Identifying primes in polynomial time: the AKS algorithm

Andrew Putman

Abstract

We prove a remarkable theorem Agrawal–Kayal–Saxena saying that there is an algorithm to determine if an integer is prime whose running time is a polynomial in the number of digits of the integer.

1 Introduction

Given an integer $n \geq 2$, the most naive way to determine if n is prime is to test whether or not n is divisible by m for all $2 \leq m \leq \lfloor \sqrt{n} \rfloor$. This has computational complexity $\sqrt{n}$; however, it is natural to want a algorithm whose computational complexity is not a polynomial in n itself, but rather in the number of digits of n, i.e. a polynomial in $\log n$. A remarkable theorem of Agrawal–Kayal–Saxena says that such an algorithm exists:

Theorem 1.1 ([1, 2])**.** *There exists an algorithm for determining if an integer $n \geq 2$ is prime with computational complexity $\log^C n$ for some constant C.*

The goal of this note is to prove Theorem 1.1. Remarkably, the proof is almost entirely elementary.

Remark 1.2. Rather than summarizing the history leading up to Theorem 1.1, we will simply refer the reader to the introduction of [1] or to the detailed survey [3].

The outline of this note is as follows. In §2, we state Theorem 2.2, which is the characterization of prime numbers underlying Theorem 1.1. We then show how to derive Theorem 1.1 from Theorem 2.2. The proof of Theorem 2.2 requires a weak form of the prime number theorem that we prove in §3. In §4, we extract from our weak form of the prime number theorem the precise statement we need. Finally, in §5, we reduce Theorem 2.2 to a statement about the arithmetic of $\overline{\mathbb{F}}_p$ and in §6 we prove this arithmetic statement.

Throughout these notes, $\log(n)$ denotes the base-2 logarithm of n and $\phi(r)$ denotes the Euler totient function.

Acknowledgments. I would like to thank Peter Shalen for some helpful corrections.

2 The algorithm

If n is prime, then in $(\mathbb{Z}/n)[x]$ it is easy to see that $(x+a)^n \equiv x^n + a$ for all $a \in \mathbb{Z}/n$. The converse to this is also true; indeed, something even stronger holds:

Lemma 2.1. *Consider some $n \geq 2$. Assume that for some $a \in \mathbb{Z}$ with $\gcd(a, n) = 1$ we have $(x + a)^n \equiv x^n + a$ in $(\mathbb{Z}/n)[x]$. Then n is prime.*

Proof. Assume that n is not prime, and let p be a prime dividing n. If p^m is the maximal power of p dividing n, then

$$\binom{n}{p} = \frac{n \cdot (n-1) \cdots (n-p+1)}{p \cdot (p-1) \cdots (1)}$$

is not divisible by p^m, so the term $\binom{n}{p} x^p a^{n-p}$ of the binomial expansion of $(x + a)^n$ does not vanish in $(\mathbb{Z}/n)[x]$. □

To test if some $n \geq 2$ is prime, therefore, it would be enough to prove that $(x+1)^n \equiv x^n + 1$ in $(\mathbb{Z}/n)[x]$. Unfortunately, $(x + 1)^n$ has $(n + 1)$ terms, so this algorithm would have computational complexity a polynomial in n. The main theorem underlying Theorem 1.1 avoids this issue by expanding binomials not in $(\mathbb{Z}/n)[x]$, but rather in $(\mathbb{Z}/n)[x]/(x^r - 1)$ for some appropriate r:

Theorem 2.2. *Consider some $n \geq 2$. For all $1 \leq r \leq \lceil \log^5 n \rceil$ and $1 \leq a \leq \lfloor \sqrt{\phi(r)} \log n \rfloor$, assume that the following hold:*

- $\gcd(a, n) = 1$, *and*
- *In $(\mathbb{Z}/n)[x]/(x^r - 1)$, we have $(x + a)^n \equiv x^n + a$.*

Then n is a power of a prime.

The remainder of this section shows how to prove Theorem 1.1 from Theorem 2.2. The following sections then prove Theorem 2.2.

