MATHEMATICS OF COMPUTATION
Volume 76, Number 257, January 2007, Pages 389–403
S 0025-5718(06)01786-8
Article electronically published on September 14, 2006

PROVING PRIMALITY IN ESSENTIALLY QUARTIC RANDOM TIME

DANIEL J. BERNSTEIN

ABSTRACT. This paper presents an algorithm that, given a prime n, finds and verifies a proof of the primality of n in random time $(\lg n)^{4+o(1)}$. Several practical speedups are incorporated into the algorithm and discussed in detail.

1. INTRODUCTION

This paper presents an algorithm that proves the primality of any prime n in random time $(\lg n)^{4+o(1)}$:

- Section 3 defines certificates and proves that n is a prime power if it has a certificate.
- Section 4 presents an algorithm that, given a prime n, finds a reasonably small certificate for n in random time $(\lg n)^{2+o(1)}$.
- Section 6 presents an algorithm to verify a reasonably small certificate in time $(\lg n)^{4+o(1)}$.

One can prove that n is not a perfect power in time $(\lg n)^{1+o(1)}$, as explained in [6] and [10], so prime-power proving is tantamount to prime proving.

Section 7 discusses verification speed in more detail. Some of the complications in the certificate definition are irrelevant to the $4+o(1)$ result but produce speedups visible at the level of detail of Section 7. A simplified proof that n is a prime power under stronger assumptions, without these complications, appears in Section 2.

Genealogy. This algorithm uses an idea that one might call "proving primality with combinatorics." This idea was introduced by Agrawal, Kayal, and Saxena in [4]. (Primitive forms of the idea were used by Fellows and Koblitz in [15], and by Konyagin and Pomerance in [20].) The Agrawal-Kayal-Saxena algorithm proves primality in polynomial time, using combinatorics in cyclotomic extensions of $\mathbf{Z}/n$.

The algorithm in this paper replaces cyclotomic extensions with random Kummer extensions, so that it can twist $x-1$ into $\zeta x-1$, $\zeta^2 x-1$, etc.; see the proof of Theorem 2.1. This idea was introduced by Berrizbeitia in [11], in the special case of Kummer extensions whose degrees are powers of 2. Berrizbeitia's algorithm proves primality in random time $(\lg n)^{4+o(1)}$ for a sparse set of primes n, namely those for which n^2-1 is divisible by some power of 2 near $(\lg n)^2$.

Received by the editor February 13, 2004 and, in revised form, December 9, 2004.
2000 *Mathematics Subject Classification.* Primary 11Y11.
The author was supported by the National Science Foundation under grant DMS–0140542, and by the Alfred P. Sloan Foundation. He used the libraries at the Mathematical Sciences Research Institute, the University of California at Berkeley, and the American Institute of Mathematics.

Cheng in [12] adapted Berrizbeitia's idea to prime degrees. Cheng's algorithm proves primality in random time $(\lg n)^{4+o(1)}$ for a larger set of primes n, namely those for which $n-1$ is divisible by a prime $e \approx (\lg n)^2$.

This paper generalizes to arbitrary positive integers $e \approx (\lg n)^2$ dividing $n^d - 1$ for any $d \in n^{o(1)}$. A standard result from analytic number theory implies that every prime n has a suitable pair (d, e); see Theorems 5.1 and 5.2. For practical purposes, the only interesting case is $d = 1$, as discussed in Section 7.

My generalization was independent of Cheng's adaptation. I read Berrizbeitia's paper on 26 January 2003 and promptly sent email to a few people saying how I expected it to generalize to any n. I was then told about Cheng's paper, which had been published on 16 January 2003. I posted a draft of this paper, with a detailed proof of Theorem 3.2, on 28 January 2003, and announced the result on the NMBRTHRY mailing list on 29 January 2003.

Mihăilescu and Avanzi realized, independently of my work, that Berrizbeitia's idea could be generalized to arbitrary positive integers e. They eventually posted their generalization; see [24]. See Section 8 for a simplified proof of a similar generalization, and a discussion of how this generalization differs from mine.

Most papers in this field have been written with a casual disregard for constant factors in run time: crude binomial-coefficient bounds and suboptimal parameter choices are typically embedded into statements of theorems. Starting a few days after [4] and continuing through this paper, I have attempted to state each theorem in its natural level of generality, to incorporate any modifications that would reduce run time, and to optimize parameters. See Sections 3 and 7 of this paper.

Competition. Another way to prove the primality of n is to exhibit a factor of the Jacobian group of a hyperelliptic curve over $\mathbf{Z}/n$. Adleman and Huang in [2] proved that every prime n has a certificate of this type that can be found in random time $(\lg n)^{O(1)}$ and verified in time at most $(\lg n)^{3+o(1)}$. The $O(1)$ here is large.

A previous algorithm of Atkin, using small-discriminant complex-multiplication elliptic curves, is *conjectured* to find a certificate of the same type in time at most $(\lg n)^{5+o(1)}$. An improved algorithm, pointed out by Shallit and reported in [21, page 711], is *conjectured* to find a certificate of the same type in time at most $(\lg n)^{4+o(1)}$. As above, the certificates can be verified in time $(\lg n)^{3+o(1)}$.

The algorithm in this paper is *proven* to find and verify certificates in random time at most $(\lg n)^{4+o(1)}$.

For readers who want to actually prove the primality of various numbers n, rather than prove theorems about how quickly one can prove primality, the impact of the new algorithm is less clear. Is the $(\lg n)^{4+o(1)}$ time for the new algorithm smaller than the $(\lg n)^{4+o(1)}$ time to find elliptic-curve certificates? My current impression is that the answer is no, but that further results along the lines of [7] or [13] could change the answer. See the end of Section 7 for further discussion of [7] and [13].

The literature contains many more methods to distinguish prime numbers from composite numbers. See my survey paper [9] for a comparison of the speed and effectiveness of these methods. For example, there are randomized compositeness-proving algorithms that reliably detect (but do not prove) primality and that take time only $(\lg n)^{2+o(1)}$.

2. The idea in a nutshell

Theorem 2.1. *Let n, d, and e be positive integers such that $2^e - 1 \geq n^{2d\lfloor\sqrt{e}\rfloor}$ and e divides $n^d - 1$. Let f be a monic polynomial in $(\mathbf{Z}/n)[y]$ of degree d. Define R as the ring $(\mathbf{Z}/n)[y]/f$. Let r be an element of R such that $r^{n^d-1} = 1$ in R, $r^{(n^d-1)/q} - 1$ is a unit in R for each prime q dividing e, and $r - 1$ is a unit in R. If $(x-1)^{n^d} = r^{(n^d-1)/e}x - 1$ in the ring $R[x]/(x^e - r)$ then n is a power of a prime.*

Theorem 2.1 improves in two ways upon the theorems of Berrizbeitia in [11] and Cheng in [12]:

- d is allowed to be any positive integer. Berrizbeitia considered only $d \in \{1, 2\}$, and Cheng considered only $d = 1$. Larger d's are important for the $(\lg n)^{4+o(1)}$ result in this paper. On the other hand, as discussed in Section 7, the case $d = 1$ is the only important case in practice.
- e is allowed to be any positive divisor of $n^d - 1$. Berrizbeitia considered only powers of 2 (although with slightly more general moduli $x^{2^i e} - r$), and Cheng considered only primes e; the proofs relied on e having only one prime divisor. Arbitrary e's are important for the $(\lg n)^{4+o(1)}$ result in this paper. Arbitrary e's also save time in practice, because they allow many more n's to be handled with $d = 1$.

Theorem 3.2 saves more time by allowing somewhat smaller e's.

Proof. If $n = 1$ then n is a power of a prime, so assume that $n > 1$.

