Computing the Galois group of a polynomial

Curtis Bright

April 15, 2013

Abstract

This article outlines techniques for computing the Galois group of a polynomial over the rationals, an important operation in computational algebraic number theory. In particular, the *linear resolvent polynomial* method of [6] will be described.

1 Introduction

An *automorphism* on a field K is a bijective homomorphism from K to itself. If L/K is a finite extension then the set of automorphisms of L which fix K form a group under composition, and this group is denoted by $\mathrm{Gal}(L/K)$.

Let f be a univariate polynomial with rational coefficients. Throughout this article we suppose that f has degree n and roots $\alpha_1, \ldots, \alpha_n$, which for concreteness we assume lie in the complex numbers, a field in which f splits into linear factors.

The *splitting field* $\mathrm{spl}(f)$ of f is the smallest field in which f splits into linear factors, and it may be denoted by $\mathbb{Q}(\alpha_1, \ldots, \alpha_n)$, a finite extension of $\mathbb{Q}$ generated by the roots of f. Then the *Galois group* of f is defined to be

$$\mathrm{Gal}(f) \coloneqq \mathrm{Gal}(\mathbb{Q}(\alpha_1, \ldots, \alpha_n)/\mathbb{Q}).$$

That is, the group of automorphisms of the splitting field of f over $\mathbb{Q}$.

1.1 The structure of Gal(f)

The elements $\sigma \in \mathrm{Gal}(f)$ may then be thought of as automorphisms of the complex numbers (technically only defined on the splitting field of f) which fix $\mathbb{Q}$. In fact, the automorphism group of $\mathbb{Q}$ is trivial, so the automorphisms of $\mathbb{C}$ fix $\mathbb{Q}$ in any case.

The automorphism group of $\mathbb{C}$, however, is uncountable. From a completely naïve perspective this could mean there is no finite way of expressing $\mathrm{Gal}(f)$, but a simple argument shows these fears are unfounded.

The first thing to note is that $\sigma \in \mathrm{Gal}(f)$ is completely determined by the values $\sigma(\alpha_1), \ldots, \sigma(\alpha_n)$, which follows by the homomorphism properties of σ and the fact that the α_i generate $\mathrm{spl}(f)$ over $\mathbb{Q}$. Secondly, note that $\sigma(\alpha_i)$ is also a root of f:

$$f(\sigma(\alpha_i)) = \sigma(f(\alpha_i)) = \sigma(0) = 0$$

Similarly $\sigma^{-1}(\alpha_i)$ (which exists since σ is a bijection) is also a root of f. In short, roots of f come from and go to other roots of f under σ.

It is often convenient to think of σ as a permutation of the α_i, and by abuse of notation we say $\sigma \in S_n$, the symmetric group of n elements.

1.2 Cubic examples

A visual way of expressing $\mathrm{Gal}(f)$ is to plot the roots of f on the complex plane and describe permutations of the roots by using arrows to specify which roots are sent where. Figure 1 demonstrates what such diagrams look like. In the case of 1(a) and 1(b) the Galois group contains six permutations, which is the most possible since $|S_3| = 3! = 6$.

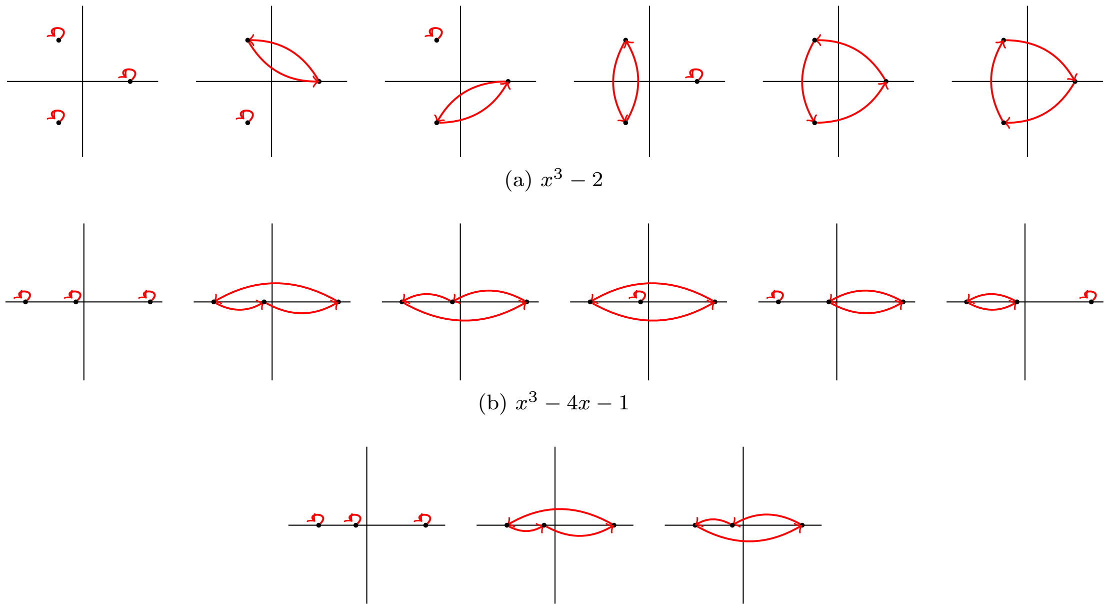

(a) $x^3 - 2$

(b) $x^3 - 4x - 1$

(c) $x^3 - 3x - 1$

Figure 1: The Galois groups of three sample irreducible cubics.

Already, the subtlety of the problem is evident. Viewing figure 1(c) completely superficially, we find that a slight perturbation of the roots from figure 1(b) causes the Galois group to lose the three transpositions of the roots. Naturally, one would like an explanation for such interesting behaviour.

