearly some right coset of Z_ξ. This equality gives a correspondence between $[\Psi]$ and $Z_\xi \backslash Z_\xi \sigma G_{\boldsymbol{a}}$. So

$$\begin{aligned} M^F_{\boldsymbol{a},\Psi}(X) = \Gamma^F_{\boldsymbol{a},\sigma}(X) &= \prod_{u \in (Z_\xi \backslash Z_\xi \sigma G_{\boldsymbol{a}})} (X - e^F_{\boldsymbol{a}}(u)) \\ &= \prod_{\tau \in [\Psi]} \Big(X - \sum_{j=1}^{n} h(a)\tau(a_j)\Big) = \Gamma_{[\Psi]}(X). \end{aligned}$$

□

Corollary 4.5. *Let $\boldsymbol{a}, \boldsymbol{b}$ be orderings of the roots of f in two different splitting fields, and let $\Psi \in G_{\boldsymbol{a}}$ and $\Phi \in G_{\boldsymbol{b}}$. If the $\Gamma_C(X)$ are pairwise coprime for different conjugacy classes of $G_{\boldsymbol{a}}$, then*

there is an isomorphism $K(\boldsymbol{a}) \to K(\boldsymbol{b})$ under which Ψ corresponds to Φ, $\iff$ $\Gamma_{[\Psi]}\big(\sum_j h(b_j)\Phi(b_j)\big) = 0.$

The condition that the Γ_C are coprime is satisfied for $h(x)$ in a Zariski dense open set in the space of all polynomials of degree at most $n-1$.

Proof. The equivalence follows from Proposition 3.8 and the lemma above. For the last assertion apply Lemma 8.2. □

Example 4.6. Take $f(x) = x^4 + 14$ over $\mathbb{Q}$. It splits completely over $L = \mathbb{Q}_5$ and $L' = \mathbb{C}$, with roots in $\mathbb{Q}_5$

$$a_1 = 1+3\cdot 5+2\cdot 5^2+\cdots, \quad a_2 = 2+2\cdot 5+0\cdot 5^2+\cdots, \quad a_3 = -a_1, \quad a_4 = -a_2,$$

and

$$b_1 = \sqrt[4]{-14}, \quad b_2 = i\sqrt[4]{-14}, \quad b_3 = -\sqrt[4]{-14}, \quad b_4 = -i\sqrt[4]{-14}$$

in $\mathbb{C}$ (with, say, $\operatorname{Arg} b_1 = \pi/4$). The Galois group of f is $G = \mathrm{D}_8$, which we view as a subgroup of S_4 via the action on the a_i. It is generated by the 4-cycle $a_1 \mapsto a_2 \mapsto a_3 \mapsto a_4 \mapsto a_1$ and the transposition $a_1 \leftrightarrow a_3$. We will illustrate how to identify the conjugacy class of complex conjugation $b_1 \leftrightarrow b_4,\ b_2 \leftrightarrow b_3$ in G, using the polynomials $\Gamma_C(x)$.

There are two conjugacy classes of double transpositions in G, namely $C_1 = \{(12)(34), (14)(23)\}$ and $C_2 = \{(13)(24)\}$. Let $h(x) = x$ and compute

$$\begin{aligned} \Gamma_{C_1}(X) &= (X - (2a_1a_2 + 2a_3a_4))(X - (2a_1a_4 + 2a_2a_3)) = X^2 - 224, \\ \Gamma_{C_2}(X) &= X - (2a_1a_3 + 2a_2a_4) = X. \end{aligned}$$

These are coprime, and Corollary 4.5 applies:

$$\sum_{j=1}^{4} b_j \bar{b}_j = 2b_1 b_4 + 2b_2 b_3 = \sqrt{224}$$

is a root of $\Gamma_{C_1}(X)$, so complex conjugation corresponds to an element of C_1.

Note that the coefficients of the $\Gamma_C(X)$ were computed as 5-adic numbers. Since they are integers and we can bound them from the (complex) absolute values of the roots of f, they can be identified exactly.

5. Frobenius elements

Now suppose K is a global field. We turn to our initial problem of computing Frobenius elements in Galois groups. We use the following remarkable property of the directed edges invariant:

Proposition 5.1. *Let $f(x) \in \mathbb{F}_q[x]$ be a polynomial with roots $a_1, \ldots, a_n \in \bar{\mathbb{F}}_q$ counted with multiplicity, and let $\phi = \mathrm{Frob}_q \in \mathrm{Gal}(\bar{\mathbb{F}}_q/\mathbb{F}_q)$. For every polynomial $h(x) \in \mathbb{F}_q[x]$,*

$$\sum_{j=1}^{n} h(a_j)\phi(a_j) = \mathrm{Tr}_{A/\mathbb{F}_q}(h(X)X^q),$$

where X is the class of x in the algebra $A = \mathbb{F}_q[x]/f$.

This is an immediate consequence of the lemma below (with $H(x) = h(x)x^q$).

Lemma 5.2. *Let k be a field and $f(x) \in k[x]$ a polynomial with roots $a_1, \ldots, a_n \in \bar{k}$ counted with multiplicity. Then for every $H(x) \in k[x]$,*

$$\sum_{j=1}^{n} H(a_j) = \mathrm{Tr}_{A/k}(H(X)),$$

where X is the class of x in $A = k[x]/f$.

Proof. Consider X as a linear map $A \to A$, $Y \mapsto XY$. Its minimal polynomial is f, since $f(X) = 0$ but no linear combination of $1, X, \ldots, X^{n-1}$ is zero. So the generalised eigenvalues of X are exactly the a_i, and those of $H(X)$ are therefore $H(a_i)$ (look at the Jordan normal form of X over $\bar{k}$). The result follows. □

Theorem 5.3 (generalised Euler's criterion). *Let K be a global field and let $f(x) \in K[x]$ be a separable polynomial with roots $a_1, \ldots, a_n$ in $\bar{K}$ and Galois group G. Fix $h(x) \in K[x]$ and for each conjugacy class C of G, set*

$$\Gamma_C(X) = \prod_{\sigma \in C} \Big(X - \sum_{j=1}^{n} h(a_j)\sigma(a_j) \Big).$$

(a) *The polynomials* $\Gamma_C(X)$ *have coefficients in* K.

