Polynomial Factorisation using Drinfeld Modules

Anand Kumar Narayanan

Polynomial Factorisation using Drinfeld Modules

ANAND KUMAR NARAYANAN

The arithmetic of Drinfeld modules has recently yielded novel algorithms for factoring polynomials over finite fields; a computational problem with applications to digital communication, cryptography and complexity theory. We offer a gentle invitation to these developments, assuming no prior knowledge of Drinfeld modules.

Factoring polynomials modulo a prime

Let $\mathbb{F}_q$ denote the finite field of integers modulo an odd prime number q. Polynomials $\mathbb{F}_q[x]$ over $\mathbb{F}_q$ share striking analogies with integers, yet we begin with an algorithmic distinction. While factoring integers remains a notoriously difficult problem, factoring polynomials in $\mathbb{F}_q[x]$ is long known to be easy, at least with access to randomness.

Polynomial factorisation over finite fields is not a mere curiosity, but has many applications. In number theory, finite fields arise as residue fields of global fields such as number fields. While determining the splitting of a prime in a number field, one factors a polynomial defining the number field modulo the prime. Several instances of polynomial factorisation appear while factoring integers using quadratic/number field sieve algorithms or while performing index calculus to compute discrete logarithms, both foundational problems in analysing the security of cryptographic systems. In digital communication, polynomial factorisation aids in the construction of certain error correcting codes (BCH and cyclic redundancy codes), structures vital to reliable transmission of information in the presence of noise.

Let us recount a simple polynomial factorisation algorithm. We are given a monic $\mathfrak{f}(x) \in \mathbb{F}_q[x]$ of degree n whose factorisation into irreducible polynomials $\mathfrak{f}(x) = \prod_{i=1}^{m} \mathfrak{p}_i(x)$ is sought. Assume $\mathfrak{f}(x)$ is square free, that is the $\mathfrak{p}_i(x)$ are distinct. This is without loss of generality for there are algorithms to rapidly reduce to this special case. It takes $n \log_2 q$ bits to write down $\mathfrak{f}(x)$ and we seek algorithms that run in time polynomial in n and $\log q$. Berlekamp was the first to show there is a randomized polynomial time algorithm, but we follow a different two step process.

Distinct degree factorisation

The first step, known as distinct degree factorisation, decomposes $\mathfrak{f}(x)$ into factors each of which is a product of irreducible polynomials of the same degree. By Fermat's little theorem, $x^q - x = \prod_{a \in \mathbb{F}_q} x - a$. To extract the product of linear factors of $\mathfrak{f}(x)$, take the greatest common divisor $\gcd(x^q - x, \mathfrak{f}(x))$. To extract products of degree two factors, degree three factors and so on iteratively, look to the succinct expression

$$x^{q^d} - x = \prod_{\mathfrak{p}:\deg(\mathfrak{p})|d} \mathfrak{p}(x)$$

for the product of monic irreducible polynomials $\mathfrak{p}(x)$ of degree dividing d. At the d^{th} iteration, with smaller degree factors already removed, $\gcd(x^{q^d} - x, \mathfrak{f}(x))$ yields the product of degree d irreducible factors.

Care in handling x^{q^d} is required for its degree is exponential in $n \log q$. All we need is $x^{q^d} \mod \mathfrak{f}(x)$, easily accomplished by a sequence of q^{th} powers modulo $\mathfrak{f}(x)$, each performed by repeated squaring. Better still, $x^{q^d} \mod \mathfrak{f}(x)$ can be rapidly computed with a fast algorithm to compose two polynomials modulo $\mathfrak{f}(x)$ in concert with q^{th} powers. Kaltofen and Shoup devised an ingenious improvement over this naive iteration resulting, in a significant speed up. The Kaltofen–Shoup algorithm implemented using the modular composition algorithm of Kedlaya–Umans performs distinct degree factorisation with run time exponent $3/2$ in the degree n and 2 in $\log q$.

