

Local Computations for Global Problems: Factoring

Claus Fieker

MPI Leipzig, August 31, 2021.

Intro/ Basics

We start by factoring $f \in \mathbb{Z}[t]$ monic and square-free, to discuss (briefly) $f \in \mathbb{Z}[t]$ non-monic, to move to $f \in (\mathbb{Z}[x]/g)[t]$ monic, square-free (for some $g \in \mathbb{Z}[x]$ monic and irreducible) to finally move to $f \in \mathbb{Z}[t, x]$ and a factorisation over \mathbb{C} .

The basic idea is always the same:

- factor modulo p (or similar)
- lift to p^k , k large
- try combinations and move back to the ground ring

Gauss-lemma shows that all divisors of the monic and integral f are monic and integral. That helps.

Assume that all polynomials mentioned are square-free. (Mostly also monic and integral for a suitable notion of integral)

Zassenhaus

Pick (any) prime p s.th. f is square-free mod p , then

$$f \equiv \prod \tilde{f}_i \pmod{p}$$

(which is effective via Berlekamp or Cantor-Zassenhaus or variations)

Use Hensel lifting to obtain a factorisation

$$f \equiv \prod f_i \pmod{p^{(2^l)}}$$

for some l to be determined by some (well known) bound on the size of coefficients of factors of f .

Note: all those bounds are actually bounds on *complex polynomials*, not integral ones!

...

Then enumerate all subsets S of $\{f_i|i\}$ by size, smallest first. For each such S test if a symmetric lift of $\prod_{i \in S} f_i$ divides f .

If yes: remove S from the set of f_i and output the factor. This will find all irreducible factors.

And of course, if f induces a Galois group of C_2^n , this has a cost of $2^{2^{n-1}}$ possibly.

van Hoeij et. al.

What is the real problem: Find as small as possible set S of factors, such that the product, after lifting to \mathbb{Z} has small coefficients. Or Find $\epsilon_i \in \{0, 1\}$ s.th. $\prod f_i^{\epsilon_i}$ is small.
Idea: make this linear.

Original idea of Mark: $f = \prod f_i$ over \mathbb{Q}_p , $f = \prod (t - \alpha_i)$ over some suitable (unramified) extension.

Now let $g|f$ for some $g \in \mathbb{Q}_p[t]$, then $g = \prod_{i \in G} (t - \alpha_i)$.

For $k > 0$ define $t_k(g) = \sum_{i \in G} \alpha_i^k$.

The main theorem on symmetric functions states that $t_k(g) \in \mathbb{Z}_p$ since $g \in \mathbb{Z}_p[t]$.

Furthermore, $t_k(gh) = t_k(g) + t_k(h)$ for all coprime $g, h|f$ and, if $g|f$, $g \in \mathbb{Z}[t]$ then

$$t_k(g) \in \mathbb{Z}$$

for all k .

We know f , hence know a bound for the absolute value of the complex roots of f , say B .

Thus for all divisors $g|f$ we have $|t_k(g)| \leq \deg g B^k \leq n B^k$ uniformly for all factors. (This is not best possible)

We also know $g = \prod f_i^{\epsilon_i}$ for $\epsilon_i \in \{0, 1\}$ suitable, hence $t_k(g) = \sum \epsilon_i t_k(f_i)$. By the bound above, independently of the precision, this needs to have a “small” lift in \mathbb{Z} .

So build a lattice. Start with $A = (t_k(f_i))_{i,k}$, then append a diagonal matrix with p^k and prepend the identity to keep track of operations. For large precision (larger than the bound), the only really short vectors in the lattice start with the 0, 1 in front, and have the lifted power-sums elsewhere. Newton relations compute the power-sums from the polynomial directly - without explicit roots, so easy.