1.3 The lack of transpositions in Gal($x^3 - 3x - 1$)

In fact, the lack of transpositions in $\mathrm{Gal}(x^3 - 3x - 1)$ has a simple explanation linked to the *discriminant*, which may be defined to be the quantity

$$\mathit{\Delta}(f) := (-1)^{n(n-1)/2} \prod_{i \neq j} (\alpha_i - \alpha_j) = \prod_{i<j} (\alpha_i - \alpha_j)^2.$$

For the polynomial in question, we find that $\mathit{\Delta}(x^3 - 3x - 1) = 81$ is a perfect square, so $\sqrt{\mathit{\Delta}} \in \mathbb{Z}$ and in particular $\sigma \in \mathrm{Gal}(f)$ fixes $\pm\sqrt{\mathit{\Delta}}$.

Therefore taking the square root of $\mathit{\Delta}$ and applying σ one finds that

$$\pm\sqrt{\mathit{\Delta}} = \prod_{i<j} (\alpha_i - \alpha_j) = \prod_{i<j} (\sigma(\alpha_i) - \sigma(\alpha_j)).$$

However, if σ was a transposition note that every factor in the final product remains constant except one, in which case a factor of -1 is introduced. For example, if σ transposes α_1 and α_2 then the factor $(\sigma(\alpha_1) - \sigma(\alpha_2))$ becomes $(\alpha_2 - \alpha_1) = -(\alpha_1 - \alpha_2)$. Thus in this case we have $\pm\sqrt{\mathit{\Delta}} = \mp\sqrt{\mathit{\Delta}}$, a contradiction since $\mathit{\Delta} \neq 0$.

In other words, if $\sigma \in S_3$ is a transposition of α_i and α_j then σ cannot be extended into a automorphism of $\mathbb{C}$, since that would require defining $\sigma(9) = -9$ in addition to $\sigma(9) = 9$. In Figure 2 this is shown visually; in the second case the action isn't actually well-defined because of the conflict as to where the point $9 \in \mathbb{C}$ is sent.

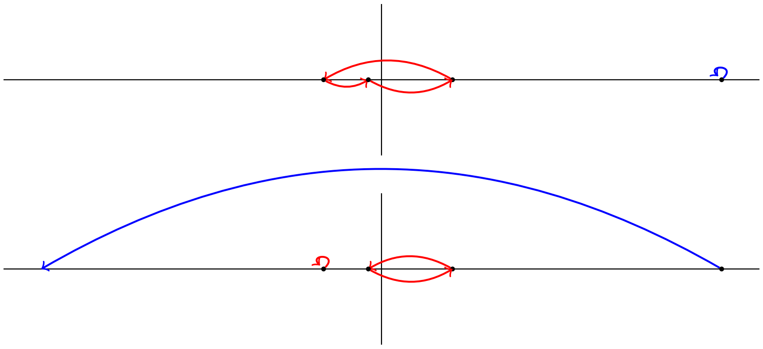

Figure 2: Two elements of $S_3 \supset \text{Gal}(x^3 - 3x - 1)$ acting on $\sqrt{\Delta}$.

1.4 Quartic examples

As a demonstration of some of the complexity that can occur in higher degree, in Figure 3 we give an example of two Galois groups which arise from quartics. In each case the Galois group has four elements, but the groups are non-isomorphic.

The roots in 4(a) are the primitive fifth roots of unity and the Galois group of their minimal polynomial is isomorphic to C_4, the cyclic group on four elements. The roots in 4(b) are the primitive eighth roots of unity and the Galois group of their minimal polynomial is isomorphic to $V_4 \cong C_2 \times C_2$, the Klein four-group.

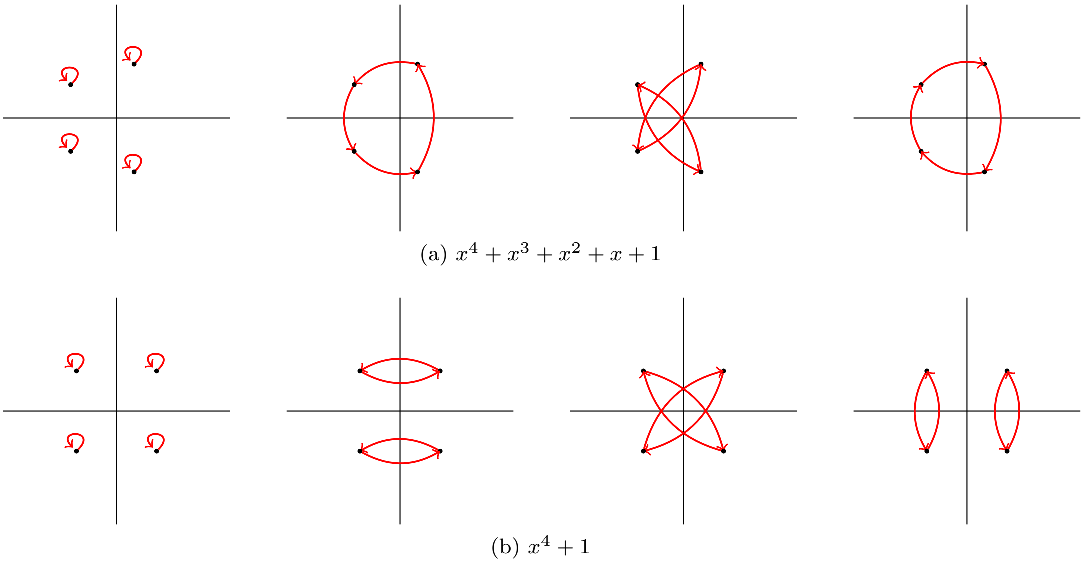

(a) $x^4 + x^3 + x^2 + x + 1$

(b) $x^4 + 1$

Figure 3: The Galois groups of two sample irreducible quartics.