We start by observing that the hypotheses of Theorem 2.2 can only be satisfied for fairly large n. Indeed, since $\gcd(n, n) \neq 1$, for the hypotheses of Theorem 2.2 to hold we must have that $a = n$ does not satisfy the inequalities in its statement. The following lemma shows that assuming that $n \geq 10000$ ensures this:

Lemma 2.3. *Consider some $n \geq 10000$. Then for all $1 \leq r \leq \lceil \log^5 n \rceil$ and $1 \leq a \leq \lfloor \sqrt{\phi(r)} \log n \rfloor$, we have $a < n$.*

Proof. Using the trivial inequality $\phi(r) \leq r - 1$, we have

$$a \leq \sqrt{\phi(\log^5(n) + 1)} \log n \leq \sqrt{\log^5 n} \log n = \log^{3.5} n.$$

It is an easy exercise to see that $\log^{3.5} n < n$ for $n \geq 10000$. □

We now prove Theorem 1.1.

Proof of Theorem 1.1, assuming Theorem 2.2. Consider some integer $n \geq 2$. We determine if n is prime via the following algorithm:

1. If $n < 10000$, then do it by brute force (which takes constant time).
2. Next, test if n is a proper power by computing $\lfloor n^{1/k} \rfloor$ for $1 < k \leq \log n$. If it is, then n is not prime. This takes $\log^{O(1)} n$ time.
3. Next, test if n satisfies the hypotheses of Theorem 2.2. By expanding out $(x+a)^n$ in $(\mathbb{Z}/n)[x]/(x^r - 1)$ via repeated squaring, this can be done in in $\log^{O(1)} n$ time. If n satisfies the hypotheses of Theorem 2.2, then since n is not a proper power the theorem implies that n is prime.
4. If n does not satisfy the hypotheses of Theorem 2.2, then we can conclude that it is composite as follows:
   - If $\gcd(n, a) \neq 1$ for some a as in the theorem, then since $n \geq 10000$ we know by Lemma 2.3 that $a < n$ and thus that n is not prime.
   - If for some a as in the theorem we have $(x+a)^n \not\equiv x^n + a$ in $(\mathbb{Z}/n)[x]/(x^r - 1)$, then n is composite; indeed, for n prime we have $(x+a)^n \equiv x^n + a$ in $(\mathbb{Z}/n)[x]$ for all a. □

3 The lower bound in Chebyshev's inequality

The deepest fact that goes into the proof of Theorem 2.2 is a weak form of the prime number theorem. Letting $\pi(n)$ be the number of primes less than n, the prime number theorem says that $\pi(n)$ is asymptotic to $n/\ln(n)$. Much easier than this is Chebyshev's inequality, which states that for some positive constants $C < D$ we have

$$C \cdot \frac{n}{\ln n} \leq \pi(n) \leq D \cdot \frac{n}{\ln n}. \tag{3.1}$$

The result we need is the key step in a remarkably simple proof of the lower bound of (3.1) due to Nair [4]:

Theorem 3.1. *Let $d_n = \operatorname{lcm}\{1, \ldots, n\}$. Then for $n \geq 7$ we have $d_n \geq 2^n$.*

Proof. We have

$$d_7 = 2^2 \cdot 3 \cdot 5 \cdot 7 = 420 \geq 2^7 = 128 \quad \text{and} \quad d_8 = 2^3 \cdot 3 \cdot 5 \cdot 7 = 840 \geq 2^8 = 256.$$

It is thus enough to prove the theorem for $n \geq 9$.

For $1 \leq m \leq n$, define

$$I_{m,n} = \int_0^1 x^{m-1}(1-x)^{n-m} dx.$$

We will calculate this in two different ways.

The first way we will calculate $I_{m,n}$ is to directly expand out the binomial and integrate:

$$\begin{aligned} I_{m,n} &= \int_0^1 x^{m-1}(1-x)^{n-m}dx = \int_0^1 x^{m-1}\sum_{k=0}^{n-m}(-1)^k\binom{n-m}{k}x^k dx \\ &= \sum_{k=0}^{n-m}(-1)^k\binom{n-m}{k}\int_0^1 x^{m+k-1}dx \\ &= \sum_{k=0}^{n-m}(-1)^k\binom{n-m}{k}\frac{1}{m+k}. \end{aligned}$$

Each of these denominators divides $d_n = \operatorname{lcm}\{1,\ldots,n\}$, so we see that $I_{m,n}d_n$ is a positive integer.

