Draft.

PROVING PRIMALITY AFTER AGRAWAL-KAYAL-SAXENA

DANIEL J. BERNSTEIN

1. Introduction

Section 2 presents three theorems concluding, under various assumptions, that n is a prime power:

- Theorem 2.1 is the original Agrawal-Kayal-Saxena theorem.
- Theorem 2.2 is a special case of Lenstra's variant.
- Theorem 2.3 is the general case of Lenstra's variant.

For the reader's convenience, the theorems are arranged in order of difficulty, and a separate proof is provided for each theorem.

Section 3 explains why PRIMES is in P. A proof based on Theorem 2.3 is much simpler than a proof based on Theorem 2.1: thanks to Lenstra, no analytic number theory is needed. Section 3 also summarizes previous results on the PRIMES-in-P problem.

The rest of the paper looks much more closely at algorithm speed, and discusses ways that the Agrawal-Kayal-Saxena computation can be accelerated:

- Section 4 summarizes the state of the art.
- Section 5 discusses the computation of $(x+b)^n$ in the ring $(\mathbf{Z}/n)[x]/(x^r-1)$.
- Sections 6, 7, and 8 discuss several "high-level" improvements.

These sections consider the speed that algorithms are reasonably conjectured to have, not the speed that we can prove. For example, it is conjectured that any integer $n \geq 2$ is a primitive root modulo r for many primes r in any reasonably large interval. This means that one can easily find numbers r suitable for Theorem 2.2, so the extra generality of Theorem 2.3 is unnecessary.

Sometimes people combine the handicaps of (1) insisting on proof and (2) looking at time in detail: they try to minimize the proven upper bound for the run time. This paper does not currently discuss these efforts.

Pedro Berrizbeitia on 2002.11.22 proposed a variant of the Agrawal-Kayal-Saxena approach, using high-power-of-2 roots of a non-square rather than rth roots of 1. I have not yet investigated this variant.

1

2. The core theorems

Theorem 2.1 (2002.08.06; Manindra Agrawal, Neeraj Kayal, and Nitin Saxena)**.** *Let n be a positive integer. Let q and r be prime numbers. Let S be a finite set of integers. Assume that q divides $r - 1$; that $n^{(r-1)/q} \bmod r \notin \{0, 1\}$; that $\gcd\{n, b - b'\} = 1$ for all distinct $b, b' \in S$; that $\binom{\#S+q-1}{\#S} \geq n^{2\lfloor \sqrt{r} \rfloor}$; and that $(x + b)^n = x^n + b$ in the ring $(\mathbf{Z}/n)[x]/(x^r - 1)$ for all $b \in S$. Then n is a power of a prime.*

The theorem actually stated in the Agrawal-Kayal-Saxena paper is that n is prime if a particular algorithm prints 1. That algorithm constrains q, r, and S much more heavily than Theorem 2.1 does, as discussed in Section 4 of this paper. Theorem 2.1 is my 2002.08.10 statement of the properties of q, r, and S used in the Agrawal-Kayal-Saxena proof. (My main contribution to the statement was the inequality $\binom{\#S+q-1}{\#S} \geq n^{2\lfloor \sqrt{r} \rfloor}$. Carl Pomerance had pointed out that varying S was useful, and that q need not be the largest prime divisor of $r - 1$.)

Proof. Find a prime divisor p of n such that $p^{(r-1)/q} \bmod r \notin \{0, 1\}$. (If every prime divisor p of n has $p^{(r-1)/q} \bmod r \in \{0, 1\}$ then $n^{(r-1)/q} \bmod r \in \{0, 1\}$.)

By hypothesis, $(x + b)^n = x^n + b$ in $\mathbf{F}_p[x]/(x^r - 1)$ for all $b \in S$. Substitute x^{n^i} for x: $(x^{n^i} + b)^n = x^{n^{i+1}} + b$ in $\mathbf{F}_p[x]/(x^{n^i r} - 1)$, hence in $\mathbf{F}_p[x]/(x^r - 1)$. By induction, $(x + b)^{n^i} = x^{n^i} + b$ in $\mathbf{F}_p[x]/(x^r - 1)$ for all $i \geq 0$. By Fermat's little theorem, $(x + b)^{n^i p^j} = (x^{n^i} + b)^{p^j} = x^{n^i p^j} + b$ in $\mathbf{F}_p[x]/(x^r - 1)$ for all $j \geq 0$.

Consider the products $n^i p^j$ with $0 \leq i \leq \lfloor \sqrt{r} \rfloor$ and $0 \leq j \leq \lfloor \sqrt{r} \rfloor$; note that $1 \leq n^i p^j \leq n^{2\lfloor \sqrt{r} \rfloor}$. There are $(\lfloor \sqrt{r} \rfloor + 1)^2 > r$ pairs (i, j), so there are distinct pairs $(i, j), (k, \ell)$ such that $n^i p^j \equiv n^k p^\ell \pmod{r}$. Write $t = n^i p^j$ and $u = n^k p^\ell$. Then $x^t = x^u$ in $\mathbf{F}_p[x]/(x^r - 1)$, so $(x + b)^t = x^t + b = x^u + b = (x + b)^u$ in $\mathbf{F}_p[x]/(x^r - 1)$ for all $b \in S$. Note that $|t - u| < n^{2\lfloor \sqrt{r} \rfloor}$.

Find an irreducible polynomial h in $\mathbf{F}_p[x]$ dividing $(x^r - 1)/(x - 1)$. A standard fact about cyclotomic polynomials is that $\deg h$ is the order of p modulo r; so $\deg h$ is a multiple of q; so $\deg h \geq q$.

Now $(x + b)^t = (x + b)^u$ in the finite field $\mathbf{F}_p[x]/h$ for all $b \in S$. Note that $x + b \in (\mathbf{F}_p[x]/h)^*$, since $\deg h \geq q \geq 2$. Define G as the subgroup of $(\mathbf{F}_p[x]/h)^*$ generated by $\{x + b : b \in S\}$; then $g^t = g^u$ for all $g \in G$.

G has at least $\binom{\#S+q-1}{\#S} \geq n^{2\lfloor \sqrt{r} \rfloor} > |t - u|$ elements: namely, all products $\prod_{b \in S} (x + b)^{e_b}$ with $\sum_b e_b \leq q - 1$. (The irreducibles $x + b$ are distinct in $\mathbf{F}_p[x]$, because each difference $(x + b) - (x + b') = b - b'$ is coprime to n by hypothesis; so these products $\prod_b (x + b)^{e_b}$ are distinct in $\mathbf{F}_p[x]$. These products have degree smaller than q, hence smaller than $\deg h$, so they remain distinct modulo h.)

Thus $t = u$. (The equation $g^t = g^u$ cannot have more than $|t - u|$ nonzero solutions in a field unless $t = u$.) In other words, $n^i p^j = n^k p^\ell$. If $i = k$ then $p^j = p^\ell$ so $(i, j) = (k, \ell)$, contradiction. Consequently n is a power of p. $\square$

This proof has been streamlined in several ways. For example, Agrawal, Kayal, and Saxena multiplied equations $x^n + b = (x + b)^n$ to obtain $g(x^n) = g(x)^n$, and then obtained $g^t = g^u$ by various substitutions; I noticed on 2002.08.10 that it is easier to substitute first and multiply later. Agrawal, Kayal, and Saxena took a generator g for the group G, and concluded from $g^t = g^u$ that $t = u$; Kiran Kedlaya pointed out on 2002.11.02 that it is easier to count solutions of $g^t = g^u$.

Theorem 2.2 (2002.08.14; Manindra Agrawal, Neeraj Kayal, Nitin Saxena, and Hendrik W. Lenstra, Jr.). *Let n and r be positive integers. Let S be a finite set of integers. Assume that n is a primitive root modulo r; that $\gcd\{n, b - b'\} = 1$ for all distinct $b, b' \in S$; that $\binom{\#S + \varphi(r) - 1}{\#S} \geq n^{2\lfloor\sqrt{\varphi(r)}\rfloor}$; and that $(x+b)^n = x^n + b$ in the ring $(\mathbf{Z}/n)[x]/(x^r - 1)$ for all $b \in S$. Then n is a power of a prime.*

Theorem 2.2 is narrower than Theorem 2.1 in one way: it requires that n be a primitive root modulo r, not merely that $n^{(r-1)/q} \bmod r \notin \{0,1\}$. On the other hand, Theorem 2.2 requires only $\binom{\#S + \varphi(r) - 1}{\#S} \geq n^{2\lfloor\sqrt{\varphi(r)}\rfloor}$ instead of $\binom{\#S + q - 1}{\#S} \geq n^{2\lfloor\sqrt{r}\rfloor}$, and it drops the requirement that r be prime.