(b) *Let* $\mathfrak{p}$ *be a prime of* K *with residue field* $\mathbb{F}_q$, *and* C *a conjugacy class of* G. *If* $\mathfrak{p}$ *does not divide the denominators of the coefficients of* f *and* h, *the leading coefficient of* f *and the resultants* $\operatorname{Res}(\Gamma_C, \Gamma_{C'})$ *for* $C' \neq C$, *then the coefficients of* $\Gamma_C(X)$ *are integral at* $\mathfrak{p}$ *and*

$$\operatorname{Frob}_\mathfrak{p} \in C \quad \Longleftrightarrow \quad \Gamma_C\big(\operatorname{Tr}_{\frac{\mathbb{F}_q[x]}{f(x)}/\mathbb{F}_q}(h(x)x^q)\big) = 0 \bmod \mathfrak{p}.$$

(c) *For all* $h(x)$ *in some Zariski dense open set in the space of polynomials of degree at most* $n-1$, *we have* $\operatorname{Res}(\Gamma_C, \Gamma_{C'}) \neq 0$ *for every pair of conjugacy classes* $C \neq C'$.

Proof. (a) This follows from Lemma 4.4, Definition 3.4 and Remark 2.8.

(b) $\Gamma_C(X)$ is clearly integral at the required primes.

$\Longrightarrow$: If $\operatorname{Frob}_\mathfrak{p} \in C$ then $\sum_{j=1}^n h(a_j)\operatorname{Frob}_\mathfrak{p}(a_j)$ is a root of $\Gamma_C(X)$ by the definition of Γ_C, and it reduces mod $\mathfrak{p}$ to $\operatorname{Tr}_{(\mathbb{F}_q[x]/f(x))/\mathbb{F}_q}(h(x)x^q)$ by Proposition 5.1.

$\Longleftarrow$: The polynomial $\Gamma_C(X)$ is distinguished from the others by any one of its roots mod $\mathfrak{p}$ by the assumption that $\mathfrak{p} \nmid \operatorname{Res}(\Gamma_C, \Gamma_{C'})$ for $C \neq C'$.

(c) Apply Lemma 8.2. □

Remark 5.4 (choice of h). If the resultants $\operatorname{Res}(\Gamma_C, \Gamma_{C'})$ are nonzero, then Theorem 5.3(b) describes the Frobenius element for all but finitely many primes $\mathfrak{p}$. If one of the resultants vanishes, or equivalently, Γ_C has a common factor with some $\Gamma_{C'}$, the statement does not apply to C for any $\mathfrak{p}$. However, this is rare and easily avoided by choosing a different h; most choices will work by Theorem 5.3(c).

Alternatively, for any fixed h with $1 < \deg h < n$ it is possible to replace f by another polynomial $\tilde{f}$ of degree n with the same splitting field so that the resulting Γ_C are coprime. To see this, consider

$$\gamma_C(X) = \prod_{\sigma \in C}\Big(X - \sum_{j=1}^n h(x_j)x_{\sigma(j)}\Big),$$

and note that they are coprime as polynomials in X over $K(x_1, \ldots, x_n)$. Now apply Lemma 8.1(b) to $F_1 = \prod_{C \neq C'} \operatorname{Res}(\gamma_C, \gamma_{C'})$ and $F_2 = 0$. We obtain a Zariski dense open set of polynomials $B(t)$ of degree at most $n-1$ for which $\tilde{f} = \prod_j (x - B(a_j))$ works.

Remark 5.5 (Euler's criterion). The classical criterion

$$a^{(p-1)/2} \equiv \left(\frac{a}{p}\right) \bmod p$$

says that $a^{(p-1)/2} = \pm 1$ determines whether $x^2 - a$ has a root modulo p. Similarly, to see whether $x^3 - a$ has a root modulo $p \equiv 1 \bmod 3$ one checks whether $a^{(p-1)/3}$ is 1 or another third root of unity in $\mathbb{F}_p^\times$, etc.

One can reformulate this as a matrix statement: Take a 2×2 matrix M with minimal polynomial $x^2 - a$ (respectively 3×3 and $x^3 - a$). Then M^{p-1} is the scalar matrix with $a^{(p-1)/2}$ or $a^{(p-1)/3}$, respectively, on the diagonal, so its trace determines whether the polynomial has a root in $\mathbb{F}_p$; for example, for $x^3 - a$ the distinction is whether $\frac{1}{3} \operatorname{Tr} M^{p-1}$ is 1 or a root of $x^2 + x + 1$.

Theorem 5.3 generalises this to arbitrary polynomials over global fields. Observe that for a polynomial

$$f(x) = x^n + c_{n-1}x^{n-1} + \cdots + c_0,$$

the trace in the theorem can be interpreted as a trace of a matrix, for instance,

$$\operatorname{Tr}_{\frac{\mathbb{F}_q[x]}{f(x)}/\mathbb{F}_q}(x^d) = \operatorname{Tr} \begin{pmatrix} & & -c_0 \\ 1 & & -c_1 \\ & \ddots & \vdots \\ & 1 & -c_{n-1} \end{pmatrix}^d \bmod q.$$

Therefore (a minor modification of) the trace $\operatorname{Tr} M^{q-1}$ for a matrix M with minimal polynomial f determines the splitting behaviour of $f \bmod \mathfrak{p}$ and the conjugacy class of Frobenius, in the same way as above. See also Sections 1 and 7.

Remark 5.6 (ramified primes). The condition that $\mathfrak{p}$ does not divide any resultant $\operatorname{Res}(\Gamma_C, \Gamma_{C'})$ excludes all primes that ramify in the splitting field of f over K. Indeed, if $\sigma \neq 1$ is an element of inertia at $\mathfrak{q}$ for some $\mathfrak{q} \mid \mathfrak{p}$, it is easy to see that $\Gamma_{[1]}$ and $\Gamma_{[\sigma]}$ have a common root mod $\mathfrak{p}$.

Remark 5.7 (extending to all $\mathfrak{p}$). In order to deal with the primes dividing the resultants, we may work over the completion $K_\mathfrak{p}$ instead of the residue field $\mathbb{F}_q$. Compute the splitting field $L/K_\mathfrak{p}$ of f and the roots $b_1, \ldots, b_n$. Choose a lift Ψ of the Frobenius element in $\operatorname{Gal}(L/K_\mathfrak{p})$ and evaluate

$$\sum_{j=1}^{n} h(b_j)\Psi(b_j).$$

This number is now a root of precisely one of the Γ_C, and this C is the conjugacy class of the chosen Frobenius lift Ψ. (See Corollary 4.5.)

Remark 5.8 (Artin L-functions). Suppose L/K is a Galois extension of number fields with Galois group G, represented as a splitting field of some polynomial $f(x) \in K[x]$. Recall that a complex representation ρ of G is called an *Artin*

representation. It has an L-series defined by the Euler product over all primes of K,

$$L(\rho, s) = \prod_{\mathfrak{p}} \frac{1}{P_{\mathfrak{p}}(q^{-s})}\,.$$

Here q is the size of the residue field at $\mathfrak{p}$ and $P_{\mathfrak{p}}(T) = \det(1 - \mathrm{Frob}_{\mathfrak{p}}\, T \mid \rho^{I_{\mathfrak{p}}})$ is the inverse characteristic polynomial of Frobenius on the subspace of ρ fixed by the inertia group $I_{\mathfrak{p}}$ at $\mathfrak{p}$.