Equal degree factorisation

Distinct degree factorization leaves us with the problem of factoring polynomials all of whose irreducible factors are of the same known degree d. All known algorithms for this task with polynomial runtime in $\log q$ are randomized. Even for the simplest case of factoring a quadratic polynomial into two linear factors, no unconditional deterministic polynomial time algorithms are known. It is closely related to the problem of finding a quadratic nonresidue modulo a given large prime.

The following randomized algorithm can be traced backed to ideas of Gauss and Legendre. For a uni-

formly random $\mathfrak{a}(x)$ of degree less than n,

$$\gcd\left(\mathfrak{a}(x)^{\frac{q^d-1}{2}} - 1\,,\ \mathfrak{f}(x)\right)$$

gives a random factorisation of degree d irreducible factors. This follows since raising $\mathfrak{a}(x)$ to the $(q^d-1)/2$-th power modulo a degree d irreducible polynomial results in either 1 or -1, depending on whether $\mathfrak{a}(x)$ reduces to a quadratic residue or not. Remarkably, this computation can be performed with run time exponent 1 in the degree n using an algorithm of von zur Gathen and Shoup implemented with the aforementioned Kedlaya–Umans modular composition.

In summary, the best known polynomial factorisation algorithms have run time exponent $3/2$ in the degree with the bottleneck being distinct degree factorisation. To lower this exponent is an outstanding problem. In fact, to lower this exponent, it suffices for there to be an algorithm that merely estimates the degree of some irreducible factor.

Drinfeld modules and polynomial factorisation

The use of Drinfeld modules to factor polynomials over finite fields originated with Panchishkin and Potemine [3]. Drawing inspiration from Lenstra's elliptic curve integer factorization, they recast the role of the group of rational points on random elliptic curves modulo primes with random finite Drinfeld modules.

We describe three Drinfeld module based algorithms for polynomial factorisation. The first two were devised in [2] and the third in [1]. The first estimates factor degrees using Euler–Poincaré characteristics in hopes of speeding up distinct degree factorisation. The second is a Drinfeld analogue of Lenstra's algorithm, closely related to the aforementioned algorithm of Panchishkin and Potemine [3]. Our exposition begins with a short account of finite Drinfeld modules followed by Euler–Poincaré characteristic and Frobenius distributions, important ingredients in the first two algorithms. The third algorithm involves Drinfeld modules with complex multiplication with an analogue of Deligne's congruence playing a vital role. It is also the fastest of the three algorithms, with runtime complexity matching the best known algorithms, in theory and practice. The hope is, the rich arithmetic of Drinfeld modules will inform new algorithms to beat the $3/2$ exponent barrier.

Finite Drinfeld modules

Drinfeld introduced the modules bearing his name as an analogue of elliptic curve complex multiplication theory. He in fact called them elliptic modules. Drinfeld modules and their generalisations have played a crucial role in the class field theory of function fields and in proving global Langlands conjecture over function fields for GL_n. We settle for a concrete simple notion of Drinfeld modules sufficient for our context. Throw the q^{th} power Frobenius σ into $\mathbb{F}_q[x]$ resulting in $\mathbb{F}_q(x)\langle\sigma\rangle$, the skew polynomial ring with the commutation rule $\sigma\,\mathfrak{u}(x) = \mathfrak{u}\,(x)^q\sigma$,, for all $\mathfrak{u}(x) \in \mathbb{F}_q[x]$. A rank-2 Drinfeld module over $\mathbb{F}_q(x)$ is (the $\mathbb{F}_q[x]$ module structure on the additive group scheme over $\mathbb{F}_q(x)$ given by) a ring homomorphism

$$\begin{aligned}\phi : \mathbb{F}_q[x] &\longrightarrow \mathbb{F}_q(x)\langle\sigma\rangle\\ x &\longmapsto x + \mathfrak{g}_\phi(x)\,\sigma + \Delta_\phi(x)\,\sigma^2\end{aligned}$$

for some $\mathfrak{g}_\phi(x) \in \mathbb{F}_q[x]$ and nonzero $\Delta_\phi(x) \in \mathbb{F}_q[x]$.