Theorem 2.2 is the special case $v = \varphi(r)$ of Theorem 2.3.

Proof. If $n = 1$ then n is a power of a prime, so assume that $n \geq 2$. Let p be a prime divisor of n.

By hypothesis, $(x + b)^n = x^n + b$ in $\mathbf{F}_p[x]/(x^r - 1)$ for all $b \in S$. Substitute x^{n^i} for x: $(x^{n^i} + b)^n = x^{n^{i+1}} + b$ in $\mathbf{F}_p[x]/(x^{n^i r} - 1)$, hence in $\mathbf{F}_p[x]/(x^r - 1)$. By induction, $(x + b)^{n^i} = x^{n^i} + b$ in $\mathbf{F}_p[x]/(x^r - 1)$ for all $i \geq 0$. By Fermat's little theorem, $(x + b)^{n^i p^j} = (x^{n^i} + b)^{p^j} = x^{n^i p^j} + b$ in $\mathbf{F}_p[x]/(x^r - 1)$ for all $j \geq 0$.

Consider the products $n^i p^j$ with $0 \leq i \leq \lfloor\sqrt{\varphi(r)}\rfloor$ and $0 \leq j \leq \lfloor\sqrt{\varphi(r)}\rfloor$; note that $1 \leq n^i p^j \leq n^{2\lfloor\sqrt{\varphi(r)}\rfloor}$. Each $n^i p^j \bmod r$ lies in a set $(\mathbf{Z}/r)^*$ of size $\varphi(r)$, and there are $(\lfloor\sqrt{\varphi(r)}\rfloor + 1)^2 > \varphi(r)$ pairs (i, j), so there are distinct pairs $(i, j), (k, \ell)$ such that $n^i p^j \equiv n^k p^\ell \pmod{r}$. Write $t = n^i p^j$ and $u = n^k p^\ell$. Then $x^t = x^u$ in $\mathbf{F}_p[x]/(x^r - 1)$, so $(x + b)^t = x^t + b = x^u + b = (x + b)^u$ in $\mathbf{F}_p[x]/(x^r - 1)$ for all $b \in S$. Note that $|t - u| + 1 \leq n^{2\lfloor\sqrt{\varphi(r)}\rfloor}$.

Find an irreducible polynomial h in $\mathbf{F}_p[x]$ dividing the rth cyclotomic polynomial Φ_r. Define G as the subset of $\mathbf{F}_p[x]/h$ generated under multiplication by $\{0\} \cup \{x + b : b \in S\}$. Then $g^t = g^u$ for all $g \in G$.

G has at least $1 + \binom{\#S + \varphi(r) - 1}{\#S} > n^{2\lfloor\sqrt{\varphi(r)}\rfloor} \geq |t - u| + 1$ elements: namely, 0, and all products $\prod_{b \in S}(x + b)^{e_b}$ with $\sum_b e_b \leq \varphi(r) - 1$. (The irreducibles $x + b$ are distinct in $\mathbf{F}_p[x]$, because each difference $(x + b) - (x + b') = b - b'$ is coprime to n by hypothesis. Thus these products $\prod_b (x + b)^{e_b}$, along with 0, are distinct in $\mathbf{F}_p[x]$; it suffices to show that they are distinct in $\mathbf{F}_p[x]/h$. Say $e = \prod_b (x + b)^{e_b}$; $f = \prod_b (x + b)^{f_b}$ or $f = 0$; and $e = f$ in $\mathbf{F}_p[x]/h$. Then

$$e^{n^a} = \prod_b (x + b)^{n^a e_b} = \prod_b (x^{n^a} + b)^{e_b} = e(x^{n^a})$$

in $\mathbf{F}_p[x]/h$ for all $a \geq 0$; similarly $f^{n^a} = f(x^{n^a})$ in $\mathbf{F}_p[x]/h$ for all $a \geq 0$; so $e(x^{n^a}) = f(x^{n^a})$ in $\mathbf{F}_p[x]/h$. Write $g = e - f$. If $c \in \{0, 1, \ldots, r - 1\}$ and $\gcd\{r, c\} = 1$ then $c \equiv n^a \pmod{r}$ for some a, so $g(x^c) = g(x^{n^a})$ in $\mathbf{F}_p[x]/(x^r - 1)$, so $g(x^c) = 0$ in $\mathbf{F}_p[x]/h$; i.e., y^c is a root of g in the field $\mathbf{F}_p[y]/h(y)$. The powers y^c are distinct in $\mathbf{F}_p[y]/h(y)$, so $\prod_c (x - y^c) = \Phi_r$ divides g in $(\mathbf{F}_p[y]/h(y))[x]$, hence in $\mathbf{F}_p[x]$. Thus $e = f$ in $\mathbf{F}_p[x]/\Phi_r$. Both e and f have degree smaller than $\varphi(r) = \deg \Phi_r$, so $e = f$ in $\mathbf{F}_p[x]$.)

Thus $t = u$. (The equation $g^t = g^u$ cannot have more than $|t - u| + 1$ solutions in a field unless $t = u$.) In other words, $n^i p^j = n^k p^\ell$. If $i = k$ then $p^j = p^\ell$ so $(i, j) = (k, \ell)$, contradiction. Consequently n is a power of p. $\qquad\square$

Theorem 2.3 (2002.08.14; Manindra Agrawal, Neeraj Kayal, Nitin Saxena, and Hendrik W. Lenstra, Jr.)**.** *Let n, r, and v be positive integers. Let S be a finite set of integers. Assume that n and r are coprime; that n has order v modulo r; that $\gcd\{n, b - b'\} = 1$ for all distinct $b, b' \in S$; that $\binom{\#S + \varphi(r) - 1}{\#S} \geq n^{2d\lfloor\sqrt{\varphi(r)/d}\rfloor}$ for every positive integer d dividing $\varphi(r)/v$; and that $(x + b)^n = x^n + b$ in the ring $(\mathbf{Z}/n)[x]/(x^r - 1)$ for all $b \in S$. Then n is a power of a prime.*

Proof. If $n = 1$ then n is a power of a prime, so assume that $n \geq 2$. Let p be a prime divisor of n.

Define $d = \#((\mathbf{Z}/r)^*/\langle n, p \rangle)$. Lift $(\mathbf{Z}/r)^*/\langle n, p \rangle$ to positive integers $m_1, \ldots, m_d$. Note that $v = \#\langle n \rangle$ divides $\#\langle n, p \rangle = \varphi(r)/d$, so d divides $\varphi(r)/v$.

By hypothesis, $(x + b)^n = x^n + b$ in $\mathbf{F}_p[x]/(x^r - 1)$ for all $b \in S$. Substitute x^{n^i} for x: $(x^{n^i} + b)^n = x^{n^{i+1}} + b$ in $\mathbf{F}_p[x]/(x^{n^i r} - 1)$, hence in $\mathbf{F}_p[x]/(x^r - 1)$. By induction, $(x + b)^{n^i} = x^{n^i} + b$ in $\mathbf{F}_p[x]/(x^r - 1)$ for all $i \geq 0$. By Fermat's little theorem, $(x + b)^{n^i p^j} = (x^{n^i} + b)^{p^j} = x^{n^i p^j} + b$ in $\mathbf{F}_p[x]/(x^r - 1)$ for all $j \geq 0$. Substitute x^{m_w} for x: $(x^{m_w} + b)^{n_i p^j} = x^{n^i p^j m_w} + b$ in $\mathbf{F}_p[x]/(x^r - 1)$.

