Future directions in algorithmic number theory

The American Institute of Mathematics

This is a hard–copy version of a web page available through http://www.aimath.org
Input on this material is welcomed and can be sent to workshops@aimath.org
Version: Wed Apr 30 11:57:16 2003

This document is an outcome of the ARCC workshop "Future directions in algorithmic number theory," held at AIM in Palo Alto, March 24-28, 2003.

Table of Contents

Chapter A: Lecture Notes

Lecture notes were TeXed in real time by John Voight.

A.1 Agrawal: A Polynomial Time Algorithm for Testing Primality

The primality testing algorithm we present is based upon the following identity: n is prime if and only if

$$(1+X)^n \equiv 1 + X^n \pmod{n},$$

where we consider this congruence as an identity in the polynomial ring $(\mathbb{Z}/n\mathbb{Z})[X]$. This easily gives a randomized algorithm which runs in polynomial time: rather than verifying it completely (which would be far too expensive), you verify it modulo a randomly chosen degree $\log n$ polynomial $Q(X)$; repeating this sufficiently often, you expect to witness a failure of this congruence to hold fairly quickly if n is composite.

Conjecture. *To prove that n is prime, it is enough to let $Q(X)$ run over the set*

$$\{X-1, X^2-1, \ldots, X^{\log^2 n}-1\}.$$

We were unable to prove this conjecture, but, instead the modified conjecture, which is then a proposition:

Proposition (Modified conjecture)**.** *To prove that n is prime, it is enough to let $Q(X)$ run over the set*

$$R = \{(X+a)^r - 1 : 1 \le r \le 16(\log n)^5, 1 \le a \le 8(\log n)^{7/2}\},$$

except when n has a 'small' prime factor ($\le (\log n)^5$).

Remark. One can replace the assumption that n does not have a small prime factor by adding to the list the polynomials X^k.

For a fixed r, this gives only that n is a prime power, which is a condition readily checked.

We now prove the modified conjecture.

If n is prime, clearly all of these conditions will hold.

Assume that n is composite and does not have a 'small' prime factor. Assume that

$$(1+X)^n \equiv 1 + X^n \pmod{n, Q(X)}$$

for every $Q(X) \in R$. First, observe that n is not a prime power. This follows as

$$(1+X)^n \equiv 1 + X^n \pmod{n, (X+1)^r - 1}$$

so substituting X for $X+1$ everywhere we get

$$(X-1)^n \equiv X^n - 1 \pmod{n, X^r - 1}$$

and substituting X^k for X we get

$$(X^k-1)^n \equiv X^{nk} - 1 \pmod{n, X^r - 1}$$

hence

$$\prod_{k<r} (X^k-1)^n \equiv \prod_{k<r} (X^{kn} - 1)$$

so $r^n \equiv r \pmod{n}$ for all $r \le 16(\log n)^5$.

Suppose $p^2 \mid n$ for p a prime. Then $r^{n-1} \equiv 1 \pmod{p^2}$ (n does not have a small prime divisor) so $r^{\gcd(n-1,p(p-1))} \equiv 1 \pmod{p^2}$ as the group is cyclic. Therefore $r^{p-1} \equiv 1 \pmod{p^2}$ for all $r \leq 16(\log n)^5$. A simple counting argument (look at all of the possible numbers whose prime divisors are all smaller than $(\log n)^2$, the identity holds for them but there are more than p of them) shows that this cannot happen.

Assume that n is composite but not a prime power. Let $p \mid n$, p prime, and fix r. Now:

Claim. We have

$$(1+X)^n \equiv 1 + X^n \pmod{n, (X+a)^r - 1}$$

for $1 \leq a \leq 8(\log n)^{7/2}$ if and only if

$$(X-a)^n \equiv X^n - a \pmod{n, X^r - 1}$$

for $1 \leq a \leq 8(\log n)^{7/2}$.

This can be seen by replacing X by $X - a$ as appropriate.

Definition. A number m is *introspective* for $g(X)$ if

$$g(X)^m \equiv g(X^m) \pmod{p, x^r - 1}.$$

Both n and p are introspective for $X - a$, $1 \leq a \leq 8(\log n)^{7/2}$. Indeed, a prime p is trivially introspective for every polynomial.

Observe that if m_1 and m_2 are introspective for $g(X)$, then so is $m_1 m_2$. This is clear as

$$g(X)^{m_1 m_2} \equiv g(X^{m_1})^{m_2} \equiv g(X^{m_1 m_2}) \pmod{p, X^r - 1}.$$

Secondly, observe trivially that if m is introspective for $g_1(X)$ and $g_2(X)$, then it is for $g_1(X)g_2(X)$.

Now let $I = \{n^i p^j : i, j \geq 0\}$ and

$$T = \left\{ \prod_{1 \leq a \leq 8(\log n)^{7/2}} (X-a)^{e_a} : e_a \geq 0 \right\}.$$

Observe that every $m \in I$ is introspective for every $g(X) \in T$.

Let t be the order of the group generated by n and p in $(\mathbb{Z}/r\mathbb{Z})^*$. There are $> t$ elements in I less than or equal to $n^{2\sqrt{t}}$.

Furthermore, there are $> n^{2\sqrt{t}}$ distinct polynomials of degree $< t$ which are distinct modulo p and $h(X)$, where $h(X)$ is an irreducible factor of the rth cyclotomic polynomial in $\mathbb{F}_p$. We show this as follows.

The number of polynomials in T of degree $< t$ is at least $2^{\min\{t, 8(\log n)^{7/2}\}}$ by simply considering products of distinct linear factors in T. Assume that $t > 4(\log n)^2$ and $t < 16(\log n)^5$. (We will come back to this: we can choose the value of r to obtain it.) Then

$$n^{2\sqrt{t}} = 2^{2\sqrt{t}\log n} < 2^t$$

and

$$n^{2\sqrt{t}} < n^{8(\log n)^{5/2}} = 2^{8(\log n)^{7/2}}.$$

Let $F = \mathbb{F}_p[X]/(h(X))$, and let $g_1(X), g_2(X)$ be of degree $< t$ in T and $g_1(X) \neq g_2(X)$. Suppose $g_1(X) = g_2(X) \pmod{p, h(X)}$. Then

$$g_1(X^m) \equiv g_1(X)^m \equiv g_2(X)^m \equiv g_2(X^m) \pmod{p, h(X)}$$

for all $m \in I$. Therefore X^m is a root in F of the polynomial $g_1(Y) - g_2(Y)$. This is a contradiction, as each of the elements X^m are distinct (they are roots of unity in the field) for the t representatives of the group generated by n and p, but the degree of $g_1(Y) - g_2(Y)$ is $< t$.

Let $m_1 \equiv m_2 \pmod{r}$, $m_1 \neq m_2 \in I$ and $m_1, m_2 \leq n^{2\sqrt{t}}$; this is possible as t is the order of the group generated by n and p in $\mathbb{Z}/r\mathbb{Z}$. Let $g(X) \in T$ be of degree $< t$. Then

$$g(X)^{m_1} \equiv g(X^{m_1}) \equiv g(X^{m_2}) \equiv g(X)^{m_2} \pmod{p, h(X)}$$

so $g(X) \pmod{h(X), p}$ is a root in F of the polynomial $Y^{m_1} - Y^{m_2}$. Since g was an arbitrary element of T, each $g(X)$ is a root of this polynomial with degree $\leq n^{2\sqrt{t}}$ but there are more than $n^{2\sqrt{t}}$ of them, contradiction.

Finally, we show that $t > 4(\log n)^2$ and $t < 16(\log n)^5$. Suppose the order of n is $\leq 4(\log n)^2$ in $(\mathbb{Z}/r\mathbb{Z})^*$. Then $r \mid \prod_{d \leq 4(\log n)^2}(n^d - 1) < 2^{16(\log n)^5}$. Now use the fact that the least common multiple of the first k numbers is at least 2^k.

It is easy to see that this algorithm has runtime $(\log n)^{10.5}$.

A.2 Agrawal: Finding Quadratic Nonresidues

Let $p - 1 = 2^\ell s$, s odd. To find a quadratic non-residue, compute continuously -1, $(-1)^{1/2}$, $(-1)^{1/4}$, $\ldots$, $(-1)^{1/2^{\ell-1}}$. One can quickly compute $(-1)^{1/2}$ from an algorithm due to Schoof, but one gets stuck at $(-1)^{1/8}$.

The idea: in primality testing, we want to know if $\mathbb{Z}/n\mathbb{Z}$ is a field, so we embed it into $(\mathbb{Z}/n\mathbb{Z})[X]/(X^r - 1)$; this ring has enough structure to pull out a nice algorithm. Assume that $\ell \geq 2$, and we try only to compute $\sqrt{-1}$. Now we embed in $\mathbb{F}_p[X]/(X^2 + 1)$. Consider $g(X) = (1 - X)^s$. Observe that $(g(X))^{2^\ell} = 1$ in $\mathbb{F}_p[X]/(X^2 + 1)$, so $g(X)$ is a 2^ℓth root of unity.

Assume that $\ell = 2$. In this case, $g(X)$ is a fourth root of unity. But X is also a fourth root, so $g(X) = X^k \pmod{X - \omega}$ where ω is the 'real' fourth root of unity. Consider $g(X) \bmod (X^2 + 1)$: observe that $g(X) \neq X^k \pmod{X^2 + 1}$ for any k. For suppose

$$(1 - X)^s \equiv X^k \pmod{X^2 + 1};$$

then

$$(1 - 1/X)^s = 1/X^k \pmod{1/X^2 + 1},$$

so

$$(1 - X)^s = -(X^s/X^k) \pmod{X^2 + 1},$$

hence $-(X^s/X^k) = X^k \pmod{X^2 + 1}$, a contradiction because s is odd. So compute $\gcd(g(X) - X^k, X^2 + 1)$, for each k, one of them will factor $X^2 + 1$.

If $\ell > 2$, then you cannot argue $g(X) \equiv X^k \pmod{X - \omega}$. If $\ell > 2$, then it is possible that $g(X)$ is an eighth root or a sixteenth root or so on. Suppose that $g(X)$ modulo $X^2 + 1$ is an eighth root, for example. Then $g(X^2) = (1 - X^2)^s \equiv X^k \pmod{X - \zeta}$ for some factor $X - \zeta$ of $X^4 + 1$ and k odd. But $(1 - X^2)^s$ is even, so it cannot be an odd power, so $\gcd((1 - X^2)^s - X^k, X^4 + 1)$ will give either a linear factor (in which case we are done) or a product of quadratic factors $(X - \zeta^2)(X + \zeta^2)$ or the products similar to $(X - \zeta)(X - \zeta^3)$ and $(X - \zeta)(X + \zeta^3)$.

Let $h(X)$ be a quadratic factor of X^4+1. Now if $(1-X)^s \not\equiv X^k \pmod{h(X)}$, then $h(X)$ can be factored and we are done. Suppose

$$(1-X)^s \equiv X^k \pmod{(X-\zeta)(X-\zeta^3)}$$

and

$$(1-X)^s \equiv X^{k'} \pmod{(X+\zeta)(X+\zeta^3)}.$$

If you replace X by X^3, then you get

$$(1-X^3)^s = X^{3k} \pmod{(X^3-\zeta)(X^3-\zeta^3)}$$

so

$$(1-X^3)^s \equiv X^{3k} = (1-X)^{3s} \pmod{(X-\zeta)(X-\zeta^3)}.$$

That means that 3 is introspective for $(1-X)^s$. The same argument applies to the other congruence, so one obtains that 3 is introspective for $(1-X)^s \bmod X^4+1$.

Now try all over again now with $X-a$ replacing $1-X$; the bad case will be when we have 3 is introspective for $(a-X)^s \bmod X^4+1$ for a large number of a. Here we get stuck, but this should be impossible.

Remarks.

A. There was a solution due to Lehmer which says for any fixed ℓ that you can solve this problem? (Cohen) But that assumes the existence of a nonresidue to begin with, which is exactly our problem. (Bernstein) But also this doesn't seem to scale well. (Cohen) There is a strategy to deal with this, and we always work modulo a degree four polynomial. (Agrawal)

B. The fact that $s > 1$ means that the same techniques as in the primality test do not seem to apply. Can you solve the problem if $s = 3$, or for other small s? (Lenstra) The case of a Fermat prime is trivial (3 is a nonresidue, and there probably aren't any above 65537). (Lenstra, Elkies) It seems as though if s is bounded, there are only a finite number of problematic a. For example, $(a-X^3)^s \equiv (a-X)^3 \pmod{h(X)}$ does not hold for 'many' a (Pomerance).

C. One strategy to solve this problem: translate this problem into rings and stare at it. (Lenstra)

D. Is there a strategy to deal with cases beyond eighth roots (where we have the special situation that every odd integer has square 1)? (Elkies) Yes, but we need to solve this problem first. (Agrawal)

E. How many values of a do you need? (Voloch) It seems unlikely you can generate the whole group with the $(X-a)^s$ without the GRH. (Bernstein)

A.3 Bernstein: Proving Primality After Agrawal-Kayal-Saxena

The lecture notes are available on the speaker's website: the talk (and more)[1], related problems[2], older paper on AKS[3], and putting AKS into context[4].

