Local Computations for Global Problems: Galois Theory

Claus Fieker

MPI-Leipzig, September 1, 2021

Intro/ Basics

We will deal with Galois theory

- how to compute Galois groups
- ... and what to do with it.

The Galois theory will be "easy" as we can re-use most of the tools from the factorisation.

Galois

Let $f \in \mathbb{Z}[x]$ be monic and irreducible (both can be relaxed). We want the Galois group of (a splitting field of) f.
But what does this mean?
$f = x^4 - 2$, then "obviously" $\mathrm{Gal}(f) = D_4$ (in Magma, D_8 in Gap), but this is a useless result, we cannot use this further as we do not have the operation of the group on anything, just the isomorphism type.

Step 0: compute the roots of f and then we can talk about the Galois group as an explicit group of permutations of the roots. Of course, we will choose an unramified p-adic splitting field...
And fix the ordering of the roots for the rest of the computation.

History

- Stauduhar, 1969, PhD with Lehmer, Berkeley
- Soicher, 1985, PhD with Conway, Cambridge
- Geier, 1993, Heidelberg, Diploma
- Eichenlaub, 1996, PhD with Olivier, Bordeaux
- Geissler, 1996, 2002, Diploma and PhD
- Hulpke, 1996
- McKay, 1995 (short cosets)
- F-Klüners, 2003

Girstmeyer, Yokoyama (97, p-adic), Valibouze, Renoult

Variants

I present Stauduhar's method, based on the numerical evaluation of *relative invariants* which are roots of *relative resolvents*..
Original: numeric $=$ complex.
There are symbolic methods based on factoring so called *resolvent polynomials* which can be absolute or relative (linked to invariants) or other (Soicher)
Galois ideals capture different aspects, splitting field as a 0-dim. ideal.
Recognition algorithms that only find the *abstract group*, not any explicit action.
Solvability can be decided in poly-time (Susan Landau).

Step 1: Stauduhar's basic tool is now:
Suppose G is known to contain $\mathrm{Gal}(f)$ and U is a maximal subgroup of G. Then we need $I \in \mathbb{Z}[x_1, \ldots, x_n]$ s.th. $\mathrm{Stab}_G(I) = U$ and s.th.

$$R := \prod_{\sigma \in G//U} (t - I^\sigma(\alpha_1, \ldots, \alpha_n)) \in \mathbb{Z}[t]$$

is square-free (if not, then replace α_i by $T(\alpha_i)$ for some suitable polynomial $T \in \mathbb{Z}[t]$. This is called a Tschirnhaus transformation. There are only finitely many bad ones)
Then if $I^\sigma(\alpha_i \ldots, \alpha_n) \in \mathbb{Z}$ is a root, then $\mathrm{Gal}(f) \subseteq U^\sigma$ and we re-start the procedure. If there are no roots for no maximal subgroup, then G is the result.

Change of perspective

We know $\mathrm{Gal}(f) \leq G$. That implies that $\mathbb{Z}[x_1, \ldots, x_n]^G \to \mathbb{Z} : x_i \mapsto \alpha_i$ is well defined.
Let $U < G$ be maximal and $I \in \mathbb{Z}[x_1, \ldots, x_n]$ such an invariant, then, as rings:

$$\mathbb{Q}[x_1, \ldots, x_n]^U = \mathbb{Q}[x_1, \ldots, x_n]^G[I]$$

ie. I is a primitive element for the invariant ring of U as an extension of the invariant ring of G.
Thus if $I(\alpha_1, \ldots, \alpha_n) \in \mathbb{Z}$, then $\mathbb{Z}[x_1, \ldots, x_n]^U$ also maps to $\mathbb{Z}$.
On termination we have explicitly constructed the invariant ring of the Galois group over the the ring generated by the elementary symmetric functions. Under evaluation this is the largest subring mapping to $\mathbb{Z}$.

How do we prove $\theta := I^\sigma(\alpha_1, \ldots, \alpha_n) \in \mathbb{Z}$? Given that we have only a p-adic approximation.
We know more: given f, we have bounds on the complex roots, this implies bounds on the possible complex size of θ if we were to choose complex roots.
Also, by assumption θ is algebraic of degree $\leq (G : U)$ and θ looks like $\mu \in \mathbb{Z}$. So fix μ.

The task is to prove that $\theta - \mu = 0$. We know (see above)

- bounds on the possible complex size of the difference
- bounds on the degree
- a lower bound on the norm: we have $\theta - \mu = 0 \bmod p^k$ for the precision that we are currently using.

And that is enough: $N(\theta - \mu)$ is divisible by p^k, thus we have lower bound (if non-zero). This implies a lower bound in $\|.\|_e$ as above (arithmetic-geometric-means), the lower bound is monotonous in k but bounded from above by the complex bound (which is independent of the precision!) Thus if k is large enough, the norm has to be 0.