Step 1: Move to a field. R is a nonzero ring, so it maps onto a field k. Explicitly: find a prime p dividing n; find an irreducible polynomial g in $\mathbf{F}_p[y]$ dividing the image of f; then $k = \mathbf{F}_p[y]/g$ is a field.

Write $N = \#R = n^d$ and $P = \#k = p^{\deg g}$. Define ζ as the image of $r^{(N-1)/e}$ in k. Then ζ has order e in k. (Indeed, $r^{N-1} = 1$ in R by hypothesis, so $\zeta^e = 1$ in k. Furthermore, if q is a prime dividing e, then $r^{(N-1)/q} - 1$ is a unit in R by hypothesis, so its image $\zeta^{e/q} - 1$ in k is a unit; hence $\zeta^{e/q} \neq 1$ in k.) Consequently, e divides $P - 1$.

Step 2: Combinatorially enumerate many powers of $x - 1$. Define A as the ring $k[x]/(x^e - r)$. By hypothesis $(x-1)^N = r^{(N-1)/e}x - 1$ in $R[x]/(x^e - r)$, so $(x-1)^N = \zeta x - 1$ in A. Substitute $\zeta^i x$ for x: $(\zeta^i x - 1)^N = \zeta^{i+1}x - 1$ in the ring $k[x]/((\zeta^i x)^e - r) = A$. Thus $(x-1)^{N^i} = \zeta^i x - 1$ in A for any integer $i \geq 0$.

There are $2^e - 1$ vectors $(a_0, a_1, \ldots, a_{e-1}) \in \{0, 1\}^e$ such that $\sum_i a_i \leq e - 1$. Any product $\prod_i (\zeta^i x - 1)^{a_i} = (x-1)^{\sum_i N^i a_i}$ is a power of $x - 1$ in A. I claim that these products are distinct, so there are at least $2^e - 1$ powers of $x - 1$ in A.

Indeed, say a, b are two such vectors with $\prod_i (\zeta^i x - 1)^{a_i} = \prod_i (\zeta^i x - 1)^{b_i}$ in A. Then $\prod_i (\zeta^i x - 1)^{a_i} = \prod_i (\zeta^i x - 1)^{b_i}$ in $k[x]$: distinct polynomials of degree at most $e - 1$ remain distinct when reduced modulo $x^e - r$. The polynomials $x - 1, \zeta x - 1, \ldots, \zeta^{e-1}x - 1$ are coprime in $k[x]$, so $a_i = b_i$ by unique factorization.

Step 3: Find colliding powers of $x - 1$. The nonzero element $r^{(P-1)/e}$ of k has eth power $r^{P-1} = 1$; but the eth roots of 1 in k are exactly the powers of ζ. Thus $r^{(P-1)/e} = \zeta^\ell$ in k for some integer ℓ. Now $x^P = x^{P-1}x = r^{(P-1)/e}x = \zeta^\ell x$ in A, so $x^{P^j} = \zeta^{j\ell}x$ in A for any integer $j \geq 0$. Thus $(x-1)^{N^i P^j} = \zeta^{i+j\ell}x - 1$ in A.

Consider the pairs (i, j) with $0 \leq i \leq \lfloor\sqrt{e}\rfloor$ and $0 \leq j \leq \lfloor\sqrt{e}\rfloor$. There are $(\lfloor\sqrt{e}\rfloor + 1)^2 > e$ pairs (i, j), and only e possible powers $\zeta^{i+j\ell}$, so there are distinct

pairs $(i,j),(i',j')$ with $\zeta^{i+j\ell} = \zeta^{i'+j'\ell}$. Define $u = N^i P^j$ and $v = N^{i'} P^{j'}$; then u and v are positive integers bounded by $N^{2\lfloor\sqrt{e}\rfloor}$, and $(x-1)^u = (x-1)^v$.

The remainder $(x^e - r) \bmod (x-1) = 1-r$ is a unit in k, so $x-1$ is a unit in A. Thus $(x-1)^{u-v} = 1$ in the unit group A^*. If $u \neq v$ then there are at most $|u-v|$ powers of $x-1$, but $|u-v| < N^{2\lfloor\sqrt{e}\rfloor} \leq 2^e - 1$; contradiction.

Hence $u = v$; i.e., $N^{i-i'} = P^{j'-j}$. If $i = i'$ then $j' = j$; contradiction. Thus a nontrivial power of n is a power of p; so n is a power of p. □

3. Certificates

Definition 3.1. Let n, d, and e be positive integers. Let c and c_- be integers. Let f be a monic polynomial in $(\mathbf{Z}/n)[y]$ of degree d. Define R as the ring $(\mathbf{Z}/n)[y]/f$. Let r be an element of R. Let S be a subset of R. Assume that

- e divides $n^d - 1$;
- $e > c \geq c_- \geq 0$;
- $r^{n^d-1} = 1$ in R;
- $r^{(n^d-1)/q} - 1$ is a unit in R for each prime q dividing e;
- s is a unit in R for all $s \in S$;
- $s^e - (s')^e$ is a unit in R for all distinct $s, s' \in S$;
- $s^e - r$ is a unit in R for all $s \in S$;
- $\binom{e\#S}{c_-}\binom{c}{c_-}\binom{e\#S-c_-+e-1-c}{e-1-c} \geq n^{d\lceil\sqrt{e/3}\rceil}$; and
- $(x-s)^{n^d} = r^{(n^d-1)/e}x - s$ in the ring $R[x]/(x^e - r)$ for all $s \in S$.

Then (d, e, c, c_-, f, r, S) is a **certificate for** n.

For example, $(1, 840, 419, 246, y, 17, \{1\})$ is a certificate for

$$3141592653589793238462643383279502884 1;$$

$(1, 2430, 1214, 928, y, 2, \{1, 2\})$ is a certificate for

$$271828182845904523536028747135266249775724709369995957496696762772407663035354759457 1;$$

and $(1, 57449, 28724, 16826, y, 2, \{1\})$ is a certificate for $2^{1024} + 643$.

Theorem 3.2. *Let n, d, and e be positive integers. Let c and c_- be integers. Let f be a monic polynomial in $(\mathbf{Z}/n)[y]$ of degree d. Define R as the ring $(\mathbf{Z}/n)[y]/f$. Let r be an element of R. Let S be a subset of R. Assume that (d, e, c, c_-, f, r, S) is a certificate for n. Then n is a power of a prime.*

Readers who warmed up by reading the proof of Theorem 2.1 will recognize the overall structure, and many components, of the proof of Theorem 3.2. However, through the definition of a certificate, Theorem 3.2 includes three features not present in Theorem 2.1:

- It uses $\sqrt{e/3}$ as suggested by Lenstra, instead of $2\sqrt{e}$. This reduces the lower bound on e by a factor of about 12. (Intermediate results by various people are not discussed here.)
- It allows $c > 0$ as suggested by Voloch, with an optimization suggested by Vaaler; see the proof to understand the role of c. This reduces the lower bound on e by an extra factor of about 6.4673912111.
- It allows $\#S$ to vary. (Berrizbeitia and Cheng considered only $\#S = 1$, or $\#S = 2^i$ for modulus $x^{2^i e} - r$.) This allows e to be chosen even smaller, at some cost; see Section 7.

Smaller e's allow faster verification of the hypotheses, i.e., faster primality proofs; on the other hand, they make the "human part" of the proof more complicated.

Additional comments on the proof appear at the end of this section.

Proof. If $n = 1$ then n is a power of a prime, so assume that $n > 1$.

Step 1: Move to a field. R is a nonzero ring, so it maps onto a field k. Explicitly: find a prime p dividing n; find an irreducible polynomial g in $\mathbf{F}_p[y]$ dividing the image of f; then $k = \mathbf{F}_p[y]/g$ is a field.