Consider the products $n^i p^j$ with $0 \leq i \leq \lfloor\sqrt{\varphi(r)/d}\rfloor$ and $0 \leq j \leq \lfloor\sqrt{\varphi(r)/d}\rfloor$; note that $1 \leq n^i p^j \leq n^{2\lfloor\sqrt{\varphi(r)/d}\rfloor}$. Each $n^i p^j \bmod r$ lies in a set $\langle n, p \rangle$ of size $\varphi(r)/d$, and there are $(\lfloor\sqrt{\varphi(r)/d}\rfloor + 1)^2 > \varphi(r)/d$ pairs (i, j), so there are distinct pairs $(i, j), (k, \ell)$ such that $n^i p^j \equiv n^k p^\ell \pmod{r}$. Write $t = n^i p^j$ and $u = n^k p^\ell$. Then $x^t = x^u$ in $\mathbf{F}_p[x]/(x^r - 1)$, so $(x^{m_w} + b)^t = x^{t m_w} + b = x^{u m_w} + b = (x^{m_w} + b)^u$ in $\mathbf{F}_p[x]/(x^r - 1)$ for all $b \in S$ and all w. Note that $|t - u| + 1 \leq n^{2\lfloor\sqrt{\varphi(r)/d}\rfloor}$.

Find an irreducible polynomial h in $\mathbf{F}_p[x]$ dividing the rth cyclotomic polynomial Φ_r. Define G as the subset of $\mathbf{F}_p[x]/h$ generated under multiplication by $\{0\} \cup \{x^{m_1} + b : b \in S\} \cup \cdots \cup \{x^{m_d} + b : b \in S\}$. Then $g^t = g^u$ for all $g \in G$.

The subset G^d of $(\mathbf{F}_p[x]/h)^d$ has at least $1 + \binom{\#S + \varphi(r) - 1}{\#S} > n^{2d\lfloor\sqrt{\varphi(r)/d}\rfloor} \geq (|t - u| + 1)^d$ elements: namely, all vectors $(\prod_{b \in S}(x^{m_1} + b)^{e_b}, \ldots, \prod_{b \in S}(x^{m_d} + b)^{e_b})$ with $\sum_b e_b \leq \varphi(r) - 1$, along with the zero vector. (The irreducibles $x + b$ are distinct in $\mathbf{F}_p[x]$, because each difference $(x + b) - (x + b') = b - b'$ is coprime to n by hypothesis. Thus the products $\prod_b (x + b)^{e_b}$, along with 0, are distinct in $\mathbf{F}_p[x]$; it suffices to show that they have distinct images in $(\mathbf{F}_p[x]/h)^d$ under the map $e \mapsto (e(x^{m_1}), e(x^{m_2}), \ldots, e(x^{m_d}))$. Say $e = \prod_b (x + b)^{e_b}$; $f = \prod_b (x + b)^{f_b}$ or $f = 0$; and $e(x^{m_w}) = f(x^{m_w})$ in $\mathbf{F}_p[x]/h$ for all w. Then

$$e(x^{m_w})^{n^a p^v} = \prod_b (x^{m_w} + b)^{n^a p^v e_b} = \prod_b (x^{n^a p^v m_w} + b)^{e_b} = e(x^{n^a p^v m_w})$$

in $\mathbf{F}_p[x]/h$ for all a, v, w; similarly $f(x^{m_w})^{n^a p^v} = f(x^{n^a p^v m_w})$ in $\mathbf{F}_p[x]/h$ for all a, v, w; so $e(x^{n^a p^v m_w}) = f(x^{n^a p^v m_w})$ in $\mathbf{F}_p[x]/h$. Write $g = e - f$. If $c \in \{0, 1, \ldots, r - 1\}$ and $\gcd\{r, c\} = 1$ then $c \equiv n^a p^v m_w \pmod{r}$ for some a, v, w, so $g(x^c) = g(x^{n^a p^v m_w})$ in $\mathbf{F}_p[x]/(x^r - 1)$, so $g(x^c) = 0$ in $\mathbf{F}_p[x]/h$; i.e., y^c is a root of g in the field $\mathbf{F}_p[y]/h(y)$. The powers y^c are distinct in $\mathbf{F}_p[y]/h(y)$, so $\prod_c (x - y^c) = \Phi_r$ divides g in $(\mathbf{F}_p[y]/h(y))[x]$, hence in $\mathbf{F}_p[x]$. Thus $e = f$ in $\mathbf{F}_p[x]/\Phi_r$. Both e and f have degree smaller than $\varphi(r) = \deg \Phi_r$, so $e = f$ in $\mathbf{F}_p[x]$.)

Thus G has more than $|t - u| + 1$ elements; so $t = u$. (The equation $g^t = g^u$ cannot have more than $|t - u| + 1$ solutions in a field unless $t = u$.) In other words, $n^i p^j = n^k p^\ell$. If $i = k$ then $p^j = p^\ell$ so $(i, j) = (k, \ell)$, contradiction. Consequently n is a power of p.   □

3. PRIMES IS IN P

What it means. "PRIMES $\in$ P" means that there is a deterministic polynomial-time algorithm A such that $A(i) = 1$ if and only if i is the decimal expansion of a prime number. PRIMES is, by definition, the set of those decimal expansions. P is the set of languages recognizable by deterministic polynomial-time algorithms.

"Deterministic" means that the operations carried out by the algorithm are determined entirely by (1) the algorithm and (2) the input. In particular, the algorithm does not make any random choices. "Polynomial time" means that there is some polynomial p such that, for every string i, the algorithm takes at most $p(\text{length of } i)$ seconds on input i.

It is important to specify the representation of numbers as strings—in this case, as decimal expansions—because the time needed for a computation depends heavily on the input representation. For example, one could represent a positive integer in factored form as a sequence of decimal expansions of primes; addition of integers would then be horribly difficult, but checking primality would be very easy.

How it is proven. The shortest known proof that PRIMES $\in$ P uses Theorem 2.3. Here is a suitable algorithm A, suggested 2002.08.13 by Lenstra:

- Check whether the input is the decimal expansion of a positive integer. If not, print 0 and stop.
- Check whether that positive integer n is a perfect power: a square, a cube, etc. If so, print 0 and stop: n is not prime.
- Compute $N = 2n(n-1)(n^2-1)(n^3-1)\cdots(n^{4\lceil \lg n\rceil^2}-1)$. Find the smallest prime number r that does not divide N. (Proof that r is bounded by a polynomial: Note that $n > 1$ so $N > 0$. There is a positive integer $k \in (8 + o(1))(\lg n)^5$ such that $N < 2^k$. Chebyshev showed that the primes $\leq 2k$ have product at least 2^k, so they cannot all divide N. Thus $r \leq 2k$.)
- Check whether n is equal to any prime below r. If so, print 1 and stop: n is prime.
- Check whether n is divisible by any prime below r. If so, print 0 and stop: n is composite.
- Check whether $(x+b)^n = x^n + b$ in the ring $(\mathbf{Z}/n)[x]/(x^r-1)$ for all $b \in S$ where $S = \{1, 2, \ldots, r\}$. If not, print 0 and stop: n is composite.
- Print 1 and stop: n is a prime power by Theorem 2.3, and therefore a prime. (By construction, r and n are coprime; all differences $\pm 1, \pm 2, \ldots, \pm(r-1)$ of elements of S are coprime to n; and the order v of n modulo r is larger than $4\lceil \lg n\rceil^2$. If d is a positive integer dividing $\varphi(r)/v$ then $d \leq \varphi(r)/v < \varphi(r)/4(\lg n)^2$, so $2d\lfloor\sqrt{\varphi(r)/d}\rfloor \leq 2d\sqrt{\varphi(r)/d} = \sqrt{4d\varphi(r)} \leq \varphi(r)/\lg n$, so $n^{2d\lfloor\sqrt{\varphi(r)/d}\rfloor} \leq 2^{\varphi(r)}$. But $\binom{\#S+\varphi(r)-1}{\#S} = \binom{2\varphi(r)}{\varphi(r)+1} \geq 2^{\varphi(r)}$ since $\varphi(r) \geq 2$.)

This takes polynomial time, using standard polynomial-time subroutines for basic arithmetic, power detection, prime enumeration, and modular exponentiation.