Problem(s)

- $(G : U)$ might be huge
- Sym(n) is not a good starting group
- the index is too large
- the invariants are hard to find
- maximal subgroups are non-trivial to compute

Subfields

Let $K = \mathbb{Q}[\alpha] = \mathbb{Q}[t]/f$ the stem field of f, and α some root of f.
If k has a subfield, $K \supset k = \mathbb{Q}[\beta]$ for some $\beta = h(\alpha)$, then this implies a partitioning of the roots of f: two roots are in the same partition iff $h(\alpha_i) = h(\alpha_j)$.
This implies that the Galois group

- is imprimitive
- has to have this partition as a block-system

Thus, we get G has to be contained in some suitable wreath-product.
If we have more than one subfield, we use an intersection. (also hard to compute)

...

Two question:

- how to get subfields without the Galois group?
- what about primitive groups?

Subfields - very naive

Subfields correspond to block-systems, hence partitions. So:

- enumerate all suitable partitions
- for each partition P test if P defines a subfield

This is Dixon's original subfield algorithm, implemented by J. Klüners.

Partition to Subfield

Let $P = (B_1, \ldots, B_k)$ be the potential block system, a partitioning of the roots fixed at the beginning. Then $|B_i| = |B_j|$ for all i, j.
Fact: if all $\beta_i = \prod_{\alpha \in B_i} \alpha$ are pairwise different, then they are the p-adic conjugates of a primitive element of the subfield. If not, then there is some $x \in \mathbb{Z}$ s.th. $\beta_i = \prod_{\alpha \in B_i} \alpha + x$ works.
(We can also try to sum rather than multiply)

From the p-adic conjugates of a primitive element we can easily (Newton) get the p-adic approximation of the minimal polynomial. From the formula for β we get bounds for the coefficients, hence the precision necessary to get the exact (potential) polynomial. Good, but does it define a subfield? Here we also need the "subfield–polynomial", we need $h \in \mathbb{Q}[x]$ s.th. $\beta = h(\alpha)$.

Subfield Polynomial

Fundamentally, finding h is an interpolation problem:

$$h(\alpha_k) = \beta_i$$

for all $\alpha_k \in B_i$, the same block.
From here we also get bounds and some algorithm, but $h \in \mathbb{Q}[x]$ so we have denominators.
If f is monic and integral, then α is an algebraic integer, hence β, hence we can use all our denominator techniques to get results.
However, we can do even better:

Co-different

Let f be monic, irreducible and integral. Then

$$\mathcal{C} = \{\gamma \in K | \, \mathrm{Tr}(\gamma\mathbb{Z}[\alpha]) \subseteq \mathbb{Z}\}$$

is the *codifferent* of $\mathbb{Z}[\alpha] = \mathbb{Z}[t]/f$.
We know

- $\mathcal{O} \subseteq \mathcal{C} = 1/f'(\alpha)\mathbb{Z}[\alpha]$
- Let $g = \sum g_i t^i = f/(t-\alpha)$, then $\mathrm{Tr}(g_i/f'(\alpha)\alpha^j) = \delta_{i,j}$

Thus we can compute the coefficients of $f'(\alpha)h$ via trace and multiplication with the dual basis.
This too gives *good* and *easy* bounds.

Klüeners - van Hoeij

All partitions? That is a lot!
New idea, due to van Hoeij again, more generically 1st: Let $f \in k[t]$ be irreducible and K any field containing a root of f.
Then

$$f = \prod f_i \in K[t]$$

wlog, $f_1 = t - \alpha$.
For each i, define $K_i = K[x]/f_i$ and

$$\phi_i : k = \mathbb{Q}[\alpha] \to K_i = K[x]/f_i : \alpha \to x$$

This is a field embedding. Then

$$L_i := \ker(\phi_i - \mathrm{Id})$$

Is a subfield. All subfields are intersections of the L_i.

...

This is slow, when directly applied ($K = k[t]/f$) due to

- factoring (over a number field)
- linear algebra (over a number field)

But K can be any field containing a root, so choose an unramified p-adic field.
Trade-off: the linear algebra is now hard:

- over $\mathbb{Q}_p$
- imprecise.

...

Klüeners and van Hoeij improved this: they they to find a *small* basis for the subfield contained in the *co-different*.
Small: small coefficients in the linear combination
Codifferent: no-denominators
Small: find using LLL, again.
From the basis (potential basis, due to errors), find the partition and use above to verify the partition.