Write $N = \#R = n^d$ and $P = \#k = p^{\deg g}$. Note that P divides N. If $N = P$ then n is a power of p, so assume that $N > P$. (Similarly, one can assume that $N \neq P^2$; this allows a variation in the Minkowski argument below.)

Define ζ as the image of $r^{(N-1)/e}$ in k. Then ζ has order e in k. (Indeed, $r^{N-1} = 1$ in R by hypothesis, so $\zeta^e = 1$ in k. Furthermore, if q is a prime dividing e, then $r^{(N-1)/q} - 1$ is a unit in R by hypothesis, so its image $\zeta^{e/q} - 1$ in k is a unit; hence $\zeta^{e/q} \neq 1$ in k.) Consequently, e divides $P - 1$.

Step 2: Combinatorially enumerate many group elements. Find an irreducible polynomial h in $k[x]$ dividing the image of $x^e - r$. Then $k[x]/h$ is a field.

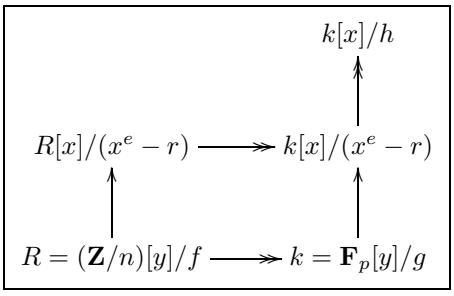

If $s \in S$ then $(x - s)^N = r^{(N-1)/e}x - s$ in $R[x]/(x^e - r)$ by hypothesis, so $(x - s)^N = \zeta x - s$ in $k[x]/(x^e - r)$. Substitute $\zeta^i x$ for x: $(\zeta^i x - s)^N = \zeta^{i+1}x - s$ in $k[x]/((\zeta^i x)^e - r) = k[x]/(x^e - r)$. Thus $(x - s)^{N^i} = \zeta^i x - s$ in $k[x]/(x^e - r)$ for any integer $i \geq 0$.

Note that $\zeta^i x - s$ is a unit in $k[x]/h$. (The remainder $(x^e - r) \bmod (\zeta^i x - s) = s^e - r$ is a unit in k, so $\zeta^i x - s$ is a unit in $k[x]/(x^e - r)$, hence in $k[x]/h$.) Note also that $\zeta^i x - s$ and $\zeta^{i'} x - s'$ are coprime in $k[x]$ unless $(\zeta^i, s) = (\zeta^{i'}, s')$. (If they are not coprime, then $s\zeta^{i'} = s'\zeta^i$ in k, so $s^e = (s')^e$ in k. If $s \neq s'$ then $s^e - (s')^e$ is a unit in R by hypothesis, so it is a unit in k; contradiction. Thus $s = s'$; so $s\zeta^{i'} = s\zeta^i$; also s is a unit in R by hypothesis, so $\zeta^{i'} = \zeta^i$.)

Consider functions $a : \{0, 1, \ldots, e - 1\} \times S \to \mathbf{Z}$ such that

- $\#\{(i, s) : a(i, s) < 0\} = c_-$;
- $\sum_{i,s} -a(i, s)[a(i, s) < 0] \leq c$; and
- $\sum_{i,s} a(i, s)[a(i, s) \geq 0] \leq e - 1 - c$.

There are $\binom{e\#S}{c_-}\binom{c}{c_-}\binom{e\#S-c_-+e-1-c}{e-1-c} \geq N^{\lceil\sqrt{e/3}\rceil} \geq N^{\sqrt{e/3}}$ such functions. I claim that the products $\prod_{i,s}(\zeta^i x - s)^{a(i,s)}$ are distinct in $(k[x]/h)^*$; so there are at least $N^{\sqrt{e/3}}$ such products.

Indeed, assume that a, b are two such functions, and that $\prod_{i,s}(\zeta^i x - s)^{a(i,s)} = \prod_{i,s}(\zeta^i x - s)^{b(i,s)}$ in $(k[x]/h)^*$. Clear denominators to obtain polynomials

$$A = \prod_{i,s}(\zeta^i x - s)^{a(i,s)[a(i,s)\geq 0]-b(i,s)[b(i,s)<0]} \in k[x],$$
$$B = \prod_{i,s}(\zeta^i x - s)^{b(i,s)[b(i,s)\geq 0]-a(i,s)[a(i,s)<0]} \in k[x]$$

with $A = B$ in $k[x]/h$. Now $A(\zeta^j x) = A^{N^j} = B^{N^j} = B(\zeta^j x)$ in $k[x]/h$, since $\zeta^{i+j}x - s = (\zeta^i x - s)^{N^j}$ in $k[x]/(x^e - r)$. Thus $A - B$ has roots $x, \zeta x, \zeta^2 x, \ldots, \zeta^{e-1}x$ in $k[x]/h$; these roots are distinct, since x is invertible in $k[x]/h$; but $A - B$ has degree at most $c + e - 1 - c < e$, so it cannot have e roots unless it is zero. Thus $A = B$ in $k[x]$; so $a(i,s)[a(i,s) \geq 0] - b(i,s)[b(i,s) < 0] = b(i,s)[b(i,s) \geq 0] - a(i,s)[a(i,s) < 0]$ by unique factorization into coprimes; so $a(i,s) = b(i,s)$.

Step 3: Find colliding powers. The nonzero element $r^{(P-1)/e}$ of k has eth power $r^{P-1} = 1$; but the eth roots of 1 in k are exactly the powers of ζ. Thus $r^{(P-1)/e} = \zeta^\ell$ in k for some integer ℓ. Now $x^P = x^{P-1}x = r^{(P-1)/e}x = \zeta^\ell x$ in $k[x]/(x^e - r)$, so $x^{P^j} = \zeta^{j\ell}x$ in $k[x]/(x^e - r)$ for any integer $j \geq 0$. Thus $(x - s)^{N^i P^j} = \zeta^{i+j\ell}x - s$ in $k[x]/(x^e - r)$.

Define L as the set of $(\alpha, \beta) \in \mathbf{Z} \times \mathbf{Z}$ such that e divides $\alpha + (\beta - \alpha)\ell$; then L is a lattice of determinant e. Define C as the set of $(\alpha, \beta) \in \mathbf{R} \times \mathbf{R}$ such that $\max\{|\alpha| \lg(N/P), |\beta| \lg P, |\alpha \lg(N/P) + \beta \lg P|\} \leq \sqrt{e/3} \lg N$; then C is a closed convex symmetric set of area $3(e/3)(\lg N)^2/(\lg P)\lg(N/P) \geq 4e$. By Minkowski's theorem there is a nonzero point $(\alpha, \beta) \in L \cap C$. Assume without loss of generality that $\alpha \geq 0$.

(Variation: $N \neq P^2$ so the area of C is larger than $4e$. Thus one can use a simpler form of Minkowski's theorem, ignoring the fact that C is closed.)

If $\beta \geq 0$ define $u = (N/P)^\alpha P^\beta$ and $v = 1$; then u and v are positive integers bounded by $N^{\sqrt{e/3}}$, and $(x - s)^{uP^\alpha} = (x - s)^{N^\alpha P^\beta} = \zeta^{\alpha+\beta\ell}x - s = \zeta^{\alpha\ell}x - s = (x - s)^{P^\alpha} = (x - s)^{vP^\alpha}$ in $k[x]/(x^e - r)$. If $\beta < 0$ define $u = (N/P)^\alpha$ and $v = P^{-\beta}$; then u and v are positive integers bounded by $N^{\sqrt{e/3}}$, and $(x-s)^{uP^\alpha} = (x-s)^{N^\alpha} = \zeta^\alpha x - s = \zeta^{(\alpha-\beta)\ell}x - s = (x - s)^{P^{\alpha-\beta}} = (x - s)^{vP^\alpha}$ in $k[x]/(x^e - r)$.