Theorem 5.3 and Remark 5.7 allow us to explicitly compute the coefficients of such L-series. For the unramified primes, they recover the conjugacy class of $\mathrm{Frob}_{\mathfrak{p}}$ in G, which determines the local polynomial $P_{\mathfrak{p}}(T)$. For the ramified primes, it suffices to find the restriction of ρ to the local Galois group $G_{\mathfrak{p}}$ at $\mathfrak{p}$ with respect to an embedding $G_{\mathfrak{p}} \hookrightarrow G$ as a decomposition group. Assuming we can find $G_{\mathfrak{p}}$, Remark 5.7 enables us to identify the conjugacy class in G of any element of $G_{\mathfrak{p}}$, under this embedding. This is sufficient to compute the character of ρ on $G_{\mathfrak{p}}$, and thus also $\rho^{I_{\mathfrak{p}}}$ and $P_{\mathfrak{p}}(T)$. Note that we have *not* actually found the decomposition group at $\mathfrak{p}$ as a *subgroup* of G, which appears to be a harder problem.

This algorithm to compute Frobenius elements and L-series of Artin representations has now been implemented in Magma [Bosma et al. 1997]. For the functional equation of the L-series one also needs to identify the conjugacy class of the complex conjugation. If G is represented as acting on the roots of f in a p-adic field, this can be done with the same method. (See Corollary 4.5 and Example 4.6.)

Remark 5.9 (complexity). From the complexity point of view, the computation of Frobenius elements for "good" primes has two steps:

One is the initial precomputation of the polynomials Γ_C, each of which takes $O(n|C|)$ operations in some field containing the a_j (for instance $\mathbb{C}$ or $\overline{\mathbb{Q}}_p$). This needs to be done for all conjugacy classes that are not determined by their cycle type.

The second step deals with a specific prime $\mathfrak{p}$ of K with residue field $\mathbb{F}_q$. We determine the cycle type of $\mathrm{Frob}_{\mathfrak{p}}$ by computing $\gcd(f, x^{q^j} - x)$ for $j \le n/2$, which takes $O(n \log q)$ multiplications of $n \times n$ matrices over $\mathbb{F}_q$. Then we evaluate the trace $\mathrm{Tr}(h(x)x^q)$ with another $O(n + \log q)$ matrix multiplications. Finally, we substitute the trace into all Γ_C corresponding to the cycle type of $\mathrm{Frob}_{\mathfrak{p}}$, which is $O(d)$ coefficient reductions and multiplications in $\mathbb{F}_q$, where d is the number of elements in G of this cycle type.

Here is as an illustration for polynomials of degree at most 11. There are 474 transitive groups G on at most 11 points, for each of which we took a polynomial $f \in \mathbb{Q}[x]$ with $\mathrm{Gal}\, f = G$ as a permutation group on the roots. (We used the database in Magma [Bosma et al. 1997, V2.16].) For each G we computed Frob_p for all $p < 100000$ with $p \nmid \Delta_f$, using Serre's trick (Example 3.9) and the algorithm

above. Together with the Galois group computation and the precomputation of the Γ_C this took under 15 seconds on a 3GHz dual-core CPU for each G, with only four exceptions: $G = \mathrm{A}_5^2 \rtimes \mathrm{C}_2$, $\mathrm{A}_5^2 \rtimes \mathrm{C}_2^2$, $\mathrm{A}_5^2 \rtimes \mathrm{C}_4$ and M_{11}. These took 17, 254, 1512 and 61 seconds respectively, with approximately 10–30 seconds taken by computing Frobenius elements and the rest by precomputing the $\Gamma_C(x)$. These four groups have large conjugacy classes of the same cycle type (the largest being the two classes of size 1800 for $\mathrm{A}_5^2 \rtimes \mathrm{C}_4$).

Remark 5.10 (additional symmetries). Suppose all conjugacy classes of elements of some order o and a fixed cycle type are closed under the power maps $g \mapsto g^k$ for k in some nontrivial subgroup $H \subset (\mathbb{Z}/o\mathbb{Z})^\times$ (for instance they are self-inverse, like in dihedral groups). Then one may replace $\Gamma_C(X)$ in Theorem 5.3 by

$$\prod_\sigma \Big(X - \sum_{j=1}^{n} h(a_j) \Big(\sum_{k \in H} \sigma^k(a_j) \Big) \Big),$$

taking the product over some representatives for C modulo the action of H, and modifying the trace accordingly. In practice, this speeds up the computation of the Γ_C, as their degree drops by a factor of $|H|$.

6. Examples: Abelian groups

If the Galois group is abelian, its conjugacy classes are of size 1, and all the Γ_C of Theorem 5.3 are linear, that is, $\Gamma_C(X) = X - r_C$ with $r_C \in K$. For a good choice of $h(x)$ and all but finitely many primes $\mathfrak{p}$, the trace $\mathrm{Tr}(h(x)x^q)$ agrees with exactly one of the r_C modulo $\mathfrak{p}$, which then determines the conjugacy class of $\mathrm{Frob}_\mathfrak{p}$.

In the examples below, ζ_n denotes a primitive n-th root of unity.

Example 6.1. Let $K = \mathbb{Q}(i)$ and

$$f(x) = x^4 + 2x^3 + (3+3i)x^2 + 4ix - 1 + i.$$

Its complex roots are

$$a_1 = -0.31795 - 0.57510i, \qquad a_2 = 0.50870 - 1.1289i,$$
$$a_3 = -1.4682 + 1.8471i, \qquad a_4 = -0.72255 - 0.14308i,$$

to 5 decimal places. The splitting field L is a C_4-extension of $\mathbb{Q}(i)$, non-Galois over $\mathbb{Q}$, and the Galois group of L/K is $\langle (1234) \rangle < \mathrm{S}_4$. Take $h(x) = x^2$. An elementary computation gives

$$\Gamma_{[\mathrm{id}]} = X - (10+6i), \qquad \Gamma_{[(1234)]} = X - (4+4i),$$
$$\Gamma_{[(13)(24)]} = X - (-2+2i), \qquad \Gamma_{[(1432)]} = X + 8.$$

For a prime $\mathfrak{p} \neq (1+i), (2-i), (3)$ (the primes dividing $r_C - r_{C'}$ for $C \neq C'$) with residue field $\mathbb{F}_q$, we deduce that the Frobenius at $\mathfrak{p}$ is determined by