To better understand the map, it is instructive to compute by hand as to where x^2, x^3 and so on, get mapped to. By design, $\mathfrak{b}(x)$ maps to a polynomial in σ with constant term $\mathfrak{b}(x)$,

$$\mathfrak{b}(x) \longmapsto \phi_\mathfrak{b} := \mathfrak{b}(x) + \sum_{i=1}^{2\deg(\mathfrak{b})} \phi_{\mathfrak{b},i}(x)\sigma^i.$$

Consider an $\mathbb{F}_q[x]$ algebra $\boldsymbol{M}$. In our algorithms to factor $\mathfrak{f}(x)$, $\boldsymbol{M}$ will often turn out to be $\boldsymbol{M} = \mathbb{F}_q[x]/(\mathfrak{f}(x))$. One way to make an $\mathbb{F}_q[x]$ algebra $\boldsymbol{M}$ into an $\mathbb{F}_q[x]$ module is to retain the addition and scalar multiplication but simply forget the multiplication. The Drinfeld module ϕ endows a new $\mathbb{F}_q[x]$ module structure to $\boldsymbol{M}$ by twisting the scalar multiplication. For $\mathfrak{b}(x) \in \mathbb{F}_q[x]$ and $a \in \boldsymbol{M}$, define the scalar multiplication

$$\mathfrak{b}(x) \star a := \phi_\mathfrak{b}(a) = \mathfrak{b}(x)a + \sum_{i=1}^{2\deg(\mathfrak{b})} \phi_{\mathfrak{b},i}(x)a^{q^i},$$

where the arithmetic on the right is performed in the $\mathbb{F}_q[x]$ algebra $\boldsymbol{M}$. Let $\phi(\boldsymbol{M})$ denote the new $\mathbb{F}_q[x]$ module structure thus endowed to $\boldsymbol{M}$.

Euler–Poincaré characterisitic

Cardinality is an integer valued measure of the size of a finite abelian group (equivalently, a finite $\mathbb{Z}$-module). A convoluted definition is to assign as the cardinality of a cyclic group of prime order the corresponding prime, and for cardinality of finite abelian groups that sit in an exact sequence to be multiplicative. The Euler–Poincaré characteristic χ is an $\mathbb{F}_q[x]$-valued

cardinality measure of a finite $\mathbb{F}_q[x]$ module defined completely analogously. For a finite $\mathbb{F}_q[x]$ module A, $\chi(A) \in \mathbb{F}_q[x]$ is the monic polynomial such that:

- If $A \cong \mathbb{F}_q[x]/(\mathfrak{p}(x))$ for a monic irreducible $\mathfrak{p}(x)$, then $\chi(A) = \mathfrak{p}(x)$.
- If $0 \to A_1 \to A \to A_2 \to 0$ is exact, then $\chi(A) = \chi(A_1)\chi(A_2)$.

For the $\mathbb{F}_q[x]$ module $\phi(\mathbb{F}_q[x]/(\mathfrak{f}(x)))$ featuring in our algorithms, the Euler–Poincaré characteristic $\chi(\phi(\mathbb{F}_q[x]/(\mathfrak{f}(x))))$ has a simple linear algebraic interpretation: the characteristic polynomial of the linear map ϕ_x on $\mathbb{F}_q[x]/(\mathfrak{f}(x))$. In particular, it is a degree n polynomial that can be computed efficiently.

