Square-free Smooth Polynomials in Residue Classes and Irreducible Polynomial Generators

Christian Bagshaw

January 18, 2023

Euclid's Proof

When Euclid's proof of the infinitude of the primes is typically presented, it goes as follows:

Proof.

Suppose there are only finitely many primes $p_1, ..., p_n$. Then by the fundamental theorem of arithmetic,

$$N = p_1 p_2 ... p_n + 1$$

must have a prime factor, but of course it cannot be any of $p_1, ..., p_n$. Thus there cannot be finitely many primes. $\square$

Euclid's Proof

But Euclid didn't frame it as a proof by contradiction! Euclid's proof was constructive - he demonstrates that this idea will produce an infinite sequence of primes.

Proof.

Let $p_1, ..., p_n$ be a finite list of primes.

We define p_{n+1} to be a prime factor of $p_1 p_2 ... p_n + 1$.

The fundamental theorem of arithmetic ensures that p_{n+1} is well-defined, and of course it cannot be a prime already in the list.

This allows us to recursively define an infinite sequence of distinct primes. $\square$

Euclid-Mullin Sequence

If we start with $p_1 = 2$ and take the smallest prime factor of $p_1 p_2 ... p_n + 1$ each time, this sequence is called the Euclid-Mullin sequence:

$$2, 3, 7, 43, 13, 53, 5, 6221671, 38709183810571, 139, ...$$

The natural question is: does every prime occur in this sequence?

Euclid-Mullin Sequence

We don't know... and in fact there are some who think this problem may be undecidable.

Most people think it should be true, and the typical thought is something like: lets suppose we have some prime p not in the sequence and all the primes less than p are in the sequence.

At each stage in the sequence, $p_1 p_2 ... p_n + 1$ takes on some value $\pmod p$. We just need it to hit 0.

For p to not appear in the sequence would mean that these values have a very serious bias away from 0 $\pmod p$. But primes should be somewhat random $\pmod p$.

Euclid-Mullin Sequence Variant

If instead we take p_{n+1} to be the largest prime factor of $p_1 p_2 ... p_n + 1$, starting with $p_1 = 2$ the sequence looks like

$$2, 3, 7, 43, 139, 50207, 340999, 2365347734339,$$
$$11128622507, 821384623, 57088979368630 1597, ...$$

Again we may ask: does every prime occur in this sequence?

Andrew Booker proved that infinitely many primes will not appear! (11 is the first prime not to appear)

Moving to $\mathbb{F}_q[T]$

We fix a prime power q, and let $\mathbb{F}_q$ denote the finite field of q elements. We can consider similar questions over $\mathbb{F}_q[T]$.

Sequences of Irreducible Polynomials

We can define an infinite sequence of monic irreducible polynomials as follows: Given a finite list of monic irreducible polynomials $P_1, ..., P_n$, recursively define P_{n+1} to be an irreducible factor of

$$P_1 P_2 ... P_n + 1.$$

Can this sequence be chosen to contain every irreducible polynomial?

At the moment we don't know...

Does taking maximums omit infinitely many irreducible polynomials?

If $q = p$ is prime, then there is a very natural way to define a total order on $\mathbb{F}_q[T]$:

$$f(T) \leq g(T) \iff f(p) \leq g(p)$$

where you consider the coefficients as integers.

Some work of Kurokawa and Satoh [Kurokawa and Satoh, 2008] shows that if you are working in $\mathbb{F}_2[T]$, start with T as the first term in the sequence and take the largest factor each time, then $T^3 + T + 1$ never appears.

If you take the largest factor each time, do you omit infinitely many irreducible polynomials?

What do we know?

Instead of adding 1 each time you could add some $a \in \mathbb{F}_p[T]$.

If you start with T as the first term in the sequence, Kurokawa and Satoh proved that there are infinitely many pairs (p, a) such that $P_1 P_2 ... P_n + a$ is always irreducible, at every stage of the sequence.

Working $(\mathrm{mod}\ P)$

Whether or not an irreducible polynomial P eventually appears in the sequence is of course very closely related to how often

$$P | P_1 P_2 ... P_n + 1$$

for some square-free product of irreducibles $P_1 ... P_n$.

Or equivalently we may want to look at the equation

$$P_1 P_2 ... P_n \equiv -1 \ (\mathrm{mod}\ P).$$

For ease of notation, we will call a polynomial k-smooth if all of it's irreducible factors have degree less than or equal to k.

Main Result

Theorem (B.)