1.5 Motivation

The following well-known theorem (e.g., [4, Theorem 14.39]) provides some motivation as to why the Galois group of a polynomial is of interest.

Theorem. *The roots of f are solvable in radicals if and only if* $\text{Gal}(f)$ *is a* solvable group, *i.e., there exists a chain of subgroups*

$$\{1\} = G_0 \trianglelefteq G_1 \trianglelefteq \cdots \trianglelefteq G_k = \text{Gal}(f)$$

where each G_i/G_{i-1} is abelian.

Actually, one only needs to know $\mathrm{Gal}(f)$ up to conjugacy in S_n to apply this theorem, i.e., up to reordering of permutation indices. Since one does not typically care about the ordering of $\alpha_1, \ldots, \alpha_n$ anyway, we only concern ourselves with computing the Galois group up to conjugacy, i.e., the Galois group under some ordering of the roots.

2 Problem simplifications

For simplicity, we will assume that f is irreducible over $\mathbb{Q}$. From this we note two immediate consequences.

First, f is a *separable* polynomial, i.e., it has distinct roots. This follows from the fact that f' has smaller degree than f and is nonzero (as $\mathrm{char}(\mathbb{Q}) = 0$) and f has no nontrivial divisors, so $\gcd(f, f') = 1$. But if α were a double root of f then it would also be a root of f', and therefore the minimal polynomial of α would divide $\gcd(f, f') = 1$, a contradiction.

Second, $\mathrm{Gal}(f)$ is a *transitive* group, i.e., for all α_i and α_j there is some $\sigma \in \mathrm{Gal}(f)$ which sends α_i to α_j. To show this, one checks that there is an embedding of $\mathbb{Q}(\alpha_i)$ in $\mathbb{C}$ with $\alpha_i \mapsto \alpha_j$ (which one then extends to $\mathrm{spl}(f)$ as necessary). Intuitively, this holds since $\mathbb{Q}(\alpha_i) \cong \mathbb{Q}(\alpha_j)$; since α_i and α_j are roots of the same minimal polynomial they are "algebraically indistinguishable"—you can replace α_i with α_j without breaking anything.

Note that while one can find automorphisms in which $\alpha_i \mapsto \alpha_j$ and $\alpha_j \mapsto \alpha_k$ *separately*, this does not mean one can enforce both *simultaneously*—after the freedom of choosing where to send α_i, one may not have any more choice in the matter.

In general, one has that $\mathrm{Gal}(gh) \subseteq \mathrm{Gal}(g) \times \mathrm{Gal}(h)$, but we do not consider the specific details of computing $\mathrm{Gal}(gh)$.

Furthermore, we assume that f is monic and has integer coefficients. This is not unreasonable since the general case reduces to this case by applying transformations of the form (for nonzero $c \in \mathbb{Q}$)

$$\begin{aligned} f(x) &\mapsto cf(x) \\ f(x) &\mapsto f(cx) \end{aligned}$$

which do not change the splitting field of f—the roots don't change in the first case and in the second case they are scaled by a nonzero rational.

If $f(x) := \frac{1}{b}\sum_{i=0}^{n} a_i x^i$ for a_i, $b \in \mathbb{Z}$ then the reduction to the special case proceeds by applying

$$f(x) \mapsto b a_n^{n-1} f(x/a_n).$$

The coefficients are scaled by b to remove denominators, and the roots are scaled by a_n as a step toward making the polynomial monic. The root scaling has the effect of removing a factor of a_n^k from the x^k coefficient, so to finish making the polynomial monic (and remove denominators again) it is also necessary to scale the coefficients by a_n^{n-1}.

3 Symmetric polynomials

A multivariate polynomial $p \in R[x_1, \ldots, x_n]$ is *symmetric* if it is fixed under all permutations of the n indeterminants, i.e.,

$$p(x_1, \ldots, x_n) = p(x_{\sigma(1)}, \ldots, x_{\sigma(n)})$$

for all $\sigma \in S_n$.

One possible ambiguity is that the symmetricity of p is dependent on the polynomial ring in which p lives. For example,

$$x_1^2 + \cdots + x_n^2$$

is a symmetric polynomial when considered as a polynomial in n indeterminants, but *not* as a polynomial in $n+1$ indeterminants. In the latter case transposing x_1 and x_{n+1} gives a new polynomial distinct from p.

The *elementary symmetric polynomials* $s_1, \ldots, s_n \in R[x_1, \ldots, x_n]$ are defined by:

$$\begin{aligned} s_1 &\coloneqq x_1 + x_2 + \cdots + x_n \\ s_2 &\coloneqq x_1x_2 + x_1x_3 + \cdots + x_{n-1}x_n \\ &\vdots \\ s_n &\coloneqq x_1x_2 \cdots x_n \end{aligned}$$

They are important since they appear (up to sign) as the coefficients of the *general polynomial of degree n*:

$$\prod_{i=1}^{n}(x - x_i) = x^n - s_1x^{n-1} + \cdots + (-1)^n s_n$$

Applying the substitution $x_i \mapsto \alpha_i$ to this gives $\prod_{i=1}^{n}(x-\alpha_i) = f(x)$, so if $f \in \mathbb{Z}[x]$ then we have that the elementary polynomials evaluated at $\alpha_1, \ldots, \alpha_n$ are in fact integers.