The original 2002.08.06 Agrawal-Kayal-Saxena PRIMES $\in$ P proof used a much more difficult 1985 theorem of Fouvry: there are many primes r such that $r-1$ has a prime divisor q exceeding $r^{2/3}$. (Pomerance pointed out that one could instead use a 1969 theorem of Goldfeld, with 2/3 replaced by something slightly larger than 1/2.) Consequently one can find reasonably small q, r with $q \geq 4\lceil\sqrt{r}\rceil\lceil\lg n\rceil$ and $n^{(r-1)/q} \bmod r \notin \{0, 1\}$. The set $S = \{1, 2, \ldots, 2\lceil\sqrt{r}\rceil\lceil\lg n\rceil\}$ then satisfies the

inequality $\binom{\#S+q-1}{\#S} \geq n^{2\lfloor\sqrt{r}\rfloor}$ in Theorem 2.1. One checks as above whether n is a perfect power, whether n has a small prime factor, and whether the equations $(x+b)^n = x^n + b$ hold in $(\mathbf{Z}/n)[x]/(x^r - 1)$.

Previous work, part 1: proving compositeness. Every composite integer n has a very short proof of its compositeness: namely, a nontrivial divisor of n. In short: "PRIMES $\in$ coNP." For example, the integer 314159265358979323 is composite, because it is divisible by 317213509. (I should be saying "certificate" rather than "proof," but the distinction doesn't matter here.)

Unfortunately, for many composite integers n, these proofs are difficult to find. The fastest known algorithm is the number field sieve, which is conjectured to take time $\exp(O((\lg n)^{1/3}(\lg\lg n)^{2/3}))$; much slower than polynomial time.

Another approach: if n does not divide $2^n - 2$ then n is composite. For almost all composite integers n, this test quickly proves the compositeness of n. Unfortunately, a few composite integers n cannot be proven composite in this way.

Artjuhov 1966, et al.: If $n \in 5 + 8\mathbf{Z}$ divides none of b, $b^{(n-1)/2} + 1$, $b^{(n-1)/4} + 1$, $b^{(n-1)/4} - 1$ then n is composite. Similar comments apply for $n \in 3+4\mathbf{Z}$, $n \in 9+16\mathbf{Z}$, etc. This test takes polynomial time.

Miller 1976, and Oesterlé 1979: If the generalized Riemann hypothesis is true, and n is composite, then Artjuhov's test proves n composite for some positive integer $b \leq 34(\lg n)^2$. In short: "GRH implies PRIMES $\in$ P." The only problem here is that we don't have a proof of GRH.

Rabin 1976, Solovay-Strassen 1977 (variant), Monier 1980, and Atkin-Larson 1982: If n is composite, then Artjuhov's test proves n composite for at least 75% of all b's in $\{1, 2, \ldots, n-1\}$. Trying Artjuhov's test for a bunch of random b's thus has an excellent chance of proving n composite. In short: "PRIMES $\in$ coRP." (It is widely believed that BPP = coRP = RP = P.)

Previous work, part 2: proving primality. Every prime has a short proof of primality using an 1876 theorem of Lucas: if $q_1, \ldots, q_t$ are prime, $n - 1 = q_1 \cdots q_t$, n divides $a^{n-1} - 1$, and n does not divide $a^{(n-1)/q_1} - 1, \ldots, a^{(n-1)/q_t} - 1$, then n is prime. In short: "PRIMES $\in$ NP." However, these proofs are difficult to find when $n - 1$ is difficult to factor.

Pocklington 1914, Morrison 1975, Brillhart-Lehmer-Selfridge 1975, Adleman-Pomerance-Rumely 1979: Improved "cyclotomic primality-proving algorithms," given any integer n, prove the primality of n in time $\exp(O(\lg\lg n \lg\lg\lg n))$. This is, unfortunately, a little slower than polynomial time.

Goldwasser-Kilian 1986, using Schoof 1985: The "elliptic-curve primality-proving algorithm" uses randomness, takes polynomial time, and is conjectured to have a good chance of finding a proof that n is prime. A variant (Adleman-Huang 1992) is proven to have a good chance. In short: "PRIMES $\in$ RP."

Given an integer n, we can try both a compositeness-proving algorithm and a primality-proving algorithm in polynomial time, and repeat until one of the methods succeeds. The result—just like the result of the Agrawal-Kayal-Saxena algorithm— is either a proof that n is composite or a proof that n is prime. ("PRIMES $\in$ ZPP"; by definition, ZPP = RP $\cap$ coRP.) However, this *could* take a huge number of repetitions if we are extremely unlucky. In contrast, the Agrawal-Kayal-Saxena algorithm uses no randomness and always finishes in polynomial time.

4. Speedups

How do we find, as quickly as possible, a proof that a given integer n is prime? Here are some of the answers before Agrawal, Kayal, and Saxena:

- First find a proof of the Riemann hypothesis. Then use the Miller-Oesterlé method, with speedups by Bach. This is not a satisfactory answer: we're still having trouble with the first step.
- Use the cyclotomic primality-proving method, with speedups by Cohen, Lenstra, Bosma, van der Hulst, and Mihailescu. This approach proves the primality of 2000-digit numbers in roughly 10^{14} clock cycles.
- Use the elliptic-curve primality-proving method, with speedups by Atkin and Morain. This approach proves the primality of 2000-digit numbers in roughly 10^{15} clock cycles.

Does the Agrawal-Kayal-Saxena idea save any time here?

Speed of the algorithm, with improvements. As discussed in Section 5, the Agrawal-Kayal-Saxena computation involves $\approx \#S \lg n$ squarings of integers having $\approx 2r \lg n$ bits. The computation can be accelerated with "low-level" speedups in integer squaring and with "high-level" improvements in the size of r and $\#S$. One way to view the progress of "high-level" improvements is to watch the conjectured asymptotics for the product $r\#S$:

- $r\#S/(\lg n)^4 \in 1024 + o(1)$, with $\#S/r \in 0.25 + o(1)$: Original algorithm. (2002.08.06, Agrawal, Kayal, Saxena.)
- $r\#S/(\lg n)^4 \in 2.2575\ldots + o(1)$, with $\#S/r \in 8.746\ldots + o(1)$: Choose r and $\#S$, subject to the condition $\binom{\#S+q-1}{\#S} \geq n^{2\lfloor\sqrt{r}\rfloor}$ in Theorem 2.1, to minimize the time spent by the algorithm. (2002.08.11, Bernstein.)
- $r\#S/(\lg n)^4 \in 0.017637\ldots + o(1)$, with $\#S/r \in 17.492\ldots + o(1)$: Increase the degree bound from q to $\varphi(r)$ (2002.08.13, Lenstra); replace p, n with $p, n/p$ (2002.08.13, Lenstra); and re-optimize parameters with the condition $\binom{\#S+\varphi(r)-1}{\#S} \geq n^{\lfloor\sqrt{\varphi(r)}\rfloor}$. See Section 7 and Section 8.
- $r\#S/(\lg n)^4 \in 0.0044093\ldots + o(1)$, with $\#S/r \in 17.492\ldots + o(1)$: Balance p and n/p (2002.08.26, Bjorn Poonen), and re-optimize with $\binom{\#S+\varphi(r)-1}{\#S} \geq n^{\lceil\sqrt{\varphi(r)/2}\rceil}$. See Section 8.
- $r\#S/(\lg n)^4 \in 0.0022046\ldots + o(1)$, with $\#S/r \in 8.746\ldots + o(1)$: Substitute $1/x$ for x (2002.08.28, Felipe Voloch), and re-optimize with $\binom{2\#S+\varphi(r)-1}{2\#S} \geq n^{\lceil\sqrt{\varphi(r)/2}\rceil}$. See Section 6.
- $r\#S/(\lg n)^4 \in 0.0011324\ldots + o(1)$, with $\#S/r \in 5.0399\ldots + o(1)$ and $d/r \in 0.4993\ldots + o(1)$: Use negative powers of $x + b$ (2002.09.13, Voloch), and re-optimize with $\binom{2\#S+d}{2\#S}\binom{2\#S+\varphi(r)-2d-1}{2\#S-d} \geq n^{\lceil\sqrt{\varphi(r)/2}\rceil}$. See Section 7.
- $r\#S/(\lg n)^4 \in 0.00050329\ldots + o(1)$, with $\#S/r \in 5.0399\ldots + o(1)$ and $d/r \in 0.4993\ldots + o(1)$: Use Minkowski's theorem (2002.12.23, Lenstra), and re-optimize with $\binom{2\#S+d}{2\#S}\binom{2\#S+\varphi(r)-1-2d}{2\#S-d} \geq n^{\lceil\sqrt{\varphi(r)/3}\rceil}$. See Section 8.
- $r\#S/(\lg n)^4 \in 0.0005026686484\ldots + o(1)$, with $\#S/r \in 4.983781039\ldots + o(1)$, $d/r \in 0.5 + o(1)$, and $i/r, j/r \in 0.4749814378\ldots + o(1)$: Use a slightly

better lower bound (2003.01.24, Jeff Vaaler) for the number of negative-power combinations, and re-optimize with $\binom{2\#S}{i}\binom{d}{i}\binom{2\#S-i}{j}\binom{\varphi(r)-1-d}{j} \geq n^{\lceil\sqrt{\varphi(r)/3}\rceil}$.