Let $q \geq 7$ be a prime power, and let $P \in \mathbb{F}_q[T]$ be an irreducible polynomial of degree at least 2. Then every residue class modulo P has a non-constant, square-free, $\deg P - 1$*-smooth representative.*

That is, for any $a \in \mathbb{F}_q[T]$ with $P \nmid a$, there exists distinct irreducible polynomials $P_1, ..., P_n$, all of degree less than $\deg P$*, such that*

$$P_1 P_2 ... P_n \equiv a \pmod{P}.$$

An integer analogue of this result was proven by Andrew Booker and Carl Pomerance [Booker and Pomerance, 2017].

Main Result

To tie this back to our original question about sequences, $a = -1$ works perfectly fine in the previous theorem. So for any irreducible polynomial P we have that

$$P | P_1 P_2 ... P_n + 1$$

for some distinct irreducible polynomials $P_1, ..., P_n$ of degree less than $\deg P$.

Some ideas from the proof

General framework is due to Booker and Pomerance, although of course the details are different in our setting. I don't want to bore anyone with all the details, so I will just outline an idea I enjoyed.

To avoid me saying the same thing over and over, lets set $\deg P = r$ and lets call a "square-free, $r-1$-smooth" representative a "nice" representative.

So when I refer to a "nice" representative for a, think of writing $a \pmod P$ as a product of distinct irreducibles that all have degree less than r.

Some ideas from the proof

We want to show that for any irreducible P, every a has a nice representative.

The idea is that we show there exists some special subgroup

$$H \subseteq (\mathbb{F}_q[T]/P(T))^*$$

such that every coset of H has a nice representative, and every element of H has a nice representative that shares no factors with the coset representatives.

Finding these coset representatives

A taste: how do we find these coset representatives?

Our subgroup is going to be

$$H_d = \{f \in (\mathbb{F}_q[T]/P(T))^* \ : \ f^{(q^r-1)/d} = 1\}$$

for some divisor $d|q^r - 1$.

Finding these coset representatives

Two quick definitions...

Let χ be a non-trivial group homomorphism

$$\chi : (\mathbb{F}_q[T]/P(T))^* \to \mathbb{C}^*$$

of order d. Just think of a map that takes things $\pmod P$ and spits out something on the unit circle, and such that $\chi^d = 1$.

By setting $\chi(0) = 0$, you can extend this to a function on all of $\mathbb{F}_q[T]$ (think Dirichlet characters, but for polynomials).

Let μ_q^2 be the characteristic function for square-free polynomials.

Finding these coset representatives

Then it turns out by the magic of orthogonality...

$$\frac{1}{d}\sum_{i=1}^{d}\sum_{\deg f\in[1,r)}\mu_q^2(f)\chi^i(f)\bar{\chi}^i(a)$$

counts the number of square-free polynomials in the coset aH_d with degree less than r.

Don't worry too much about how this works, just think that we have a function that counts nice representatives in cosets. If you show that the value of this is positive, then you are done.

Finding these coset representatives

We manipulate a bit (separate out the case $i = d$) and it turns out we just have to show

$$\sum_{\deg f \in [1,r)} \mu_q^2(f) > \left| \sum_{i=1}^{d-1} \sum_{\deg f \in [1,r)} \mu_q^2(f) \chi^i(f) \bar{\chi}^i(a) \right|.$$

You can just apply a counting argument to get an exact value for the LHS, and for the RHS you apply some fancy "character sum" estimates.

Why $q \geq 7$?

We believe the result holds for $q \geq 2$, except for the cases

$$P = T^3 + T + 1 \text{ and } P = T^3 + T^2 + 1$$

over $\mathbb{F}_2[T]$.

That is, you could manually check that 1 has no nice representatives modulo either of these polynomials.

But it turns out the analytic stuff from before left a lot of remaining cases...

Why $q \geq 7$?

The main ingredient to bound that stuff from before was utilizing the following [Han, 2020]: for any positive integer t,

$$\left| \sum_{\substack{f \in \mathbb{F}_q[T] \\ \deg f = t \\ f \text{ monic}}} \chi(f) \right| \leq q^{t/2} \binom{\deg P - 1}{t}.$$

The triangle inequality trivially bounds the left hand side by q^t.
But the binomial coefficients on the right can get close to 2^t, so for small q the right can be close to $q^{t/2}2^t$, which could even be worse than trivial. The difficulty here is we need something completely explicit, not asymptotic.

A good improvement upon this bound for small prime powers would certainly lead to progress here.