Pth powering is invertible on the powers of $x - s$, since $(x - s)^{P^e} = x - s$; so $(x - s)^u = (x - s)^v$ in $k[x]/(x^e - r)$. Consequently $(x - s)^{u-v} = 1$ in $(k[x]/h)^*$. Take N^ith powers: $(\zeta^i x - s)^{u-v} = 1$ in $(k[x]/h)^*$. Thus each of the aforementioned products $\pi = \prod_{i,s}(\zeta^i x - s)^{a(i,s)}$ in $(k[x]/h)^*$ satisfies $\pi^{u-v} = 1$.

On the other hand, if $u \neq v$ then the number of $|u - v|$th roots of 1 in the field $k[x]/h$ is at most $|u - v| < N^{\sqrt{e/3}}$; contradiction. Thus $u = v$; i.e., $N^\alpha = P^{\alpha-\beta}$. If $\alpha = 0$ then $P^{-\beta} = 0$ so $\beta = 0$, but (α, β) was nonzero by construction; contradiction. Hence n is a power of p. □

Notes on the proof. Consider the subgroup G of $(k[x]/(x^e - r))^*$ generated by $\{x - s : s \in S\}$. The proof may be summarized as follows: G is large; G is cyclic; if n is not a power of p, then G has small exponent; contradiction.

There are three different ways to prove that G is cyclic:

- Choose e and r so that $k[x]/(x^e - r)$ is forced to be a field. This is the approach used by Berrizbeitia and Cheng; it is also the reason for their restrictions on e.
- Used in Theorem 2.1: Choose $S = \{1\}$, so that G is cyclic by definition. This is the simplest approach.
- Used in Theorem 3.2: The available equations for $x - s$ imply that G is always isomorphic to its image in $(k[x]/h)^*$. This idea was first published by Macaj in [23]; it was discovered independently by Agrawal.

The original approach of Agrawal, Kayal, and Saxena in [4] was to work instead with the cyclic image of G in $(k[x]/h)^*$ and to force the degree of h to be fairly large. A subsequent improvement by Lenstra was to work with an isomorphic image of G in a product of copies of $(k[x]/h)^*$. Both approaches are quantitatively worse than proving that G is cyclic.

Except in the simple cyclic-by-definition case, all of these approaches rely on the fact that a finite multiplicative subgroup of a field is cyclic. The proof of that fact starts from the observation that a large subgroup cannot have small exponent, and then does some extra work to construct a generator. The extra work is unnecessary for this application: the only reason to prove that G is cyclic is to prove that it does not have small exponent. This simplification was pointed out to me by Kiran Kedlaya (who was preparing to explain [4] to high-school students); it is used in the proof of Theorem 3.2.

The proof of Theorem 3.2 would still work if $n^{d\lceil\sqrt{e/3}\rceil}$ were replaced by $n^{d\sqrt{e/3}}$ in the definition of a certificate. However, this change would complicate certificate testing and would have very little benefit.

4. Finding a certificate: algorithm

Every prime n has a certificate of the form $(d, e, 0, 0, f, r, \{1\})$ with $d \in (\lg n)^{o(1)}$ and $e \in (\lg n)^{2+o(1)}$. Furthermore, this certificate can be found in random time $(\lg n)^{2+o(1)}$. This section discusses the construction of d and e, then the construction of f, and finally the construction of r.

As discussed in Section 6, one can then verify that this is a certificate for n in time $(\lg n)^{4+o(1)}$. As discussed in Section 7, one can reduce the $o(1)$ by choosing certificates more carefully.

Finding d and e. Theorem 5.1 states that there is a positive integer d such that $n^d - 1$ has a divisor $e \geq 6$ between $d^2 \lceil \lg n \rceil^2$ and $(d+1)d^2 \lceil \lg n \rceil^2$. Theorem 5.2 states that the smallest such d is in $\exp(O(\lg \lg \lg n \lg \lg \lg \lg n))$, hence in $(\lg n)^{o(1)}$.

To compute the smallest d, one can try $d = 1$, then $d = 2$, etc.; success will occur within $(\lg n)^{o(1)}$ tries. For each d there are $(\lg n)^{2+o(1)}$ possible divisors e between $d^2 \lceil \lg n \rceil^2$ and $(d+1)d^2 \lceil \lg n \rceil^2$, each e having $(\lg n)^{o(1)}$ bits. One can compute $n^d - 1$ modulo all these e's simultaneously in time $(\lg n)^{2+o(1)}$; see, e.g., [8, Section 18].

Finding f. For every prime number n and positive integer d, there is a monic irreducible polynomial $f \in (\mathbf{Z}/n)[y]$ of degree d.

One standard way to find f is to generate a uniform random monic polynomial f of degree d, see if it is irreducible, and try again if not. There are many choices of f that work: the expected number of trials is approximately d. If d is chosen as above, then the expected number of trials is in $(\lg n)^{o(1)}$.

One standard way to check the irreducibility of a single f is to see whether f has factors in common with $x^n - x, x^{n^2} - x, \ldots, x^{n^{d-1}} - x$. Each nth powering in $(\mathbf{Z}/n)[y]/f$ takes time $(\lg n)^{2+o(1)}$ if $d \in (\lg n)^{o(1)}$, so the total time for an irreducibility test is $(\lg n)^{2+o(1)}$.

I do not mean to suggest that this is the state of the art in constructing irreducible polynomials.

Finding r. For every prime number n, positive integer d, positive integer e dividing $n^d - 1$, and monic irreducible polynomial $f \in (\mathbf{Z}/n)[y]$ of degree d, there is an element r of the field $R = (\mathbf{Z}/n)[y]/f$ such that $r^{(n^d-1)/e}$ has order e; for example, any generator r of R^*. Furthermore, if $e \geq 6$ and $e \geq d^2 \lceil \lg n \rceil^2$, as in Theorems 5.1 and 5.2, then $(d, e, 0, 0, f, r, \{1\})$ is a certificate for n by Theorem 5.3.

One standard way to find r is to generate a uniform random element r of $R - \{0\}$, to see if $r^{(n^d-1)/e}$ has order e, and to try again if not. There are many choices of r that work: the expected number of trials is the product of $q/(q-1)$ for primes q dividing e, which is in $(\lg n)^{o(1)}$ if $e \in (\lg n)^{2+o(1)}$.

One standard way to check whether $r^{(n^d-1)/e}$ has order e is to check that $r^{n^d-1} = 1$ and that $r^{(n^d-1)/q} \neq 1$ for each prime q dividing e. There are only $(\lg n)^{o(1)}$ such primes q if $e \in (\lg n)^{2+o(1)}$, and all of them are easy to find since e is small; the main work is to compute $r^{(n^d-1)/e}$ in the first place, which takes time $(\lg n)^{2+o(1)}$.

As above, I do not mean to suggest that this is the state of the art in finding elements of specified order.

5. Finding a certificate: theorems

Theorem 5.1. *Let n be an integer with $n \geq 2$. Then there exists a positive integer d such that $n^d - 1$ has a divisor $e \geq 6$ with $d^2 \lceil \lg n \rceil^2 \leq e < (d+1)d^2 \lceil \lg n \rceil^2$.*

Proof. Observe that $n^2 - 1 \geq \lceil \lg n \rceil^2$ and $n^6 + n^4 + n^2 + 1 \geq 64$. Thus $e \geq 64 \lceil \lg n \rceil^2$ where $e = n^8 - 1$. Define d as the largest multiple of 8 with $d^2 \leq e/\lceil \lg n \rceil^2$. Then $d \geq 8$, and e divides $n^d - 1$. Finally, $(d^2 - 16)d \geq 64$, so $d^3 + d^2 \geq d^2 + 16d + 64 = (d+8)^2 > e/\lceil \lg n \rceil^2$. □

Theorem 5.2. *There are constants n_0 and α such that, for every prime number $n \geq n_0$, there is a positive integer $d \leq \exp(\alpha \log(3 \log \lg n) \log \log(3 \log \lg n))$ such that $n^d - 1$ has a divisor $e \geq 6$ with $d^2 \lceil \lg n \rceil^2 \leq e < (d+1)d^2 \lceil \lg n \rceil^2$.*

This is a typical application of a well-known theorem of Odlyzko and Pomerance; see [3, Theorem 3]. The point is that the product of the small primes dividing $n^d - 1$ grows, at a minimum, almost exponentially with d. Older theorems suffice for the bound $d \in (\lg n)^{o(1)}$.