Thus the algorithm has been sped up by a factor of $2037127.2\ldots+o(1)$. Presumably the $o(1)$ here is positive and fairly large for typical input sizes: reducing r has slightly more than a linear impact on integer-multiplication time. Of course, "two million times faster" does not mean "fast."

How the improved algorithm works. The following theorem incorporates all of the above improvements into Theorem 2.2:

Theorem 4.1. *Let n and r be positive integers. Let d, i, and j be nonnegative integers. Let S be a finite set of integers with $0, 1, -1 \notin S$. Assume that n is a primitive root modulo r; that $\varphi(r) \geq 2$; that $\gcd\{n, bb'-1\} = 1$ for all $b, b' \in S$; that $\gcd\{n, b-b'\} = 1$ for all distinct $b, b' \in S$; that $\binom{2\#S}{i}\binom{d}{i}\binom{2\#S-i}{j}\binom{\varphi(r)-1-d}{j} \geq n^{\lceil\sqrt{\varphi(r)/3}\rceil}$; that $b^{n-1} = 1$ in the ring $\mathbf{Z}/n$ for all $b \in S$; and that $(x+b)^n = x^n + b$ in the ring $(\mathbf{Z}/n)[x]/(x^r - 1)$ for all $b \in S$. Then n is a power of a prime.*

One uses this theorem as follows, given a positive integer n. Check whether n is a perfect power; if so, n is composite. Select an integer $r_0 \geq 3$. Find the smallest prime $r \geq r_0$ such that n is a primitive root modulo r. Select an integer d between 0 and $\varphi(r) - 1$. Select an integer i between 0 and d. Select an integer j between 0 and $\varphi(r) - 1 - d$. Select an integer $s \geq 0$ such that

$$\binom{2s}{i}\binom{d}{i}\binom{2s-i}{j}\binom{\varphi(r)-1-d}{j} \geq n^{\lceil\sqrt{\varphi(r)/3}\rceil}.$$

Define $S = \{2, 3, \ldots, s+1\}$. Check whether $\gcd\{n, b\} = 1$ for all $b \in S$; if not, the primality of n is easily determined, since n has a small factor. Check whether $\gcd\{n, bb'-1\} = 1$ for all $b, b' \in S$; if not, the primality of n is easily determined. Check whether $\gcd\{n, b-b'\} = 1$ for all distinct $b, b' \in S$; if not, the primality of n is easily determined. Check whether $b^{n-1} = 1$ in $\mathbf{Z}/n$ for all $b \in S$; if not, n is composite. Check whether $(x+b)^n = x^n + b$ in $(\mathbf{Z}/n)[x]/(x^r - 1)$ for all $b \in S$; if not, n is composite. Finally, n is prime by Theorem 4.1.

The best choice of r_0 depends on the speed of integer arithmetic. The choice $r_0 \approx 0.01(\lg n)^2$ is close to asymptotically optimal, but somewhat smaller choices are better for reasonable sizes of n. Having found r, one can select $d \approx 0.5\varphi(r)$, select $i = j \approx 0.475\varphi(r)$, and then find the smallest possible s by binary search; there is ample time to compute the binomial coefficients.

One can precompute a table of reasonable choices (r, d, i, j, s, n_0) such that

$$\binom{2s}{i}\binom{d}{i}\binom{2s-i}{j}\binom{\varphi(r)-1-d}{j} \geq n_0^{\lceil\sqrt{\varphi(r)/3}\rceil}.$$

If n is a primitive root modulo r, and $n \leq n_0$, then one can use (r, d, i, j, s) for n.

5. Exponentiation

The bottleneck in the Agrawal-Kayal-Saxena computation is calculating $(x+b)^n$ in the ring $(\mathbf{Z}/n)[x]/(x^r - 1)$ for each $b \in S$. The standard way to do this involves $\approx \#S \lg n$ squarings in $(\mathbf{Z}/n)[x]/(x^r - 1)$ and a very small amount of additional work.

How, then, do we square an element of $(\mathbf{Z}/n)[x]/(x^r - 1)$? One answer: Define $k = \lceil \lg rn^2 \rceil$. Lift to $\mathbf{Z}[x]/(x^r - 1)$, obtaining a polynomial whose coefficients are between 0 and $n - 1$. Note that the square has coefficients smaller than $rn^2 \leq 2^k$. (One can decrease k a bit by using coefficients between $-n/2$ and $n/2$.) Map to $\mathbf{Z}/(2^{kr}-1)$ by $x \mapsto 2^k$. Multiply in $\mathbf{Z}/(2^{kr}-1)$. Recover the product in $\mathbf{Z}[x]/(x^r-1)$. Reduce each coefficient modulo n to obtain the product in $(\mathbf{Z}/n)[x]/(x^r - 1)$.

So the computation boils down to $\approx \#S \lg n$ squarings of big integers modulo $2^{kr} - 1$, and $\approx r\#S \lg n$ reductions modulo n of integers below 2^k. One can do each squaring in time $O(kr \lg kr \lg \lg kr) = O(r \lg n \lg(r \lg n) \lg \lg(r \lg n))$, and one can do each reduction in time $O(\lg n \lg \lg n \lg \lg \lg n)$, by standard fast-arithmetic techniques. Total time: $O(r\#S(\lg n)^2 \lg(r \lg n) \lg \lg(r \lg n))$.

The widely available GMP library takes roughly 100 clock cycles per output bit for big-integer squaring. Every improvement in this time produces a corresponding improvement in the speed of the Agrawal-Kayal-Saxena computation.

Factoring the modulus. If r is even, one can factor $x^r - 1$ into $x^{r/2} - 1$ and $x^{r/2} + 1$, and perform separate computations modulo $x^{r/2} - 1$ and $x^{r/2} + 1$. Various other factorizations are possible, depending on r and n; but most factorizations are a waste of time.

6. Substitution

This section identifies a finite coprime set T of monic non-constant polynomials $f \in (\mathbf{Z}/n)[x]$ such that $f^n = f(x^n)$ in $(\mathbf{Z}/n)[x]/(x^r - 1)$. Here **coprime** means that any distinct $f, g \in T$ satisfy $f(\mathbf{Z}/n)[x] + g(\mathbf{Z}/n)[x] = (\mathbf{Z}/n)[x]$. The next section will consider products of elements of T; coprimality will ensure that these products are all different.

Select a finite set S of integers; for example, the Agrawal-Kayal-Saxena algorithm chooses $S = \{-1, -2, \ldots, -s\}$. Compute $(x+b)^n$ in $(\mathbf{Z}/n)[x]/(x^r - 1)$ for each $b \in S$, as explained in Section 5. Check that $(x + b)^n = x^{n \bmod r} + b$ in $(\mathbf{Z}/n)[x]/(x^r - 1)$; if this fails then n is composite. Now $(x + b)^n = x^n + b$.

The Agrawal-Kayal-Saxena algorithm chooses T as $\{x + b : b \in S\}$. Checking the coprimality of T means checking that n is coprime to the differences $b - b'$ of distinct elements $b, b' \in S$. If this fails then the primality of n is easily determined, provided that the elements of S are reasonably small. The same proviso will be assumed throughout the following discussion.

Substituting $1/x$ for x. Voloch on 2002.08.28 suggested doubling the size of T by choosing T as $\{x + b : b \in S\} \cup \{x + 1/b : b \in S\}$.