The computational side

Because this bound gets worse as q gets small, the analytic proof DID NOT work for any irreducible polynomial P over $\mathbb{F}_q[T]$ satisfying

$$
\begin{aligned}
&q = 2... \\
&q = 3, \quad \deg P \leq 48 \\
&q = 4, \quad \deg P \leq 22 \\
&q = 5, \quad \deg P \leq 12 \\
&q = 7, \quad \deg P \leq 5 \\
&q = 8, \quad \deg P \leq 5 \\
&q = 9, \quad \deg P \leq 4 \\
&q = 11, \quad \deg P \leq 4 \\
&\vdots \\
&q = 61, \quad \deg P = 2
\end{aligned}
$$

The computational side

$$
\begin{aligned}
q = 7, &\quad \deg P \leq 5 \\
q = 8, &\quad \deg P \leq 5 \\
q = 9, &\quad \deg P \leq 4 \\
q = 11, &\quad \deg P \leq 4 \\
&\vdots \\
q = 61, &\quad \deg P = 2
\end{aligned}
$$

All these cases were manually checked with a close-to-trivial algorithm in Sage - just a bit of work to keep track of square-free polynomials already found.
Took my Macbook Air about 5 hours...

The computational side

$$q = 3, \quad \deg P \leq 48$$
$$q = 4, \quad \deg P \leq 22$$
$$q = 5, \quad \deg P \leq 12$$

A very efficient method to compute smooth square-free representatives for each equivalence class of many irreducible polynomials could close up this problem for $q = 3, 4, 5$.

Applying this to generate irreducibles

Lets move on from that stuff, and talk about how we can use this result to generate some polynomials in a similar fashion to Euclid.

Theorem (B.)

Let $p \geq 7$ be a prime. Starting with T as the first term in the sequence, recursively define a sequence of monic irreducible polynomials in $\mathbb{F}_p[T]$ where if N is the product of the terms generated so far take

$$P = \min\{g \text{ monic and irreducible } : \ g \nmid N, \ g | h+1 \text{ for some } h | N\}$$

as the next term in the sequence. This sequence produces all monic irreducible polynomials in $\mathbb{F}_p[T]$ in order.

Applying this to generate irreducibles

How does this relate to the original problem?
If our sequence contains $P_1, ..., P_n$, then instead of taking the smallest irreducible factor of $P_1P_2...P_n + 1$, we are taking the smallest irreducible factor of

$$P_{e_1}...P_{e_k} + 1$$

where we consider all possible subsets

$$\{P_{e_1}, ..., P_{e_k}\} \subseteq \{P_1, ..., P_n\}.$$

Instead of multiplying all the terms in the sequence together and adding 1, we just multiply some of them together and add 1.

Another way to generate irreducibles

Theorem (B.)

Let $q \geq 7$ be a prime power. Starting with all degree 1 polynomials as the first terms in the sequence, recursively define a sequence of irreducible polynomials in $\mathbb{F}_q[T]$ where if N is the product of the terms generated so far, take as the next term some irreducible factor of $h + N/h$ where $h|N$. Such a sequence can be chosen to contain every irreducible polynomial.

Another way to generate irreducibles

How does this relate to the original problem?
If our sequence contains $P_1, ..., P_n$, then instead of taking the smallest irreducible factor of $P_1P_2...P_n + 1$, we are looking at irreducible factors of

$$\prod_{P_i \in S} P_i + \prod_{P_j \notin S} P_j$$

where we consider all possible subsets

$$S \subseteq \{P_1, ..., P_n\}.$$

Instead of multiplying all the terms in the sequence together and adding 1, we just split them into two sets and add the products.

Why does this work?

Again the technique is due to Booker and Pomerance.
The idea is: suppose P is an irreducible not in the sequence, and all irreducible polynomials of degree less than P are in the sequence. Let N be the product of the terms in the sequence so far.

If $-N$ is square modulo P, then there exists some non-zero a such that

$$-N \equiv a^2 \pmod{P} \implies a + N/a \equiv 0 \pmod{P}.$$

But the main result implies that there exists some $h|N$ such that $a \equiv h \pmod{P}$.
So

$$P|h + N/h.$$

Why does this work?

If $-N$ is not a square modulo P, then lets show you can pick the next term in the sequence (say Q) to also not be a square modulo P. Thus at the next stage, the product of all the terms is QN, and $-QN$ is a square modulo P, so we are back in the first case.

Why does this work?

You can use the famous Hasse bound to show that there is some $x \pmod P$ that doesn't lie on the elliptic curve

$$y^2 \equiv x^3 + xN \equiv x^2(x + N/x) \pmod P$$

This means that

$$x + N/x$$

is not a square modulo P.
Let h be a nice representative for x modulo P, and then let Q be an irreducible factor that divides $h + N/h$ that is not a square modulo P. Take this as the next term in the sequence.