The second way is to inductively calculate to obtain an exact formula. The base of the induction is

$$I_{1,n} = \int_0^1 (1-x)^{n-1}dx = \frac{1}{n}.$$

For $2 \le m \le n$, we can use integration by parts to see that

$$I_{m,n} = \int_0^1 x^{m-1}(1-x)^{n-m}dx = \frac{m-1}{n-m+1}\int_0^1 x^{m-2}(1-x)^{n-m+1}dx = \frac{m-1}{n-(m-1)}I_{m-1,n}.$$

Combining these two facts, we see that

$$\begin{aligned} I_{m,n} &= \frac{m-1}{n-(m-1)}I_{m-1,n} = \frac{(m-1)(m-2)}{(n-(m-1))(n-(m-2))}I_{m-2,n} \\ &= \cdots = \frac{(m-1)(m-2)\cdots(1)}{(n-(m-1))(n-(m-2))\cdots(n-1)}I_{1,n} \\ &= \frac{(m-1)(m-2)\cdots(1)}{(n-(m-1))(n-(m-2))\cdots(n-1)(n)} \\ &= \frac{1}{m}\frac{m!(n-m)!}{n!} = \frac{1}{m\binom{n}{m}}. \end{aligned}$$

Since $I_{m,n}d_n$ is a positive integer, we deduce that $m\binom{n}{m}$ divides d_n for all $1 \le m \le n$.

In particular, $n\binom{2n}{n}$ divides d_{2n} and

$$(n+1)\binom{2n+1}{n+1} = (2n+1)\binom{2n}{n}$$

divides d_{2n+1}. Since d_{2n} divides d_{2n+1} and $\gcd(n, 2n+1) = 1$, these two facts imply that $n(2n+1)\binom{2n}{n}$ divides d_{2n+1}. In particular,

$$\begin{aligned} d_{2n+1} &\ge n(2n+1)\binom{2n}{n} \\ &\ge n\sum_{k=0}^{2n}\binom{2n}{k} \\ &= n\cdot(1+1)^{2n} = n\cdot 2^{2n}. \end{aligned}$$

We conclude that for $n \geq 4$ we have

$$d_{2n+2} \geq d_{2n+1} \geq n2^{2n} \geq 4 \cdot 2^{2n} = 2^{2n+2}.$$

Thus $d_n \geq 2^n$ for $n \geq 9$, as desired. □

Though we will not need it, at this point we might as well prove the lower bound in Chebyshev's inequality:

Theorem 3.2. *For $n \geq 4$, we have $\pi(n) \geq \frac{n}{\log n}$.*

Proof. The theorem is easily checked for $4 \leq n \leq 6$, so we only need to prove it for $n \geq 7$. As in Theorem 3.1, let $d_n = \operatorname{lcm}\{1, \ldots, n\}$. The primes that divide d_n are precisely the primes $p \leq n$, and the exponent of p is the largest m such that $p^m \leq n$, i.e. $m = \lfloor \log(n)/\log(p) \rfloor$. From this, we see that

$$d_n \leq \prod_{p \leq n} p^{\log(n)/\log(p)}.$$

Theorem 3.1 says that $d_n \geq 2^n$ for $n \geq 7$, so for these values of n we see that

$$2^n \leq \prod_{p \leq n} p^{\log(n)/\log(p)}.$$

Taking the base-2 logarithms of both sides of this, we see that

$$n \leq \sum_{p \leq n} \frac{\log n}{\log p} \cdot \log p = \log(n)\pi(n).$$

Dividing both sides by $\log n$ now gives the desired inequality. □

4 Making n have high order modulo r

Our main use of Theorem 3.1 will be to prove the following lemma, which will play a key role in the proof of Theorem 2.2. If $\gcd(n, r) = 1$, then write $\operatorname{ord}_r(n)$ for the order of n in $\mathbb{Z}/r$.