It is overkill to assume that n is prime; what matters is that n has no tiny prime divisors.

Proof. Choose α such that d below always exists, and choose $n_0 > 8$ such that H below always exists.

Given $n \geq n_0$, select a real number $H > 16$ such that $H \leq (\lg n)^3$, $D + 1 < n$, and $H/D^2 \geq \lceil \lg n \rceil^2$, where $D = \exp(\alpha \log \log H \log \log \log H)$. Asymptotically one can take H in $(\lg n)^{2+o(1)}$, and thus D in $(\lg n)^{o(1)}$, satisfying $H/D^2 \geq \lceil \lg n \rceil^2$, so

the extra constraints $H \leq (\lg n)^3$ and $D+1 < n$ are automatically satisfied for n large enough. Note that $D \leq \exp(\alpha \log(3 \log \lg n) \log \log(3 \log \lg n))$.

By [3, Theorem 3], there is a positive integer $d \leq D$ such that $H \leq \pi$, where π is the product of the primes q with $q-1$ dividing d.

Now $d^2 \lceil \lg n \rceil^2 \leq D^2 \lceil \lg n \rceil^2 \leq H \leq \pi$. Find the smallest positive integer $e \geq d^2 \lceil \lg n \rceil^2$ dividing π; note that $e \geq 6$ since $n > 8$. Each prime q is at most $d+1$, so e must be smaller than $(d+1)d^2 \lceil \lg n \rceil^2$.

Finally, e divides $n^d - 1$. Indeed, each prime q is at most $d+1 \leq D+1 < n$, so q does not divide n, so q divides $n^{q-1}-1$, hence $n^d - 1$. □

Theorem 5.3. *Let n be a prime number. Let d be a positive integer. Let $e \geq 6$ be a divisor of $n^d - 1$ such that $d^2 \lceil \lg n \rceil^2 \leq e$. Let f be a monic irreducible polynomial in $(\mathbf{Z}/n)[y]$ of degree d. Let r be an element of the ring $(\mathbf{Z}/n)[y]/f$ such that $r^{(n^d-1)/e}$ has order e. Then $(d, e, 0, 0, f, r, \{1\})$ is a certificate for n.*

Proof. Write $R = (\mathbf{Z}/n)[y]/f$. Observe that R is a field.

By hypothesis, n, d, and e are positive integers; e divides $n^d - 1$; $r^{n^d-1} = (r^{(n^d-1)/e})^e = 1$; if q is a prime dividing e, then $r^{(n^d-1)/q} - 1 = (r^{(n^d-1)/e})^{e/q} - 1 \neq 0$, so $r^{(n^d-1)/q} - 1$ is a unit; and $e > 1$, so $r^{(n^d-1)/e} \neq 1$, so $r \neq 1$, so $1^e - r = 1 - r$ is a unit.

Furthermore, $e \geq 6$, so $\binom{2e-1}{e-1} \geq 2^e$ and $e \geq (\sqrt{e/3}+1)^2$. Thus $(\lg \binom{2e-1}{e-1})^2 \geq e^2 \geq (\sqrt{e/3}+1)^2 e \geq (\sqrt{e/3}+1)^2 d^2 \lceil \lg n \rceil^2 \geq \lceil \sqrt{e/3} \rceil^2 d^2 (\lg n)^2$; i.e., $\binom{2e-1}{e-1} \geq n^{d\lceil\sqrt{e/3}\rceil}$.

Finally, $(x-1)^{n^d} = x^{n^d} - 1 = x^{n^d-1}x - 1 = r^{(n^d-1)/e}x - 1$ in $R[x]/(x^e - r)$. □

6. Checking a certificate

This section presents an algorithm that decides whether (d, e, c, c_-, f, r, S) is a certificate for n, given positive integers n, d, e, integers c and c_-, a monic degree-d polynomial $f \in (\mathbf{Z}/n)[y]$, an element r of $R = (\mathbf{Z}/n)[y]/f$, and a subset S of R.

This algorithm takes time $(\lg n)^{4+o(1)}$ for reasonably small inputs. "Reasonably small" means that d is in $(\lg n)^{o(1)}$, $\#S$ is in $(\lg n)^{o(1)}$, e is at most $(\lg n)^{2+o(1)}$, and $e > c \geq c_- \geq 0$. Note that the certificates $(d, e, 0, 0, f, r, \{1\})$ constructed in Section 4, with $d \in (\lg n)^{o(1)}$ and $e \in (\lg n)^{2+o(1)}$, are reasonably small.

The reader is assumed to be familiar with fast multiplication. See, e.g., [8].

The basic conditions. Computing $n^d - 1$, and checking that it is divisible by e, takes time $(\lg n)^{1+o(1)}$. Checking that $e > c \geq c_- \geq 0$ takes time $(\lg n)^{o(1)}$.

Multiplying in $\mathbf{Z}/n$ takes time $(\lg n)^{1+o(1)}$. Thus multiplying in R takes time $(\lg n)^{1+o(1)}$. Computing the $n^d - 1$ power of r in R takes $(\lg n)^{1+o(1)}$ multiplications in R, hence time $(\lg n)^{2+o(1)}$.

The units. There are $(\lg n)^{o(1)}$ primes q dividing e; finding them by trial division takes time $(\lg n)^{1+o(1)}$. Computing the $(n^d-1)/q$ power of r in R takes time $(\lg n)^{2+o(1)}$. Checking whether $r^{(n^d-1)/q} - 1$ is a unit in R takes time $(\lg n)^{1+o(1)}$.

Computing s^e in R for each $s \in S$ takes time $(\lg n)^{1+o(1)}$. Checking all the remaining units takes time $(\lg n)^{1+o(1)}$.

The binomial coefficients. Computing $(e\#S)!$ takes time at most $(\lg n)^{2+o(1)}$, since $e\#S$ is at most $(\lg n)^{2+o(1)}$. Similarly, computing $c_-!$ and $(e\#S-c_-)!$ takes time at most $(\lg n)^{2+o(1)}$. Thus computing the binomial coefficient $\binom{e\#S}{c_-}$ takes time at most $(\lg n)^{2+o(1)}$. Similar comments apply to the other binomial coefficients.

Computing $n^{d\lceil\sqrt{e/3}\rceil}$ takes time $(\lg n)^{2+o(1)}$. Checking whether $n^{d\lceil\sqrt{e/3}\rceil} \le \binom{e\#S}{c_-}\binom{c}{c_-}\binom{e\#S-c_-+e-1-c}{e-1-c}$ takes time $(\lg n)^{2+o(1)}$.

The big exponentiation. Multiplying in $R[x]/(x^e-r)$ takes time $(\lg n)^{3+o(1)}$. Computing each $(x-s)^{n^d}$ in $R[x]/(x^e-r)$ takes $(\lg n)^{1+o(1)}$ multiplications in $R[x]/(x^e-r)$, hence time $(\lg n)^{4+o(1)}$.

7. Optimizations and practical performance

This section looks at verification speed more closely in the important case $d=1$.

Why $d=1$ in practice. A substantial fraction of primes n have suitable divisors e of $n-1$: divisors slightly above the lower bound in Theorem 3.2. What about the other primes n?