$\mathrm{Tr}_{(\mathbb{F}_q[x]/f(x))/\mathbb{F}_q}(x^{q+2}) \equiv$	$10+6i$	$4+4i$	$-2+2i$	-8
$\mathrm{Frob}_{\mathfrak{p}} \quad =$	id	(1234)	(13)(24)	(1432)

Example 6.2 (Kummer extensions). Suppose $\zeta = \zeta_n \in K$ and $L = K(\sqrt[n]{s})$ is a Kummer extension of degree n. It is abelian with Galois group C_n whose elements are determined by

$$\sigma_i : \sqrt[n]{s} \mapsto \zeta^i \sqrt[n]{s} \quad \text{for } i = 1, \dots, n.$$

Take $f(x) = x^n - s$ and $h(x) = x^{n-1}$. Then

$$\Gamma_{[\sigma_i]}(X) = X - \sum_{j=1}^{n} h(\zeta^j \sqrt[n]{s})\sigma_i(\zeta^j \sqrt[n]{s}) = X - ns \cdot \zeta^i.$$

For a prime $\mathfrak{p}$ of K with residue field $\mathbb{F}_q$, because $n \,|\, q-1$, we have

$$\begin{aligned}\mathrm{Tr}_{(\mathbb{F}_q[x]/f(x))/\mathbb{F}_q}(h(x)x^q) &= \mathrm{Tr}_{(\mathbb{F}_q[x]/x^n-s)/\mathbb{F}_q}(x^{q+n-1}) \\ &= \mathrm{Tr}_{(\mathbb{F}_q[x]/x^n-s)/\mathbb{F}_q}(s^{(q-1)/n+1}) = ns \cdot s^{(q-1)/n}.\end{aligned}$$

So Theorem 5.3 says that for $\mathfrak{p} \nmid ns$,

$$\mathrm{Frob}_{\mathfrak{p}} = \sigma_i \quad \Longleftrightarrow \quad s^{(q-1)/n} \equiv \zeta^i \mod \mathfrak{p},$$

which is the classical criterion for Kummer extensions.

Example 6.3 ($\mathbb{Q}(\zeta_p)/\mathbb{Q}$). Let $\zeta = \zeta_p$ for some prime $p > 2$, and take

$$K = \mathbb{Q}, \quad L = \mathbb{Q}(\zeta), \quad f(x) = x^{p-1} + \cdots + x + 1.$$

Thus $\mathrm{Gal}(L/K) \cong (\mathbb{Z}/p\mathbb{Z})^\times$, with elements $\sigma_i : \zeta \mapsto \zeta^i$ for $i = 1, \dots, p-1$. For $h(x) = x^2$ we have $\Gamma_{[\sigma_i]}(X) = X - r_i$ with $r_i \in \mathbb{Q}$ given by

$$r_i = \sum_{j=1}^{p-1} (\zeta^j)^2 \sigma_i(\zeta^j) = \sum_{j=1}^{p-1} \zeta^{j(2+i)} = \begin{cases} -1 & \text{if } i \neq p-2, \\ p-1 & \text{if } i = p-2. \end{cases}$$

For a prime q of $\mathbb{Q}$,

$$\begin{aligned}\mathrm{Tr}_{(\mathbb{F}_q[x]/f(x))/\mathbb{F}_q}(h(x)x^q) &= \mathrm{Tr}_{(\mathbb{F}_q[x]/f(x))/\mathbb{F}_q}(x^{q+2}) \equiv \mathrm{Tr}_{(\mathbb{Z}[x]/f(x))/\mathbb{Z}}(x^{q+2}) \mod q \\ &\equiv \mathrm{Tr}_{F/\mathbb{Q}}(\zeta^{q+2}) \equiv \begin{cases} -1 \mod q & \text{if } p \nmid q+2, \\ p-1 \mod q & \text{if } p \,|\, q+2. \end{cases}\end{aligned}$$

Hence Theorem 5.3(b) shows that for all $q \neq p$,

$$\mathrm{Frob}_q = \sigma_{p-2} \quad \Longleftrightarrow \quad q \equiv -2 \mod p.$$

The same computation with $h(x) = x^{p-k}$ for varying k yields the classical criterion

$$\mathrm{Frob}_q = \sigma_k \quad \Longleftrightarrow \quad q \equiv k \mod p.$$

Note that none of these $h(x)$ work for all conjugacy classes simultaneously, because the $\Gamma_{[\sigma_j]}$ are not coprime. This tends to happen when the roots of f are "too nice" and $h(x)$ is "too simple". By Lemma 8.2, most h do work. In our example, a general polynomial

$$h(x) = \lambda_1 x^{p-1} + \cdots + \lambda_{p-1} x + \lambda_p \quad \text{has } \Gamma_{[\sigma_i]}(X) = X + h(1) - p\lambda_i,$$

and these are distinct if and only if $\lambda_1, \ldots, \lambda_{p-1}$ are. The primes to which the theorem then applies are those not dividing $p \prod(\lambda_i - \lambda_j)$.

Example 6.4 (cyclotomic extensions). In general, suppose $L = K(\zeta_n)$ is some cyclotomic extension, and $f(x)$ is the minimal polynomial of ζ_n over K. As in the previous example, $G = \mathrm{Gal}(L/K) \hookrightarrow (\mathbb{Z}/n\mathbb{Z})^\times$, and we write σ_i for the automorphism with $\sigma_i(\zeta_n) = \zeta_n^i$. We do the same computation as above: For $h(x) = x^k$ and $\mathfrak{p}$ a prime of K with residue field $\mathbb{F}_q$,

$$\Gamma_{[\sigma_i]}(X) = X - \sum_{g \in G} g(\zeta_n)^k \sigma_i(g(\zeta_n)) = X - \sum_{g \in G} g(\zeta_n)^{k+i} = X - \mathrm{Tr}_{L/K}(\zeta_n^{k+i}),$$
$$\mathrm{Tr}_{(\mathbb{F}_q[x]/f(x))/\mathbb{F}_q}(x^{k+q}) \equiv \mathrm{Tr}_{L/K}(\zeta_n^{k+q}) \mod \mathfrak{p}.$$

Because $\mathrm{Tr}_{L/K}(\zeta_n^j)$ is $|G|$ precisely when $n \mid j$, the polynomial $\Gamma_{[\sigma_{n-k}]}$ differs from all the other $\Gamma_{[\sigma_j]}$, and we find that

$$\mathrm{Frob}_\mathfrak{p} = \sigma_{n-k} \quad \Longleftrightarrow \quad q \equiv n - k \mod n$$

for almost all $\mathfrak{p}$. (One may improve "almost all" to "all $\mathfrak{p} \nmid n$" by taking several h.)