Frobenius distribution of Drinfeld modules

Let us put our newly defined $\mathbb{F}_q[x]$ modules and cardinality measure χ to use. Take an elliptic curve E over the rational numbers and reduce it at a prime p. The $\mathbb{F}_p$-rational points $E(\mathbb{F}_p)$ famously form a finite abelian group with cardinality $p + 1$ up to an error determined by the Frobenius trace $t_{E,p}$. The Hasse bound, considered the Riemann hypothesis for elliptic curves over finite fields, asserts that $|t_{E,p}| \leq 2\sqrt{p}$. Thereby,

$$\#(E(\mathbb{Z}/(p))) = p + 1 - \underbrace{t_{E,p}}_{-2\sqrt{p} \leq \ \leq 2\sqrt{p}}.$$

Gekeler established the following Drinfeld module analogue. Take a Drinfeld module ϕ, a monic irreducible polynomial $\mathfrak{p}(x)$, and consider the resulting $\mathbb{F}_q[x]$ module $\phi(\mathbb{F}_q[x]/(\mathfrak{p}(x)))$. Its Euler–Poincaré characteristic equals

$$\chi\left(\phi\left(\mathbb{F}_q[x] \Big/ (\mathfrak{p}(x))\right)\right) = \mathfrak{p}(x) + \underbrace{\mathrm{t}_{\phi,\mathfrak{p}}(x)}_{\leq \deg(\mathfrak{p})/2},$$

which is $\mathfrak{p}(x)$ plus an error determined by the Frobenius trace $t_{\phi,\mathfrak{p}}(x)$ of degree at most half that of $\mathfrak{p}(x)$. The analogy with the Hasse bound is striking. The error in each case takes roughly at most half the number of bits as the estimate.

Factor degree by Euler–Poincaré characteristic

Gekeler's bound concerns Drinfeld modules ϕ at an irreducible $\mathfrak{p}(x)$. What happens at $\mathfrak{f}(x) = \prod_i \mathfrak{p}_i(x)$, our polynomial to factor? The multiplicativity of the Euler–Poincaré characteristic implies

$$\begin{aligned}\chi\left(\phi\left(\mathbb{F}_q[x] \Big/ (\mathfrak{f}(x))\right)\right) &= \prod_i \chi\left(\phi\left(\mathbb{F}_q[x] \Big/ (\mathfrak{p}_i(x))\right)\right) \\ &= \prod_i \left(\mathfrak{p}_i(x) + \mathrm{t}_{\phi,\mathfrak{p}_i}(x)\right) = \mathfrak{f}(x) + \mathrm{t}_{\phi,\mathfrak{f}}(x),\end{aligned}$$

for some $\mathrm{t}_{\phi,\mathfrak{f}}(x)$ of degree at most $s_f/2$, where $s_{\mathfrak{f}}$ denotes the degree of the smallest degree factor of $\mathfrak{f}(x)$. Thus, we have an extension of Gekeler's bound to reducible polynomials

$$\chi\left(\phi\left(\mathbb{F}_q[x] \Big/ (\mathfrak{f}(x))\right)\right) = \mathfrak{f}(x) + \underbrace{\mathrm{t}_{\phi,\mathfrak{f}}(x)}_{\leq s_f/2},$$

implying $\mathfrak{f}(x)$ and $\chi(\phi(\mathbb{F}_q[x]/(\mathfrak{f}(x))))$ agree at the high degree coefficients! The number of agreements tells us information about the smallest factor degree.

For a randomly chosen ϕ, $\mathrm{t}_{\phi,\mathfrak{f}}(x)$ likely has degree exactly $\lfloor s_f/2 \rfloor$ (with probability at least $1/4$). The number of agreements not merely bounds but determines the degree of the smallest factor. To claim this probability, one needs to prove for a randomly chosen ϕ, the Frobenius traces corresponding to the irreducibles of smallest degree do not conspire yielding cancellations. To this end, we seek equidistribution formulae for the Frobenius traces. Analogous to elliptic curves, there is a correspondence between the number of isomorphism classes of Drinfeld modules with a given trace and Gauss class numbers in certain imaginary quadratic orders. The latter can be computed using analytic class number formulae.

An algorithm to estimate the degree of the smallest degree factor of a given $\mathfrak{f}(x)$ is now apparent. Pick a Drinfeld module ϕ (by choosing $\mathfrak{g}_\phi, \Delta_\phi$ at random of degree less than n). Compute the Euler–Poincaré characteristic $\chi(\phi(\mathbb{F}_q[x]/(\mathfrak{f}(x))))$ and count the number of high degree coefficients it agrees in with $\mathfrak{f}(x)$.