Lemma 4.1. *For $n \geq 3$, there exists some $r \leq \lceil \log^5 n \rceil$ such that $\gcd(n, r) = 1$ and $\operatorname{ord}_r(n) > \log^2 n$.*

Proof. Set $B = \lceil \log^5 n \rceil$, and let r be the smallest integer that does not divide the integer

$$A = n^{\lfloor \log B \rfloor} \cdot \prod_{k=1}^{\lfloor \log^2 n \rfloor} (n^k - 1).$$

We first prove that $r \leq B$. Observe that

$$A < n^{\log B} \cdot \prod_{k=1}^{\lfloor \log^2 n \rfloor} n^k = n^{\log B + \sum_{k=1}^{\lfloor \log^2 n \rfloor} k}$$
$$\leq n^{\log B + \frac{1}{2}\log^2(n)\cdot(\log^2(n)-1)} \leq n^{\log^4 n} = 2^{\log^5 n} \leq 2^B.$$

If every element of $\{1, \ldots, B\}$ divided A, then their least common multiple d_B would also divide A. Since $n \geq 3$, we have $B = \lceil \log^5 n \rceil \geq 11$ and thus Theorem 3.1 says that $d_B \geq 2^B$, so this would contradict the above (strict) upper bound on A. We conclude that some element of $\{1, \ldots, B\}$ does not divide A, so $r \leq B$, as claimed.

We now prove that $\gcd(n, r) = 1$. Assume for the sake of contradiction that $\gcd(n, r) > 1$. Write $r = st$ where s is a product of prime powers for primes that divide n and t is a product of prime powers for primes that do not divide n. Since $\gcd(n, r) > 1$, it follows that $s > 1$, so $t < r$. Since r is the smallest integer that does not divide

$$A = n^{\lfloor \log B \rfloor} \cdot \prod_{k=1}^{\lfloor \log^2 n \rfloor} (n^k - 1),$$

it follows that t divides A. Since t is coprime to n, this implies that

$$t \text{ divides } \prod_{k=1}^{\lfloor \log^2 n \rfloor} (n^k - 1). \tag{4.1}$$

If p^k is a prime power dividing s, then since $p^k \leq r \leq B$ we must have

$$k \leq \lfloor \log(B)/\log(p) \rfloor \leq \lfloor \log(B) \rfloor.$$

This implies that p^k divides $n^{\lfloor \log B \rfloor}$. Since this is true for all prime powers dividing s, we deduce that

$$s \text{ divides } n^{\lfloor \log B \rfloor}. \tag{4.2}$$

Combining (4.1) and (4.2), we deduce that

$$r = st \text{ divides } A = n^{\lfloor \log B \rfloor} \cdot \prod_{k=1}^{\lfloor \log^2 n \rfloor} (n^k - 1),$$

contradicting the fact that r was chosen to not divide A. We conclude that in fact $\gcd(n, r) = 1$, as claimed.

Since $\gcd(n, r) = 1$, it makes sense to talk about $\mathrm{ord}_r(n)$, and we conclude by proving that $\mathrm{ord}_r(n) > \log^2 n$. Setting $\ell = \mathrm{ord}_r(n)$, by definition we have $n^\ell \equiv 1$ in $\mathbb{Z}/r$, i.e. that r divides $n^\ell - 1$. By construction, r cannot divide $n^k - 1$ for $1 \leq k \leq \lfloor \log^2 n \rfloor$. We conclude that $\ell > \log^2 n$, as desired. □

5 Reduction to the arithmetic of $\overline{\mathbb{F}}_p$

We now prove Theorem 2.2 (or, rather, reduce it to a slightly different statement we will prove in the next section).

Proof of Theorem 2.2. We start by recalling what we must prove. Consider some $n \geq 2$. For all $1 \leq r \leq \lceil \log^5 n \rceil$ and $1 \leq a \leq \lfloor \sqrt{\phi(r)} \log n \rfloor$, assume that the following hold:

- $\gcd(a, n) = 1$, and
- In $(\mathbb{Z}/n)[x]/(x^r - 1)$, we have $(x + a)^n \equiv x^n + a$.