It is clear that the elementary symmetric polynomials are symmetric, since any permutation simply rearranges the terms of each s_i. What is more interesting is that *all* symmetric polynomials can be expressed in terms of the elementary symmetric polynomials. The following is known as the *fundamental theorem of symmetric polynomials.* [4, Corollary 14.31]

Theorem. *Every symmetric polynomial in $R[x_1, \ldots, x_n]$ can be written as a polynomial in $s_1, \ldots, s_n$ with coefficients in R.*

In fact, the decomposition can be effectively computed using a generalization of the Euclidean algorithm. The idea is to define an ordering on the monomials $x_1^{a_1} \cdots x_n^{a_n}$ by, for example, lexicographic ordering on the exponent vectors $(a_1, \ldots, a_n)$. If the symmetric polynomial p has leading term $cx_1^{a_1} \cdots x_n^{a_n}$ under this ordering then one checks that

$$p - cs_1^{a_1-a_2}s_2^{a_2-a_3} \cdots s_n^{a_n} \tag{1}$$

is a symmetric polynomial with smaller leading term. Applying this recursively one eventually expresses p in terms of the s_i (it eventually finishes since $\mathbb{N}^n$ does not contain an infinite strictly decreasing sequence under lexicographic ordering). Also, so that (1) is actually a *polynomial* it is necessary to show that $a_i \geq a_{i+1}$. This holds since otherwise p, being symmetric, would contain the term $cx_1^{a_1} \cdots x_i^{a_{i+1}} x_{i+1}^{a_i} \cdots x_n^{a_n}$, a contradiction to $cx_1^{a_1} \cdots x_n^{a_n}$ being the leading term of p.

3.1 The orbit of a polynomial

The *orbit* of a polynomial $p \in R[x_1, \ldots, x_n]$ under S_n is the set of polynomials that p can be sent to by permuting the x_i, and this will be denoted by $\operatorname{orb}(p)$. This can be thought of as measuring "how close" a polynomial is to being symmetric.

For example, if $\operatorname{orb}(p)$ is as small as possible, i.e., $\operatorname{orb}(p) = \{p\}$, then p is fixed any permutation of the x_i so it is symmetric. The other extreme is when $\operatorname{orb}(p)$ is as large as possible, i.e., $|\operatorname{orb}(p)| = n!$; in this case p can be thought of as being "as far from symmetric as possible". Of course, possibilities between these two extremes are also possible, e.g., the orbit of $x_1 + x_2$ under S_3 is $\{x_1 + x_2, x_1 + x_3, x_2 + x_3\}$.

4 The resolvent polynomial

The most important definition for what follows is that of a *resolvent polynomial.* Intuitively, the resolvent polynomial is defined as a polynomial whose roots are "combinations" of the roots $\alpha_1, \ldots, \alpha_n$ of f; the manner in which the α_i are combined is determined by a multivariate polynomial p. That is, the resolvent polynomial is defined in terms of two polynomials $f \in \mathbb{Z}[x]$ and $p \in \mathbb{Z}[x_1, \dots, x_n]$ to be the new univariate polynomial

$$R_{p,f}(y) \coloneqq \prod_{p_i \in \operatorname{orb}(p)} \big(y - p_i(\alpha_1, \dots, \alpha_n)\big).$$

For example, if $p \coloneqq x_1 + x_2$ then the roots of $R_{p,f}$ are all sums of the roots of f; since f has n roots the resolvent polynomial will have $\binom{n}{2}$ roots, as demonstrated by the following examples:

$f(x)$	$R_{p,f}(y)$
$x^3 - 2$	$y^3 + 2$
$x^4 + 1$	$y^6 - 4y^2$
$x^4 + x^3 + x^2 + x + 1$	$y^6 + 3y^5 + 5y^4 + 5y^3 - 2y - 1$

A useful example is to take $p \coloneqq \prod_{i<j}(x_i - x_j)$. In this case p is *nearly* symmetric, but not quite, since as previously noted if σ is a transposition then $\sigma(p) = -p$. Note that every permutation can be decomposed into a product of transpositions (e.g., using bubblesort) and, furthermore, the parity of the number of transpositions in the decomposition is invariant. The *even* permutations induce a even number of factors of -1 and therefore leave p fixed, while the *odd* permutations induce an odd number of factors of -1 and therefore negate p, so we have $\operatorname{orb}(p) = \{p, -p\}$. Then

$$R_{p,f}(y) = \Big(y - \prod_{i<j}(\alpha_i - \alpha_j)\Big)\Big(y + \prod_{i<j}(\alpha_i - \alpha_j)\Big) = y^2 - \operatorname{disc}(f).$$

4.1 The coefficients of the resolvent

Note that since the resolvent is defined with respect to the orbit of p it is symmetric in the α_i, i.e., permuting the α_i will permute the roots of $R_{p,f}$ but does not in fact change $R_{p,f}$. In other words, the coefficients of $R_{p,f}$ are symmetric polynomials in α_1, …, α_n, and by the fundamental theorem of symmetric polynomials can be written in terms of the elementary symmetric polynomials in α_1, …, α_n. However, as previously noted the elementary symmetric polynomials in α_1, …, α_n are (up to sign) exactly the coefficients of f, and therefore integers.

This shows that $R_{p,f} \in \mathbb{Z}[y]$ when $p \in \mathbb{Z}[x_1, \dots, x_n]$ and $f \in \mathbb{Z}[x]$, and furthermore gives an algorithm for computing the coefficients of $R_{p,f}$. In practice, another way of computing $R_{p,f}$ is to approximate the roots of f via numerical root-finding methods, form all combinations of the roots as specified by p, and then expand the product from the definition to find approximations of the coefficients of $R_{p,f}$. Since the coefficients are known to be integers, if the approximations are known with sufficient accuracy (absolute error less than 0.5) then the approximations may simply be rounded to the nearest integer. It is perhaps less elegant than the symbolic method, but is simple to implement in practice.