Assume that $0 \notin S$, $1 \notin S$, and $-1 \notin S$. Check that $b^n = b$ in $\mathbf{Z}/n$ for all $b \in S$; if this fails then n is composite. Check that b (which is nonzero) is coprime to n for all $b \in S$, that $b - b'$ (which is nonzero) is coprime to n for all distinct $b, b' \in S$, and that $bb' - 1$ (which is nonzero) is coprime to n for all $b, b' \in S$; if any of this fails then the primality of n is easily determined.

Fix $b \in S$. Check that $(x + b)^n = x^n + b$ in $(\mathbf{Z}/n)[x]/(x^r - 1)$. Substitute $1/x$ for x, multiply by x^n, and divide by $b^n = b$, to see that $(x + 1/b)^n = x^n + 1/b$.

The polynomials $x + b$ and $x + b'$ are coprime for all distinct $b, b' \in S$: their difference $b - b'$ is a unit. The polynomials $x + 1/b$ and $x + 1/b'$ are coprime for all distinct $b, b' \in S$: their difference $(b' - b)/bb'$ is a unit. The polynomials $x + b$ and $x + 1/b'$ are coprime for all $b, b' \in S$: their difference $(bb' - 1)/b'$ is a unit.

Voloch comments that one can often skip the tests of $b^n = b$ in $\mathbf{Z}/n$: substituting $x = 1$ in $(x + b)^n = x^n + b$ produces $(1 + b)^n = 1 + b$. But these tests are too fast to worry about.

Substituting x^2 for x. On 2002.08.31, I suggested including $x^2 + b$ and $x^2 + 1/b$, and perhaps $x^3 + b$ etc. One has to exclude squares from S, check that $b^2 - b'$ (which is nonzero) is coprime to n for all $b, b' \in S$, and check that $b^2 b' - 1$ (which is nonzero) is coprime to n for all $b, b' \in S$. Further checks are necessary with $x^3 + b$ etc.

This can be combined with the negative-powers idea discussed in Section 7. It should save a small percentage in the numbers in Section 4; the importance of a polynomial in T drops dramatically as its degree grows, but there's still some benefit. I haven't done the necessary calculations yet.

Substituting $-x$ for x. If r is even then $(x+b)^n = x^n + b$ implies $(x-b)^n = x^n - b$, so one can double the size of T. However, in the usual case that $r/2$ is odd, it is simpler and more flexible to chop r in half and multiply $\#S$ by 2.

Using x. As several people have observed, one can include x in T even without 0 in S, since $(x + 0)^n = x^n + 0$. This is an extremely small speedup.

7. COMBINATION

Consider the polynomials $g \in (\mathbf{Z}/n)[x]$ such that $g^n = g(x^n)$ in $(\mathbf{Z}/n)[x]/(x^r-1)$.

The goal of this section is to count the number of different possibilities for the image of g in $\mathbf{F}_p[x]/h$. Here p is a prime divisor of n, and h is an irreducible factor of Φ_r in $\mathbf{F}_p[x]$, as in the proofs of Theorem 2.1 and Theorem 2.2. What matters, eventually, is that the number of images is larger than the number $|t - u| + 1$ constructed in these proofs.

Section 6 explained how to compute a finite coprime set T of monic non-constant polynomials $f \in (\mathbf{Z}/n)[x]$ such that $f^n = f(x^n)$ in $(\mathbf{Z}/n)[x]/(x^r - 1)$. Agrawal, Kayal, and Saxena build many polynomials g by multiplying elements of T. Any function $e : T \to \{0, 1, \dots\}$ with $\sum_{f \in T} e(f) \deg f < \deg h$ determines a product $g = \prod_{f \in T} f^{e(f)}$ in $(\mathbf{Z}/n)[x]$ with $\deg g < \deg h$ and $g^n = g(x^n)$ in $(\mathbf{Z}/n)[x]/(x^r-1)$. The elements of T are coprime in $(\mathbf{Z}/n)[x]$, hence in $\mathbf{F}_p[x]$, so different e's produce different g's in $\mathbf{F}_p[x]$, hence in $\mathbf{F}_p[x]/h$. Furthermore, all these products are nonzero in $\mathbf{F}_p[x]$, hence in $\mathbf{F}_p[x]/h$; one obtains another polynomial $g = 0$.

The number of images of g in $\mathbf{F}_p[x]/h$ obtained this way is 1 plus the number of choices of e. This is the same as the sum of coefficients of $y^0, y^1, \dots, y^{\deg h-1}$ in the power series $1 + \prod_{f \in T}(1 + y^{\deg f} + y^{2 \deg f} + \cdots)$. One can easily compute these coefficients, and thus compute a lower bound for the number of images of g, given a lower bound for $\deg h$.

The original Agrawal-Kayal-Saxena paper used an extremely crude lower bound for the number of images, and compensated by choosing r and $\#S$ unnecessarily large. It is important to use a good lower bound—such as the exact number—so that r and $\#S$ can be reduced.

In the extreme case that $\deg f = 1$ for all $f \in T$, the series is $1 + (1 - y)^{-\#T} = 1 + \sum_{k \geq 0}(-1)^k \binom{-\#T}{k} y^k$. The coefficient sum is $1 + \binom{\#T+\deg h-1}{\#T}$.

If T is the union of T_1 and T_2 where $\deg f = 1$ for all $f \in T_1$ and $\deg f = 2$ for all $f \in T_2$, then the series is

$$1 + (1 - y)^{-\#T_1}(1 - y^2)^{-\#T_2} = 1 + \sum_{k \geq 0}(-1)^k \binom{-\#T_1}{k} y^k \sum_{j \geq 0}(-1)^j \binom{-\#T_2}{j} y^{2j}.$$

Note that $\binom{-T_2}{j}$ is relatively small compared to $\binom{-T_1}{2j}$, unless T_2 is much larger than T_1; degree-2 polynomials are much less valuable than degree-1 polynomials.

Increasing the degree bound to $\varphi(r)$. Lenstra pointed out on 2002.08.13 that, when n is a primitive root modulo r, these polynomials g are distinct in $\mathbf{F}_p[x]/h$ if they are distinct in $\mathbf{F}_p[x]/\Phi_r$. Thus one can add coefficients of $y^0, y^1, \dots, y^{\varphi(r)-1}$, not just $y^0, y^1, \dots, y^{\deg h-1}$.

This is the basic difference between Theorem 2.1 and Theorem 2.2. Here's the general proof: By hypothesis, $g^n = g(x^n)$ in $(\mathbf{Z}/n)[x]/(x^r - 1)$. Substitute x^{n^i} for x: $g(x^{n^i})^n = g(x^{n^{i+1}})$ in $(\mathbf{Z}/n)[x]/(x^r - 1)$. By induction, $g^{n^i} = g(x^{n^i})$ in $(\mathbf{Z}/n)[x]/(x^r - 1)$ for all $i \geq 0$. If $c \in \{0, 1, \dots, r - 1\}$ and $\gcd\{r, c\} = 1$ then $c \equiv n^a \pmod{r}$ for some a, so $g(x^c) = g(x^{n^a}) = g^{n^a}$ in $(\mathbf{Z}/n)[x]/(x^r - 1)$, so $g(x^c) = g^{n^a}$ in $\mathbf{F}_p[x]/h$. If f is another such polynomial with $f = g$ in $\mathbf{F}_p[x]/h$, then $f(x^c) = f^{n^a} = g^{n^a} = g(x^c)$ in $\mathbf{F}_p[x]/h$, so $(f - g)(x^c) = 0$ in $\mathbf{F}_p[x]/h$; this is true for all c, so $f - g$ is divisible by Φ_r.

Using negative powers. Voloch pointed out on 2002.09.13 that one could use negative powers of the elements of T.

Choose a nonnegative integer $d \leq \varphi(r)$. Voloch took $d = \lfloor 0.5\varphi(r) \rfloor$. An obvious generalization is to allow any d.

Assume for simplicity that $\deg f \in \{1, 2\}$ for all $f \in T$, assume as above that n is a primitive root modulo r, and assume that $\varphi(r) \geq 3$. Then n must have a prime factor p that, modulo r, is neither 1 nor -1. The degree of an irreducible factor of Φ_r in $\mathbf{F}_p[x]$ is the order of p modulo r, which is at least 3; thus every $f \in T$ is invertible in $\mathbf{F}_p[x]/\Phi_r$.