A single elliptic-curve-primality-proving step conjecturally takes time $(\lg n)^{3+o(1)}$ to reduce the problem of proving the primality of n to the problem of proving the primality of an auxiliary prime n'. Here n' is slightly shorter than n and "looks random." One can find several choices for n' at similar speed.

Consequently one can—conjecturally—parlay a fast algorithm for a substantial fraction of primes n into a fast algorithm for all primes n. This was suggested by Cheng in [12].

Consider, for example, a prime n that does not have any suitable divisors of $n-1$, but that does have suitable divisors of n^2-1. Here are two ways to prove that n is prime:

- Apply Theorem 3.2 to n with $d=2$.
- Use an elliptic-curve-primality-proving step to locate an auxiliary n' that has suitable divisors of $n'-1$. Apply Theorem 3.2 to n' with $d=1$.

Experiments support the conjecture that the savings in moving from $d=2$ to $d=1$ is far above the cost of locating a suitable n'.

Of course, the same argument might mean that even the $d=1$ case is of no practical interest: perhaps Theorem 3.2 will never be faster than a series of elliptic-curve-primality-proving steps. On the other hand, perhaps future improvements to Theorem 3.2 will make certificate verification so fast that the $d=2$ case is worth considering again.

Choosing c and c_-. The choice $c=c_-=0$ in Section 4 is far from optimal. If $c\approx\alpha e$ and $c_-\approx\beta e$ then the product $\binom{e\#S}{c_-}\binom{c}{c_-}\binom{e\#S-c_-+e-1-c}{e-1-c}$ is approximately $\exp(e\gamma)$ where $\gamma=(\#S-\beta+1-\alpha)\log(\#S-\beta+1-\alpha)+\#S\log\#S+\alpha\log\alpha-2(\#S-\beta)\log(\#S-\beta)-2\beta\log\beta-(\alpha-\beta)\log(\alpha-\beta)-(1-\alpha)\log(1-\alpha)$.

One can, either with a computer program or by hand, easily find α and β that maximize γ for any given $\#S$. Any choice of $c\approx\alpha e$ and $c_-\approx\beta e$ is reasonable; a small amount of additional searching will locate the optimal c and c_-.

It turns out that the optimal α and β have simple expressions: $\alpha=1/2$ and $\beta=(\#S+1-\sqrt{\#S^2+1})/2$. For example, say $\#S=1$. The product of binomial

coefficients is about $5.828427\ldots^e$ if one takes $c \approx e/2$ and $c_- \approx (2-\sqrt{2})e/2 = (0.2928932\ldots)e$. For comparison: The product of binomial coefficients is about 4^e if one takes $c = 0$ and $c_- = 0$.

Choosing e and $\#S$. Say there are many possibilities for $(e, \#S)$—or, in the elliptic-curve context, many possibilities for an auxiliary $(n, e, \#S)$—such that the maximized product of binomial coefficients exceeds $n^{\lceil\sqrt{e/3}\rceil}$. One should choose the possibility that minimizes verification time.

As a first approximation, this means minimizing $e\#S$: verification time can be crudely modeled as $(\lg n)^2 e\#S$. The following table shows $e\#S/(\lg n)^2$ as a function of $e/(\lg n)^2$, when $\#S$ is chosen as small as possible:

$\#S$	works for $e/(\lg n)^2$ between about	and about	so $e\#S/(\lg n)^2$ is between about	and about
1	$0.051540\ldots$	∞	$0.051540\ldots$	∞
2	$0.027664\ldots$	$0.051540\ldots$	$0.055328\ldots$	$0.103081\ldots$
3	$0.020415\ldots$	$0.027664\ldots$	$0.061247\ldots$	$0.082992\ldots$
4	$0.016832\ldots$	$0.020415\ldots$	$0.067328\ldots$	$0.081663\ldots$
5	$0.014653\ldots$	$0.016832\ldots$	$0.073269\ldots$	$0.084160\ldots$
6	$0.013169\ldots$	$0.014653\ldots$	$0.079017\ldots$	$0.087923\ldots$
7	$0.012082\ldots$	$0.013169\ldots$	$0.084575\ldots$	$0.092187\ldots$
8	$0.011244\ldots$	$0.012082\ldots$	$0.089958\ldots$	$0.096658\ldots$

If e drops substantially below $0.01(\lg n)^2$, then $e\#S$ explodes: $\#S = 100$ works for $e/(\lg n)^2$ down to about $0.004037\ldots$; $\#S = 1000$ works for $e/(\lg n)^2$ down to about $0.002164\ldots$; $\#S = 10000$ works for $e/(\lg n)^2$ down to about $0.001347\ldots$; and so on.

A more precise model of verification time includes logarithmic factors that grow with e but not with $\#S$. Reducing e at the expense of $\#S$ often saves time even if it increases $e\#S$.

Multiplying quickly. One can square an element of $(\mathbf{Z}/n)[x]/(x^e - r)$ as follows:

- Lift to $\mathbf{Z}[x]$, obtaining polynomials of degree at most $e-1$ with coefficients in $\{0, 1, \ldots, n-1\}$.
- Choose p so that $2^p > en^2$, and map to $\mathbf{Z}[x]/(x - 2^p) \cong \mathbf{Z}$, obtaining an integer with approximately $2e \lg n$ bits. If integers are represented in the usual form, then this operation—polynomial evaluation at 2^p—is a simple matter of copying bytes.
- Square in $\mathbf{Z}$.
- Recover the square in $\mathbf{Z}[x]$. This is a simple matter of copying bytes.
- Reduce modulo $x^e - r$. This is particularly easy if r is small.
- Reduce each coefficient modulo n.

In the C programming language, for example, one can represent an element of $(\mathbf{Z}/n)[x]/(x^e - r)$ as an array of e integers in $\{0, 1, \ldots, n-1\}$, each integer being represented in turn as an `mpz_t` variable using Granlund's GMP 4.1.3 library in [19]. Specifically, the `mpz_t` variables `poly[0]`, `poly[1]`, ..., `poly[`$e-1$`]` represent the polynomial `poly[0]` + `poly[1]`$x + \cdots +$ `poly[`$e-1$`]`x^{e-1} in $(\mathbf{Z}/n)[x]/(x^e - r)$. The following function then replaces `poly` with its square in $(\mathbf{Z}/n)[x]/(x^e - r)$, using GMP's `mpz_import` and `mpz_export` functions to copy bytes:

```
void square(mpz_t *poly,int e,mpz_t n,int r)
{
  char *s;
  size_t unused;
  mpz_t t;
  int p8;
  int j;

  mpz_init(t);
  mpz_mul(t,n,n);
  mpz_mul_ui(t,t,e); /* t = en^2 */
  p8 = mpz_sizeinbase(t,256); /* en^2 fits into p8 bytes */

  s = calloc(e,p8); if (!s) exit(111);
  /* GMP convention: if no memory, exit */
  for (j = 0;j < e;++j)
    mpz_export(s + j * p8,&unused,-1,1,0,0,poly[j]);
    /* warning: overflows buffer if poly[j] is too big */
  mpz_import(t,e * p8,-1,1,0,0,s); /* t = poly(256^p8) */
  free(s);

  mpz_mul(t,t,t); /* t = poly^2(256^p8) */

  s = calloc(2 * e,p8); if (!s) exit(111);
  mpz_export(s,&unused,-1,1,0,0,t);
  for (j = 0;j < e;++j) {
    mpz_import(poly[j],p8,-1,1,0,0,s + (j + e) * p8);
    mpz_mul_si(poly[j],poly[j],r);
    mpz_import(t,p8,-1,1,0,0,s + j * p8);
    mpz_add(poly[j],t,poly[j]);
    mpz_mod(poly[j],poly[j],n);
  }
  free(s);

  mpz_clear(t);
}
```

A straightforward certificate-verification program, using this `square` function, takes $1.18 \cdot 10^{11}$ clock cycles on a 1300MHz Pentium M to verify the aforementioned certificate $(1, 2430, 1214, 928, y, 2, \{1, 2\})$ for the prime $\lfloor 10^{84} \exp 1 \rfloor$. About 88% of that time is spent in GMP's `mpz_mul` function for integer squaring, so there is little benefit in attempting to streamline any other operations.