Drinfeld module analogue of Lenstra's algorithm

It is instructive to begin with Lenstra's elliptic curve integer factorisation algorithm before seeing its Drinfeld module incarnation. Pollard designed his *p-1* algorithm to factor an integer that has a prime factor modulo which the multiplicative group has smooth order. But this smoothness condition is rarely met. Lenstra recast the role of the multiplicative group with the additive group associated with a random elliptic curve. If the integer has a prime factor modulo which the randomly chosen elliptic curve has

smooth order, the algorithm succeeds in extracting that factor. For a random elliptic curve, this smoothness condition is met with a probability depending sub-exponentially on the size of the smallest prime factor. Consequently, it is among the most popular algorithms for integer factorisation, particularly as an initial step to extract small factors.

Pollard's $p-1$ algorithm

Fix a positive integer B as the smoothness bound and denote by m, the product of all prime powers at most B. Given an N to factor, choose a positive integer $a < N$ at random. Assume a is prime to N for otherwise $\gcd(a, N)$ is a nontrivial factor of N. If N has a prime factor p with every prime power factor of $p-1$ at most B,

$$a^m - 1 = (a^{p-1})^{m/(p-1)} - 1 \cong 0 \mod p \implies p \mid a^m - 1$$

and $\gcd(a^m - 1, N)$ is likely a nontrivial factor of N. The running time is exponential in the size of B. For typical N, B needs to be as big as the smallest factor of N and thus the running time is typically exponential in the size of the smallest factor of N.

Lenstra's algorithm

Lenstra's elliptic curve factorization algorithm factors every N in (heuristic) expected time sub-exponential in the size of the smallest factor p of N. A key insight of Lenstra was to substitute the multiplicative group $(\mathbb{Z}/p\mathbb{Z})^\times$ in Pollard's $p-1$ algorithm with the group $E(\mathbb{F}_p)$ of $\mathbb{F}_p$ rational points of a random elliptic curve E over $\mathbb{F}_p$. The running time depends on the smoothness of the group order $|E(\mathbb{F}_p)|$ for a randomly chosen E. The Hasse bound guarantees that $||E(\mathbb{F}_p)| - (p+1)| \leq 2\sqrt{p}$ and Lenstra proved that his algorithm runs in expected time sub-exponential in the size of p assuming a heuristic on the probability that a random integer in the interval $[p+1-2\sqrt{p}, p+1+2\sqrt{p}]$ is smooth.

Drinfeld module analogue

In the ensuing Drinfeld module version, random elliptic curves will be recast with random Drinfeld modules to factor polynomials. As before, let $\mathfrak{f}(x) \in \mathbb{F}_q[x]$ denote the polynomial to factor. Pick a Drinfeld module ϕ at random and a random element α in the $\mathbb{F}_q[x]$ module $\phi(\mathbb{F}_q[x]/(\mathfrak{f}(x)))$. The order, $\mathrm{Ord}(\alpha)$, of α is the smallest degree polynomial $\mathfrak{b}(x) \in \mathbb{F}_q(x)$ that annihilates α, that is $\phi_{\mathfrak{b}}(\alpha) = 0$. Extract and divide away the linear factors of $\mathrm{Ord}(\alpha)$,

$$\mathfrak{r}(x) := \mathrm{Ord}(\alpha)/\gcd(\mathrm{Ord}(\alpha), x^q - x)$$

and apply the Drinfeld action $\phi_{\mathfrak{r}}(\alpha)$ on α. It is likely that $\phi_{\mathfrak{r}}(a)$ is a zero at some but not all factors $\mathfrak{p}_i(x)$ of $\mathfrak{f}(x)$ and $\gcd(\phi_{\mathfrak{r}}(a), \mathfrak{f})$ gives a random factorisation.