Consider functions $e : T \to \mathbf{Z}$ such that $\sum_{f \in T}[e(f) > 0]e(f) \deg f < \varphi(r) - d$ and $\sum_{f \in T} -[e(f) < 0]e(f) \deg f \leq d$. Each function determines a product $g = \prod_{f \in T} f^{d+e(f)}$ in $(\mathbf{Z}/n)[x]$.

Say e' is another such function, determining a product g'. I claim that g and g' are distinct in $\mathbf{F}_p[x]/\Phi_r$ unless $e = e'$. Indeed, say $\prod_{f \in T} f^{d+e(f)} = \prod_{f \in T} f^{d+e'(f)}$ in $\mathbf{F}_p[x]/\Phi_r$. Divide by $\prod_{f \in T} f^d$, multiply by $\prod_{f \in T} f^{-[e(f)<0]e(f)}$, and multiply by $\prod_{f \in T} f^{-[e'(f)<0]e'(f)}$, to obtain the equation

$$\prod_{f \in T} f^{[e(f)>0]e(f)-[e'(f)<0]e'(f)} = \prod_{f \in T} f^{[e'(f)>0]e'(f)-[e(f)<0]e(f)}$$

between two polynomials of degree below $\varphi(r) - d + d = \varphi(r)$. The exponents must match: $[e(f) > 0]e(f) - [e'(f) < 0]e'(f) = [e'(f) > 0]e'(f) - [e(f) < 0]e(f)$, so $e(f) = e'(f)$.

In the special case $\deg f = 1$, there are at least $\binom{\#T+d}{\#T}\binom{\#T+\varphi(r)-1-2d}{\#T-d}$ functions e. This lower bound was used for the figures $0.0011324\ldots$ and $0.00050329\ldots$ in Section 4.

Vaaler pointed out a slightly better lower bound (reported to me on 2003.01.24), namely $\binom{\#T}{i}\binom{d}{i}\binom{\#T-i}{j}\binom{\varphi(r)-1-d}{j}$ for any i, j: choose i positions where e will be negative, choose negative values for those i positions adding up to at most d, choose j remaining positions where e will be positive, and choose positive values for those j positions adding up to at most $\varphi(r) - d - 1$. The exact number of functions e is the sum of this quantity over all i between 0 and d and all j between 0 and $\varphi(r) - d - 1$, so it is reasonably well approximated by the maximum of this quantity.

Increasing the degree bound further with ABC. The strategy used above to count products in $\mathbf{F}_p[x]/\Phi_r$ is to count products in $\mathbf{F}_p[x]$ of degree below $\varphi(r)$. Voloch pointed out another strategy on 2002.08.28, using Mason's ABC theorem. I improved Voloch's result on 2002.08.31, and improved it further on 2003.01.24.

Assume for simplicity that $\deg f = 1$ for all $f \in T$, assume as above that n is a primitive root modulo r, and assume that $\varphi(r) \geq 2$. Then n must have a prime factor p that is not 1 modulo r. Every $f \in T$ is invertible in $\mathbf{F}_p[x]/\Phi_r$. Also assume that n has no prime divisors $\leq 3\varphi(r)$; then $p > 3\varphi(r)$.

Consider products $\prod_{f \in T} f^{e(f)}$ with $e(f) \geq 0$. First step: Assume that *three* distinct products A, B, C in $\mathbf{F}_p[x]$, with $\gcd \{A, B, C\} = 1$, are all the same modulo Φ_r; I claim that $\max \{\deg A, \deg B, \deg C\} > 2\varphi(r) - \deg \operatorname{rad} ABC$. Here $\operatorname{rad} X$ means the product of all the irreducibles that divide X.

Assume without loss of generality that $\deg A = \max \{\deg A, \deg B, \deg C\}$. Also assume that $\deg A \leq 2\varphi(r)$; otherwise there is nothing to prove.

Define $U = (A - B)/\Phi_r$ and $V = (A - C)/\Phi_r$. Then $U \neq 0$; $V \neq 0$; $V - U \neq 0$; $U, V, V - U$ have degree at most $\deg A - \varphi(r)$; and $V(A - B) = U(A - C)$, i.e., $(V - U)A + UC = VB$.

Define $D = \gcd\{(V - U)A, VB, UC\}$. Then $(V - U)A/D + UC/D = VB/D$; $\gcd\{(V - U)A/D, UC/D, VB/D\} = 1$; and $(V - U)A/D, UC/D, VB/D$ all have degree at most $2 \deg A - \varphi(r) \leq 3\varphi(r) < p$. Apply the ABC theorem to see that $\deg(V - U)A/D < \deg \operatorname{rad}(((V - U)A/D)(VB/D)(UC/D))$.

Fact: $D \operatorname{rad}(((V - U)A/D)(VB/D)(UC/D))$ divides $UV(V - U) \operatorname{rad} ABC$. Thus $\deg(V - U)A < \deg UV(V - U) \operatorname{rad} ABC$.

(In other words: If $d = \min\{w + a, v + b, u + c\}$ and $\min\{a, b, c\} = 0$ then $d + [u + v + w + a + b + c > 3d] \leq u + v + w + [a + b + c > 0]$. Proof: Without loss of generality assume $a = 0$. Then $d \leq w \leq u + v + w$. If $d < u + v + w$ then $d + [\cdots] \leq d + 1 \leq u + v + w \leq u + v + w + [\cdots]$ as claimed. If $a + b + c > 0$ then $d + [\cdots] \leq u + v + w + 1 = u + v + w + [\cdots]$ as claimed. Otherwise $u + v + w + a + b + c = d \leq 3d$ so $d + [u + v + w + a + b + c > 3d] = d \leq u + v + w \leq u + v + w + [\cdots]$ as claimed.)

Hence $\deg A < \deg UV \operatorname{rad} ABC \leq 2(\deg A - \varphi(r)) + \deg \operatorname{rad} ABC$; i.e., $\deg A > 2\varphi(r) - \deg \operatorname{rad} ABC$ as claimed.

(In particular, $\deg A > 2\varphi(r) - \#T$. If $\#T \leq \varphi(r)$ then this bound improves on the obvious bound $\deg A \geq \varphi(r)$. See below for results when $\#T > \varphi(r)$. Voloch's result was $\deg A > 1.2\varphi(r) - 0.2\#T$, using the weaker observation that D divides $\gcd\{UV(V - U)A, UV(V - U)B, UV(V - U)C\} = UV(V - U)\gcd\{A, B, C\} = UV(V - U)$.)

Next step: Assume that three distinct products A, B, C in $\mathbf{F}_p[x]$ are all the same modulo Φ_r. Note that the assumption $\gcd\{A, B, C\} = 1$ has been dropped. Define $G = \gcd\{A, B, C\}$; then G is coprime to Φ_r, so $A/G, B/G, C/G$ are all the same modulo Φ_r, so

$$\max\left\{\deg \frac{A}{G}, \deg \frac{B}{G}, \deg \frac{C}{G}\right\} > 2\varphi(r) - \deg \operatorname{rad} \frac{ABC}{GGG} \geq 2\varphi(r) - \deg \operatorname{rad} ABC.$$

But $\deg G \geq \deg \operatorname{rad} G = \deg \operatorname{rad} ABC - \deg \operatorname{rad} A - \deg \operatorname{rad} B - \deg \operatorname{rad} C + \deg \operatorname{rad} \gcd\{A, B\} + \deg \operatorname{rad} \gcd\{A, C\} + \deg \operatorname{rad} \gcd\{B, C\}$ by inclusion-exclusion. Thus $\max\{\deg A, \deg B, \deg C\} > 2\varphi(r) - \deg \operatorname{rad} A - \deg \operatorname{rad} B - \deg \operatorname{rad} C + \deg \operatorname{rad} \gcd\{A, B\} + \deg \operatorname{rad} \gcd\{A, C\} + \deg \operatorname{rad} \gcd\{B, C\}$.