[1] http://cr.yp.to/papers.html#quartic
[2] http://cr.yp.to/papers.html#abccong
[3] http://cr.yp.to/papers.html#aks
[4] http://cr.yp.to/primetests.html

A.4 Edixhoven: About Point Counting over Arbitrary Finite Fields

Consider a system of equations $f_1(x_1,\dots,x_n)=0,\ \dots,\ f_m(x_1,\dots,x_n)=0$ given by polynomials $f_i \in \mathbb{F}_q[x_1,\dots,x_n]$. This is not essentially different than the case of a single hypersurface (every variety is birational to a hypersurface, or, also, one can use the inclusion-exclusion principle). We let $q=p^r$.

Question. For fixed n, is there an algorithm that computes the number of solutions in $\mathbb{F}_q$ in time polynomial in the quantities: $\log q$ (or r and $\log p$), $d=\max_i \deg f_i$, and m?

If you fix p, and $m=1$, then the answer is yes (Lauder-Wan) using p-adic methods. We discuss the case where p is not fixed. If p is not fixed, then one knows that the answer is yes for elliptic curves (Schoof) using ℓ-torsion points and curves of a given genus via the ℓ-torsion of their jacobian (Pila).

Conceptually, all methods use Lefschetz fixed points formula:

$$\#X(\mathbb{F}_q)=\textstyle\sum_{i=0}^{2\dim X}(-1)^i \operatorname{Tr}(\operatorname{Frob}_q \,|H_c^i(X))$$

where we denote $H_c^i(X)$ cohomology with compact support. This is true for X a scheme of finite type which is separated over $\mathbb{F}_q$. For these cohomology groups, one can take a p-adic approach using a de Rham-type cohomology, lifting X to a p-adic ring R and take the hypercohomology of the de Rham sequence, quite explicit and computable but the complexity is worse than linear in p. Instead, one can also use mod ℓ methods ($\ell \neq p$); here one takes the groups $H_c^i(X_{\overline{\mathbb{F}_q},et},\mathbb{F}_\ell)$ which is a lot less explicit. The H^i derive from injective resolutions on the etale topology. In this setup, there is the advantage that you can choose ℓ. For an elliptic curve, one has

$$E(\overline{\mathbb{F}_q})[\ell]^\vee = H^1(E_{\overline{\mathbb{F}_q},et},\mathbb{F}_\ell).$$

What is the simplest interesting case where we want to but cannot yet compute an H^i with $i\geq 2$? We think of surfaces, or modular forms of weight ≥ 3 (to generalize the case of elliptic curves, which correspond to eigenforms of weight 2). We assume that there is a cohomology group of dimension ≥ 2 (if it is of dimension 1, Frobenius acts as a power of the cyclotomic character).

We consider as an example the modular form $\Delta = q\prod_{n\geq 1}(1-q^n)^{24}=\sum_{n\geq 1}\tau(n)q^n$ is an eigenform of weight 12, viewed as a function on the upper half-plane. Then $\Delta(dq/q)^{\otimes 6}$ is an $SL_2(\mathbb{Z})$-invariant on $\mathfrak{H}$, so it descends to $\mathfrak{H}/SL_2(\mathbb{Z})$. The variety we work with is the ten-fold product of the universal elliptic curve E; we find that $H^{11}(E^{10})$ has dimension 2. For all p, and $\ell\neq p$, $\tau(p) \bmod \ell$ is the trace of Frob_p on $H^{11}(E^{10}_{\overline{\mathbb{F}_p},et},\mathbb{F}_\ell)$, which is also the trace of Frob_p on $H^{11}(E^{10}_{\overline{\mathbb{Q}},et},\mathbb{F}_\ell)$ with $\operatorname{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ acting on it: it is the two-dimensional Galois representation (modulo ℓ) associated to Δ. The action factors through $\operatorname{Gal}(K_{\Delta,\ell}/\mathbb{Q})$ acting faithfully where $K_{\Delta,\ell}/\mathbb{Q}$ is a finite extension.

Compute explicitly the extension $K_{\Delta,\ell}$. One gets as a byproduct a computation of the actual representation, which we cannot easily compute now. A bit of work yields that $K_{\Delta,\ell}$ is the Galois closure of the field definition of a suitable element $x\in J_1(\ell)(\overline{\mathbb{Q}})[\ell]$, where $J_1(\ell)$ is the jacobian of $X_1(\ell)$, if $X_1(\ell)(C)=\mathfrak{H}/\Gamma_1(\ell)$ together with te cusps, and $\Gamma_1(\ell)$ is the set of matrices $\begin{pmatrix} a & b \\ c & d\end{pmatrix}$ with $a\equiv 1 \pmod{\ell}$, $c\equiv 0 \pmod{\ell}$. Still: need to compute this efficiently. This is not so easy because the genus g of $X_1(\ell)$ is quadratic in ℓ.

We have the following strategy for finding $\mathbb{Q}(x)$ (based a suggestion of Jean-Marc Couveignes). We have a surjection

$$X_1(\ell)(\mathbb{C})^g \to J_1(\ell)(\mathbb{C}) = \mathbb{C}^g/\Lambda$$

where $\Lambda = H_1(X_1(\ell)(\mathbb{C}), \mathbb{Z})$. We have $\mathbb{C}^g/\Lambda \supset (1/\ell)\Lambda/\Lambda \ni x$. (For r prime, $T_r \cdot x = \tau(r) \cdot x$.) This map on complex points is given by $(Q_1, \ldots, Q_g) \in X_1(\ell)(\mathbb{C})$ maps to the point $[Q_1 + \cdots + Q_g - gP_0] = \sum_{i=1}^g \int_{P_0}^{Q_i} (\omega_1, \ldots, \omega_g) \in \mathbb{C}^g/\Lambda$. Here, for P_0 one can choose a $\mathbb{Q}$-rational cusp.

Generically, there is a unique point $(Q_1, \ldots, Q_g)$ up to permutation which gives the point x, namely, $\alpha = \sum_i j(Q_i) \in \mathbb{Q}(x)$. Now one estimates the height of α, approximates α in $\mathbb{C}$ by lifting (numerically) the straight line path from 0 to x (possible because the map $X_1(\ell)^g \to J_1(\ell)$ is generically unramified). It will probably be a good idea to replace the divisor gP_0 by a sum of g distinct points $P_1, \ldots, P_g$, with small height, defined over a small and solvable extension of $\mathbb{Q}$. As x and the P_i determine the Q_i (up to permutation), and as x is a torsion point (Néron-Tate height zero) one expects that the height of α is not much bigger than that of the P_i. So one hopes that the required number of correct digits of the approximation of α grows polynomially in ℓ. The tool to be used for estimating the height of α is Arakelov geometry. A good indication that this proposed strategy works is that the height estimate works well in the function field case (a nice application of the Grothendieck-Riemann-Roch theorem).

A.5 Gao: Factoring Polynomials under GRH

We consider the following problem: Given a prime p and $f \in \mathbb{F}_p[x]$, where $\deg f = n$, f is separable, and f splits completely, find a proper factor of f (in deterministic polynomial time).

Berlekamp algorithm reduces general polynomials in $\mathbb{F}_q[x]$ (q a power of p) to polynomials of the above type. Without GRH, we are stuck already at $x^2 - a$. So throughout we assume GRH.

Ronyai (1988) shows that this can be done in time $(n^n \log p)^{O(1)}$, or more precisely $(n^r \log p)^{O(1)}$ whenever $r \mid n$, $r > 1$, so in particular if n is even f can be split in deterministic polynomial time. Bach, von zur Gathen and Lenstra (2001) give an algorithm with polynomial time if $\phi_k(p)$ is smooth for some k where $\phi_k(x)$ is the kth cyclotomic polynomial. Evdokimov (1994) shows for any n and p, f can be factored in time $(n^{\log n} \log p)^{O(1)}$. We discuss work in Cheng and Huang (2000), and Gao (2001) plus some unpublished results. GRH will be needed only to compute an rth nonresidue in $\mathbb{F}_p$ or in its extensions, for $1 \le r \le n$.

Definition. An algebra $R/\mathbb{F}_p$ is called *elementary* if $R \cong (\mathbb{F}_p)^{\oplus m}$ for some m.

Let R be elementary over $\mathbb{F}_p$. Then we write $R = \mathbb{F}_p\epsilon_1 + \cdots + \mathbb{F}_p\epsilon_m$ where the ϵ_i are primitive idempotents, which are unique in R.

Fact.

A. If $f, g \in R[x]$, then $\gcd(f, g)$ can be defined properly and can be computed in deterministic polynomial time for any elementary algebra R.

B. If $f \in R[x]$ is monic, separable (i.e. $(f, f') = (1)$), and f splits completely, then $R_1 = R[x]/(f(x))$ is also elementary. Given a zerodivisor in R_1, one can compute a proper factor of f or a zerodivisor of R.

C. Given a nontrivial ring endomorphism of R_1 over R, one can find a proper factor of f or a zerodivisor of R. (Need GRH to get an rth nonresidue, for $1 \le r \le n$, where $n = \deg f$.)

D. Given a quadratic nonresidue in $\mathbb{F}_p$, there is a deterministic polynomial time algorithm for computing square roots in R. More precisely we have a function $\sigma : R \to R$, such that (i) $\sigma(A)$ is a square root of A if A is a square in R, (ii) if $A = \sum_{i=1}^{m} a_i\epsilon_i \in R$, $a_i \in \mathbb{F}_p$, then

$$\sigma(A) = \sum_{i=1}^{m} \sigma(a_i)\epsilon_i,$$

and (iii) for $a \in \mathbb{F}_p$, $\sigma(a^2) = \pm a$. For example, if $p \equiv 3 \pmod 4$, then we can take $\sigma(A) = A^{(p+1)/4}$, and for $a \in \mathbb{F}_p$,

$$\sigma(a^2) = \begin{cases} a, & \text{if } a \text{ is a square,} \\ -a, & \text{otherwise.} \end{cases}$$

Now let $f \in \mathbb{F}_p[x]$ be separable, so that

$$f = \prod_{i=1}^{n}(x - a_i), \quad a_i \in \mathbb{F}_p.$$

To factor f, define

$$R_2 = \mathbb{F}_p[z_1, z_2]/(f(z_1), f(z_2)),$$

which is the tensor product of $\mathbb{F}_p[z]/(f(z))$ with itself. Then R_2 is an elementary algebra. In the following, we will identity z_1 and z_2 with their images in R_2. Let

$$\epsilon_i = \frac{\prod_{j\ne i}(z_1 - a_j)}{\prod_{j\ne i}(a_i - a_j)}, \quad 1 \le i \le n,$$

and

$$\eta_i = \frac{\prod_{j\ne i}(z_2 - a_j)}{\prod_{j\ne i}(a_i - a_j)}, \quad 1 \le i \le n.$$

Then $\epsilon_i\eta_j$, $1 \le i, j \le n$, are all the primitive idempotents of R_2. They have the following properties:

$$\sum_{i=1}^{n} \epsilon_i = 1, \quad \sum_{j=1}^{n} \eta_j = 1,$$

and

$$z_1 = a_1\epsilon_1 + \cdots + a_n\epsilon_n = \sum_{i,j} a_i\epsilon_i\eta_j,$$

$$z_2 = a_1\eta_1 + \cdots + a_n\eta_n = \sum_{i,j} a_j\epsilon_i\eta_j.$$

Let

$$A = \frac{1}{2}(z_1 + z_2 + \sigma((z_1 - z_2)^2) \in R_2,$$

which can be computed in deterministic polynomial time. Then

$$A = \sum_{i,j} \frac{1}{2}(a_i + a_j + \sigma((a_i - a_j)^2)\epsilon_i\eta_j$$

where

$$\frac{1}{2}(a_i + a_j + \sigma((a_i - a_j)^2)) = \begin{cases} a_i, & \text{if } \sigma((a_i - a_j)^2) = a_i - a_j \\ a_j, & \text{if } \sigma((a_i - a_j)^2) = a_j - a_i. \end{cases}$$

Hence A encodes information about the "squareness" of the differences of the roots of f. By using characteristic polynomials and gcd technique, one can extract the factors of f.

Definition. For $1 \le i \le n$, define

$$\Delta_i = \{1 \le j \le n : j \ne i, \sigma((a_i - a_j)^2) = -(a_i - a_j)\}.$$

Theorem. *We can always find a proper factor of f in $\mathbb{F}_p[x]$ except when*

$$\#\Delta_i = (n-1)/2, \quad 1 \le i \le n, \tag{1}$$

and in this exception case f can be factored over $\mathbb{F}_p[z_1]$ as

$$f(x) = (x - z_1) f_0 f_1$$

where

$$f_0 = \sum_{i=1}^{n} \prod_{j \in \Delta_i} (x - a_j)\epsilon_i, \quad f_1 = \sum_{i=1}^{n} \prod_{j \notin \Delta_i, j \ne i} (x - a_j)\epsilon_i.$$

To further factor f_0 over $R_1 = \mathbb{F}_p[z_1]/(f(z_1))$, we compute in the ring

$$R_3 = R_1[z_2, z_3]/(f_0(z_2), f_0(z_3)),$$

and similarly for f_1.