A brief outline of why this is the case follows. As with groups, the order of an element divides the cardinality of an $\mathbb{F}_q[x]$ module. That is, $\mathrm{Ord}(\alpha)$ divides the Euler–Poincaré characteristic of $\phi(\mathbb{F}_q[x]/(\mathfrak{f}(x)))$. In fact, with high probability, $\mathrm{Ord}(\alpha)$ equals the Euler–Poincaré characteristic

$$\mathrm{Ord}(\alpha) = \prod_i \left(\mathfrak{p}_i(x) + \mathfrak{t}_{\phi,\mathfrak{p}_i}(x)\right).$$

If the factors on the right have/don't have a linear factor independently and roughly uniformly at random, then the algorithm indeed yields a random factorisation. This is indeed the case!

The factors on the right lie in the short intervals I_i centred at $\mathfrak{p}_i(x)$ with interval degree bounded by $\deg(\mathfrak{p}_i)/2$. The Frobenius trace distribution assures a certain semi-circular equidistribution of $\mathfrak{p}_i(x) + \mathfrak{t}_{\phi,\mathfrak{p}_i}(x)$ in the short interval I_i. Remarkably, unlike over integers, factorisation patterns in the short intervals I_i are unconditionally proven to be random enough. In summary, for random ϕ, the factorisation patterns of each $\mathfrak{p}_i(x) + \mathfrak{t}_{\phi,\mathfrak{p}_i}(x)$ is like that of a random degree $\deg(\mathfrak{p}_i)$ polynomial. Further, they are independent.

The computation of $\mathrm{Ord}(\alpha)$ dominates the runtime and can be performed efficiently through linear algebra. This is in stark contrast to integers, where finding the order of an element in the multiplicative group modulo a composite appears hard.

Drinfeld Modules with Complex Multiplication:

Our last algorithm is distinguished in that it samples from Drinfeld modules with complex multiplication. A Drinfeld module ϕ has complex multiplication if its endomorphism ring

$$\mathrm{End}_{\mathbb{F}_q(x)}(\phi) \otimes_{\mathbb{F}_q[x]} (\mathbb{F}_q(x))$$

is isomorphic to a quadratic extension of $\mathbb{F}_q(x)$. A typical Drinfeld module has an endomorphism ring only isomorphic to $\mathbb{F}_q(x)$. Complex multiplication is the rare case where the Drinfeld module is more symmetric than typical. As with reducing elliptic curves over rational numbers at primes, Drinfeld modules can be reduced at irreducible polynomials. The reduction

is deemed supersingular if the endomorphism ring is noncommutative, and ordinary otherwise. Every Drinfeld module with complex multiplication has the remarkable feature that the density of irreducible polynomials where it is supersingular is roughly half.

To factor a given polynomial $\mathfrak{f}(x)$, the strategy is to choose a random Drinfeld module with complex multiplication. Using explicit formulae, construct a Drinfeld module ϕ with complex multiplication by the quadratic extension $\mathbb{F}_q(x)(\sqrt{x-c})$ with $c \in \mathbb{F}_q$ chosen at random. Then attempt to separate out the irreducible factors of $\mathfrak{f}(x)$ where ϕ is supersingular from the ordinary. This likely results in a random factorisation of $\mathfrak{f}(x)$, which is recursively factored to obtain the complete factorisation.

To separate the supersingular factors, we look to the Hasse invariant, an indicator of supersingularity. The Hasse invariant $h_{\phi,\mathfrak{p}} \in \mathbb{F}_q[x]/(\mathfrak{p}(x))$ of ϕ at an irreducible $\mathfrak{p}(x)$ vanishes if and only if ϕ is supersingular at $\mathfrak{p}(x)$. For the chosen ϕ, we construct a polynomial that is a simultaneous lift of Hasse invariants at all irreducible polynomials of degree at most that of $\mathfrak{f}(x)$. The efficient construction of this lift relies critically on a Drinfeld module analogue of Deligne's congruence due to Gekeler. The common irreducible factors of this lift and $\mathfrak{f}(x)$ are precisely the irreducible factors of $\mathfrak{f}(x)$ where ϕ is supersingular. The GCD of $\mathfrak{f}(x)$ and the lift separates out the supersingular factors from the ordinary, as desired.