Future improvements. It would be surprising for the group discussed in Section 3 to have size much smaller than p^e. Theorem 2.1 uses a lower bound near 2^e. Theorem 3.2 uses a lower bound near $5.828427\ldots^e$. Can one do better?

An improvement from 2^e to $2^{\gamma e}$ means that the lower bound on e drops by a factor of about γ^2. A lower bound close to p^e, with $\#S$ small, would reduce the lower bound on e to about $(4 \lg(n/p))/3 \lg p < (4/3) \lg n$, saving a factor of $(\lg n)^{1+o(1)}$. Mihăilescu has pointed out that in this case one can further reduce the

lower bound on e by using unit-group factors to quickly increase the lower bound on primes p dividing n.

Voloch in [27] suggested considering products in $\mathbf{F}_p[x]$ of degree somewhat larger than e, and applying the ABC theorem. I showed in [7] that, if $\#S = 1$ and n does not have any tiny factors, then four distinct products of degree at most $1.1e$ cannot be congruent modulo $x^e - r$, so the group size is at least $\frac{1}{3}\binom{\lfloor 2.1e \rfloor}{e} \approx 4.2768947738\ldots^e$. Perhaps there are better results along these lines. (I suspect that any such results will work with multiple derivatives directly, rather than applying the ABC theorem.) If 1000 distinct products of degree at most $2e$ cannot be congruent modulo $x^e - r$, then the group size is at least $\frac{1}{999}\binom{3e}{e} \approx 6.75^e$.

Cheng and Wan in [14, Section 4] pointed out a connection between the group size and list decoding of Reed-Solomon codes. Cheng in [13, Theorem 3] used this connection to prove that the group size is at least about 5.17736^e. Perhaps there are better results along these lines.

8. Smaller powers

Theorem 8.1. *Let n, d, and e be positive integers such that $2^e - 1 \geq n^{2\lfloor\sqrt{de}\rfloor}$ and e divides $n^d - 1$. Let f be a monic polynomial in $(\mathbf{Z}/n)[y]$ of degree d. Define R as the ring $(\mathbf{Z}/n)[y]/f$. Assume that $y^{n^d-1} = 1$ in R, that $y^{(n^d-1)/q} - 1$ is a unit in R for each prime q dividing e, that $y - 1$ is a unit in R, and that $f(z) = (z-y)(z-y^n)(z-y^{n^2})\cdots(z-y^{n^{d-1}})$ in the ring $R[z]$. If $(x-1)^n = x^n - 1$ in the ring $R[x]/(x^e - y)$ then n is a power of a prime.*

The big difference between Theorem 2.1 and Theorem 8.1 is that Theorem 2.1 uses n^dth powers in $R[x]/(x^e - y)$ while Theorem 8.1 uses nth powers. This requires extra effort in the proof, and requires some extra (easily tested) hypotheses on f, but it has two advantages: first, checking nth powers is about d times faster than checking n^dth powers; second, e can be chosen about d times smaller, because $n^{2d\lfloor\sqrt{e}\rfloor}$ has been replaced with $n^{2\lfloor\sqrt{de}\rfloor}$.

The impact of this change depends on the reader's perspective:

- For readers who care about proven performance and—stepping beyond the scope of this paper—want a more precise bound than $(\lg n)^{4+o(1)}$: Theorem 8.1 provides better speed than Theorem 2.1. An analogous modification in Theorem 3.2 would save even more time.
- For readers who care about proven performance at the level of detail of $(\lg n)^{4+o(1)}$: The speedup is not visible; there is no improvement in the limiting exponent 4.
- For readers who care about practical performance: There is no speedup; d is always 1 in practice, as discussed in Section 7.

Berrizbeitia's theorem in [11] is the case $d \in \{1, 2\}$ and $e \in \{1, 2, 4, 8, \ldots\}$ of Theorem 8.1 rather than Theorem 2.1. Mihăilescu and Avanzi in [24, Theorem 6] pointed out that nth powers could be used for arbitrary values of e.

Proof. If $n = 1$ then n is a power of a prime, so assume that $n > 1$.

Step 1: Move to a field. Find a prime p dividing n. Find an irreducible polynomial g in $\mathbf{F}_p[y]$ dividing the image of f. Then $k = \mathbf{F}_p[y]/g$ is a field. Define ζ as the image of $y^{(n^d-1)/e}$ in k.

Step 2: Combinatorially enumerate many powers of $x-1$. Consider the ring morphism $h \mapsto h(y^n)$ from $(\mathbf{Z}/n)[y]$ to R. This morphism takes f to $f(y^n) = 0$, so it induces a ring morphism $\sigma : R \to R$. Note that $\sigma(y) = y^n$.

By hypothesis $(x-1)^n - (x^n-1)$ is a multiple of $x^e - y$ in $R[x]$. Apply σ to all coefficients: $(x-1)^n - (x^n-1)$ is a multiple of $x^e - \sigma(y) = x^e - y^n$. Repeat to see that $(x-1)^n - (x^n-1)$ is a multiple of $x^e - y^{n^i}$ for each $i \geq 0$. Substitute x^{n^i} for x: $(x^{n^i}-1)^n - (x^{n^{i+1}}-1)$ is a multiple of $x^{n^i e} - y^{n^i}$, hence a multiple of $x^e - y$.

Consequently, $(x-1)^{n^i} = x^{n^i} - 1$ in $R[x]/(x^e - y)$, hence in the ring $A = k[x]/(x^e - y)$. In particular, $(x-1)^{n^d} = x^{n^d} - 1 = y^{(n^d-1)/e}x - 1 = \zeta x - 1$ in A. Substitute $\zeta^i x$ for x: $(\zeta^i x - 1)^{n^d} = \zeta^{i+1}x - 1$ in $k[x]/((\zeta^i x)^e - y) = A$. Thus $(x-1)^{n^{id}} = \zeta^i x - 1$ in A.

There are $2^e - 1$ vectors $(a_0, a_1, \ldots, a_{e-1}) \in \{0,1\}^e$ such that $\sum_i a_i \leq e-1$. Any product $\prod_i (\zeta^i x - 1)^{a_i} = (x-1)^{\sum_i n^{id} a_i}$ is a power of $x-1$ in A. These products are distinct, so there are at least $2^e - 1$ powers of $x-1$ in A.

Step 3: Find colliding powers of $x-1$. By hypothesis, in k, the product $(y^p - y)(y^p - y^n) \cdots (y^p - y^{n^{d-1}})$ equals $f(y^p)$, so it is a multiple of $g(y^p) = g(y)^p = 0$, so it is equal to 0. Hence $y^p = y^{n^\ell}$ in k for some $\ell \in \{0, 1, \ldots, d-1\}$. By induction, $y^{p^j} = y^{n^{j\ell}}$ in k for any integer $j \geq 0$.

I claim that, over all integers $i \geq 0$ and $j \geq 0$, there are at most de possibilities for $x^{n^i p^j}$ in A. Indeed, $x^{n^i p^j}$ has eth power $y^{n^i p^j} = y^{n^{i+j\ell}}$ in A. By hypothesis $y^{n^d} = y$, so $y^{n^{i+j\ell}} \in \{y, y^n, y^{n^2}, \ldots, y^{n^{d-1}}\}$. There are at most e powers of x whose eth power is y, at most e powers of x whose eth power is y^n, etc.