Theorem. *We can always split f_0 or f_1 except when*

$$\#(\Delta_i \cap \Delta_j) = (n-3)/4, \quad 1 \le i < j \le n, \tag{2}$$

and in this exception case, f_0 is factored over $R_1[z_2]/(f_0(z_2))$ as

$$f_0(x) = (x - z_2) f_{00} f_{01}$$

where $f_{00}, f_{01} \in R_1[z_2][x]$ both of degree $(n-3)/4$; similarly over the ring $R_1[z_2]/(f_1(z_2))$,

$$f_1(x) = (x - z_2) f_{10} f_{11}.$$

A system of subsets satisfying (1) and (2) is called an Hadamard design. When $n \equiv 3 \bmod 4$ is a prime, there is always an Hadamard design.

To further split f_{00}, let $R_2 = R_1[z_2]/(f_0(z_2))$, and we compute in the ring

$$R_4 = R_2[z_3, z_4]/(f_{00}(z_3), f_{00}(z_4)).$$

Similarly for f_{01}, f_{10}, and f_{11}.

Theorem. *We can always split f_{00}, f_{01}, f_{10}, or f_{11} except when*

$$\#(\Delta_i \cap \Delta_j \cap \Delta_k) = (n-7)/8, \quad 1 \le i < j < k \le n. \tag{3}$$

It can be proved, however, that (3) is impossible. **It remains open how to explore this information to obtain a proper factor of f in $\mathbb{F}_p[x]$!**

A.6 Kedlaya: Counting Points using p-adic Cohomology

We introduce a different p-adic setup for counting points (or equivalently computing zeta functions). If X is a variety over $\mathbb{F}_q$, we want to count points using 'de Rham' cohomology. We will demonstrate Monsky-Washnitzer cohomology (a kind of rigid cohomology for smooth affine varieties). We restrict to the case where X is an affine curve, since in higher dimensions other methods will be faster. Then $X = \operatorname{Spec} \mathbb{F}_q[t_1, \dots, t_n]/(f_1, \dots, f_m)$.

Let $W = W(\mathbb{F}_q)$ be the Witt vectors, and let $K = W[1/p]$ be the fraction field of W. Define

$$W\langle t_1, \dots, t_n\rangle^\dagger = \left\{ \textstyle\sum_I c_I t^I : I = (i_1, \dots, i_n) \in \mathbb{Z}^n_{\geq 0}, c_I \in W, \right. \\ \left. \text{convergent for } t_i \in \overline{K}, |t_i| \leq 1 + \epsilon \text{ for some } \epsilon > 0 \right\}.$$

In other words, $v_p(c_I) \geq r\epsilon_I - c$ for some r, c with $r > 0$. Modulo p, $W\langle t_1, \dots, t_n\rangle^\dagger$ reduces to the polynomial ring $\mathbb{F}_q[t_1, \dots, t_n]$, since all but finitely many coefficients are divisible by p. We define $K\langle t_1, \dots, t_n\rangle^\dagger = W\langle t_1, \dots, t_n\rangle^\dagger[1/p]$. Let

$$A^{\text{int}} = W\langle \widetilde{t_1}, \dots, \widetilde{t_n}\rangle^\dagger / (\widetilde{f_1}, \dots, \widetilde{f_n})$$

where $\widetilde{f_i}$ is a lift of f_i. Then

$$A^{\text{int}}[1/p] = K\langle \widetilde{t_1}, \dots, \widetilde{t_n}\rangle^\dagger / (\widetilde{f_1}, \dots, \widetilde{f_n}).$$

There exists a lift so that A^{int} is flat over W.

Let Ω^1_A be the A-module generated by symbols $dt_1, \dots, dt_n$ modulo the submodule generated by $d\widetilde{f_1}, \dots, d\widetilde{f_m}$. Then there is a K-linear derivation $d : A \to \Omega^1_A$. Letting $\Omega^i_A = \bigwedge^i_A \Omega^i_A$; you get the de Rham complex

$$A = \Omega^0_A \xrightarrow{d} \Omega^1_A \xrightarrow{d} \dots$$

and you 'define'

$$H^i_{MW}(X) = \frac{\ker(\Omega^i_A \to \Omega^{i+1}_A)}{\operatorname{img}(\Omega^{i-1}_A \to \Omega^i_A)}.$$

It turns out that $H^q_{MW}(X)$ is independent of choices (A is unique up to noncanonical isomorphism, funny automorphisms are homotopic to the identity) and given $X \to Y$, there is a map $A_Y \to A_X$ and the induced maps $H^i_{MW}(Y) \to H^i_{MW}(X)$ also do not depend on choices.

The spaces $H^i(X)$ are finite-dimensional, but it is not obvious; it relies upon relating this cohomology to rigid cohomology for proper varieties, namely, crystalline cohomology which we know is finite-dimensional for other reasons. Moreover, they satisfy the Lefschetz trace formula: if $F : X \to X$ is the q-power Frobenius, then (Monsky)

$$\#X(\mathbb{F}_{q^i}) = \textstyle\sum_j (-1)^j \operatorname{Tr}((qF^{-1})^i | H^j(X)).$$

The idea: try to compute $H^i(X)$ and the map induced by F (find $A^{\text{int}} \to A^{\text{int}}$ lifting q-power Frobenius).

Example. Look at $X = \operatorname{Spec} \mathbb{F}_q[x,y,z]/(y^2 - f(x), yz - 1)$ with $\operatorname{char} \mathbb{F}_q = p$ odd. Let $\deg f = 2g+1$, f monic. Lift it to $A = K\langle x,y,z\rangle^\dagger/(y^2 - P(x), yz-1)$, where $P(x)$ is monic, degree $2g+1$ over W. It is easy to compute that $H^0(X)$ is one-dimensional. Now $H^1(X)$ is generated by $x^i\,dx/y$ for $i = 0,\dots,2g-1$, and $x^i\,dx/y^2$, $i = 0,\dots,2g$. Note $H^1(X)$ splits under $y \mapsto -y$ into plus and minus eigenspaces.

You need to find relations in $H^1(X)$ that $d(x^i/y^j) = 0$. (This is a special situation: all relations are 'algebraic'.) Lift the p-power Frobenius by $W \to W$ by the Witt vector Frobenius, $x \mapsto x^p$, and $y \mapsto y^p\sqrt{F(P(x))/P(X)^p}$. Compose this map with itself n times to get a q-power Frobenius lift, and this allows us to compute the zeta function of a genus g hyperelliptic curve over $\mathbb{F}_{p^n}$ in time $\widetilde{O}(g^4n^3p)$.

A.7 Lauder: Counting Solutions to Equations in Many Variables over Finite Fields

We present an algorithm which allows us to count solutions to a homogeneous equation $f(X_1,\dots,X_n) \in \mathbb{F}_q[X_1,\dots,X_n]$ of degree d (for simplicity we assume $d \geq 2$, $n \geq 2$) with running time which does not increase exponentially in number of variables. In other words, we are interested in computing the number of projective solutions

$$N_k = \#\{(x_1 : \cdots : x_n) \in \mathbb{P}^{n-1}_{\mathbb{F}_{q^k}} : f(x_1,\dots,x_n) = 0\}$$

for every $k \geq 1$. We encode these numbers in the generating function

$$Z(f,T) = \exp(\textstyle\sum_k N_k T^k/k) \in \mathbb{Q}[[T]]$$

which, by a theorem of Dwork, is in fact a rational function. We assume that f is nonsingular, i.e. f and $\partial f/\partial x_i$ for $i = 1,\dots,n$ have no common projective solution. In this situation, we know that

$$Z(f,T) = \frac{P(T)^{(-1)^{n+1}}}{(1-T)(1-qT)\dots(1-q^{n-2}T)}$$

where $\deg P = (1/d)((d-1)^n + (-1)^n(d-1))$.

If we compute N_k naively for $k = 1,\dots,\deg P$ then we can of course recover the polynomial $P(T)$; the time required to do this, however, requires $(q^{\deg P})^n \approx 2^{d^{n-1}\log q}$ evaluations of f. The input is given by $\binom{d+n-1}{n-1} \leq d^{n-1}$ terms of size $\log q$, and the output size is approximately $(d^{n-1}\log q)^{O(1)}$. We would like the running time to be polynomial in this quantity.

If $n = 2$, we are counting solutions of a univariate polynomial, and this can be done in time $(d\log q)^{O(1)}$. For $n = 3$, we have an algorithm of Schoof-Pila for curves which has run time $(\log q)^\Delta$ where Δ depends on d exponentially. In general, there is an algorithm (due to L. and Wan) which runs in time $(pd^n\log q)^{O(n)}$. Notice the n in the exponent—we would like to lose this dependence.

The new result: If f is 'sufficiently generic' (we exclude a Zariski closed set which is efficiently computable), $p \neq 2$, and $p \nmid d$, then we can find $P(T)$ using $(pd^n\log q))^{O(1)}$ bit operations. As a corollary, we see that if $f \in \mathbb{Z}[X_1,\dots,X_n]$ is sufficiently generic, then there exists an algorithm which takes as input a prime p, outputs the number of solutions $f \bmod p = 0$, and has run time $O(p^{2+\epsilon})$.

Recall that $P(T) = \det(I - T\operatorname{Frob}_q | H^{n-2}(X))$, where we write X for the projective variety defined by the equation $f = 0$. The action of Frob_q can be represented by a matrix

with entries in a field of characteristic zero, and we find that

$$N_k = (-1)^n \operatorname{Tr}(\operatorname{Frob}_q^k | H^{n-2}(X)) + 1 + q^k + \cdots + (q^k)^{n-2}.$$

For curves, for example, the dimension is $n-2=1$, and $H^1(X)$ is a $\mathbb{Z}_\ell$-module, for $\ell \neq p$.

Instead, we work with the p-adic theory, where $H^{n-2}(X)$ is a R-module for a ring $R \supset \mathbb{Z}_p$. We compute instead $\operatorname{Frob}_q = \operatorname{Frob}_p^{\log_p(q)}$, and compute the matrix of $\operatorname{Frob}_p | H^{n-2}(X)$. (Specifically, $R = \mathbb{Q}_q(\pi)$ where $\mathbb{Q}_q$ is the unramified extension of $\mathbb{Q}_p$ of degree $\log_p q$ and $\pi^{p-1} = -p$. Also our $H^{n-2}(X)$ is actually the primitive part of the cohomology space.) Consider the family

$$f_Y = \sum_{i=1}^{n} a_i X_i^d + Y h(X_1, \ldots, X_n),$$

and assume that $a_1 \ldots a_n \neq 0$. Then $f_1 = f$ and f_0 is a diagonal form, for which it is easy to count the number of solutions. Now $\operatorname{Frob}_p(Y)$ is a p-adic analytic function with the property that $\operatorname{Frob}_p(Y)$ evaluated at a Teichmuller lift of $y^{1/p}$ is exactly the Frob_p associated to f_y.

We see that $\operatorname{Frob}_p(Y) = C(Y^p)^{-1} \operatorname{Frob}_p(0) C^{\tau-1}(Y)$ where $\tau \bmod p : \alpha \mapsto \alpha^p$, and $C(Y)$ is a matrix of power series around the origin satisfying the differential equation $dC/dY = C(Y)B(Y)$ with initial condition $C(0) = I$, where $B(Y)$ is easily computed. This gives a way to compute $\operatorname{Frob}_p(Y)$ in a radius around the origin, and we need to extend this to the closed disc of radius 1.

The entries of the matrices are p-adic holomorphic functions, so to compute these modulo a power of p we find rational functions with denominators corresponding to the values of Y where the variety becomes singular and then recover the numerator from the power series. Using overconvergence we get a bound on the degree of these rational functions. We evaluate the rational functions at $Y = 1$. One gets nice complexity because we work with univariate power series, with decay in the coefficients on the order of $1/p$, and one needs to take the power of p approximately on the order of $d^n \log q$.

In the old algorithm, we would work in $H^{n-2}(X)$, the ring of power series in $X_1, \ldots, X_n$ over an $R = \mathbb{Q}_p^{(m)}(\pi)$ where $\mathbb{Q}_p^{(m)}$ is the unramified extension of $\mathbb{Q}_p$ of degree m (if $q = p^m$) and $\pi^{p-1} = -p$ modulo an infinite subspace; the power series we must work with have on the order of $(pd^n \log q)^n$ terms.

Problem. *Find an algorithm which counts the number of points on a curve in time* $(d \log q)^{O(1)}$.

A.8 Lenstra: Primality Testing with Pseudofields

This is joint work with Carl Pomerance.

If f, g are real-valued functions on a set X and $g > 0$, then we say $f = \widetilde{O}(g)$ if there exists $c \in \mathbb{R}_{>0}$ such that for all $x \in X$, $|f(x)| \leq g(x) \max\{2, \log g(x)\}^c$.

Theorem. *There is a deterministic algorithm that given* $n \in \mathbb{Z}$, $n > 1$, *correctly decides whether or not* n *is prime and that has runtime* $\widetilde{O}((\log n)^6)$.