4.2 Gal(f) acting on the roots of the resolvent

If $\sigma \in \operatorname{Gal}(f)$ then as previously noted σ fixes $R_{p,f}$ but permutes its roots—we say that σ *acts on* the roots of $R_{p,f}$. As before, this permutation may be expressed visually by plotting the roots of $R_{p,f}$ in the complex plane and using arrows between the roots to represent which root is sent where. Figure 4 demonstrates this for three example polynomials f we've previously seen, using $p \coloneqq x_1 + x_2$.

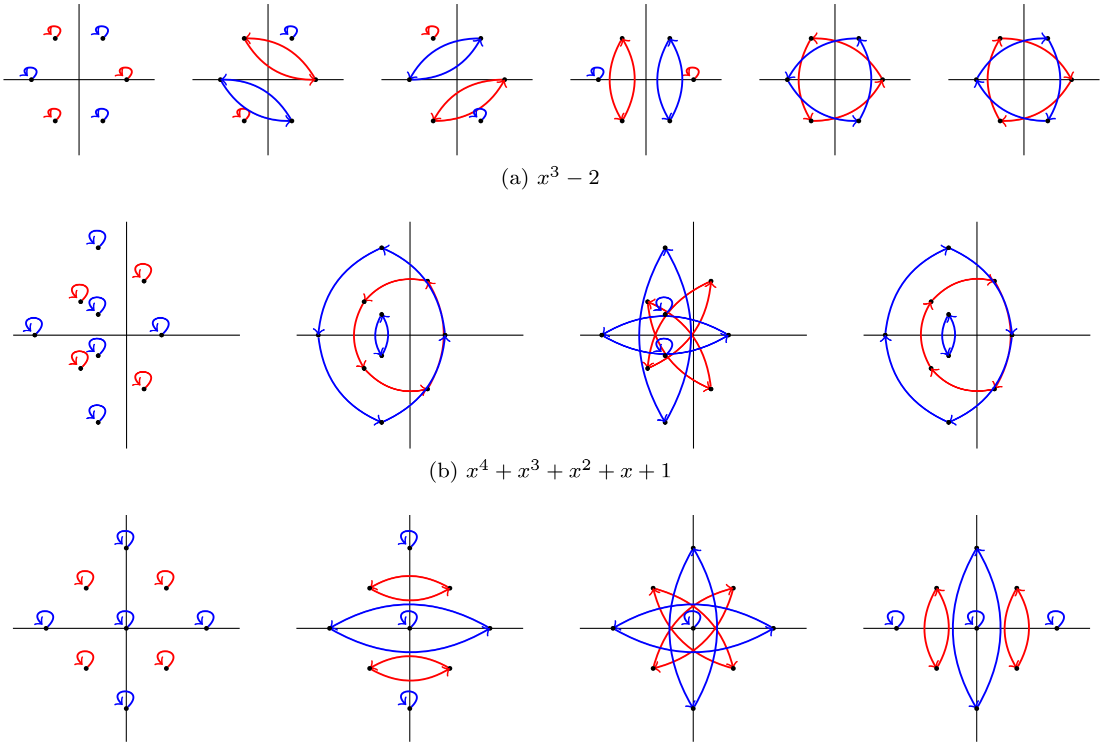

(a) $x^3 - 2$

(b) $x^4 + x^3 + x^2 + x + 1$

(c) $x^4 + 1$

Figure 4: $\sigma \in \mathrm{Gal}(f)$ acts on the roots of $R_{p,f}$.

The key observation here is that the action by $\sigma \in \mathrm{Gal}(f)$ on the roots of $R_{p,f}$ actually gives $\mathrm{Gal}(R_{p,f})$. More precisely, let $\phi\colon \mathrm{Gal}(f) \to \mathrm{Gal}(R_{p,f})$ be defined so that $\phi(\sigma)$ is the action by σ on the roots of $R_{p,f}$. Alternatively, if $R_{p,f}$ has m roots one can think of ϕ as a partial function $\phi\colon S_n \to S_m$, although the roots of $R_{p,f}$ should be distinct for ϕ to be unambiguous. For example, if $R_{p,f}$ has a double root $\beta_1 = \beta_2$ with $\phi(\sigma)(\beta_1) = \beta_1$ then when viewed as a permutation $\phi(\sigma)$ could either fix or transpose 1 and 2. Formally, we have [2, Theorem 6.3.3] the following:

Theorem. *If the roots of* $R_{p,f}(y)$ *are distinct then* $\mathrm{Gal}(R_{p,f}) = \phi(\mathrm{Gal}(f))$.

Intuitively this holds since by definition of ϕ we have that

$$\phi(\mathrm{Gal}(f)) \subseteq \mathrm{Gal}(R_{p,f}),$$

i.e., if σ is an automorphism then $\phi(\sigma)$ is also an automorphism. For the other direction, note that since $R_{p,f}$'s roots are built out of f's roots, we have $\mathrm{spl}(R_{p,f}) \subseteq \mathrm{spl}(f)$. If we think of the automorphisms of $\mathrm{spl}(R_{p,f})$ as being automorphisms of $\mathrm{spl}(f)$ (after extending them as necessary) then we have $\mathrm{Gal}(R_{p,f}) \subseteq \mathrm{Gal}(f)$, or by "projecting" down into $\mathrm{spl}(R_{p,f})$ we have

$$\mathrm{Gal}(R_{p,f}) \subseteq \phi(\mathrm{Gal}(f)).$$

4.3 Tschirnhausen transformation

As mentioned, to use the above theorem we require that $R_{p,f}$ have distinct roots, however will not always be the case; see for example Figure 4(c). However, by applying a simple transformation we can find a new irreducible monic polynomial $g \in \mathbb{Z}[y]$ with the same splitting field as f and with $R_{p,g}$ separable. Although the algorithm is nondeterministic and in theory might not terminate, in practice it usually finishes rather quickly. It proceeds as follows:

1. Choose a "random" polynomial $A \in \mathbb{Z}[x]$ of degree less than that of f. For example, chose the coefficients uniformly at random from some finite range.