We must prove that n is a power of a prime.

This is trivial if $n = 2$, so we can assume that $n \geq 3$. The first step is to use Lemma 4.1 to find some $r_0 \leq \lceil \log^5 n \rceil$ such that $\gcd(n, r_0) = 1$ and $\mathrm{ord}_{r_0}(n) > \log^2 n$. Choosing some prime p dividing n, let $\mu \in \overline{\mathbb{F}}_p$ be a primitive r_0-th root of unity. We then have a ring homomorphism

$$\Psi \colon (\mathbb{Z}/n)[x]/(x^r - 1) \to \overline{\mathbb{F}}_p$$

taking $\mathbb{Z}/n$ to $\mathbb{F}_p$ and taking x to μ. Using our assumptions, for $1 \leq a \leq \lfloor \sqrt{\phi(r_0)} \log n \rfloor$ we have

$$0 \equiv \Psi((x - a)^n - (x^n - a)) \equiv (\mu - a)^n - (\mu^n - a),$$

so in $\overline{\mathbb{F}}_p$ we have

$$(\mu - a)^n \equiv \mu^n - a.$$

The theorem now follows from Theorem 6.1 below, which says that under precisely the circumstances we are considering, these identities in $\overline{\mathbb{F}}_p$ imply that n is a power of p. □

6 Completing the proof

The following theorem concerning $\overline{\mathbb{F}}_p$ was promised at the end of the previous section.

Theorem 6.1. *Let $n \geq 2$ and let p be a prime dividing n. For some $r \geq 2$ satisfying $\gcd(n, r) = 1$ and $\mathrm{ord}_r(n) > \log^2 n$, let $\mu \in \overline{\mathbb{F}}_p$ be a primitive r-th root of unity. For all integers $1 \leq a \leq \lfloor \sqrt{\phi(r)} \log n \rfloor$, assume that the following hold:*

- $\gcd(a, n) = 1$, *and*
- $(\mu + a)^n \equiv \mu^n + a$ *in* $\overline{\mathbb{F}}_p$.

Then n is a power of p.

Remark 6.2. If n is a power of p, then $(\mu + a)^n \equiv \mu^n + a$ holds for all $a \in \mathbb{F}_p$. The point of Theorem 6.1 is that for the converse to hold, we only need to check this for a certain range of values of a.

Proof of Theorem 6.1. For the sake of contradiction, make the following assumption:

$$\text{The integer } n \text{ is not a power of } p. \tag{6.1}$$

For $e \geq 1$ and $f \in \mathbb{Z}[x]$, write $e \sim f$ if

$$f(\mu)^e \equiv f(\mu^e) \quad \text{in} \quad \overline{\mathbb{F}}_p.$$

The hypotheses of the theorem say that $n \sim x + a$ for all integers $0 \leq a \leq \lfloor \sqrt{\phi(r)} \log n \rfloor$. The first step is to enlarge the number of exponents e and polynomials f for which this holds. Define

$$\mathcal{E} = \left\{ (\frac{n}{p})^k \cdot p^{k'} \mid k, k' \geq 0 \right\} \subset \mathbb{Z}$$

and

$$\mathcal{P} = \left\{ \prod_{k=1}^{m} (x - a_i) \mid m \geq 1 \text{ and } 0 \leq a_i \leq \lfloor \sqrt{\phi(r)} \log n \rfloor \text{ for all } 1 \leq i \leq m \right\} \subset \mathbb{Z}[x].$$

The factors in elements of $\mathcal{P}$ are allowed to repeat. We then have the following.