The theorem of AKS begins (in brief) as follows. Let $n \in \mathbb{Z}_{>1}$ be a positive integer, r a prime number with $r \nmid n$, and

$$(X + a)^n \equiv X^n + a \pmod{n, X^r - 1}$$

for a small set of a, and the multiplicative order of n modulo r is at least $c(\log n)^2$. Together with other conditions, we conclude n is a prime. This has runtime $\widetilde{O}(r^{3/2}(\log n)^3)$. For $r = O((\log n)^5)$, we get $\widetilde{O}((\log n)^{21/2})$ with effective constant. For an ineffective constant, one can get $\widetilde{O}((\log n)^{15/2})$. It would be optimal to have $\widetilde{O}((\log n)^6)$, which would require $r = \widetilde{O}((\log n)^2)$ but here we run into Artin's conjecture on primes with prescribed primitive root. To get around this, we generalize the situation slightly to give more parameters.

In fact, we replace $X^r - 1$ by $(X^r - 1)/(X - 1) = X^{r-1} + \cdots + X + 1$; the proof remains essentially unchanged. Instead of considering congruences, we instead think of equalities in the ring

$$A = \mathbb{Z}[X]/(n, X^{r-1} + \cdots + X + 1) \supset \mathbb{Z}/n\mathbb{Z}.$$

This is a Galois extension of $\mathbb{Z}/n\mathbb{Z}$ of degree $r - 1$. We now replace the polynomial $X^{r-1} + \cdots + X + 1$ with other polynomials $f(X)$ for which the same proof techniques apply. In this case, A is a field if and only if n is a prime number and n is a primitive root modulo r. We are led to introduce rings that 'try very hard' to be fields.

Definition. A *pseudofield* is a pair A, α where A is a ring (commutative with 1) and $\alpha \in A$ such that there exists $n \in \mathbb{Z}_{>1}$ and $d \in \mathbb{Z}_{>0}$ satisfying:

- $\mathbb{Z}/n\mathbb{Z}$ is a subring of A with $\#A \leq n^d$, and there exists $\sigma \in \operatorname{Aut} A$ with:
  - $\sigma^d \alpha = \alpha$;
  - for all $q \mid d$ prime, $\sigma^{d/q}(\alpha) - \alpha \in A^*$; and
  - $\sigma\alpha = \alpha^n$;
- Equivalently, $A \cong (\mathbb{Z}/n\mathbb{Z})[X]/(f)$ as a ring with $X \mapsto \alpha$ for some monic polynomial $f \in (\mathbb{Z}/n\mathbb{Z})[X]$ of degree d satisfying:
  - $f(X) \mid f(X^n)$;
  - $f(X) \mid (X^{n^d} - X)$; and
  - for all primes $q \mid d$, $(f, X^{n^{d/q}} - X) = (1)$.

Also equivalently, a pseudofield is characterized by the conditions that $A \supset \mathbb{Z}/n\mathbb{Z}$ be Galois with cyclic group generated by σ, α generates A as a ring, and $\sigma\alpha = \alpha^n$.

Theorem. *There is a deterministic algorithm that given $n \in \mathbb{Z}_{>1}$ and $f \in (\mathbb{Z}/n\mathbb{Z})[X]$ decides if $(\mathbb{Z}/n\mathbb{Z})[X]/(f)$ is a pseudofield in time*

$$\widetilde{O}((d + \log n) d \log n).$$

It is routine to verify this using the second set of conditions.

Example.

A. If r is a prime with $r \nmid n$, then $A = (\mathbb{Z}/n\mathbb{Z})[X]/(X^{r-1} + \cdots + 1)$ is a pseudofield if and only if the order of n modulo r is $r - 1$.
B. If n is prime, then A, α is a pseudofield if and only if A is a field and $A = \mathbb{F}_n[\alpha]$.
C. If $A = \mathbb{Z}/n\mathbb{Z}$ (so that $d = 1$) and $\alpha = (a \bmod n)$, then A, α is a pseudofield if and only if $a^n \equiv a \pmod{n}$—in other words, n is a pseudoprime to base a.

Theorem. *Let A, α be a pseudofield of characteristic n and degree d such that $d > (\log n/\log 2)^2$, n has no prime factor $\leq k = \lfloor \sqrt{d}(\log n/\log 2) \rfloor$, and such that*

$$(\alpha + a)^n = a^n + a$$

for $a = 1, 2, \ldots, k \pmod n$. *Then* n *is a power of a prime number.*

The proof uses: for each prime $p \mid n$, there exists a unique $\tau \in \langle\sigma\rangle$ such that for all $\beta \in A$, $\tau(\beta) \equiv \beta^p \pmod{pA}$. This is a result coming from Galois theory for rings.

This leads to a deterministic primality test with runtime equal to the time to construct the pseudofield plus the time to check the conditions; the latter takes time $\widetilde{O}(d^{3/2}(\log n)^3)$.

In the context of primality testing, there is a procedure which converts any (honest) method for constructing finite fields to a method for checking primality. If the algorithm on input n crashes, then n was not prime; if it returns a polynomial f, then one checks (efficiently) if this gives rise to a pseudofield, which then verifies that n is prime, and otherwise produces a proof that n is not prime.

Therefore we look for algorithms for constructing finite fields. Our construction relies on the following theorem:

Theorem (Kummer 1846)**.** *For* r *prime and* $q \mid (r-1)$*, put*

$$f_{q,r} = \prod_{\substack{i \in \mathbb{F}_r \\ i^q = 1}} \left(X - \sum_{j^{(r-1)/q} = i} \zeta_r^j \right) \in \mathbb{Z}[X]$$

where ζ_r *is a primitive* r*th root of unity in* $\mathbb{C}$*. The polynomial* f *is monic and irreducible of degree* q*. If* p *is prime,* $p \neq r$*, then* $(f_{q,r} \bmod p) \in \mathbb{F}_p[X]$ *is irreducible if and only if the order of* $p^{(r-1)/q}$ *modulo* r *is equal to* q*.*

Fact. Let A_i, α_i be a pseudofield of characteristic n and degree $d_i > 1$ for $i = 1, 2$ such that $\gcd(d_1, d_2) = 1$; then $A_1 \otimes_{\mathbb{Z}} A_2, \alpha_1 \otimes \alpha_2$ is a pseudofield of characteristic n and degree $d_1 d_2$.

Theorem. *There exists an effective computable constant* c *such that there is a deterministic algorithm that given* $n \in \mathbb{Z}_{>1}$ *finds a finite sequence of pairs* $(r_1, q_1), \ldots, (r_k, q_k)$ *such that:*

- $q_i > 1$*,* q_i *pairwise coprime;*
- r_i *prime,* $q_i \mid (r_i - 1)$*;*
- *The order of* $n^{(r_i-1)/q_i}$ *modulo* r_i *is* q_i*; and*
- $d = \prod q_i$ *satisfies*

$$(\log n / \log 2)^2 < d < c(\log n / \log 2)^2$$

and $\max r_i < d$*.*

This algorithm runs in time $\widetilde{O}((\log n)^{24/11})$*, with an effective constant.*

A.9 Pomerance and Bleichenbacher: Constructing Finite Fields

Consider the following problem: Given a prime p and an integer $d > 1$, find an irreducible polynomial $f \in \mathbb{F}_p[X]$ of degree d. And do so in time polynomial in d and $\log p$. There is a randomized algorithm which attacks this problem by picking a polynomial at random (approximately 1 out of every d polynomials will be irreducible), and testing each for irreducibility (which is fast), continuing this procedure until you find one, then stop. But we are interested here in a deterministic algorithm. Already for $d = 2$, this is a difficult problem, equivalent to finding a quadratic nonresidue modulo p.

Assuming the ERH, Adleman and Lenstra have a solution to this problem. Unconditionally, they also find an irreducible polynomial of degree d' with $d \leq d' < cd \log p$, where c is an effectively computable number. Letting $d = (\log n)^2$ and $n = p$ (where you do not know

a priori if n is prime), the polynomial produced has degree $O(\log^3 n)$, and the runtime of the AKS algorithm becomes $\widetilde{O}((\deg f)^{3/2}\log^3 n) = \widetilde{O}((\log n)^{15/2})$. We improve this theorem to the following:

Theorem. *Such a polynomial can be produced with $d \le d' \le 4d$ for p sufficiently large and $d \ge (\log p)^{11/6+\epsilon}$.*

The bound for the 'sufficiently large' part depends effectively on the choice of ϵ.

Theorem. *There is an effectively computable function N_ϵ and a deterministic algorithm such that if $\epsilon > 0$, $n > N_\epsilon$, and $D > (\log n)^{11/6+\epsilon}$, the algorithm produces pairs $(q_1, r_1), \ldots, (q_k, r_k)$ such that for each i, r_i is prime, $r_i < D$, $q_i \mid r_i - 1$, the order of $n^{(r_i-1)/q_i}$ modulo r_i is q_i. Further, the $q_1, \ldots, q_k$ are pairwise coprime, and $D \le \prod_i q_i \le 4D$. This algorithm runs in time $\widetilde{O}_{\text{eff}}(D^{12/11})$.*

Let η_i be the Gaussian period of degree q_i in $\mathbb{Q}(\zeta_{r_i})$ (the trace of ζ_{r_i} into the unique subfield of degree q_i over $\mathbb{Q}$). The element $\eta = \eta_1 \ldots \eta_k$ has degree $q_1 \ldots q_k$ over the rationals (by coprimality). If n is prime, and if $f(x)$ is the minimal polynomial of η then $f \bmod n$ is irreducible over $\mathbb{F}_n$. Checking if $f(X) \mid f(X^n)$ in $(\mathbb{Z}/n\mathbb{Z})[X]$ (together with some other conditions), we see that f gives rise to a pseudofield.

Let $x = D^{6/11-\epsilon/4}$. Throughout, we assume that n is 'sufficiently large' with the bound being effectively computable, depending only on the choice of ϵ.

Proposition. *All but $O(x/\log^3 x)$ primes $r \le x$ have a prime $q \mid (r-1)$ with $q > x^{1/(\log\log x)^2}$ and the order of $n^{(r-1)/q}$ modulo r is equal to q.*

Therefore up to x, almost all of the primes are useful in the context of our theorem. This proposition is a natural extension of the argument in the original AKS paper, together with an added ingredient about the distribution of primes r such that $r-1$ is smooth due to Pomerance and Shperlinski.

Proposition. *Let Q be a set of primes q with $x^{1/(\log\log x)^2} < q \le x^{1/2}$ and $\sum_{q\in Q} 1/(q-1) < (3-\epsilon)/11$. Then there are $> \delta x/\log^2 x$ primes $r \le x$ such that $r-1$ is free of primes from $Q \cup (x^{1/2}, x)$.*

This proposition follows from a method of Balog, together with some effective estimates on the distribution of primes in residue classes. Together, these two propositions give the corollary:

Corollary. *Let Q be the set of primes q in the first proposition satisfying $q \le \sqrt{x}$. Then*

$$\sum_{q\in Q} \frac{1}{q-1} \ge \frac{3-\epsilon}{11}.$$

Proof. If this inequality did not hold, then by the first and second propositions, there must be primes $r \le x$ having the properties of both propositions. By the first proposition, $r-1$ has a prime factor $q > x^{1/)\log\log x)^2}$ and the order of $n^{(r-1)/q}$ modulo r is equal to q. By one of the properties in the second proposition, $q \le x^{1/2}$. Then $q \in Q$. This contradicts the second proposition. □

Proposition. *There exists a subset of Q in the corollary with product in the interval $[D, 4D]$.*

The proof of this theorem relies upon combinatorial number theory (essentially, you can solve a bin packing problem using the primes q). It relies upon:

Theorem (Continuous Frobenius theorem)**.** *If S is an open subset of $\mathbb{R}_{>0}$, S is closed under addition, and $1 \notin S$, then for any t, $0 < t \leq 1$, the du/u measure of $S \cap (0,t)$ is $\leq t$, i.e.*

$$\int_0^t \chi_S(u)\frac{du}{u} \leq t$$

where $\chi_S(u)$ is the characteristic function of S.

Now a remark about effectivity. It was proven by de la Vallée Poussin in 1896 that

$$\pi(x,k,a) = \#\{p \leq x : x \equiv a \pmod{k}\} \sim \pi(x)/\phi(k)$$

as $x \to \infty$. The Siegel-Walfisz theorem states that this is true for $k < (\log x)^A$ for any fixed A. This theorem is inherently ineffective because it depends on the existence or nonexistence of Siegel zeros. The Siegel-Walfisz theorem is ubiquitous in analytic number theory, being used in the Bombieri-Vinogradov theorem, Fouvry's theorem, and much else. The analytic number theory we use is an effective version of the Bombieri-Vinogradov theorem that does not rely on the Siegel-Walfisz theorem, and we replace the Fouvry theorem by a (weaker) result of Deshouillers-Iwaniec.