2. Compute
$$g(y) \coloneqq R_{A,f}(y) = \prod_{i=1}^{n} \bigl(y - A(\alpha_i)\bigr).$$
This may be efficiently computed as the resultant of $f(x)$ and $y - A(x)$ with respect to x.

3. If g and $T \coloneqq R_{p,g}$ are squarefree (i.e., $\gcd(g, g') = \gcd(T, T') = 1$) then output g; otherwise choose a new $A \in \mathbb{Z}[x]$ and start over.

If the algorithm terminates then g is squarefree and all of its roots $A(\alpha_i)$ are conjugate, so g is irreducible. Also, T is squarefree over $\mathbb{Z}$, and therefore over $\mathbb{Q}$ by Gauss' Lemma. If α was a double root of T then $h \coloneqq \operatorname{minpoly}(\alpha)$ divides T (say $T = hk$) as well as $T' = h'k + k'h$, so it must also divide $h'k$. However, h is irreducible so as noted before $\gcd(h, h') = 1$, and it follows that h divides k, a contradiction to T being squarefree. Thus T has distinct roots.

Furthermore, since the roots of g are built out of the roots of f we have that
$$\operatorname{spl}(g) \subseteq \operatorname{spl}(f). \qquad (2)$$
Now, if one considers the action ϕ of $\operatorname{Gal}(f)$ on the roots of g one sees that it is invertible: since the roots of g are distinct and are constructed from f's roots via the *univariate* polynomial A one can compute from which root of f each root of g was derived from and thereby find σ from $\phi(\sigma)$. Thus ϕ is injective and $|\operatorname{Gal}(f)| \leq |\operatorname{Gal}(g)|$. Since the splitting field of any polynomial over $\mathbb{Q}$ is a Galois extension of $\mathbb{Q}$, one concludes that
$$[\operatorname{spl}(f) : \mathbb{Q}] \leq [\operatorname{spl}(g) : \mathbb{Q}].$$
Combining this with (2) we have $\operatorname{spl}(g) = \operatorname{spl}(f)$ as required.

5 The orbit-length partition

The idea behind the method for determining the Galois group of f is to make use of the previous theorem which says that $\operatorname{Gal}(R_{p,f}) = \phi(\operatorname{Gal}(f))$. Of course, although one can compute $R_{p,f}$ one can't necessarily compute $\operatorname{Gal}(R_{p,f})$ since this is the problem we're trying to solve; in fact it is likely even harder since typically the degree of the resolvent will be larger than the degree of f. However, it is not necessary to know the full Galois group of $R_{p,f}$ to be able to make use of the theorem; in particular we will see how knowledge of $\operatorname{Gal}(R_{p,f})$ can be used to limit the possibilities for $\operatorname{Gal}(f)$.

5.1 'Local' transitivity

By way of motivation, consider the Galois group of $R_{p,f}$ from Figure 4(b), where $f \coloneqq x^4 + x^3 + x^2 + x + 1$ and $p \coloneqq x_1 + x_2$:

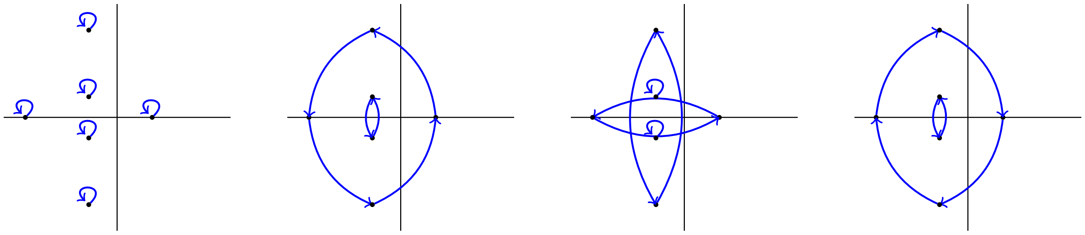

Note that the inner two roots and outer four roots are ‘locally’ transitive, i.e., there is an automorphism which sends any inner root to any other inner root, as well as an automorphism which sends any outer root to any other outer root.

Since automorphisms permute the roots of any minimal polynomial it follows that locally transitive roots will share a minimal polynomial. Conversely, if roots of $R_{p,f}$ share a minimal polynomial then there will be an automorphism which sends any one to any other, so the roots will be locally transitive.

In other words, the local transitivity of the roots of $R_{p,f}$ under the action by $\mathrm{Gal}(f)$ can be determined by factoring $R_{p,f}$ over $\mathbb{Z}$ to find the minimal polynomials of its roots. Since $R_{p,f}$ is monic it is primitive and therefore it is equivalent to factor it over $\mathbb{Q}$, which can be done in polynomial time [5].

For the above example, we have

$$R_{p,f} = y^6 + 3y^5 + 5y^4 + 5y^3 - 2y - 1 = (y^4 + 2y^3 + 4y^2 + 3y + 1)(y^2 + y - 1)$$

so while this doesn’t give us the Galois group of $R_{p,f}$ it does tell us its local transitivity, i.e., the orbits of the action by $\mathrm{Gal}(f)$ on the roots of $R_{p,f}$. Note that the orbits of an action on a set S form a partition of S; the sizes of the orbits are known as the *orbit-length partition*. In the above case the orbit-length partition is $(4, 2)$, corresponding to the degrees of the irreducible factors of $R_{p,f}$.