Final step: Assume that $m \geq 3$ products in $\mathbf{F}_p[x]$ are all the same modulo Φ_r. Let Z be the set of those products. Define $d = \max\{\deg A : A \in Z\}$. Then $d > 2\varphi(r) - \deg \operatorname{rad} A - \deg \operatorname{rad} B - \deg \operatorname{rad} C + \deg \operatorname{rad} \gcd\{A, B\} + \deg \operatorname{rad} \gcd\{A, C\} + \deg \operatorname{rad} \gcd\{B, C\}$ for any distinct $A, B, C \in Z$. Average this inequality over all choices of A, B, C to see that

$$d > 2\varphi(r) - \frac{3}{m}\sum_A \deg \operatorname{rad} A + \frac{6}{m(m-1)}\sum_{A \neq B} \deg \operatorname{rad} \gcd\{A, B\}.$$

On the other hand, $\#T \geq \deg \operatorname{rad} \bigcup_A A \geq \sum_A \deg \operatorname{rad} A - \sum_{A \neq B} \deg \operatorname{rad} \gcd\{A, B\}$ by inclusion-exclusion, so

$$d + \frac{3}{m}\#T > 2\varphi(r) - \left(\frac{3}{m} - \frac{6}{m(m-1)}\right)\sum_{A \neq B} \deg \operatorname{rad} \gcd\{A, B\}.$$

Note that $3/m - 6/m(m-1) \geq 0$. One can bound each term $\deg \operatorname{rad} \gcd\{A, B\}$ by the simple observation that $A/\gcd\{A, B\}$ and $B/\gcd\{A, B\}$ are distinct congruent

products of degree at most $d - \deg \gcd \{A, B\}$: thus $d - \deg \gcd \{A, B\} \geq \varphi(r)$, so $\deg \operatorname{rad} \gcd \{A, B\} \leq d - \varphi(r)$. Hence $d + (3/m)\#T > 2\varphi(r) - (d - \varphi(r))(3m - 9)/2$; i.e., $d > ((3m - 5)/(3m - 7))\varphi(r) - (6/m(3m - 7))\#T$.

Summary: Take any integer $m \geq 3$. There are $\binom{\#T+d}{\#T}$ products of degree at most d, where $d = \lfloor ((3m - 5)/(3m - 7))\varphi(r) - (6/m(3m - 7))\#T \rfloor$. It is impossible for m of them to be congruent modulo Φ_r. Thus there are at least

$$\frac{1}{m-1}\binom{\#T + d}{\#T} = \frac{1}{m-1}\binom{\left\lfloor \frac{3m-5}{3m-7}\varphi(r) + \left(1 - \frac{6}{m(3m-7)}\right)\#T \right\rfloor}{\#T}$$

different products in $\mathbf{F}_p[x]/\Phi_r$. For example, the number of different products is at least $\frac{1}{2}\binom{2\varphi(r)}{\#T}$; at least $\frac{1}{3}\binom{\lfloor 1.4\varphi(r)+0.7\#T \rfloor}{\#T}$; at least $\frac{1}{4}\binom{\lfloor 1.25\varphi(r)+0.85\#T \rfloor}{\#T}$; and so on.

Can this idea be combined with using negative powers to improve the speed of AKS? The difficulty is that clearing the denominator in (A, B, C) has twice as much of an impact as clearing the denominator in (A, B).

Can the bound $d > ((3m - 5)/(3m - 7))\varphi(r) - (6/m(3m - 7))\#T$ be improved? Is there a better way to handle m products than to combinatorially merge the information that ABC provides for sets of 3 products?

Using leading coefficients. All the nonzero polynomials constructed above are monic. One can also use non-monic polynomials. This is an extremely small speedup.

8. CONGRUENCE

Section 7 explained how to count images in $\mathbf{F}_p[x]/h$ of polynomials $g \in (\mathbf{Z}/n)[x]$ such that $g^n = g(x^n)$ in $(\mathbf{Z}/n)[x]/(x^r - 1)$. Here p is a prime divisor of n, and h is an irreducible factor of Φ_r in $\mathbf{F}_p[x]$.

Theorems 2.1 and 2.2 use these images as follows. Repeated substitution implies that $g^t = g(x^t)$ in $\mathbf{F}_p[x]/(x^r - 1)$ for any product t of powers of n and p. If u is another product with different exponents, and if $t \equiv u \pmod{r}$, then $g^t = g(x^t) = g(x^u) = g^u$ in $\mathbf{F}_p[x]/(x^r - 1)$, hence in $\mathbf{F}_p[x]/h$. The number of images is larger than $|t - u| + 1$, but the equation $g^t = g^u$ has at most $|t - u| + 1$ roots in a field, so $t = u$, so n is a power of p.

Why is the number of images larger than $|t - u| + 1$? Answer: r, S are chosen so that the number of images exceeds $n^{2\lfloor \sqrt{r} \rfloor}$; there exist t, u with $|t - u| + 1 \leq n^{2\lfloor \sqrt{r} \rfloor}$.

Focusing on units. Lenstra pointed out on 2002.08.13 that one could replace $\lfloor \sqrt{r} \rfloor$ with $\lfloor \sqrt{\varphi(r)} \rfloor$. This is incorporated into Theorem 2.2. This is an extremely small improvement if r is prime.

Replacing p, n with $p, n/p$. Lenstra pointed out on 2002.08.13 that one could take t and u as products of powers of n/p and p, not just n and p. (Indeed, tp^k and up^k are products of powers of n and p for some k, so $g^{tp^k} = g^{up^k}$ in $\mathbf{F}_p[x]/h$; pth powering is invertible in $\mathbf{F}_p[x]/h$, so $g^t = g^u$.)

Consider the products $(n/p)^i p^j$ with $0 \leq i \leq \lfloor \sqrt{\varphi(r)} \rfloor$ and $0 \leq j \leq \lfloor \sqrt{\varphi(r)} \rfloor$; note that $1 \leq (n/p)^i p^j \leq n^{\lfloor \sqrt{\varphi(r)} \rfloor}$. There are $(\lfloor \sqrt{\varphi(r)} \rfloor + 1)^2 > \varphi(r)$ pairs (i, j), so there are distinct pairs $(i, j), (k, \ell)$ such that $(n/p)^i p^j \equiv (n/p)^k p^\ell \pmod{r}$. Bottom line: There exist t, u with $|t - u| + 1 \leq n^{\lfloor \sqrt{\varphi(r)} \rfloor}$.

Balancing p and n/p. Poonen pointed out on 2002.08.26 that there exist t, u with $|t - u| + 1 \leq n^{\lceil \sqrt{\varphi(r)/2} \rceil}$.

Consider the triangle of real numbers (x, y) with $0 \leq x$, $0 \leq y$, and $(n/p)^x p^y \leq n^{\sqrt{\varphi(r)/2}}$. Define $i = \lfloor x \rfloor$ and $j = \lfloor y \rfloor$; then i, j are integers with $0 \leq i$, $0 \leq j$, and $(n/p)^i p^j \leq (n/p)^x p^y \leq n^{\sqrt{\varphi(r)/2}}$. Furthermore, (x, y) is contained in the rectangle $[i, i + 1] \times [j, j + 1]$. The number of pairs (i, j) is exactly the total area of those rectangles, which is more than the area of the triangle, which in turn exceeds $2(\sqrt{\varphi(r)/2})^2 = \varphi(r)$.

Using Minkowski's theorem. Lenstra pointed out to me on 2002.12.23 that there exist t, u with $|t - u| + 1 \leq n^{\lceil \sqrt{\varphi(r)/3} \rceil}$.

Consider the triangle T of real numbers (x, y) with $0 \leq x$, $0 \leq y$, and $(n/p)^x p^y \leq n^{\sqrt{\varphi(r)/3}}$. Define C as the set of differences $(x, y) - (x', y')$ of points in T. Then the area of C is 6 times the area of T; hence the area of C is more than $4\varphi(r)$.

The lattice of integers (i, j) such that $(n/p)^i p^j \equiv 1 \pmod{r}$ has determinant $\varphi(r)$, so it has a nonzero point $(i, j) \in C$ by Minkowski's theorem. Assume without loss of generality that $i \geq 0$. If $j \geq 0$ then (i, j) is in T; define $t = (n/p)^i p^j$ and $u = 1$. Otherwise $(i, 0)$ and $(0, -j)$ are in T; define $t = (n/p)^i$ and $u = p^{-j}$.

DEPARTMENT OF MATHEMATICS, STATISTICS, AND COMPUTER SCIENCE (M/C 249), THE UNIVERSITY OF ILLINOIS AT CHICAGO, CHICAGO, IL 60607–7045
E-mail address: djb@cr.yp.to