Now we give an outline of the proof of the Frobenius theorem. First, it is sufficient to prove the case where $S_t = S \cap (0,t) = \bigcup_{i=1}^n (a_i, b_i)$ (i.e. S_t contains only finitely many intervals). Second, since $1 \notin S$, for all $(h_1, \ldots, h_n) \in \mathbb{N}_{\geq 0}^n$ either $\sum_{i=1}^n h_i a_i \geq 1$ or $\sum_{i=1}^n h_i b_i \leq 1$ (*). Now fix $b_1, \ldots, b_n$ such that $b_1 > b_2 > \cdots > b_n$ and consider all sets $\bigcup_{i=1}^n (a_i, b_i)$ satisfying $b_1 \geq a_1 \geq b_2 \geq \cdots \geq b_n \geq a_n$ (**) as well as condition (*). Under these conditions, there exists a maximum to $\sum_{i=1}^n (\log(b_i) - \log(a_i))$. We may thus assume that $S_t = \bigcup_{i=1}^n (a_i, b_i)$ is a maximum. We show in the paper that we can assume that $b_1 > a_1 > b_2 > \cdots > b_n > a_n$.

Let $U = \{h \in \mathbb{N}_{\geq 0}^n : ha = 1\}$ with $a = (a_1, \ldots, a_n)$. Let $a_{n+1} = b_{n+1} = 0$. For all $h = (h_1, \ldots, h_n) \in U$ and $1 \leq k \leq n$,

$$h_k \left(\sum_{i=1}^n (b_i - a_i) h_i \right) \leq h_k (b_k - b_{k+1}).$$

This is trivial if $h_k = 0$, and otherwise, $\sum_{i=1}^n a_i h_i = 1$ which implies

$$\sum_{i=1}^n a_i h_i - a_k + a_{k+1} < 1.$$

and therefore by assumption (*)

$$\sum_{i=1}^n b_i h_i - b_k + b_{k+1} \leq 1.$$

Let $v = (v_1, \ldots, v_n) \in \mathbb{R}^n$ such that $vh \geq 0$ for all $h \in U$. Then there exists $\epsilon > 0$ such that for all $0 \leq x \leq \epsilon$,

$$\bigcup_{i=1}^n (a_i + v_i x, b_i)$$

satisfies (*) and (**). By assumption

$$\sum_{i=1}^{n}(\log b_i - \log(a_i + v_i x))$$

is maximal for $x = 0$, so $\sum_{i=1}^{n} v_i/a_i \geq 0$. Then a theorem by Farkas (or the dual theorem of linear programming) implies that there exists $p_j \geq 0$ such that

$$\sum_{i=1}^{\ell} h_{ij} p_j = \frac{1}{a_i}$$

where $U = \{h_1, \ldots, h_\ell\}$ and $h_j = (h_{1j}, \ldots, h_{nj})$. Now multiply the equation

$$h_k \left(\sum_{i=1}^{n} (b_i - a_i) h_i \right) \leq h_k(b_k - b_{k+1}).$$

by $a_k p_j$ and sum up

$$\sum_{k=1}^{n}\sum_{j=1}^{\ell} a_k p_j h_{kj} \left(\sum_{i=1}^{n} (b_i - a_i) h_{ij} \right) \leq \sum_{k=1}^{n}\sum_{j=1}^{\ell} a_k p_j h_{kj}(b_k - b_{k+1}).$$

After some simple arithmetic, we find that

$$\sum_{i=1}^{n}(\log b_i - \log a_i) \leq t.$$

A.10 Silverberg: Applications of Algebraic Tori to Crytography

In this lecture we discuss 'torus-based cryptography', and counterexamples to conjectures in an article entitled *Looking Beyond XTR*, and compare TBC with Lucas-based cryptosystems and XTR, and understand LUC, XTR, and *Beyond* in terms of algebraic tori. This is joint work with Karl Rubin, and inspired by XTR.

The cryptosystem XTR (due to A. Lenstra and E. Verheul) concerns the extension $\mathbb{F}_{p^6}/\mathbb{F}_{p^2}$: they consider the subgroup of $\mathbb{F}_{p^6}^*$ of order $p^2 - p + 1$ with generator g; the public knowledge is $\mathrm{Tr}_{\mathbb{F}_{p^6}/\mathbb{F}_{p^2}}(g)$, what is shared is $\mathrm{Tr}_{\mathbb{F}_{p^6}/\mathbb{F}_{p^2}}(g^{ab})$, where $\mathrm{Tr}(g^a)$ and $\mathrm{Tr}(g^b)$ are transmitted. In this setup, you get the security of $\mathbb{F}_{p^n}^*$ while transmitting only $\phi(n)$ elements of $\mathbb{F}_p$.

Let L/k be a finite cyclic extension with intermediate field F. Let $g \in L \setminus F$, and denote by C_g the $\mathrm{Gal}(L/F)$-conjugacy class of g, so the characteristic polynomial of g over F is $\prod_{h \in C_g}(X - h)$. For $L = \mathbb{F}_{p^6}$, $F = \mathbb{F}_{p^2}$, $k = \mathbb{F}_p$, the polynomial is $x^3 - s_1 x^2 + s_2 x - s_3$, where $s_1 = \mathrm{Tr}_{L/F}(g)$, $s_3 = N_{L/F}(g)$, and $s_2 = \mathrm{Tr}_{L/F}(gg^\sigma)$, where $\langle\sigma\rangle = \mathrm{Gal}(L/F)$. If g is in the subgroup of L^* of order $p^2 - p + 1$, then $s_3 = 1$ and $s_2 = \mathrm{Tr}_{L/F}(g)^p$. Therefore knowing $\mathrm{Tr}_{L/F}(g)$ is equivalent to knowing all the symmetric polynomials on C_g which is equivalent to knowing C_g as a set, so you know C_{g^a} and this is equivalent to knowing $\mathrm{Tr}_{L/F}(g^a)$. In other words, you can exponentiate, but you cannot multiply: $\mathrm{Tr}(g)$ and $\mathrm{Tr}(h)$ do not determine $\mathrm{Tr}(gh)$; i.e. knowing C_g and C_h does not allow you to know C_{gh}.

Bosma-Hutton-Verheul conjecture that for all n, there exists a divisor $d \mid n$ such that $d \mid \phi(n)$ and for $L = \mathbb{F}_{p^n}$ and $F = \mathbb{F}_{p^d}$, you can recover all the coefficients $s_1, \ldots, s_{n/d}$ of the

characteristic polynomial of g over F from the first $\phi(n)/d$ of them for all g in the subgroup of L^* of order $\Phi_n(p)$ and not in any proper subfield.

We show that this is in fact false. In particular, when $n = 30$, it is false; for $p = 7$, $d = 1$, no 10 symmetric polynomials determine all of them, and no 8 determine any of the others (except the ones determined by the symmetry of the characteristic polynomial). For $p = 7$, $d = 2$, no 4 symmetric polynomials determine all of them.

Fact. The order $\Phi_n(p)$ subgroup of $\mathbb{F}_{p^n}^*$ is

$$\{\alpha \in \mathbb{F}_{p^n}^* : N_{\mathbb{F}_{p^n}/M}(\alpha) = 1 \text{ for all } M \subsetneq \mathbb{F}_{p^n}\}.$$

Let L/k be an abelian degree n extension of fields. Let

$$T_{L/k} = \ker\left(\mathrm{Res}_{L/k}(\mathbb{G}_m) \xrightarrow{\oplus N_{L/M}} \bigoplus_{k\subset M\subsetneq L} \mathrm{Res}_{M/k}(\mathbb{G}_m)\right);$$

recall that $\mathbb{G}_m(k) = k^*$, and $(\mathrm{Res}_{L/k}\,\mathbb{G}_m)(k) = L^*$. If L/k is not cyclic, $\dim(T_{L/k}) = 0$. If L/k is cyclic, then $T_{L/k}$ is an algebraic torus over k of dimension $\phi(n)$, i.e. $T_{L/k}$ is isomorphic over $\overline{k}$ to $\mathbb{G}_m^{\phi(n)}$. Here, $T_{L/k} \cong_L \mathbb{G}_m^{\phi(n)}$.

Assume from now on that L/k is cyclic. Then

$$T_{L/k}(k) = \{\alpha \in L^* : N_{L/M}(\alpha) = 1 \text{ for all } k \subset M \subsetneq L\}.$$

Conjecture (Voskresenskii)**.** *$T_{L/k}$ is rational, i.e. there exists a birational map $T_{L/k} \to \mathbb{A}^{\phi(n)}$.*

This is true if $n = p^a$ or p^aq^b (Klyachko 1988). It is not known for $n = pqr$. Let us look at the case when $n = 2$, say, $\mathrm{char}\, k \neq 2$. Then $L = k(\sqrt{d})$, and $T_{L/k} = \ker(N_{L/k})$ which is a conic which can be parameterized, and we obtain

$$\begin{aligned}\psi : \mathbb{P}^1 &\xrightarrow{\sim} T_{L/k}\\ a &\mapsto (a+\sqrt{d})/(a-\sqrt{d})\\ \infty &\mapsto 1\end{aligned}$$

We have $\psi(a)\psi(b) = \psi((ab+d)/(a+b))$. (This is also just Hilbert's theorem 90.) Writing T_n for $T_{L/k}$ when $k = \mathbb{F}_q$, this induces a way to do the multiplication in T_2 in $\mathbb{P}^1(k)$.

We also give an explicit example when $n = 6$ for $\mathrm{char}(k) \neq 3$, $[k(\zeta_9) : k] = 6$ (e.g. $k = \mathbb{F}_q$ with $q \equiv 2, 5 \pmod 9$). $T_{L/k}$ is dimension 2 and contained in $\mathrm{Res}_{M/k}(T_{L/M}) = T' \sim_k \mathbb{A}^3$ where M/k is the subextension of degree 3. But $T_{L/M} = \ker(N_{L/M})$ is dimension 1, and $N_{L/F} = 1$ defines a hypersurface in $T' \sim \mathbb{A}^3$. Therefore $T_6 \sim \mathbb{A}_2$, so we can use the multiplication in T_6 but represent elements of T_6 by 2 elements of $\mathbb{F}_8$. This gives rise to the cryptosystem CEILIDH.

Open problems:

A. Improve the efficiency of multiplication and exponentiation for the system CEILIDH.
B. Repeat this analysis for $n = 30$, i.e.
    1. Find explicit birational isomorphisms between T_{30} and $\mathbb{A}^8$,
    2. Find prime powers q of size $1024/30 \approx 35$ bits such that $\Phi_{30}(q)$ has a 160-bit prime factor,
    3. Are there special attacks on DL in $\mathbb{F}_{q^{30}}^*$?

Now we look to understand LUC, XTR, and *Beyond* in terms of algebraic tori. For H a subgroup of $\mathrm{Gal}(L/k) = G$ which is a direct factor, write Σ_H for the group of permutations of H. Then Σ_H acts on

$$\bigoplus_{\sigma \in G} \mathbb{A}^1 \xrightarrow{\sim}_L \mathrm{Res}_{L/k}\, \mathbb{A}^1 \supset \mathrm{Res}_{L/k}(\mathbb{G}_m).$$

Let

$$X_F = \mathrm{img}\left(T_{L/k} \to \mathrm{Res}_{L/k}(\mathbb{G}_m)/\Sigma_{\mathrm{Gal}(L/F)}\right).$$

For LUC and XTR, you look at

$$\{\mathrm{Tr}_{L/F}(\alpha) : \alpha \in T_{L/k}(k)\}$$

which is the image of $T_{L/k}(k)$ under $T_{L/k} \to X_F \to \mathrm{Res}_{F/k}(\mathbb{A}^1)$, where the latter map is a birational isomorphism.

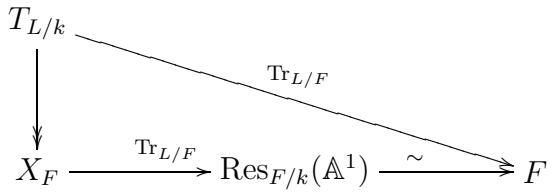

Assume that n is squarefree. Write $\mathrm{Gal}(L/F) = H_1 \times \cdots \times H_t$ where H_i are cyclic of prime order.

Theorem. *The action of* Σ_{H_i} *on* $\mathrm{Res}_{L/k}(\mathbb{G}_m)$ *preserves* $T_{L/k}$ *and* X_F *is birational to* $T_{L/k}/(\Sigma_{H_1} \times \cdots \times \Sigma_{H_t})$.

In XTR, we obtain X_F birational to T_6/S_3; in the cases $n = 30$ and $d = 1, 2$, one gets $T_{30}/(S_2 \times S_3 \times S_5)$, $T_{30}/(S_3 \times S_5)$, which are not groups.