Claim 1. *For all $e \in \mathcal{E}$ and $f \in \mathcal{P}$, we have $e \sim f$.*

Proof of claim. It is enough to prove the following three facts:

(i) For all integers $0 \leq a \leq \lfloor \sqrt{\phi(r)} \log n \rfloor$, we have $p \sim x - a$ and $\frac{n}{p} \sim x - a$.
(ii) For all $e \in \mathcal{E}$ and $f, f' \in \mathcal{P}$ such that $e \sim f$ and $e \sim f'$, we have $e \sim f \cdot f'$.
(iii) For all $e, e' \in \mathcal{E}$ and $f \in \mathcal{P}$ such that $e \sim f$ and $e' \sim f$, we have $e \cdot e' \sim f$.

We will prove each in turn.

We start with (i). Consider an integer $0 \leq a \leq \lfloor \sqrt{\phi(r)} \log n \rfloor$. Since the Frobenius is an element of the Galois group of $\overline{\mathbb{F}}_p$, we have

$$(\mu - a)^p \equiv \mu^p - a^p \equiv \mu^p - a,$$

so $p \sim x - a$. As for n/p, by assumption we have $n \sim x - a$, so

$$(\mu^p)^{n/p} - a \equiv \mu^n - a \equiv (\mu - a)^n \equiv ((\mu - a)^p)^{n/p} \equiv (\mu^p - a)^{n/p}.$$

Applying the inverse of the Frobenius automorphism to both sides of this identity replaces μ^p by μ and fixes a, and thus yields

$$\mu^{n/p} - a \equiv (\mu - a)^{n/p},$$

so $n/p \sim x - a$, as claimed.

We now prove (ii). Consider $e \in \mathcal{E}$ and $f, f' \in \mathcal{P}$ such that $e \sim f$ and $e \sim f'$. We then have

$$(f(\mu) \cdot f'(\mu))^e \equiv f(\mu)^e \cdot f'(\mu)^e \equiv f(\mu^e) f'(\mu^e),$$

proving that $e \sim f \cdot f'$.

We conclude by proving (iii). Consider $e, e' \in \mathcal{E}$ and $f \in \mathcal{P}$ such that $e \sim f$ and $e' \sim f$. We thus have

$$f(\mu)^{e'} \equiv f(\mu^{e'}). \tag{6.2}$$

Since $\gcd(n, r) = 1$ and $e \in \mathcal{E}$, we also have $\gcd(e, r) = 1$. Since $\mu \in \overline{\mathbb{F}}_p$ is a primitive r-th root of unity, this implies that $\mu^e \in \overline{\mathbb{F}}_p$ is another primitive r-th root of unity, so there is an element of the Galois group of $\overline{\mathbb{F}}_p$ taking μ to μ^e. Hitting (6.2) with this element of the Galois group, we see that

$$f(\mu^e)^{e'} \equiv f((\mu^e)^{e'}).$$

We thus have

$$f(\mu)^{ee'} = f(\mu^e)^{e'} = f(\mu^{ee'}),$$

so $e \cdot e' \sim f$, as claimed. □

We now make the following two definitions:

- Let $\overline{\mathcal{E}} \subset \mathbb{Z}/r$ be the image of $\mathcal{E} \subset \mathbb{Z}$ under the map $\mathbb{Z} \to \mathbb{Z}/r$.
- Let $\overline{\mathcal{P}} = \{f(\mu) \mid f \in \mathcal{P}\} \subset \overline{\mathbb{F}}_p$.

With these definitions, we have the following inequality. We remark that this inequality is where we use our assumption (6.1).

Claim 2. *We have* $|\overline{\mathcal{P}}| \leq n^{\sqrt{|\overline{\mathcal{E}}|}}$.

Proof of claim. Consider the following subset of $\mathcal{E}$:

$$\mathcal{E}' = \left\{ (\frac{n}{p})^k \cdot p^{k'} \mid 0 \leq k, k' \leq \lfloor \sqrt{|\overline{\mathcal{E}}|} \rfloor \right\}.$$

Our assumption (6.1) asserting that n is not a power of p implies that

$$(\frac{n}{p})^{k_1} \cdot p^{k_1'} \neq (\frac{n}{p})^{k_2} \cdot p^{k_2'}$$

for all $k_1, k_1', k_2, k_2' \geq 0$ such that $(k_1, k_1') \neq (k_2, k_2')$. This implies that $\mathcal{E}'$ has at least $(\lfloor \sqrt{|\overline{\mathcal{E}}|} \rfloor + 1)^2 > |\overline{\mathcal{E}}|$ elements. We thus must be able to find distinct $m_1, m_2 \in \mathcal{E}' \subset \mathbb{Z}$ that project to the same element of $\overline{\mathcal{E}} \subset \mathbb{Z}/r$.