Problem: this is a large lattice, both in terms of the number of power sums as well as the precision. Prohibitively large. The solution: gradual feeding, even removing rows as you go (Ill-with-removal, Novacin et. al)

There is a second form of linearisation: logarithmic derivatives. For $g|f$ map

$$L : g \rightarrow \frac{g'}{g} f$$

This too satisfies $L(gh) = L(g) + L(h)$. Instead of the power sums above, experimentally using trailing and leading terms of $L(f_i)$ allows to use smaller dimensional lattices. There are sophisticated bound on the coefficients of $L(g)$ also available.

Next observation: thinking p -adic numbers as power-series in p , we are looking for linear combinations (of $t_k(g)$ or $L(f_i)$) with few non-zero terms, the series needs to be finite. We do not care about the initial terms, so we can try to divide the coefficients by a large power of p to lower the bound to 1-ish.

This can be tuned even further, of course, and combined with Zassenhaus. (See Hart, Novocin: ISSAC)

Intro

Let $K = \mathbb{Q}[t]/\mu$ for some monic, integral, irreducible $\mu \in \mathbb{Z}[t]$ and $f \in K[x]$ monic. Wlog (?) assume $f \in (\mathbb{Z}[t]/\mu)(x)$ to avoid all denominator.

Problem: for $g|f$, $g \in K[x]$, g monic we do not have $g \in (\mathbb{Z}[t]/\mu)[x]$

$$x^3 - 2 - \sqrt{5} = \left(x - \frac{1 + \sqrt{5}}{2}\right) \left(x^2 + x \frac{1 + \sqrt{5}}{2} + \frac{3 + \sqrt{5}}{2}\right)$$

Problem: what is a symmetric lift when dealing with extensions?

Let $\mathcal{O} = \text{IntCls}(\mathbb{Z}, K)$ the maximal order, then at least for $g|f$, we have $g \in \mathcal{O}[x]$.

However, the computation of \mathcal{O} is expensive, sub-exponential in the discriminant, so too expensive.

Possible solutions

- \mathcal{O} might be already known, so use it
- $\mathcal{O} \subseteq \frac{1}{d}\mathbb{Z}[t]/\mu$ for $d = \text{disc}(\mu)$
- $\mathcal{O} \subseteq \frac{1}{\mu'(\alpha)}\mathbb{Z}[\alpha]$ for $\mu(\alpha) = 0$ (co-different, Kronecker representation)
- use rational reconstruction and compute the actual denominators

If possible, we use 1, but 3 is better (much better) than 2.

Bounds

K embeds into \mathbb{C} in n -different ways. Let $g|f$, $g \in K[x]$, then for all embeddings $g^{(i)}|f^{(i)}$ and the same bounds as in the \mathbb{Q} case can be used to obtain bounds for (the embeddings of the coefficients of) g .

Let b be a coefficient of the (unknown) g . Then, by above, $|b^{(i)}| \leq B$ for some B . Also $b = \sum a_i \alpha^i$ for $a_i \in \mathbb{Q}$ as the powers form a basis. We want

- $a_i \in \mathbb{Z}$
- bounds for $|a_i|$

There is a generic trick for the bounds: on K as a n -dimensional vector space we have 2 different norms (even more) to measure size $\|\alpha\|_e^2 = \sum |\alpha^{(i)}|^2$ and $\|\alpha\|_c^2 = \sum a_i^2$. The 1st is canonical, usually called T_2 -norm, while the 2nd is what we need. As K is fin. dim. there are constants c_1 and c_2 s.th.

$$c_1 \|\alpha\|_e \leq \|\alpha\|_c \leq c_2 \|\alpha\|_e$$

The complex bounds immediately give bounds on $\|\alpha\|_e$, hence we get bounds on $\|\alpha\|_c$ hence on a_i . (Fieker, Friedrichs ANTS-2000) c_1 and c_2 can be computed or estimated easily. Of course, this can be done for any basis not just the power basis.

Zassenhaus

Pick a nice prime ideal \mathfrak{p} , unramified, not dividing any obvious denominator and s.th. f is square-free over the residue field. Then factor f over the residue field and lift modulo a suitable power of \mathfrak{p} . Of course, this means factoring f over some unramified extension of \mathbb{Q}_p .