We have maps $L \to F$ for every symmetric function $s_1, \ldots, s_{[L:F]}$. We have a surjection $T_{L/k} \to X_F$ and an injection $X_F \hookrightarrow F^{[L:F]}$ by the direct sum of these functions. A BHV conjecture implies that for the subset consisting of the first $\lceil \phi(n)/d \rceil$ functions (where $d = [F : k]$), the map remains injective. Further, a BHF conjecture implies that every n has a divisor d so that d also divides $\phi(n)$, and the map $X_F \to \mathbb{A}^{\phi(n)}$ (induced by the first $\phi(n)/d$ symmetric functions) is a birational isomorphism. This is true for $(n, d) = (1, 1)$ (DH), $(2, 1)$ (LUC), $(6, 2)$ (XTR), and $(\ell, 1)$ and $(2\ell, 2)$ where ℓ is prime (see *Doing more with fewer bits*, by Brouwer-Pellikaan-Verheul), but:

Theorem. *This is false for* $n = 30$ *(*$d = 1, 2$*) if* $\mathrm{char}(k)$ *lies outside a finite set.*

To prove this, we first do a computer search for 2 elements of $T_{30}(\mathbb{F}_7)$ with the same image $a \in (\mathbb{F}_7)^8$ but different images in X_F. Using Hensel's Lemma, every lift of a to $\mathbb{Z}_7^8$ has at least 2 inverse images in $X_F(\mathbb{Q}_7)$. Therefore the map is not generically one-to-one over $\mathbb{Q}_7$, so it is not generically one-to-one over $\mathbb{Q}$, hence over any field of characteristic 0. Then reduce modulo p to get it over $\mathbb{F}_p$ and therefore all fields of characteristic p (outside of a finite set).

A.11 Stein: Modular Forms Database

The lecture notes[5] are available on the speaker's website. The tables[6] are also available there.

[5] http://modular.fas.harvard.edu/mfd/talks/mfd1/

[6] http://modular.fas.harvard.edu/Tables/

A.12 Voloch: Multiplicative Subgroups of a Finite Field

The lecture notes[7] are available on the speaker's website.

A.13 Wan: Partial Counting of Rational Points over Finite Fields

We are motivated by the following problem. Let $\mathbb{F}_{q^d} = \mathbb{F}_q[\alpha]/h(\alpha)$, where h is irreducible of degree $d > 1$ over $\mathbb{F}_q$. We look at the group

$$G = \langle a - \alpha : a \in \mathbb{F}_q \rangle = (\mathbb{F}_{q^d}^*)^I \subset \mathbb{F}_{q^d}^*,$$

where $I = [\mathbb{F}_{q^d}^* : G]$. When does $I = 1$, for example?

Let $D \mid (q^d - 1)$, and let

$$N = \#\{(x, y) : x - \alpha = y^D, x \in \mathbb{F}_q, y \in \mathbb{F}_{q^d}\}.$$

By a character sum argument counting, you can write this as

$$N = \sum_{\substack{\phi : \mathbb{F}_{q^d}^* \to \mathbb{C}^* \\ \phi^D = 1}} \sum_{x \in \mathbb{F}_q} \phi(x - \alpha).$$

By the Riemann hypothesis (Weil), we have $|N - q| \leq (D-1)(d-1)\sqrt{q}$. Therefore we have seen:

Proposition. *Let $S \subset \mathbb{F}_q$. Let $G_S = \langle a - \alpha : a \in S \rangle = (\mathbb{F}_{q^d}^*)^{I_S}$. Then*

$$(\#S) I_S \leq N \leq q + (I_S - 1)(d-1)\sqrt{q}.$$

If $(\#S) > (d-1)\sqrt{q}$, then

$$I_S \leq \frac{q - (d-1)\sqrt{q}}{(\#S) - (d-1)\sqrt{q}}.$$

In particular, if $(\#S) = q$, then $I_S = 1$ and $G = \mathbb{F}_{q^d}^*$.

Therefore we consider the problem: Can we compute N in time polynomial in d, D, and $\log q$?

The general setup: Let $f(x_1, \ldots, x_n) \in \overline{\mathbb{F}_q}[x_1, \ldots, x_n]$, and $d_1, \ldots, d_n \geq 1$. We want to count

$$N_{d_1,\ldots,d_n}(f) = \#\{(x_1, \ldots, x_n) : f(x_1, \ldots, x_n) = 0,\ x_i \in \mathbb{F}_q^{d_i}\}.$$

Can we compute $N_{d_1,\ldots,d_n}(f)$, or at least estimate it? How does this quantity vary when the d_i vary?

For example, we consider the Artin-Schreier hypersurface. Let

$$f(x_1, \ldots, x_n, y_1, \ldots, y_{n'}) \in \mathbb{F}_q[x_1, \ldots, x_n, y_1, \ldots, y_{n'}],$$

where $n, n' \geq 1$. For each $d \geq 1$, we consider

$$N_d(f) = \#\{(x_0, \ldots, x_n, y_1, \ldots, y_{n'}) : x_0^p - x_0 = f(x_1, \ldots, x_n, y_1, \ldots, y_{n'}),$$
$$x_i \in \mathbb{F}_{q^d}, y_j \in \mathbb{F}_q\}.$$

Heuristically (for suitable f), we expect

$$N_d(f) = q^{dn+n'} + O(q^{(dn+n')/2})$$

where the constant depends on p, f, and d.

[7] http://www.ma.utexas.edu/users/voloch/preprint.html

Theorem (Deligne)**.** *Write $f = f_m + f_{m-1} + \cdots + f_0$, where f_i are homogeneous of degree i. Assume f_m defines a smooth projective hypersurface in $\mathbb{P}^{n+n'-1}_{\mathbb{F}_q}$, and that $p \nmid m$, $d = 1$. Then*

$$|N_1(f) - q^{n+n'}| \leq (p-1)(m-1)^{n+n'} q^{(n+n')/2}.$$

What about $d > 1$?

Definition. If $d \geq 1$, we define the *dth fibred sum of f* to be

$$\bigoplus_y^d f = f(x_{11}, \ldots, x_{1n}, y_1, \ldots, y_{n'}) + \cdots + f(x_{d1}, \ldots, x_{dn}, y_1, \ldots, y_{n'}).$$

Theorem (Fu-W)**.** *Write $f = f_m + \cdots + f_0$, and assume that $\bigoplus_y^d f_m$ is smooth in $\mathbb{P}^{dn+n'-1}_{\mathbb{F}_q}$ and $p \nmid m$. Then*

$$|N_d(f) - q^{dn+n'}| \leq (p-1)(m-1)^{dn+n'} q^{(dn+n')/2}.$$

Example. In the case that we can write

$$f(x, y) = f_{1m}(x_1, \ldots, x_n) + f_{2m}(y_1, \ldots, y_{n'}) + f_{\leq m-1}(x, y),$$

and f_{1m} is smooth in $\mathbb{P}^{n-1}_{\mathbb{F}_q}$, f_{2m} is smooth in $\mathbb{P}^{n'-1}_{\mathbb{F}_q}$. Then $\bigoplus_y^d f_m$ is smooth in $\mathbb{P}^{dn+n'-1}_{\mathbb{F}_q}$ if and only if $p \nmid d$.

Since the condition that the fibred sum be smooth is Zariski open, we have shown it is nonempty if $p \nmid d$ and therefore there exist many examples of such f to which the theorem applies.

Definition. Let M_d be the set of f over $\overline{\mathbb{F}_q}$ such that $\bigoplus_y^d f_m$ is smooth. Then M_d is Zariski open in the set of all f over $\overline{\mathbb{F}_q}$ with $\deg f \leq m$.

Theorem (Gao-W)**.** *M_d is Zariski dense if and only if $p \nmid d$. In fact,*

$$\bigcap_{\substack{d=1 \\ p \nmid d}}^{\infty} M_d = \bigcap_{\substack{d=1 \\ p \nmid d}}^{p^{(m-1)^n}} M_d$$

and this intersection is Zariski open and dense.

Problem. *What about Kummer hypersurfaces*

$$x_0^D = f(x_1, \ldots, x_n, y_1, \ldots, y_n')$$

where $x_i \in \mathbb{F}_{q^d}$ and $y_j \in \mathbb{F}_q$?

Remark. We expect $|N_d - q^{dn+n'}| = O(q^{(dn+n')/2})$, but one can get the weaker estimate $O(q^{dn/2+n'-1/2})$ in many cases (Katz).

Now we consider partial zeta functions over $\mathbb{F}_q$. Let $f(x_1, \ldots, x_n) \in \mathbb{F}_q[x_1, \ldots, x_n]$, $d_1, \ldots, d_n \geq 1$. Define

$$Z_{d_1,\ldots,d_n}(f, T) = \exp\left(\sum_{k=1}^{\infty} N_{d_1,\ldots,d_n,k} \frac{T^k}{k}\right)$$

where

$$N_{d_1,\ldots,d_n,k} = \#\{(x_1, \ldots, x_n) : f(x_1, \ldots, x_n) = 0, x_i \in \mathbb{F}_{q^{d_i k}}\}.$$

Without loss of generality, we may assume $\gcd(d_1, \ldots, d_n) = 1$, since otherwise we can just enlarge the ground field $\mathbb{F}_q$.

Proposition.

A. *If* $d_1 \mid \cdots \mid d_n$, *then* $Z_{d_1,\ldots,d_n}(f,T) \in \mathbb{Q}(T)$.

B. (Faltings) $Z_{d_1,\ldots,d_n}(f,T) = \prod_{i=1}^{d} P_i(T)^{\zeta_d^i}$, *where* $d = \operatorname{lcm}(d_1,\ldots,d_n)$ *and* ζ_d *is a primitive* d*th root of unity,* $P_i(T) \in \overline{\mathbb{Q}}(T)$, *and* $P_i(0) = 1$.

By exponentiation to a root of unity, we mean the formal binomial expansion. From a counting point of view, this is 'as good as rational'.

Theorem. *In all cases,* $Z_{d_1,\ldots,d_n}(f,T) \in \mathbb{Q}(T)$.

Proof. Let $X^d = \underbrace{X \times \cdots \times X}_{d}$. We have a map

$$\sigma : X^d \to X^d$$
$$(x^{(1)},\ldots,x^{(d)}) \mapsto (x^{(d)},x^{(1)},\ldots,x^{(d-1)})$$

Faltings constructs a large subvariety $Y_{d_1,\ldots,d_n} \hookrightarrow X^d$ which is stable under σ. Then

$$N_{d_1,\ldots,d_n}(f) = \#\mathrm{Fix}(\sigma \circ \mathrm{Frob}_q \,|Y(\overline{\mathbb{F}_q})) = \sum(-1)^i \operatorname{Tr}(\sigma \circ \mathrm{Frob}\,|H_c^i(\overline{Y})).$$

Note $\sigma \circ \mathrm{Frob}_q = \mathrm{Frob}_q \circ \sigma$. This implies Faltings' 'near' rationality as in the proposition.

Now $Z_{d_1,\ldots,d_n}(f,T) \in 1 + T\mathbb{Q}[[T]]$. We refine the above argument as follows. First, for $\gcd(a,d) = 1$, you have

$$N_{d_1,\ldots,d_n}(f) = \#\,\mathrm{Fix}(\sigma^a \circ \mathrm{Frob}_q \,|Y(\overline{\mathbb{F}_q})).$$

We consider $Y \to Y/G$, where $G = \langle\sigma\rangle \cong \mathbb{Z}/d\mathbb{Z}$. We have a character $\chi : G \to \mathbb{C}^*$, and define the L-function

$$L(\chi,T) = \exp\left(\sum_{k=1}^{\infty} \frac{T^k}{k}\left(\frac{1}{d}\sum_{\tau\in G}\chi(\tau^{-1})\#\,\mathrm{Fix}(\tau \circ \mathrm{Frob}_q^k \,|Y(\overline{\mathbb{F}_q}))\right)\right).$$

By Grothendieck, $L(\chi,T) \in \mathbb{Q}(\zeta_d)(T)$. This implies that

$$Z_{d_1,\ldots,d_n}(f,T)^{\phi(d)} = \prod_{\chi\in\widehat{G}} L(\chi,T)^{\sum_{\gcd(a,d)=1}\chi(\sigma)^a} \in \mathbb{Q}(\zeta_d)(T)$$

so $Z_{d_1,\ldots,d_n}(f,T) \in \mathbb{Q}(T)$ (essentially by unique factorization). □

Open problem: can you bound the total degree of $Z_{d_1,\ldots,d_n}(f,T)$? The best bound we have is $3 \cdot 2^{d+1}(3+dm)^{d_1+\cdots+d_n+1}$, $m = \deg f$. Can this be improved to $O(d)^{O(1)}$? Yes, if $d_1 = \cdots = d_r = d$ and $d_{r+1} = \cdots = d_n = 1$ (Fu-W).

Chapter B: Problems

Problem/Question 1. If

$$(X-1)^n \equiv X^n - 1 \pmod{n, X^r - 1}$$

and $\gcd(r,n) = 1$ does this imply that

$$n^2 \equiv 1 \pmod{r}$$

when n is composite? (AKS)

Remarks.