Do we need to start with all linear polynomials?

The way I stated this result was that we need to start with all the degree 1 (linear) polynomials already in the sequence.

> Do we need to start with all degree 1 polynomials in the sequence for this to work? Can we just start with "T" in the sequence, for example, and still generate all the irreducible polynomials this way?

With my limited insight, this doesn't seem to follow from our main result.

Do we need to start with all linear polynomials?

Even to computationally check this for a given prime power is difficult. Working in $\mathbb{F}_3[T]$, again my MacBook Air tells me there are a total of 1397132 possibilities for the first 7 terms in the sequence, none of which contain more than 3 degree 1 polynomials (of course there are 6 linear polynomials in $\mathbb{F}_3[T]$).

Another question

Lets say $\deg P = r$ for simplicity. Given some a, we could also ask:

How many non-constant, nice representatives does a have modulo P?

Lets additionally only count those such that

$$\deg(P_1 P_2 ... P_n) < k$$

for some integer k.

So we are really counting solutions (say $N(a, k, P)$) to

$$P_1 ... P_n \equiv a \pmod{P}, \quad \begin{array}{l} \deg P_i < r \\ \deg(P_1 ... P_n) < k \\ P_1 ... P_n \text{ square-free.} \end{array}$$

Another question

What might we expect $N(a, k, P)$ to be?

If we had to guess, maybe we would say that when $\deg P$ is really big, square-free polynomials should distribute themselves somewhat uniformly across all the residue classes modulo P. So a has its fair share, but not too much less.
So maybe something like

$$\begin{aligned} N(a, k, P) &\approx E(r)\left(\frac{\text{number of polynomials of degree less than } k}{\text{number of residue classes modulo } P}\right) \\ &= E(r)q^{k-r} \end{aligned}$$

where $E(r)$ is some error term, taking into account that not all polynomials are square-free and $r-1$-smooth. And maybe if k isn't too big nor small, then $E(r)$ isn't too big.

Another question

Theorem (B. and Shparlinski)

If $k \in [3r/2, 3r]$ *then*

$$N(a, k, P) \geq q^{k-r+o(r)}$$

as $r \to \infty$.

As we let $\deg P \to \infty$, $N(a, k, P)$ is at least about q^{k-r} as long as k isnt too big nor too small, and the error is something like $q^{-\epsilon r}$.

Another another problem

Given some a, we could ask:

What is the smallest (in terms of degree) nice representative for a modulo P?

Lets think just a little bit more generally - so let's ask what's the smallest square-free, α-smooth representative for a modulo P.
So we are looking for the smallest value of $\deg(P_1 \ldots P_n)$ out of all solutions to

$$P_1 \ldots P_n \equiv a \pmod{P}, \quad \deg P_i < \alpha$$
$$P_1 \ldots P_n \text{ square-free.}$$

Lets call it $M(\alpha, a, P)$.

Another another problem

Theorem (B. and Shparlinski)

If $r = O(\alpha)$ *then*

$$M(\alpha, a, P) \leq r(2 + o(1))$$

as $r \to \infty$, *independently of* a.

For any fixed $\epsilon > 0$, we can let $\alpha = \epsilon r$ and this holds as $r \to \infty$.

Is the asymptotic needed?

Lets just think about "nice" representatives again (so $\alpha = r - 1$).

The proportion of polynomials that are square-free is $1 - 1/q$.

We want to find a nice representative for every a with $\deg a < r$.

If a is square-free then we are done.

If not then maybe we would expect one of

$$a + cP, \ c \text{ constant}$$

to be square-free (and probably also $r - 1$-smooth), since there are $q - 1$ of these.
If not, then maybe

$$a + cP, \ c \text{ has degree } 1$$

is square-free (and maybe also $r - 1$-smooth), since there are $q^2 - q$ of these.

Is the asymptotic needed?

This (combined with a few limited computations) makes me think that reality is something like

$$M(r-1, a, P) \leq r + C$$

for sufficiently large q, for some small constant C that is independent of q.

Thank you!

References


Booker, A. and Pomerance, C. (2017).
Squarefree smooth numbers and Euclidean prime generators.
Proc. Am. Math. Soc, 145:5035–5042.


Han, D. (2020).
A note on character sums in function fields.
Finite Fields and Their Applications, 68:101734.


Kurokawa, N. and Satoh, T. (2008).
Euclid prime sequences over unique factorization domains.
Experimental Mathematics, 17:145–152.