5.2 Limiting the possibilities for Gal(f)

Suppose that $p \coloneqq x_1 + x_2$ and f is of degree 4. Then $R_{p,f}$ has the following six roots:

$$\alpha_1 + \alpha_2 \qquad \alpha_1 + \alpha_3 \qquad \alpha_1 + \alpha_4 \qquad \alpha_2 + \alpha_3 \qquad \alpha_2 + \alpha_4 \qquad \alpha_3 + \alpha_4$$

Suppose that $\mathrm{Gal}(f) = V_4$, the Klein four-group. Up to relabeling, we can express V_4 using cycle notation as

$$V_4 \coloneqq \{1, (12)(34), (13)(24), (14)(23)\}.$$

Now consider the action of each permutation of V_4 on the roots of $R_{p,f}$:

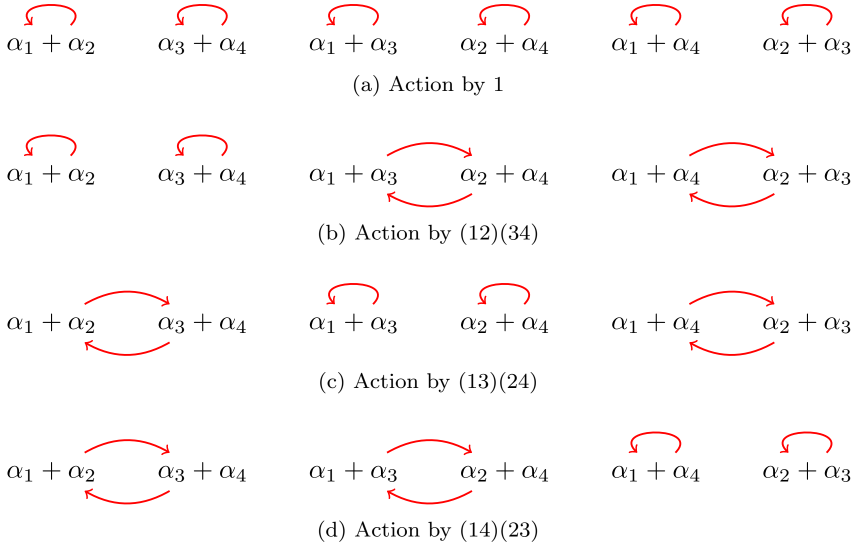

$\alpha_1 + \alpha_2$ $\alpha_3 + \alpha_4$ $\alpha_1 + \alpha_3$ $\alpha_2 + \alpha_4$ $\alpha_1 + \alpha_4$ $\alpha_2 + \alpha_3$
$\alpha_1 + \alpha_2$ $\alpha_3 + \alpha_4$ $\alpha_1 + \alpha_3$ $\alpha_2 + \alpha_4$ $\alpha_1 + \alpha_4$ $\alpha_2 + \alpha_3$
$\alpha_1 + \alpha_2$ $\alpha_3 + \alpha_4$ $\alpha_1 + \alpha_3$ $\alpha_2 + \alpha_4$ $\alpha_1 + \alpha_4$ $\alpha_2 + \alpha_3$
$\alpha_1 + \alpha_2$ $\alpha_3 + \alpha_4$ $\alpha_1 + \alpha_3$ $\alpha_2 + \alpha_4$ $\alpha_1 + \alpha_4$ $\alpha_2 + \alpha_3$

(a) Action by 1

(b) Action by (12)(34)

(c) Action by (13)(24)

(d) Action by (14)(23)

Figure 5: The actions by permutations of V_4 on the roots of $R_{p,f}$.

From this we see that the orbits of V_4 acting on the roots of $R_{p,f}$ are $\{\alpha_1+\alpha_2, \alpha_3+\alpha_4\}$, $\{\alpha_1+\alpha_3, \alpha_2+\alpha_4\}$, and $\{\alpha_1 + \alpha_4, \alpha_2 + \alpha_3\}$, i.e., V_4 gives rise to the orbit-length partition $(2, 2, 2)$. Continuing with the example from 5.1, this is a contradiction since $\mathrm{Gal}(f)$ actually gives rise to the orbit-length partition $(4, 2)$. Thus $\mathrm{Gal}(f) \neq V_4$.

One can similarly compute the orbit-length partitions for the other transitive subgroups of S_4, which results in the following table:

S_4	A_4	D_4	V_4	C_4
(6)	(6)	$(4,2)$	$(2,2,2)$	$(4,2)$

From this we conclude that in our previous example $\mathrm{Gal}(f)$ must be either D_4 or C_4. To distinguish between these cases, we try a new resolvent polynomial.

With the choice $p := x_1 - x_2$ one sees that there are 4 possibilities for where to send x_1 and 3 remaining possibilities where to send x_2, and each choice gives rise to a new polynomial since the coefficients of x_1 and x_2 are distinct. Thus $R_{p,f}$ will have degree 12, and we can compute

$$R_{p,f}(y) = y^{12} + 5y^{10} + 15y^8 + 25y^6 - 50y^4 + 125 = (y^4 + 5y^2 + 5)(y^4 + 5y + 5)(y^4 - 5y + 5),$$

so the orbit-length partition of the roots of $R_{p,f}$ under $\mathrm{Gal}(f)$ is $(4, 4, 4)$. As before, one can also compute the orbit-length partition of the roots of $R_{p,f}$ under each transitive subgroup of S_4, resulting in the following table:

S_4	A_4	D_4	V_4	C_4
(12)	(12)	$(8,4)$	$(4,4,4)$	$(4,4,4)$

Thus $\mathrm{Gal}(f)$ is either V_4 or C_4. By the process of elimination, we conclude that in this example $\mathrm{Gal}(f) = C_4$, which is correct as seen in Figure 3(a).