(i) In the simplest case, $r = 5$ and $n \equiv 2 \pmod 5$, what do the heuristics say? This feels like the traditional series of pseudoprime tests, and although they were first thought to be sufficient, counterexamples exist. (Bernstein) It was thought that the old pseudoprime tests $2^n \equiv 2 \pmod n$ combined with the quadratic test with least discriminant d with the Jacobi symbol $(d/n) = -1$ (constructing a Lucas sequence of discriminant d) would be enough to show that n is prime. There should be infinitely many composite numbers which pass both of these tests. There is no known example, and there is a \$620 prize to find an example of a number that pass the various tests. (Pomerance) There are no heuristic reasons yet to believe this. (AKS)

(ii) For $r \geq 5$, there are 5000 pairs (n, r) that all satisfy these conditions. (AKS)

(iii) The most naive heuristic (looking at 2 as a random element modulo n) for the first claim relies upon the fact that $\sum 1/n$ diverges, whereas in this case we are looking at $\sum 1/n^r$ which converges. (Lenstra) These heuristics need to take into account smoothness. (Bernstein)

(iv) Is there a reason why $n^2 \equiv 1 \pmod r$ comes into the play? For all counterexamples with $r \geq 7$, $p \mid n$ implies $p^2 \equiv 1 \pmod r$. (Lenstra, AKS)

(v) In these conditions we see that $r^n \equiv r \pmod n$. (Lenstra) So if $p \mid n$, the numbers $p - 1$, $p+1$, p^2+1 must be very smooth (for $r = 5$). The heuristics show that there should be 'lots' of primes p satisfying these three numbers; create many n from these p, and as in the case of Carmichael numbers, then perhaps for some n this should fail. (Pomerance) If $p - 1$ is smooth for all $p \mid n$ and n is squarefree, then the multiplicative group of n has smooth order. The maximal order of any element in that group can be made small, so it would not be unusual for $n^2 \equiv 1 \pmod r$. Much of this can be found in Grantham's thesis. (Pomerance)

(vi) These heuristics are compatible with the AKS primality test because there we have r growing with n. (Bernstein)

(vii) What about $r \gg (\log \log n)^2$? The largest r found was 97, and for $r \geq 13$ we found only approximately 40 such elements. (AKS)

(viii) Are prime powers special? (Lenstra) For $r \geq 5$, all n found were squarefree. (AKS) The search was done for $n \leq 10^{11}$, $r < 100$.

Problem/Question 2. Consider $\mathbb{F}_p$ and $d \in \mathbb{Z}_{>0}$, and $S \subset \mathbb{F}_p$ with $\#S = d$. Let $h(X) \in \mathbb{F}_p[X]$ with $h(s) \neq 0$ for all $s \in S$. Let G be the group generated by $X - s$ for $s \in S$. In the original AKS paper, we have $\#G \geq 2^d$. There are techniques for getting $\#G \geq (5.82)^d$ (using lattice point counting); there seems to be much more room for improvement (perhaps using ABC). (Bernstein, Voloch) Perhaps

$$(\log \#G)/d \to \infty$$

as $d \to \infty$ uniformly over h.

Remarks.

(i) It is hard to find examples with $\#G$ smaller than p^d. Every improvement speeds up by a constant factor the AKS by a factor related to square of the base of the improvement. (Bernstein)

(ii) Does it help to know if S is an interval? (Voloch, Lenstra) Not used thus far.

(iii) For Kummer extensions, the extensions will look like multiplicative cosets of a root of unity times a single number—can this be used? (Bernstein)

(iv) Is this problem independent of h? (Cohen) For example, $h(X) = X^d - 1$ with group generated by X has small order. (Bernstein)

(v) Instead of looking at $\mathbb{F}_p[X]/(h)$, look at an elliptic curve E over $\mathbb{F}_q$ and the subgroup of $E(\mathbb{F}_q)$ generated by simple x coordinates. (Elkies) Using Weil restriction of scalars, you are looking at a curve C inside of an abelian variety over $\mathbb{F}_p$ and look at the subgroup of points generated by $C(\mathbb{F}_p)$. This is then amenable to class field theory techniques. (Voloch) The interesting case is the analogous case with AKS: $q = p^d$ and $\#S = d$, $\dim A = d$.

Problem/Question 3. Find a deterministic polynomial time algorithm for recognizing perfect numbers. (Lenstra)

Given two monic polynomials $f, g \in \mathbb{Z}[X]$ with no common irreducible factor, and two integers $n, m \in \mathbb{Z}_{>0}$, find in polynomial time all $x \in \mathbb{Z}$ such that $f(x) \mid m$ and $g(x) \mid n$, with $|x| < H(f, g, m, n)$ for some function H.

Remarks.

(i) The first problem would be more interesting than perfect numbers themselves. You are given the numbers m, n in binary. (Lenstra) A solution to the second problem gives a solution to the first: If x is prime and $x^k \parallel n$, and n is perfect, then $1 + x + \cdots + x^k \mid 2n$.

(ii) There is one solution for very small H which is to just try out the necessary values of x.

(iii) There are some results on when $m \mid f(x)$, which has typographical similarity to our problem. (Coppersmith)

(iv) If n is a perfect number, there is only one choice of x^k where the factor 2 is irrelevant; therefore you are reduced to the case where $m = n$. (Pomerance)

(v) By Gary Miller's thesis, if you are given a multiple of $\phi(n)$, you can factor n using a randomized algorithm or deterministically under the GRH. (Pomerance) You can replace $\phi(n)$ by $\sigma(n)$. (Lenstra)

(vi) If you allow randomization, the first problem should be doable. (Lenstra) There is a paper of Bach-Shallit.

(vii) This algorithm will recognize perfect numbers but it may not recognize imperfect numbers. (Lenstra)

(viii) There is a more general notion (*multiply perfect*) where $\sigma(n)/n = k$ has small height. Does this affect the problem? (Elkies) No, because you try out each k one at a time, and for fixed k this is virtually identical to the original problem.

(ix) A different problem is to just consider $f = g$, or to look at rational functions which are integer-valued.

(x) Are there heuristics for the number of such x which are related to heuristics for the largest prime divisor of $f(x)$? (Pomerance)

(xi) Look at $f(x) = (x^4 + x^5 + \cdots + x^8) \mid n$. Choose $h \in \mathbb{R}$ and $g \approx e^{\sqrt{8 \log n \log h}}$. Then you can find the set of all x with $|x| \leq h$ such that $\gcd(f(x), n) > g$. This can be done 'reasonably fast'. (Bernstein) This is the state of the art due to LLL, but it completely fails to solve the problem.

Problem/Question 4. These questions are motivated by the questions posed by AKS concerning finding quadratic nonresidues modulo a prime p.

(a) It is known that finding a single quadratic nonresidue for a given prime is polynomially equivalent to solving all quadratic equations. How far can one do this for finding a single bit of data (or few bits of data) for higher degrees? (Elkies)

(b) We do not know yet that there is a deterministic polynomial time algorithm for finding quadratic nonresidues. Is there a subexponential algorithm?

Remarks.

(i) The best known deterministic algorithm is due to Burgess and Vinogradov which runs in time $p^{1/(4\sqrt{e})}$. (Pomerance) This is sheer enumeration. If you allow exponential time, is there something better than enumeration? (Bernstein)

(ii) For cubic extensions, you can use Cardano's formula. (Gao) Is it equivalent to finding a quadratic nonresidue in the cubic extension given by a cubic nonresidue? (Elkies) If cubics includes quadratics, then you can make a quadratic extension. Then there is a trick due to Berlekamp which allows you to solve cubics in the extension by solving them in the ground field. (Lenstra)

(iii) If you know a kth nonresidue in the appropriate extension, then you can factor polynomials up to degree six (unpublished). (Gao) The Galois group is cyclic, so solvability by radicals applies.

(iv) Given a quadratic nonresidue, any even degree polynomial f with $f \mid (X^p - X)$ can be split nontrivially deterministically in polynomial time, due to Ronyai. (Lenstra) But this does not determine all solutions. You can specify quadratic conditions that must be satisfied by the factors. (Gao) There is a certain combinatorial structure on systems of roots which must be attended to, and you run into difficulties at degree 7. Conjecturally, you should be able to go higher.

(v) If you have GRH then you can deterministically in polynomial time solve quadratic extensions. Can you solve higher degree equations? (Elkies) On the GRH, you can construct appropriate nonresidues (since you can construct finite fields). (Lenstra) You can do it for fixed degree in time $n^{\log n}(\log p)^{O(1)}$ with the GRH due to Ronyai, Evdokimov. (Cheng)

Problem/Question 5. Suppose that you have a nonconstant family of elliptic curves E_λ over $\mathbb{F}_p$ (e.g. the Frey curves), $\lambda \in \mathbb{P}^1(\mathbb{F}_p)$. Can you find deterministically λ such that E_λ and $E_{\lambda+1}$ are not isogenous (i.e. $\#E_\lambda(\mathbb{F}_p) \neq \#E_{\lambda+1}(\mathbb{F}_p)$). (Elkies)

Remarks.

(i) You want p sufficiently large and the degree small. (Elkies) If you can do this, then you should also be able to do it for λ and $1-\lambda$, and then you can 'separate them apart' by applying Schoof's algorithm and obtain a deterministic square root.

(ii) Over $\mathbb{Q}$, there are only finitely many isogeny classes, so you just need to check that λ is not a root of a finite list of polynomial equations. (Elkies)

(iii) Can you deterministically find λ, μ such that E_λ, E_μ are not isogeneous? (Coppersmith) What happens to Schoof's algorithm if you just run it with λ a variable? (Lenstra)

(iv) There are bounds $p^{1/2+\epsilon}$ on the number of isogeny classes (Hasse interval) over $\mathbb{F}_p$.

(v) You can also ask the question for the family of all elliptic curves over $\mathbb{F}_p$, deterministically. (Pomerance) This you can do by looking at if -1 and -2 are both squares. (Elkies)

Problem/Question 6. Let $f(X,Y) \in \mathbb{F}_q[X,Y]$ be irreducible. Let $g(Y) = f(aY+b, Y)$. Count (or estimate) the number of pairs a, b over $\mathbb{F}_q$ such that g is irreducible over $\mathbb{F}_q$. (Gao)

Remarks.

(i) Is it > 0 when $q > d^4$, where d is the total degree of f? (Gao) Or perhaps some other bound on q?

(ii) It is equivalent to count the number of reducible ones, so use the Schoof-Pila algorithm. (Elkies) But it is slightly different because you are counting points on a surface.

(iii) The polynomial $f(X,Y) = X^q - X + Y$ is always divisible by Y under such a substitution. (Lenstra) Therefore we need $q > d$.

(iv) Do you need to assume that f is irreducible over $\overline{\mathbb{F}_q}$? (Pomerance)

(v) What does Hilbert's irreducibility theorem say in this case? (Lenstra)

(vi) Applying the Chebotarev density theorem, this number is $rq^2 + O(q^{3/2})$, where $r \in \mathbb{Q}_{\geq 0}$. (Wan) How big is the constant?

(vii) View g as a polynomial in 3 variables, Y, a, b. The surface $g = 0$ is a cover of the affine plane given by the variables a, b. What is the Galois group of this cover (over $\mathbb{F}_q(a,b)$)? Is it possibly the full symmetric group? (Lenstra) Then $r(\#G)$ is equal to the number of elements in the Galois group that are a full d-cycle. Note this extension is separable whenever f is irreducible. (Edixhoven)

(viii) If $r = 0$ then all elements of the set are reducible, since any one irreducible element will have Frobenius which is a full d cycle. (Lenstra)

Problem/Question 7. Let $m \geq n \in \mathbb{Z}$. Prove there is a polynomial $g \in \mathbb{F}_q[x]$ of degree $\leq 2\log n$ so that $x^m + g(x)$ has an irreducible factor of degree n. (Gao)

Remarks.

(i) If q is large compared to fixed m, n, this seems provable. (Gao)

(ii) For $q = 2$, there is an application to computing discrete logs in $\mathbb{F}_{2^n}^*$. (Coppersmith)

(iii) If m is the smallest power of q which is $\geq n$, then α is a root of the irreducible factor, then the order of α is at least $\geq n^{(\log n)/(\log\log n)}$, so this group has large order. (Gao)

(iv) For all q, and $m = n$, it is proven if $\deg g \leq n/2$.

(v) Consider the variant where instead of restricting the degree consider restricting the number of nonzero terms of $x^m + g(x)$ to a fixed number (such as 7) (Bernstein), or at least $\leq 2\log n$ (Pomerance).

(vi) Alternatively, find a trinomial of degree $m \leq 2n$ over $\mathbb{F}_q$ with a primitive irreducible factor of degree n. (Gao) For small q this may not be possible.

(vii) For fixed n, q, what is the sparsest polynomial of degree $\leq 2n$ with a primitive and irreducible factor of degree n? (Pomerance) This should be uniform in q, surely 5 but perhaps 7. (Bernstein) Trivially, sparsity n works by taking an irreducible polynomial of degree n (excluding $(n,q) = (2,2)$). (Lenstra) Sparsity $(1-\epsilon)n$ should be possible. (Wan)

Problem/Question 8. Look at $\sum n^\rho$ over the zeros ρ of the Riemann zeta function with the imaginary part of ρ the interval $[T, 2T]$. This is $\approx T/(2\pi)\Lambda(n)$. Can you make a primality test out of this? (Conrey)

The sum $\sum n^\rho \zeta'(\rho)/\zeta''(\rho) \approx T/(2\pi)\log p \log q$ if $n = pq$. Can you make a factoring algorithm out of this?

Problem/Question 9. Compute the zeta function of $\operatorname{Spec}\mathbb{Z}/n\mathbb{Z}$, namely

$$\zeta(s) = \prod_{p|n} \frac{1}{1-p^{-s}},$$

without knowing the prime factorization of n. Computing the special value $s = 1$ gives you $\phi(n)$ which is enough to factor n. (Wan)

Replace n by a polynomial $f(x) \in \mathbb{F}_p[x]$, so given

$$\zeta(s) = \prod_{\substack{g|f \\ g \text{ monic,irreducible}}} \frac{1}{1-g}$$

can you recover the factorization of f in $\mathbb{F}_q[x]$ in deterministic polynomial time? (Wan)

Remarks.