Remark 6.5. The fact that we obtained a simple formula for Frobenius elements for cyclotomic and Kummer extensions relied on the existence of a universal expression for the trace $\mathrm{Tr}(h(x)x^q) \bmod \mathfrak{p}$. It follows from class field theory that there are such formulas in all abelian extensions.

For instance, consider Example 6.1 of a C_4-extension of $K = \mathbb{Q}(i)$ from the point of view of class field theory. There the conductor of L/K is

$$N = (1+i)^4(2-i) = 8 - 4i,$$

and the group $(\mathbb{O}_K/N)^\times$ is $\mathrm{C}_4 \times \mathrm{C}_4 \times \mathrm{C}_2$, with generators i, 7 and $3-2i$, respectively. For a prime $\mathfrak{p} = (\alpha) \subset \mathbb{Z}[i]$ not dividing N, if $\alpha \equiv i^a 7^b (3-2i)^c \mod N$, then $\mathrm{Frob}_\mathfrak{p} = (1234)^b$.

Now compare this with the description of Frobenius in Example 6.1. Writing $\mathbb{F}_q = \mathbb{Z}[i]/\mathfrak{p}$ and Tr for $\mathrm{Tr}_{(\mathbb{F}_q[x]/f(x))/\mathbb{F}_q}$, we get 4 congruences for the traces:

$$\begin{aligned}
&\mathfrak{p}=(\alpha), \alpha \equiv i^a 7^0 (3-2i)^c \mod N &&\iff \mathrm{Tr}(x^{q+2}) \equiv 10+6i \mod \mathfrak{p},\\
&\mathfrak{p}=(\alpha), \alpha \equiv i^a 7^1 (3-2i)^c \mod N &&\iff \mathrm{Tr}(x^{q+2}) \equiv 4+4i \mod \mathfrak{p},\\
&\mathfrak{p}=(\alpha), \alpha \equiv i^a 7^2 (3-2i)^c \mod N &&\iff \mathrm{Tr}(x^{q+2}) \equiv -2+2i \mod \mathfrak{p},\\
&\mathfrak{p}=(\alpha), \alpha \equiv i^a 7^3 (3-2i)^c \mod N &&\iff \mathrm{Tr}(x^{q+2}) \equiv -8 \mod \mathfrak{p}
\end{aligned}$$

for $\mathfrak{p} \neq (1+i), (2-i), (3)$.

Note that if one had a way to prove these congruences directly, one would have a proof of Artin reciprocity in the extension L/K.

7. Examples: Nonabelian groups

We continue with examples to Theorem 5.3. When G is nonabelian, the only difference is that the Γ_C are no longer linear.

Example 7.1. Let $K = \mathbb{Q}$ and $f(x) = x^3 - 2$. It has Galois group S_3 and roots $a_1 = \sqrt[3]{2}$, $a_2 = \zeta\sqrt[3]{2}$ and $a_3 = \zeta^2\sqrt[3]{2}$, where ζ is a primitive cube root of unity. Take $h(x) = x^2/6$ (the factor $\frac{1}{6}$ is only chosen for convenience) and compute the polynomials Γ_C for the three conjugacy classes:

$$\begin{aligned}
\Gamma_{[\mathrm{id}]} &= X - \tfrac{1}{6}(a_1^2 a_1 + a_2^2 a_2 + a_3^2 a_3)\\
&= X - 1,\\
\Gamma_{[(12)]} &= (X - \tfrac{1}{6}(a_1^2 a_2 + a_2^2 a_1 + a_3^3))(X - \tfrac{1}{6}(a_1^2 a_3 + a_2^3 + a_3^2 a_1))\\
&\qquad\qquad \cdot (X - \tfrac{1}{6}(a_1^3 + a_2^2 a_3 + a_3^2 a_2))\\
&= (X - \tfrac{1}{3}(\zeta + \zeta^2 + 1))(X - \tfrac{1}{3}(\zeta^2 + 1 + \zeta))(X - \tfrac{1}{3}(1 + \zeta + \zeta^2))\\
&= X^3,\\
\Gamma_{[(123)]} &= (X - \tfrac{1}{6}(a_1^2 a_2 + a_2^2 a_3 + a_3^2 a_1))(X - \tfrac{1}{6}(a_1^2 a_3 + a_2^2 a_1 + a_3^2 a_2))\\
&= (X - \tfrac{1}{3}(\zeta + \zeta + \zeta))(X - \tfrac{1}{3}(\zeta^2 + \zeta^2 + \zeta^2)) = (X - \zeta)(X - \zeta^2)\\
&= X^2 + X + 1.
\end{aligned}$$

On the other hand, for a rational prime $q = 3m + k$ with $k = 1$ or 2,

$$\begin{aligned}
\mathrm{Tr}_{(\mathbb{F}_q[x]/x^3-2)/\mathbb{F}_q}\big(\tfrac{1}{6}x^{q+2}\big) &= \mathrm{Tr}\big(\tfrac{1}{6}2^{m+1}x^{k-1}\big)\\
&= \begin{cases} 2^m & \text{if } k = 1,\\ 0 & \text{if } k = 2 \end{cases}\\
&= \begin{cases} 2^{(q-1/3)} & \text{if } q \equiv 1 \bmod 3,\\ 0 & \text{if } q \equiv 2 \bmod 3. \end{cases}
\end{aligned}$$

The conclusion of Theorem 5.3 is that, as expected, for $q \neq 2, 3$,

$$\begin{aligned} q \equiv 1 \bmod 3, 2 \in (\mathbb{F}_q)^{\times 3} &\implies \mathrm{Frob}_q = \mathrm{id}, \\ q \equiv 1 \bmod 3, 2 \notin (\mathbb{F}_q)^{\times 3} &\implies \mathrm{Frob}_q \in [(123)], \\ q \equiv 2 \bmod 3 &\implies \mathrm{Frob}_q \in [(12)]. \end{aligned}$$

Clearly, an identical computation goes through for $f(x) = x^3 - c$ (with $h(x) = x^2/3c$) over any global field K with $\zeta \notin K$.