Primitive Groups

If there are no subfields, then...
we can try to use

$$R = \prod (t - \alpha_i - \alpha_j)$$

and its factorisation as a starting point.
Fact: take "any" polynomial constructed from the roots that is rational. Then any non-trivial factor comes from the stabilizer of the roots in this factor.
Thus we get potentially better starting groups than S_n / A_n this way.
(Alt(n), Sym(n) are usually detected while trying to find a nice p-adic field. They are separated by the discriminant)

The large index problem

Here only a partial solution:
If we use an unramified p-adic extension, then we have a non-trivial automorphism σ, the Frobenius for this extension.
This is automatically an element of the Galois group.
Hence we need to only consider groups containing σ.
Translate: we need coset representatives τ s.th. $\sigma \in U^\tau$ as only those are possible for the Galois group.
This is a *huge* reduction.
However, the *requirement* for Stauduhar's method was to *know* that R is square-free. This can now not be verified...

Invariants

We need, for $U < G$ maximal, some G-relative U-invariant, ie. some I s.th. $\mathrm{Stab}_G(I) = V$.
This is trivial:

$$I = \sum_{\sigma \in U} (x_1 x_2^2 \cdots x_n^n)^\sigma$$

is such an invariant.
This is useless.
Example:

$$\sqrt{\mathrm{disc}} = \prod_{i<j} (x_i - x_j)$$

This is a $\mathrm{Sym}(n)$-relative $\mathrm{Alt}(n)$-invariant.
In this (factored) form, we can evaluate in $O(n^2)$ multiplications.
Expanded, we need $O(2^n)$ operations...
So the goal is to find invariants in some form that can be

Resolvents

Let I be a G-relative U-invariant, then

$$R := \prod_{\sigma \in G//U} (t - I^\sigma(\alpha_1, \ldots, \alpha_n)) \in \mathbb{Z}[t]$$

is called a G-relative resolvent polynomial.
Suppose R has a factor g, then, since G acts on the roots of R, the roots of g are a single G-orbit., thus the Galois group has to contain the stabiliser of this orbit.
In the simple case of g being linear, hence "a root", this is Stauduhar: the stabilizer is just U^σ by construction.
But ...
... what about other factors?

...

Other factors?

Problem(s):

- we actually need R, so far we've only looked for roots
- then we need to factor, we will always have superfluous local factors. Recall $\deg R = G : V$ is large
- then we need the roots and the stabilizer. This is hard group theory

But:

This can help a lot. In fact, entire algorithms are designed around

- (symbolically) computing resolvents
- factoring
- intersecting stabilizers.

For moderate degree, this works well.

How to use it?

We can work in the normal closure!
Let U be any subgroup of G and I s.th. $\mathrm{Stab}_G(I) = U$, then I is a primitive element for $\mathbb{Z}[x_1, \ldots, x_n]^U$ over $\mathbb{Z}[x_1, \ldots, x_n]^G$, hence, generically, $I(\alpha_1, \ldots, \alpha_n)$ is a primitive element for the fixed field of U.
Since we have the explicit G-action we can

- test primitivity (compute all conjugates (I^σ for $\sigma \in G//U$) and check there different)
- get bounds for the coefficients of the min. poly
- compute the min. poly

At worst, we might have to do a Tschirnhaus transformation to ensure the primitivity.

Applications

- change the "representation": the same abstract group can be realised as a permutation group of different degrees (e.g. $\mathrm{Sym}(3)$ can be on 3 points (cubic polynomial) or 6 points.
- compute any subfield of the normal closure explicitly
- for any chain of subgroups, obtain the corresponding tower of fields
- explicitly solve by radicals

Function Fields

Now, try the same for $K = \mathbb{Q}(t)[x]/f$ or $\mathbb{Q}[t,x]/f$. Hilbert irreducibility states that for almost all choices for $t_0 \in \mathbb{Z}$, the Galois group of K and of $k = \mathbb{Q}[x]/f(t_0, x)$ are "the same", (this then also includes subfields!)
So:

- choose a p-adic splitting field for k and compute the subfields/ Galois group
- lift the p-adic roots for k to $\mathbb{Q}_q[[t]]$ roots for K
- verify that the subfields (block systems) for k work as well for K
- verify that the invariants used later also evaluate to s.th. "integral"

...

Problem(s):

- now the bounds have to come from $\mathrm{complex}$ power series
- the integrality (for the subfield polynomial) is difficult: the codifferent shows that we are in

$$\mathsf{IntCls}(\mathbb{Q}[t], K) \subseteq \frac{1}{f'(x)}\mathbb{Q}[t][x]/f$$

we would need a bound for

$$\mathsf{IntCls}(\mathbb{Z}[t], K) \subseteq \frac{1}{f'(x)}\mathbb{Z}[t][x]/f$$

which is hard: $\mathbb{Z}[t]$ is no PID, the module theory collapses.