5.3 The effectiveness of this method

In [6] this method is examined for polynomials up to degree 7. They show that $\mathrm{Gal}(f)$ can be uniquely distinguished in all but four cases by taking p to be small linear polynomials, as well as $p_\Delta := \prod_{i<j}(x_i - x_j)$. As previously noted, p_Δ leads to the resolvent $y^2 - \mathrm{disc}(f)$ which splits if and only if $\mathrm{disc}(f)$ is a perfect square. The following table gives some possible choices for p which together will completely solve each case $n \leq 7$:

degree 3:	p_Δ or $x_1 - x_2$
degree 4:	$p_\Delta, x_1 + x_2, x_1 - x_2$
degree 5:	$p_\Delta, x_1 - x_2, (x_1 + x_2 - x_3 - x_4)^2$
degree 6:	$p_\Delta, x_1 + x_2, x_1 + x_2 + x_3, x_1 - x_2, x_1 + x_2 + x_3 + p_\Delta$
degree 7:	$p_\Delta, x_1 + x_2 + x_3$

6 A general algorithm

Although this gives a complete algorithm for small degree, the choices for p are somewhat *ad hoc* and one would like assurance the problem can always be solved, at least in principle. In fact, for any $G \subseteq S_n$ one can test if $\mathrm{Gal}(f) \subseteq G$ by selecting p so that p is fixed by exactly the permutations in G, i.e., $G = \mathrm{stab}(p)$. Although it may not be the simplest choice, one explicit possibility for p which satisfies this condition is

$$p := \sum_{\sigma \in G} x_{\sigma(1)} x_{\sigma(2)}^2 x_{\sigma(3)}^3 \cdots x_{\sigma(n)}^n.$$

For each fixed $\tau \in G$, we have $\tau(G) = G$, so τ permutes the terms of p without changing p. On the other hand, if $\tau \notin G$ then the term $x_{\tau(1)} x_{\tau(2)}^2 \cdots x_{\tau(n)}^n$ does not appear in p, so τ doesn't fix p.

Now, $R_{p,f}$ has a simple integer root if and only if

$$\mathrm{Gal}(f) \subseteq \mathrm{stab}(p(x_1, \ldots, x_n))$$

for some ordering of the x_i. [3, Proposition 13.3.2]

For suppose $R_{p,f}$ had an integer root β; say $\beta = p(\alpha_1, \ldots, \alpha_n)$. If $\mathrm{Gal}(f) \nsubseteq \mathrm{stab}(p(x_1, \ldots, x_n))$ then there is some $\sigma \in \mathrm{Gal}(f)$ with $\sigma(p(x_1, \ldots, x_n)) \neq p(x_1, \ldots, x_n)$, so these both occur in the product expansion of $R_{p,f}$ as two distinct roots. However, $p(\alpha_1, \ldots, \alpha_n) = \beta$ and $\sigma(p(\alpha_1, \ldots, \alpha_n)) = \sigma(\beta) = \beta$, a contradiction.

Conversely, say every $\sigma \in \mathrm{Gal}(f)$ fixes $p(x_1, \ldots, x_n)$. Then in particular it fixes the root $p(\alpha_1, \ldots, \alpha_n)$ of $R_{p,f}$ and it therefore has a linear minimal polynomial (which divides $R_{p,f}$).

Therefore, after a separable $R_{p,f}$ is constructed we can easily test if $\mathrm{Gal}(f) \subseteq \mathrm{stab}(p)$ up to conjugacy. Using the above construction for p we can set $\mathrm{stab}(p)$ to be any subgroup of our choosing, and in particular run the test with all transitive subgroups of S_n. Enumerating these is a problem of interest in its own right, and has already been solved for $n \leq 32$ [1].

Once all transitive subgroups $G \subseteq S_n$ are known, one can work their way through the subgroup lattice by testing if $\mathrm{Gal}(f) \subseteq G$ as required. For example, the subgroup lattice of S_4 is:

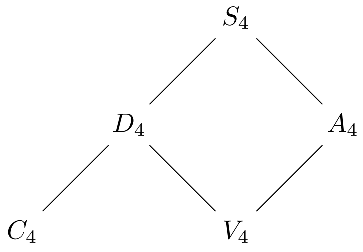

In this case, one can determine $\mathrm{Gal}(f)$ by using at most 3 containment tests.

Although this may be theoretically fulfilling, this algorithm is not practical for large n, since as one moves down the subgroup lattice the degree of the resultant becomes unwieldy. Using the p from above one has $|\mathrm{orb}(p)| = n!/|G|$, so to determine if $\mathrm{Gal}(f) \subseteq C_n$ would require $R_{p,f}$ to have degree $(n-1)!$—even writing down this polynomial is takes longer than exponential time in the size of the input, which is linear in n.

References

[1] J. J. Cannon and D. F. Holt. The transitive permutation groups of degree 32. *Experimental Mathematics*, 17(3):307–314, 2008.

[2] H. Cohen. *A Course in Computational Algebraic Number Theory.* Graduate Texts in Mathematics. Springer, 1993.

[3] D. Cox. *Galois Theory.* Pure and applied mathematics. Wiley, 2004.

[4] D. Dummit and R. Foote. *Abstract Algebra.* Wiley, 2004.

[5] A. K. Lenstra, H. W. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982.

[6] L. Soicher and J. McKay. Computing galois groups over the rationals. *Journal of Number Theory*, 20(3):273 – 281, 1985.