We can also take a general cubic polynomial and obtain an analogue of Euler's criterion for its factorisation modulo primes:

Theorem 7.2. *Let $f(x) = x^3 + bx + c$ be a separable cubic polynomial over a global field K, and $\mathfrak{p}$ a prime of K with residue field $\mathbb{F}_q$. Write*

$$T = \mathrm{Tr}_{(\mathbb{F}_q[x]/f(x))/\mathbb{F}_q}(x^{q+1}) = \mathrm{Tr} \begin{pmatrix} 0 & 0 & -c \\ 1 & 0 & -b \\ 0 & 1 & 0 \end{pmatrix}^{q+1} \bmod \mathfrak{p}.$$

If $\mathfrak{p}$ does not divide $3b(4b^3 + 27c^2)$ and the denominators of b and c, then

$$\begin{aligned} T \equiv -2b \bmod \mathfrak{p} &\iff f(x) \text{ has 3 roots mod } \mathfrak{p}, \\ T \equiv \quad b \bmod \mathfrak{p} &\iff f(x) \text{ is irreducible mod } \mathfrak{p}, \\ T \text{ is a root of } x^3 - 3b^2x - 2b^3 - 27c^2 &\iff f(x) \text{ has 1 root mod } \mathfrak{p}. \end{aligned}$$

Proof. We compute the polynomials Γ_C for $G = \mathrm{S}_3$, $h(x) = x$ by expressing their coefficients in terms the elementary symmetric functions $a_1 + a_2 + a_3 = 0$, $a_1a_2 + a_2a_3 + a_3a_1 = b$ and $a_1a_2a_3 = -c$:

$$\begin{aligned} \Gamma_{[\mathrm{id}]} &= X - (a_1^2 + a_2^2 + a_3^2) = X - (a_1 + a_2 + a_3)^2 + 2(a_1a_2 + a_1a_3 + a_2a_3) \\ &= X + 2b, \\ \Gamma_{[(12)]} &= (X - (a_1a_2 + a_2a_1 + a_3^2))(X - (a_1a_2 + a_2a_1 + a_3^2)) \\ &\qquad \cdot (X - (a_1a_2 + a_2a_1 + a_3^2)) \\ &= X^3 - 3b^2X - 2b^3 - 27c^2, \\ \Gamma_{[(123)]} &= (X - (a_1a_2 + a_2a_3 + a_3a_1))(X - (a_1a_3 + a_2a_1 + a_3a_2)) \\ &= (X - b)^2. \end{aligned}$$

The least common multiple of their pairwise resultants is $3b(4b^3 + 27c^2)$, which completes the proof by Theorem 5.3. □

An identical computation can be done for polynomials of higher degree, as long as one has the patience to work out the coefficients of the Γ_C. Here is the corresponding result for quartics:

Theorem 7.3. *Let $f(x)=x^4+bx^2+cx+d$ be a separable quartic polynomial over K, and $\mathfrak{p}$ a prime of K with residue field $\mathbb{F}_q$. Then the value*

$$\mathrm{Tr}_{\frac{\mathbb{F}_q[x]}{f(x)}/\mathbb{F}_q}(x^{q+1})$$

is a root of one of the polynomials

$$\begin{aligned}
\Gamma_{[\mathrm{id}]} &= X+2b,\\
\Gamma_{[(12)(34)]} &= X^3-2bX^2-16dX+32bd-8c^2,\\
\Gamma_{[(12)]} &= X^6+4bX^5+(2b^2+8d)X^4+(-12b^3+48bd-26c^2)X^3\\
&\quad -(23b^4-120b^2d+108bc^2+112d^2)X^2\\
&\quad -(16b^5-128b^3d+138b^2c^2+256bd^2+216c^2d)X-4b^6\\
&\quad +48b^4d-56b^3c^2-192b^2d^2-288bc^2d-27c^4+256d^3,\\
\Gamma_{[(123)]} &= X^4+(-2b^2+8d)X^2-8c^2X+b^4-8b^2d+8bc^2+16d^2,\\
\Gamma_{[(1234)]} &= X^3-2bX^2+(b^2-4d)X+c^2.
\end{aligned}$$

If $\mathfrak{p}$ does not divide the denominators of b, c and d and the pairwise resultants of the Γ_C, then this determines the degrees in the factorisation of f mod $\mathfrak{p}$: They are the cycle lengths of the permutation in the index of Γ.

A theorem of Brumer (see [Jensen et al. 2002, Theorem 2.3.5]) states that any Galois extension L/K with Galois group $G=\mathrm{D}_{10}$ is a splitting field of

$$f_{a,b}(x)=x^5+(a-3)x^4+(b-a+3)x^3+(a^2-a-1-2b)x^2+bx+a$$

for some $a, b\in K$. Using a similar argument to $G=\mathrm{S}_3$ and S_4, we find:

Theorem 7.4. *Suppose L/K is the splitting field of $f_{a,b}(x)$ as above, with*

$$G=\mathrm{Gal}(L/K)\cong \mathrm{D}_{10}.$$

If $\mathfrak{p}$ is a prime of K with residue field $\mathbb{F}_q$, not dividing $3a-b+1$ and the denominators of a and b and such that $f \bmod \mathfrak{p}$ is irreducible, then

$$\mathrm{Tr}_{\frac{\mathbb{F}_q[x]}{f(x)}/\mathbb{F}_q}(x^{q+1})$$

is either $-2a+b+1$ or $a+2$ modulo $\mathfrak{p}$. This determines which of the two conjugacy classes of 5-cycles contains $\mathrm{Frob}_{\mathfrak{p}}$.

Remark 7.5. In this setting, if $\mathrm{Frob}_{\mathfrak{p}}$ is not a 5-cycle, it is either the identity or an element of order 2. In the former case,

$$\mathrm{Tr}_{\frac{\mathbb{F}_q[x]}{f(x)}/\mathbb{F}_q}(x^{q+1})=a^2-4a-2b+3 \bmod \mathfrak{p};$$

in the latter the trace is a root of

$$\begin{aligned}
&\Gamma_{[(23)(45)]}\\
&= X^5 - (a-3)^2X^4 + (31 - 2a^3 + 4b - 3b^2 + a^2(11+2b) - 2a(21+2b))X^3\\
&\quad + (12a^3(3+2b) - a^2(137+44b) + a(114+6b-28b^2)\\
&\qquad - 51 + 7a^4 - 4a^5 - 20b + 14b^2 - 2b^3)X^2\\
&\quad + (40 + 16a^5 - 8a^6 + 32b - 17b^2 - 4b^3 + a^4(58+42b) + a^2(182+18b-52b^2)\\
&\qquad + 4a^3(-49 - 21b + b^2) - 2a(65 + 13b - 17b^2 + 6b^3))X\\
&\quad + 8a^6 - 4a^7 + 4a^5(7+5b) - 4a^4(32+17b) + 2a^3(123+85b+4b^2)\\
&- a^2(245+218b+24b^2) - 2a(-30-6b+51b^2+22b^3) + 2(-6-8b+3b^2+b^3-4b^4).
\end{aligned}$$