We need a fact: As $k \rightarrow \infty$, the shortest vector in \mathfrak{p}^k grows as well. Proof: $\beta \in \mathfrak{p}^k$ has norm $|N(\beta)| \geq N(\mathfrak{p})^k$. The arithmetic-geometric means inequality shows

$$|N(\beta)| \leq \left(\frac{\|\beta\|_e}{n}\right)^n$$

so $\|\beta\|_e \geq nN(\mathfrak{p})^{k/n}$ as required.

Note: this can (and will) be used in conjunction to the norm change constants to estimate the precision.



Lift

Problem: lift back to $\mathbb{Z}[t]/\mu$. If \mathfrak{p} is a prime of degree 1, then for all $\beta \in \mathbb{Z}[t]/\mu$ and all k there is a $b \in \mathbb{Z}$ s.th. $\beta \equiv b\mathfrak{p}^k$.

So, given $b \in \mathbb{Z}$, find a (small) $\beta \in \mathbb{Z}[t]/\mu$ s.th. $\beta \equiv b \pmod{\mathfrak{p}^k}$.

LLL again: if γ_i form a \mathbb{Z} -basis of \mathfrak{p}^k , put the coefficients into a matrix and LLL-reduce, call the resulting matrix C and set D as the 1st row of C^{-1} . Then $b - \text{round}(Db)C$ is a small(ish) representative. If \mathfrak{p} is not of degree 1, we need more than one row of C^{-1}

Armed with this, the Zassenhaus method works “exactly as before”.

van Hoeij

Also “the same”. 1st approach: use the gradual feeding, the logarithmic derivatives, and, in the matrix with the coefficients replace each coefficient in \mathbb{Z} by a vector of length n .

The p^k diagonal matrix is replaced by a block diagonal matrix with a basis of \mathfrak{p}^k .

Otherwise: business as usual.

Improvement: use a LLL basis of \mathfrak{p}^k

Improvement: scale each column by the inverse of the LLL basis matrix, then clear denominators. Result: the block diagonal matrix becomes again a diagonal matrix with p^k .

Improvement: scale the coefficients by p^l again as, same as before, we only want combinations where the trailing coefficients are 0.

We are not interested in the initial terms, only the end.

Remarks

Any p^k based algorithm in $\mathbb{Z}/$ number fields can be turned into a p^k algorithm using this lifting technique. Problem is that the LLL becomes really expensive, this is the bottleneck. Damien Stehle looked into this and told me that the input is close to worst possible for LLL.

We use this for

- n -th roots
- roots
- modular Groebner bases
- determinants
- gcd and resultants

In certain ranges of the parameters.

$\mathbb{Q}[t, x]$

We will only look into the bivariate case as the general case is “just” a question of lifting the other variables back in. The “base-case” is the bivariate one.

Idea: factor in $\mathbb{Q}(t)[x]$...

Assume $f \in \mathbb{Z}[t, x]$ is square-free and monic as a polynomial in $\mathbb{Q}(t)[x]$ (can be weakened). Pick a evaluation point t_0 s.th.

$f(t_0)(x) \in \mathbb{Z}[x]$ is square-free, (wlog $t_0 = 0$), then a prime p s.th. $f(0, x)$ is square-free modulo p .

From here, we compute a factorisation over $\mathbb{Q}_p[[t]]$ to recover the one in $\mathbb{Q}(t)$.

We have a 2-dim lifting problem: we need t and p^k ! So far we reduced f modulo $\langle t - t_0, p \rangle$ and have a factorisation modulo this ideal. We also have the co-factors to prove the factors to be coprime over the residue field.

We can lift this factorisation to any precision in t and p desirable!

Any factor $g|f$ will have $\deg_t(g) \leq \deg_t(f)$, so let's write $f = \prod f_i \bmod \langle t^d, p \rangle$ for $d > \deg_t(f)$ suitable.