Hence there are at most de possibilities for $x^{n^i p^j} - 1 = (x^{n^i} - 1)^{p^j} = (x-1)^{n^i p^j}$ in A.

Consider the pairs (i, j) with $0 \leq i \leq \lfloor \sqrt{de} \rfloor$ and $0 \leq j \leq \lfloor \sqrt{de} \rfloor$. There are $(\lfloor \sqrt{de} \rfloor + 1)^2 > de$ pairs (i, j), so there are distinct pairs $(i, j), (i', j')$ with $(x-1)^u = (x-1)^v$, where $u = n^i p^j$ and $v = n^{i'} p^{j'}$. Note that u and v are positive integers bounded by $n^{2\lfloor \sqrt{de} \rfloor}$.

The remainder $(x^e - y) \bmod (x-1) = 1 - y$ is a unit in k, so $x-1$ is a unit in A. Thus $(x-1)^{u-v} = 1$ in the unit group A^*. If $u \neq v$ then there are at most $|u-v|$ powers of $x-1$, but $|u-v| < n^{2\lfloor \sqrt{de} \rfloor} \leq 2^e - 1$; contradiction.

Hence $u = v$; i.e., $n^{i-i'} = p^{j'-j}$. If $i = i'$ then $j' = j$; contradiction. Thus a nontrivial power of n is a power of p; so n is a power of p. □

References

[1] Leonard M. Adleman, Ming-Deh A. Huang, *Proceedings of the 18th annual ACM symposium on theory of computing*, Association for Computing Machinery, New York, 1986. ISBN 0-89791-193-8. MR0990047 (89j:69001)

[2] ———, *Primality testing and abelian varieties over finite fields*, Lecture Notes in Mathematics, 1512, Springer-Verlag, Berlin, 1992. ISBN 3-540-55308-8. MR1176511 (93g:11128)

[3] Leonard M. Adleman, Carl Pomerance, Robert S. Rumely *On distinguishing prime numbers from composite numbers*, Annals of Mathematics **117** (1983), 173–206. ISSN 0003-486X. MR0683806 (84e:10008)

[4] Manindra Agrawal, Neeraj Kayal, Nitin Saxena, *PRIMES is in P* (2002). URL: `http://www.cse.iitk.ac.in/news/primality.html`.

[5] A. O. L. Atkin, Francois Morain, *Finding suitable curves for the elliptic curve method of factorization*, Mathematics of Computation **60** (1993), 399–405. ISSN 0025-5718. MR1140645 (93k:11115)

[6] Daniel J. Bernstein, *Detecting perfect powers in essentially linear time*, Mathematics of Computation **67** (1998), 1253–1283. ISSN 0025-5718. URL: `http://cr.yp.to/papers.html`. MR1464141 (98j:11121)

[7] Daniel J. Bernstein, *Sharper ABC-based bounds for congruent polynomials*, to appear, Journal de Théorie des Nombres de Bordeaux. ISSN 1246-7405. URL: `http://cr.yp.to/papers.html#abccong`. `ID 1d9e079cee20138de8e119a99044baa3`.

[8] Daniel J. Bernstein, *Fast multiplication and its applications*, to appear in Buhler-Stevenhagen *Algorithmic number theory* book. URL: `http://cr.yp.to/papers.html#multapps`. `ID 8758803e61822d485d54251b27b1a20d`.

[9] Daniel J. Bernstein, *Distinguishing prime numbers from composite numbers: the state of the art in* 2004, submitted. URL: `http://cr.yp.to/papers.html#prime2004`. `ID d72f09ae5b05f41a53e2237c53f5f276`.

[10] Daniel J. Bernstein, Hendrik W. Lenstra, Jr., Jonathan Pila, *Detecting perfect powers by factoring into coprimes*. URL: `http://cr.yp.to/papers.html#powers2`. `ID bbd41ce71e527d3c06295aadccf60979`.

[11] Pedro Berrizbeitia, *Sharpening PRIMES is in P for a large family of numbers*, (2002). URL: `http://arxiv.org/abs/math.NT/0211334`.

[12] Qi Cheng, *Primality proving via one round in ECPP and one iteration in AKS*, (2003). URL: `http://www.cs.ou.edu/~qcheng/pub.html`.

[13] Qi Cheng, *On the bounded sum-of-digits discrete logarithm problem in finite fields*, (2004). URL: `http://www.cs.ou.edu/~qcheng/pub.html`.

[14] Qi Cheng, Daqing Wan, *On the list and bounded distance decodibility of Reed-Solomon codes*, (extended abstract), (2004). URL: `http://www.cs.ou.edu/~qcheng/pub.html`.

[15] Michael R. Fellows, Neal Koblitz *Self-witnessing polynomial-time complexity and prime factorization*, Designs, Codes and Cryptography **2** (1992), 231–235. ISSN 0925–1022. URL: `http://cr.yp.to/bib/entries.html#1992/fellows`. MR1181730 (93e:68032)

[16] Shafi Goldwasser, Joe Kilian, *Almost all primes can be quickly certified*, in [1], (1986), 316–329; see also newer version [17].

[17] Shafi Goldwasser, Joe Kilian, *Primality testing using elliptic curves*, Journal of the ACM **46** (1999), 450–472; see also older version [16]. ISSN 0004-5411. MR1812127 (2002e:11182)

[18] Ronald L. Graham, Jaroslav Nešetřil, *The mathematics of Paul Erdős. I*, Algorithms and Combinatorics, **13**, Springer-Verlag, Berlin, 1997. ISBN 3-540-61032-4. MR1425172 (97f:00032)

[19] Torbjorn Granlund, *GMP* 4.1.3 : *GNU multiple precision arithmetic library*, (2004). URL: `http://www.swox.com/gmp/`.

[20] Sergei Konyagin, Carl Pomerance, *On primes recognizable in deterministic polynomial time*, in [18] (1997), 176–198. URL: `http://cr.yp.to/bib/entries.html#1997/konyagin`. MR1425185 (98a:11184)

[21] Arjen K. Lenstra, Hendrik W. Lenstra, Jr., *Algorithms in number theory*, in [26] (1990), 673–715. URL: `http://cr.yp.to/bib/entries.html#1990/lenstra-survey`. MR1127178

[22] Hendrik W. Lenstra, Jr., *Galois theory and primality testing*, in [25] (1985), 169–189. MR0812498 (87g:11171)

[23] Martin Macaj, *Some remarks and questions about the AKS algorithm and related conjecture*, (2002). URL: `http://thales.doa.fmph.uniba.sk/macaj/aksremarks.pdf`.

[24] Preda Mihăilescu, Roberto M. Avanzi, *Efficient "quasi"-deterministic primality test improving AKS*, URL: `http://www-math.uni-paderborn.de/~preda/`.

[25] I. Reiner, K. W. Roggenkamp (editors), *Orders and their applications: proceedings of the conference held in Oberwolfach, June* 3–9, 1984, Lecture Notes in Mathematics, 1142, Springer-Verlag, Berlin, 1985. ISBN 3-540-15674-7. MR0812486 (86g:16003)

[26] Jan van Leeuwen (editor), *Handbook of theoretical computer science, volume A: algorithms and complexity*, Elsevier, Amsterdam, 1990. ISBN 0-444-88071-2. MR1127166 (92d:68001)

[27] José Felipe Voloch, *On some subgroups of the multiplicative group of finite ring*, to appear, Journal de Théorie des Nombres de Bordeaux. ISSN 1246-7405. URL: `http://www.ma.utexas.edu/users/voloch/preprint.html`.

Department of Mathematics, Statistics, and Computer Science (M/C 249), The University of Illinois at Chicago, Chicago, Illinois 60607–7045

E-mail address: `djb@cr.yp.to`