Example 7.6. Here is another example, to illustrate what the Γ_C look like in general. Take $K = \mathbb{Q}$ and $L = \mathbb{Q}(E[3])$, the 3-torsion field of the elliptic curve $E : y^2 + y = x^3 - x^2$. Then $\mathrm{Gal}(L/K) \cong \mathrm{GL}_2(\mathbb{F}_3)$, and L is the splitting field of

$$f(x) = x^8 - 9x^7 + 18x^6 + 33x^5 - 93x^4 - 15x^3 - 23x^2 - 36x - 27.$$

The Γ_C for $h(x) = x^2$ are

$$\begin{aligned}
\Gamma_{[\mathrm{id}]} &= X - 144,\\
\Gamma_{[(13)(24)(56)(78)]} &= X - 3,\\
\Gamma_{[(24)(57)(68)]} &= X^{12} - 699X^{11} + 204666X^{10} - 32922129X^9 + 3212225793X^8\\
&\quad - 196600821903X^7 + 7340079612456X^6 - 145234777501584X^5\\
&\quad + 566948224573848X^4 + 26747700562448082X^3\\
&\quad - 187604198442957555X^2 - 2946247136394353892X\\
&\quad - 24290099658154516203,\\
\Gamma_{[(148)(273)]} &= X^8 - 546X^7 + 120102X^6 - 14088342X^5 + 989228043X^4\\
&\quad - 43566817716X^3 + 1248800990265X^2 - 21583664066961X\\
&\quad + 167939769912993,\\
\Gamma_{[(1432)(5768)]} &= X^6 - 258X^5 + 26448X^4 - 1344378X^3 + 34859664X^2\\
&\quad - 445164021X + 2926293624,\\
\Gamma_{[(174382)(56)]} &= X^8 - 264X^7 + 29292X^6 - 1698042X^5 + 51288993X^4\\
&\quad - 654852960X^3 + 3360584547X^2 - 277935306777X + 7299371089503,\\
\Gamma_{[(15473628)]} &= X^6 - 258X^5 + 26250X^4 - 1336755X^3 + 35700471X^2\\
&\quad - 477465444X + 2707751520,\\
\Gamma_{[(16483527)]} &= X^6 - 258X^5 + 28230X^4 - 1674048X^3 + 57362760X^2\\
&\quad - 1097286921X + 9616023198.
\end{aligned}$$

Example 7.7. As an indication of the kind of Artin L-series that may be numerically computed, we give an example with a big Galois group over $\mathbb{Q}$. We take $G = \mathrm{PGSp}(4, \mathbb{F}_3)$ of order 51840, realised through the Galois action on the 3-torsion of the Jacobian of a genus 2 curve, and evaluate the Artin L-series of an irreducible 6-dimensional representation of G.

Specifically, G is the unique double cover of the simple group $\mathrm{Sp}(4,\mathbb{F}_3)/\mathbb{F}_3^\times$ in $\mathrm{PGL}(4,\mathbb{F}_3)=\mathrm{GL}(4,\mathbb{F}_3)/\mathbb{F}_3^\times$. To obtain it as a Galois group, take the hyperelliptic curve

$$\mathscr{C}/\mathbb{Q}: y^2-(x^2+1)y=x^5-x^4+x^3-x^2.$$

Consider the field $\mathbb{Q}(J[3])$ obtained by adjoining to $\mathbb{Q}$ the coordinates of the 3-torsion points of its Jacobian $J/\mathbb{Q}$. Then $\mathrm{Gal}(\mathbb{Q}(J[3])/\mathbb{Q})$ is $\mathrm{GSp}(4,\mathbb{F}_3)$. The group we want is $G=\mathrm{GSp}(4,\mathbb{F}_3)/\{\pm1\}$, and it can be obtained from the Galois action on the 40 lines through the origin in $J[3]$. Specifically, if $(P)+(Q)-2(O)\in J[3]$ is a nonzero point with $P=(x_P,y_P)$, $Q=(x_Q,y_Q)$, the minimal polynomial f of x_Px_Q over $\mathbb{Q}$ has Galois group G;

$$f=x^{40}+27x^{39}+39x^{38}-61x^{37}+\cdots+2259x^3+3471x^2+1057x+69.$$

In its action on the roots of f, the group has several conjugacy classes of the same cycle type, and the largest Γ_C that we need has degree 2160 (using Remark 5.10).

The group has two irreducible 6-dimensional representations, ρ and ρ' (whose traces on elements of order 10 in G are $+1$ and -1, respectively). The curve $\mathscr{C}$ has good reduction outside 2 and 3, so $L/\mathbb{Q}$ is unramified at all primes $p\neq 2,3$. The conductor of ρ is $2^{10}3^{17}$ and we used our machinery to compute the local polynomials for the Artin L-series $L(\rho,s)$ for primes up to 410203. Using Magma [Bosma et al. 1997], we then evaluate

$$L(\rho,1)\approx 1.852529796,\qquad L(\rho,2)\approx 1.119877506,$$

to 10-digit precision. This computation relies implicitly on the validity of Artin's conjecture for ρ. It took half an hour on a 3GHz dual-core CPU to compute $\mathrm{Gal}(f/\mathbb{Q})$, 5 hours for the Γ_C, 3 hours for the local information at $p=2$, and 3 (ramification groups, conductor exponents etc.), 3 minutes for the Frobenius elements and the local polynomial computation and half an hour for each of the L-values.

8. Appendix: Two lemmas on Zariski density

Lemma 8.1. *Suppose K is an infinite field, $f\in K[t]$ is a separable polynomial of degree n and $a_1,\dots,a_n$ are its roots in some splitting field L.*

(a) *If $F,G\in K[x_1,\dots,x_n]$ take the same values on*

$$x_1=\beta_0+\beta_1a_1+\cdots+\beta_{n-1}a_1^{n-1},\dots,x_n=\beta_0+\beta_1a_n+\cdots+\beta_{n-1}a_n^{n-1}$$

for all $[\beta_1,\dots,\beta_n]\in K^n$, then $F=G$.

(b) *Suppose $F_1,\dots,F_d\in K[x_1,\dots,x_n]$ are distinct. There exists a polynomial $B(t)=\beta_0+\cdots+\beta_{n-1}t^{n-1}\in K[t]$ such that $B(a_1),\dots,B(a_n)$ generate L and*

the F_i take distinct values on $[B(a_1), \dots, B(a_n)]$. The set of such B contains a Zariski dense open subset of $K \oplus Kt \oplus \cdots \oplus Kt^{n-1}$.