Now, as before, we are looking for subsets of the f_i s.th. the $\sum \epsilon_i t_k(f_i)$ are small, which means: $\deg_t \sum \epsilon_i t_k(f_i) \leq \deg_t f$. As there is no carry in the addition of polynomials, this is a pure linear equation for each coefficient in degree $> \deg_t(f)$: all of them need to be 0.

Instead of try-and-error (Zassenhaus) or LLL (van-Hoeij) we use Gauss over \mathbb{F}_p to find the combination! We can apply gradual feeding as before but lin. algebra over finite fields is “fast”.

Caveat: occasionally a coefficient in \mathbb{Q}_p will be zero in \mathbb{F}_p , so we might need more t . Statistically, this is rare.

Then we lift the factorisation to p^k (and the minimal t -precision) to obtain the factors. There are bounds of the coefficients of factors of multivariate polynomials known as well.

$$K[t, x]$$

For number fields we follow the same idea as before: take the \mathbb{Q} method and replace p by a nice ideal \mathfrak{p} . The bounds from above are for complex polynomials, so we get bounds for the coefficients of the coefficients as before. The factors of an integral monic f are again in the maximal order, hence the same discussions as before about the denominators apply.

Non-Monic

If f is not monic in x , life is more interesting: we can

- do a substitution to make it monic
- factor, as we do, over $\mathbb{F}_p[[t]]$ into monic factors and apply rational reconstruction to write the coefficients as rational functions in t . The denominators, when cleared, give the matching factorisation of the leading coefficient. This too can be lifted...

$\mathbb{C}[t, x]$ - the absolute factorisation

Given $f \in \mathbb{Z}[t, x]$ (monic in x , irreducible) we want to factorise over \mathbb{C} . Trivial example: $x^2 - 2t^2$ is irreducible over \mathbb{Q} , but splits over $\mathbb{Q}[\sqrt{2}][t, x]$.

Again, we have bounds on the coefficients (complex bounds), we know the integrality, but now we're missing the field. So step 1: make the field as large as possible:

Find an evaluation point t_0 s.th. $f(t_0, x)$ is irreducible, wlog $t_0 = 0$ and a prime p s.th. $f(0, x)$ is square-free with a “small degree” splitting field over \mathbb{F}_p . Essentially we are computing the roots of f as power-series over an unramified p -adic extension of small degree...

Step 2: compute the roots over the residue field: $f = \prod (x - \alpha_i)$
for $\alpha_i \in \mathbb{F}_q[[t]] \bmod t^k$

As above, use power-sums to find the combinations of α_i that are polynomials (and not power-series). As before, this are linear equations over the finite field.

Note: we know divisibility conditions on the size of the partitioning of the roots.

Step 3: compute the “true” factors over the residue field. In general, the coefficients will live in a smaller field, so find this. Also: comparing the factors obtained we find

- the degree of the number field we are looking for
- a combination of the coefficients that will be a primitive element of the field
- as above, the factorisation of the leading coefficient

Step 4: compute a lifting over the smaller unramified extension up to a sufficient precision of the true factors.

This gives the \mathbb{Q}_q -adic conjugates of a primitive element of the target number field as well.

Using power-sums and Newton-relations, write down the minimal polynomial. If this fails, we need to increase the precision. Since we have bounds on the coefficients of the factors, we have bounds on the primitive element and the minimal polynomial as well.

Step 5: the coefficients:

Once we have the primitive element, write all the other coefficients as linear combinations of powers (this involves, via coefficient comparison, linear algebra over \mathbb{Q}_p or \mathbb{Z}_p . If we make the denominators disappear as before, we just symmetric lift the coefficients back to \mathbb{Z})

Final Step 6: test if the 1st factor is a divisor over the number field. If not, increase the precision or choose a different p if precision is too high (we have bounds on the coefficients!)