(i) Can you compute the latter using Drinfeld modules? (Lenstra) Take $A = \mathbb{F}_p[x]/(f)$. Define an $\mathbb{F}_p$-linear map $T : A \to A$ where $T(\alpha) = \alpha^p - x\alpha$ which defines the Carlitz module, and makes A into an $\mathbb{F}_p[T]$-module. If you write $f = \prod_g g_i^{e_i}$, and define $\phi(f) = \prod_i g_i^{e_i-1}(g_i - 1)$, then $\phi(f)(T)$ kills A. But also $f\prod_i(g_i - 1)$ also kills A. Mimic the factorization of integers using upper bounds on $\phi(n)$ to factor f probabilistically using this 'exponent' of the multiplicative group.

Problem/Question 10. To speed up the elliptic curve factorization algorithm, pick $E/\mathbb{Q}$ with large torsion group so that its reduction modulo n is more likely to be smooth. Mazur's results show that this torsion subgroup can be no larger than 16. There is a result of Kamienny-Mazur-Merel: there is a bound on $\#E(K)_{\text{tors}}$ in terms of $[K : \mathbb{Q}]$. So try to find a number field K where p splits completely and such that $E(K)_{\text{tors}}$ large. (Pomerance)

Remarks.

(i) Whatever advantage you gain might be swamped by the expense of working in the larger number field. (Pomerance)

(ii) One problem is that for any given number field, there are only finitely many curves over that number field with a fixed torsion subgroup (since then modular curves have genus ≥ 2).

(iii) If $n = p^2q$, choose a discriminant D such that D is a nonsquare modulo q, and find an elliptic curve $E/\mathbb{Q}(\sqrt{D})$. Then $E(\mathbb{Q}(\sqrt{D}))_{\text{tors}}$ injects into $E(\mathbb{F}_{q^2})$. (Bleichenbacher)

Problem/Question 11. Given a (random) number n of 10000 digits, it may be impossible to find 5000 digit primes p, q such that $n = pq$. Much easier: find some 10000 digit primes p, q such that n is the first 10000 digits of pq. Instead, find 7500 digit primes p, q such that n is the first 10000 digits of pq. (Coppersmith) Therefore, do this for 6000 digit primes. (Bernstein)

Remarks.

(i) Coppersmith's algorithm works as follows. Pick at random p_0 of 7500 digits. Pick $q_0 = \lfloor (n/p_0)10^{5000} \rfloor$. Add x to p_0 and y to q_0 where x, y have 2500 digits a piece to fix this up. We want

$$(p_0 + x)(q_0 + y) = p_0q_0 + p_0y + q_0x + xy \approx n10^{5000}$$

so use lattice reduction. At the end, check to make sure p_0 and q_0 are prime.

(ii) This has applications in cryptography. (Bernstein)

Problem/Question 12. Let p be prime, write $p - 1 = 2^\ell m$ where m is odd, and assume $\ell \geq 3$. Find in deterministic polynomial time $x \in \mathbb{F}_p$ such that $x^{2^{\ell-1}} + 1$ is not a $2^{\ell-1}$th power. (Kedlaya)

Remarks.

(i) This should be a sufficient condition for the deterministic nonresidue algorithm of Agrawal to work.

Problem/Question 13.

(a) Do the p-adic point counting methods of Lauder which allow you to go from one curve from another in a family apply to p-adic methods without the linear factor? (Edixhoven)

(b) Do Fesenko's methods for proving good properties of Hasse-Weil L-functions of curves over number fields provide anything useful for computations? (Edixhoven)

Remarks.

(i) Does Riemann-Roch provide anything useful? (Apparently not; involves linear algebra over matrices of size the number of points.) (Wan)

Problem/Question 14. Pila's method for computing roots of unity is exponential in g. Is there any reason it can't be made linear? (Elkies)

Remarks.

(i) Huang has a randomized algorithm (depends on factoring polynomials) with exponent $g^{O(1)}$. (Lauder) They avoid using a full projective model. (Pila)

(ii) For any prime ℓ, you work in a group of size ℓ^{2g}, so shouldn't be much worse. See also Edixhoven's talk. The calculation there (on modular curves) can also be done for Drinfeld modular curves. (Elkies) Is there an analogue of the point-counting problem for Drinfeld modules? (Kedlaya)

(iii) Has anyone tried doing Schoof-Pila in genus 2? (Kedlaya) Gaudry and Schost have applied AGM. (Couveignes) They also did small torsion. (Edixhoven)

Problem/Question 15. Given a variety over a finite field, can you verify that a given function is its zeta function (i.e., is the question in NP)? (Lenstra, Edixhoven)

Remarks.

(i) Yes (for g fixed), because you can verify the orders of the Jacobian over the first g extension fields. (Elkies)

Problem/Question 16. How can you compute with points as in Edixhoven's talk? In that application, you can avoid writing down explicit ℓ-torsion points (only the fields of definition), but can you write them down for other transformation? (Edixhoven)

Remarks.

(i) Can one work out instances of the passage from H^1 to H^2? E.g., K3 surface of Néron-Severi rank 19 (the rank 20 case is standard)? (Elkies)

(ii) More comments on ℓ-adic computation of zeta functions? (Pila)

Problem/Question 17.

(a) Can deformation theory be applied ℓ-adically? (Various)

(b) Under what circumstances do Betti numbers stay the same under reduction modulo p? (Various)

(c) Is finding the genus of a plane curve polynomial time (in the degree)? (Wan)

Problem/Question 18. Can you compute roots modulo p of a fixed polynomial (à la Schoof-Pila) in polynomial time, like $x^3 - 2$? (Schoof)

Remarks.

(i) A problem is realizing a given polynomial as the characteristic polynomial of an endomorphism, if its roots do not lie in a CM field. Does it help to consider modular forms? (Pila)

B.1 Remarks on Agrawal's Conjecture

These notes concern Agrawal's conjecture, the first problem in the problem session:

Conjecture. *Let n and r be two coprime positive integers. If*

$$(X-1)^n \equiv X^n - 1 \pmod{n, X^r - 1}$$

then either n is prime or

$$n^2 \equiv 1 \pmod{r}.$$

(If Agrawal's conjecture were true, this would improve the polynomial time complexity of the AKS primality testing algorithm from $\widetilde{O}((\log n)^{7.5})$ to $\widetilde{O}((\log n)^3)$.)

The contents are due to Lenstra and Pomerance and suggest strongly that this conjecture is false.

Proposition (Lenstra)**.** *Let $p_1, \ldots, p_k$ be k pairwise distinct prime integers, and let $n = p_1 \ldots p_k$. Suppose that:*

(i) $k \equiv 1 \pmod 4$*;*
(ii) $p_i \equiv 3 \pmod{80}$ *for all i;*
(iii) $(p_i - 1) \mid (n-1)$ *for all i; and*
(iv) $(p_i + 1) \mid (n+1)$ *for all i.*

Then

$$(X-1)^n \equiv X^n - 1 \pmod{n, X^5 - 1}$$

and $n^2 \not\equiv 1 \pmod 5$.

Remark. This result is also true for $k \equiv 3 \pmod 4$.

Proof. By assumption we get $n = 3^k \equiv 3 \pmod{80}$ because $3^4 \equiv 1 \pmod{80}$. So $n \equiv 3 \pmod 5$ and then $n^2 \not\equiv 1 \pmod 5$.

We also have the following identity:

$$(X - 1, X^4 + X^3 + X^2 + X + 1, n) = (1)$$

in the polynomial ring $\mathbb{Z}[X]$. Hence, in order to prove the identity

$$(X-1)^n \equiv X^n - 1 \pmod{n, X^5 - 1}$$

it suffices to prove that

$$(X-1)^n \equiv X^n - 1 \pmod{n, X^4 + \cdots + X + 1}.$$

The Chinese remainder theorem gives the following isomorphism:

$$\mathbb{Z}[X]/(n, X^4 + \cdots + X + 1) \cong \prod_{i=1}^{k} \mathbb{F}_{p_i}[X]/(X^4 + \cdots + X + 1).$$

Each ring factor $R_i = \mathbb{F}_{p_i}[X]/(X^4 + \cdots + X + 1)$ is actually a field since each prime p_i is prime to 5 and the 5th cyclotomic polynomial is irreducible in $\mathbb{F}_p[X]$ so that R_i is nothing but the splitting field of $\mathbb{F}_{p_i}[\zeta_5]$ for a primitive 5th root of unity ζ_5.

It therefore suffices to prove that each prime $p_i = p$ satisfies

$$(\zeta_5 - 1)^n = \zeta_5^n - 1$$

in the field $\mathbb{F}_p[\zeta_5]$. We see from (ii) that

$$(\zeta_5 - 1)^{p^2} = \zeta_5^{p^2} - 1 = \zeta_5^{-1} - 1$$

(since $p \equiv 3 \pmod 5$, we have $p^2 \equiv -1 \pmod 5$). Thus

$$(\zeta_5 - 1)^{p^2} = -\zeta_5^{-1}(\zeta_5 - 1).$$

Hence the order of $(\zeta_5 - 1)$ in $\mathbb{F}_p[\zeta_5]$ divides $10(p^2 - 1)$.

It remains to check the residue class of n modulo $10(p^2 - 1)$; more precisely, it suffices to show that

$$n \equiv p \pmod{10(p^2 - 1)}.$$

We can factor $10(p^2 - 1)$ into 4 pairwise coprime factors:

$$10(p^2 - 1) = 5(2^4)\left(\frac{p-1}{2}\right)\left(\frac{p+1}{4}\right)$$

so it suffices to verify this modulo each factor. Since $n, p \equiv 3 \pmod{80}$ by assumption, the first follows. Assumption (iii) implies that

$$n \equiv 1 \pmod{(p-1)/2}$$

and so

$$n = p \pmod{(p-1)/2}$$

since $p \equiv 1 \pmod{(p-1)/2}$, and

$$n \equiv p \pmod{(p+1)/4}$$

similarly. This completes the proof. □

By this proposition, we have a heuristic which suggests the existence of many counterexamples to the Agrawal conjecture. This argument taken from analytic number theory is very similar to the one already used by Pomerance to find counterexamples to the Baillie-PSW primality testing algorithm which can be found at `http://www.pseudoprime.com/dopo.pdf`.

Fix some arbitrarily large integer m and let T be very large. Let $P = P_m(T)$ denote the set of primes p in the interval $[T, T^m]$ such that:

A. $p \equiv 3 \pmod{80}$;
B. $(p-1)/2$ is squarefree and divisible only by primes $q \leq T$ with $q \equiv 3 \pmod 4$;
C. $(p+1)/4$ is squarefree and divisible only by primes $r \leq T$ with $r \equiv 1 \pmod 4$.

Both smoothness conditions (2) and (3) are rather restrictive: heuristically, the cardinality of the set P is asymptotically ($T \to \infty$)

$$\#P \sim c_m \frac{T^m}{(\log T^m)^2}$$

for some positive constant c_m that depends on the choice of m. In particular, we can take a sufficiently large integer T such that

$$\#P > \frac{T^m}{(\log T^m)^3}$$

which we assume from now on.

Also choose an odd integer $k \equiv 1 \pmod 4$ such that $k < T^2/(\log T^m)$. We consider the squarefree numbers n that run over products of k distinct primes of the set P. Obviously such an integer n satisfies $n < e^{T^2}$. The number of choices for n is exactly given by the binomial coefficient $\binom{\#P}{k}$, and we get the lower bound:

$$\begin{aligned}\binom{\#P}{k} &\geq \left(\frac{T^m}{(\log T^m)^3(T^2/\log T^m)}\right)^{(T^2/\log T^m)-4} \\ &> (T^{m-3})^{(T^2/\log T^m)-4} = e^{(1-3/m)T^2-4(m-3)\log T} \\ &> e^{1-(4/m)T^2}.\end{aligned}$$

for large T and fixed m.

Let Q denote the product of primes $q \leq T$ with $q \equiv 3 \pmod 4$, and let R denote the product of primes $r \leq T$ with $r \equiv 1 \pmod 4$. Then Q and R are coprime and asymptotically the product QR equals $e^{(1+o(1))T}$ as $T \to \infty$, so that $QR < e^{2T}$ for some large T. Thus, the number of choices for the numbers n that satisfy in addition $n \equiv 1 \pmod Q$ and $n \equiv -1 \pmod R$ should be asymptotically

$$e^{(1-4/m)T^2}e^{-2T} > e^{T^2(1-5/m)}.$$

But any such n is a counterexample to Agrawal's conjecture by Lenstra's proposition. We see therefore that for fixed m and for all large T, there should be at least $e^{T^2(1-5/m)}$ counterexamples to Agrawal's conjecture below e^{T^2}. That is, if we let $x = e^{T^2}$, this argument implies that the number of counterexamples $\leq x$ is expected to be $\gg x^{1-\epsilon}$ for any $\epsilon > 0$.