(c) *Let F be a T-invariant for some $T < \mathrm{S}_n$. There is a Zariski dense open set of polynomials $B(t) \in K \oplus Kt \oplus \cdots \oplus Kt^{n-1}$ for which $\boldsymbol{a}' = [B(a_1), \dots, B(a_n)]$ generate L and $e^F_{\boldsymbol{a}'} : T \setminus \mathrm{S}_n \to L$ is injective.*

Proof. (a) Let $U = K(t_1, \dots, t_n)$. As a first step, we observe that K^n is Zariski dense in $\mathbb{A}^n_U = U^n$: This is clear for $n = 1$ as K is infinite; generally, if K^n were not Zariski dense, it would be contained in a (not necessarily irreducible) hypersurface of some degree d, so it would contain at most d hyperplanes. But, by induction, it contains all $\{r\} \times U^{n-1}$ for all $r \in K$, which gives a contradiction.

Therefore, as F and G are continuous in the Zariski topology, they agree on all of U^n, that is, on all the combinations above with $[\beta_1, \dots, \beta_n] \in U^n$. Now solve the system of equations $\sum_{j=0}^{n-1} a_i^j \beta_j = t_j$ for $\beta_1, \dots, \beta_n$. (This is possible because $a_i \neq a_k$ for $i \neq k$, so the Vandermonde matrix is invertible.) Using this solution we find that $F(t_1, \dots, t_n) = G(t_1, \dots, t_n)$, so $F = G$ as polynomials.

(b) Put $F(x_1, \dots, x_n) = \prod_{i<j}(x_i - x_j)(F_i - F_j)$ and $G = 0$ and apply (a). This gives a polynomial $B(t) = \beta_0 + \cdots + \beta_{n-1}t^{n-1} \in K[t]$ that clearly satisfies the "distinct values" condition. Furthermore, $B(a_i) \neq B(a_j)$ guarantees the "generate L" condition as well: The Galois action permutes the $B(a_i)$ in the same way as the a_i, so the Galois group has the same order. Finally, consider $F(B(a_1), \dots, B(a_n))$ as a polynomial in $\beta_0, \dots, \beta_{n-1}$. Its zero set is Zariski closed in $\mathbb{A}^n$ and we proved that its complement is nonempty. This proves the last claim.

(c) Apply (b) to the set of polynomials $\{F^\sigma\}_{\sigma \in T \setminus \mathrm{S}_n}$, using that, by definition, $e^F_{\boldsymbol{a}'}(\sigma^{-1}) = F((\boldsymbol{a}')^{\sigma^{-1}}) = F^\sigma(\boldsymbol{a}')$. □

Lemma 8.2. *Suppose K is an infinite field, $f \in K[t]$ is a separable polynomial of degree n and $a_1, \dots, a_n$ are its roots in some splitting field L. Then on a Zariski dense open set of polynomials $h(x)$ in $K \oplus Kx \oplus \cdots \oplus Kx^{n-1} \cong \mathbb{A}^n_K$, the values*

$$v_h(\sigma) = \sum_{j=1}^{n} h(a_j)\sigma(a_j)$$

for $\sigma \in G = \mathrm{Gal}(L/K)$ are distinct.

Proof. For any $\sigma \in G$, the map $E_\sigma : h \mapsto v_h(\sigma)$ is K-linear $K^n \to L$. So E_σ agrees with E_τ on a K-linear subspace for every $\sigma, \tau \in G$. If none of these subspaces is all of K^n, then the complement of their union is the desired set (nonempty since K is infinite). It remains to prove that $E_\sigma \neq E_\tau$ for $\sigma \neq \tau$.

Suppose $E_\sigma = E_\tau : K^n \to L$. Then their extensions by linearity to maps $L^n \to L$ agree as well. In other words, $v_h(\sigma) = v_h(\tau)$ for all h in $L \oplus Lx \oplus \cdots \oplus Lx^{n-1}$. In

particular, taking

$$h(x) = \prod_{j \neq i} (x - a_j)$$

we get that $\sigma(a_i) = \tau(a_i)$. As this holds for all i, it follows that $\sigma = \tau$. □

Acknowledgements

Tim is supported by a Royal Society University Research Fellowship. We would like to thank Robinson College and Gonville & Caius College, Cambridge, where most of this research was carried out.

References

[Booker 2005] A. R. Booker, "Numerical tests of modularity", *J. Ramanujan Math. Soc.* **20**:4 (2005), 283–339. MR 2006k:11090 Zbl 1122.11032

[Bosma et al. 1997] W. Bosma, J. Cannon, and C. Playoust, "The Magma algebra system, I: The user language", *J. Symbolic Comput.* **24**:3-4 (1997), 235–265. MR 1484478 Zbl 0898.68039

[Buhler 1978] J. P. Buhler, *Icosahedral Galois representations*, Lecture Notes in Mathematics **654**, Springer, Berlin, 1978. MR 58 #22019 Zbl 0374.12002

[Jensen et al. 2002] C. U. Jensen, A. Ledet, and N. Yui, *Generic polynomials: Constructive aspects of the inverse Galois problem*, Mathematical Sciences Research Institute Publications **45**, Cambridge University Press, 2002. MR 2004d:12007 Zbl 1042.12001

[Roberts 2004] D. P. Roberts, "Frobenius classes in alternating groups", *Rocky Mountain J. Math.* **34**:4 (2004), 1483–1496. MR 2005m:12004 Zbl 1138.12306

[Zagier 2008] D. Zagier, "Elliptic modular forms and their applications", pp. 1–103 in *The 1-2-3 of modular forms*, edited by K. Ranestad, Springer, Berlin, 2008. MR 2010b:11047 Zbl 05808162

Communicated by Bjorn Poonen
Received 2011-08-04 Revised 2012-05-09 Accepted 2012-06-07

tim.dokchitser@bristol.ac.uk *Department of Mathematics, Bristol University, University Walk, Bristol, BS8 1TW, United Kingdom*
http://www.maths.bris.ac.uk/~matyd/

v.dokchitser@dpmms.cam.ac.uk *Emmanuel College, University of Cambridge, Cambridge, CB2 3AP, United Kingdom*
http://www.dpmms.cam.ac.uk/~vd209/

Algebra & Number Theory

msp.org/ant

See inside back cover or msp.org/ant for submission instructions.

The subscription price for 2013 is US $200/year for the electronic version, and $350/year (+$40, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscribers address should be sent to MSP.

Algebra & Number Theory (ISSN 1944-7833 electronic, 1937-0652 printed) at Mathematical Sciences Publishers, 798 Evans Hall #3840, c/o University of California, Berkeley, CA 94720-3840 is published continuously online. Periodical rate postage paid at Berkeley, CA 94704, and additional mailing offices.

ANT peer review and production are managed by EditFLOW® from Mathematical Sciences Publishers.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

http://msp.org/

Algebra & Number Theory

Volume 7 No. 6 2013

1937-0652(2013)7:6;1-7