Hasse Invariants and Deligne's congruence:

For a Drinfeld module ϕ with defining coefficients $(\mathfrak{g}_\phi, \Delta_\phi)$, we now construct the aforementioned lift of Hasse invariants. Consider the sequence $\mathfrak{r}_{\phi,k}(x) \in \mathbb{F}_q[x]$ of polynomials indexed by k starting with $\mathfrak{r}_{\phi,0}(x) := 1$, $\mathfrak{r}_{\phi,1}(x) := \mathfrak{g}_\phi(x)$ and for $m > 1$,

$$\mathfrak{r}_{\phi,m}(x) := \left(\mathfrak{g}_\phi(x)\right)^{q^{m-1}} \mathfrak{r}_{\phi,m-1}(x) - (x^{q^{m-1}} - x)\left(\Delta_\phi(x)\right)^{q^{m-2}} \mathfrak{r}_{\phi,m-2}(x).$$

Gekeler showed that $\mathfrak{r}_{\phi,m}(x)$ is the value of the normalized Eisenstein series of weight $q^m - 1$ on ϕ and established Deligne's congruence for Drinfeld modules, which ascertains for any $\mathfrak{p}$ of degree $k \geq 1$ with $\Delta_\phi(x) \neq 0 \mod \mathfrak{p}$ that $h_{\phi,\mathfrak{p}} = \mathfrak{r}_{\phi,k}(x) \mod \mathfrak{p}(x)$.

Hence $\mathfrak{r}_{\phi,k}(x)$ is a lift to $\mathbb{F}_q[x]$ of the Hasse invariants of ϕ at not just one but all irreducible polynomials of degree k. Further, $\mathfrak{r}_{\phi,k}(x), \mathfrak{r}_{\phi,k+1}(x)$ are both zero precisely modulo the supersingular $\mathfrak{p}(x)$ of degree at most k. Since a factor of $\mathfrak{f}(x)$ is of degree at most n, take

$$\gcd(\mathfrak{r}_{\phi,n}(x), \mathfrak{r}_{\phi,n+1}(x))$$

as the Hasse invariant lift.

To claim the algorithm indeed works, it remains to demonstrate that with constant probability, our random choice of ϕ with complex multiplication yielding a random factorization of $\mathfrak{f}(x)$. By complex multiplication theory and Carlitz reciprocity, this probability is identified with splitting probabilities in a certain hyperelliptic extension of $\mathbb{F}_q(x)$, and duly bounded. The overall runtime is dictated by the time taken to compute the Hasse invariant lift. An intricate algorithm for this task is devised in [1] using a fast procedure to compute its defining recurrence. Remarkably, the runtime exponent matches the best known factorisation algorithm and is comparable in practice to the fastest existing implementations. In light of this, thorough further investigation of Drinfeld module inspired polynomial factorisation is warranted!

Acknowledgements

This work was supported by the European Union's H2020 Programme (grant agreement #ERC-669891).

FURTHER READING

[1] J. Doliskani, A. K. Narayanan, É. Schost, Drinfeld Modules with Complex Multiplication, Hasse Invariants and Factoring Polynomials over Finite Fields, J. Symbolic Comput., to appear.

[2] A. K. Narayanan, Polynomial factorization over finite fields by computing Euler–Poincaré characteristics of Drinfeld modules, Finite Fields Appl. 54 (2018) 335–365.

[3] A. Panchishkin and I. Potemine, An algorithm for the factorization of polynomials using elliptic modules, Constructive methods and algorithms in number theory, p. 117. Mathematical Institute of AN BSSR, Minsk, 1989 (Russian).

Anand Kumar Narayanan

Anand is a research scientist at the Laboratoire d'Informatique de Paris 6, Sorbonne Université. His research interests are in the algorithmic aspects of number theory guided by applications in cryptography, coding theory and complexity theory. Anand was born in India and in his spare time enjoys wandering jungles and mountains.