Order the m_i such that $m_1 > m_2$, and write $m_1 = (\frac{n}{p})^k p^{k'}$. Consider the polynomial $\Psi = x^{m_1} - x^{m_2} \in \mathbb{Z}[x]$. This polynomial has degree

$$m_1 = (\frac{n}{p})^k p^{k'} \leq (\frac{n}{p})^{\lfloor \sqrt{|\overline{\mathcal{E}}|} \rfloor} p^{\lfloor \sqrt{|\overline{\mathcal{E}}|} \rfloor} \leq n^{\sqrt{|\overline{\mathcal{E}}|}}.$$

It follows that Ψ has at most $n^{\sqrt{|\overline{\mathcal{E}}|}}$ roots in $\overline{\mathbb{F}}_p$. It is therefore enough to prove that each element of $\overline{\mathcal{P}} \subset \overline{\mathbb{F}}_p$ is a root of Ψ.

Consider some $f \in \mathcal{P}$. Our goal is to prove that $f(\mu)$ is a root of Ψ. Since $m_1, m_2 \sim f$, we have

$$\Psi(f(\mu)) \equiv f(\mu)^{m_1} - f(\mu)^{m_2} \equiv f(\mu^{m_1}) - f(\mu^{m_2}) \quad \text{in} \quad \overline{\mathbb{F}}_p. \tag{6.3}$$

Since μ is a primitive r-th root of unity and m_1 equals m_2 modulo r, we have $\mu^{m_1} \equiv \mu^{m_2}$, so (6.3) is 0, as desired. □

To establish our contradiction and prove the theorem, therefore, it is enough to prove the opposite inequality

$$|\overline{\mathcal{P}}| > n^{\sqrt{|\overline{\mathcal{E}}|}}.$$

This will be Claim 7 below, which is preceded by four preliminary claims.

Claim 3. *We have* $|\overline{\mathcal{E}}| > \log^2 n$.

Proof of claim. Recall from the hypotheses of the theorem we are proving that $\mathrm{ord}_r(n) > \log^2 n$. Since $\mathcal{E} \subset \mathbb{Z}$ contains all powers of n and $\overline{\mathcal{E}}$ is the reduction of $\mathcal{E}$ modulo r, it follows that $|\overline{\mathcal{E}}| > \log^2 n$. □

Claim 4. *We have* $|\overline{\mathcal{E}}| \leq \phi(r)$.

Proof of claim. Since $\gcd(n, r) = 1$, the two generators $\frac{n}{p}$ and p of $\mathcal{E}$ both map to units in $\mathbb{Z}/r$. It follows that $|\overline{\mathcal{E}}| \leq |(\mathbb{Z}/r)^\times| = \phi(r)$. □

Claim 5. *The mod-p reduction map* $\mathbb{Z}[x] \to \mathbb{F}_p[x]$ *restricts to an injection* $\mathcal{P} \to \mathbb{F}_p[x]$.

Proof of claim. Recall that $\mathcal{P}$ is multiplicatively generated by elements of the form $x - a \in \mathbb{Z}[x]$ with a an integer satisfying $0 \leq a \leq \lfloor \sqrt{\phi(r)} \log n \rfloor$. Aside from $a = 0$, the assumptions of our theorem say that such integers a all satisfy $\gcd(a, n) = 1$. Since p divides n, it follows that $a = 0$ is the only one that is divisible by p, and hence that $p > \lfloor \sqrt{\phi(r)} \log n \rfloor$. We conclude that all the multiplicative generators of $\mathcal{P}$ map to distinct linear polynomials in $\mathbb{F}_p[x]$. This implies that the mod-p reduction map $\mathbb{Z}[x] \to \mathbb{F}_p[x]$ restricts to an injection from $\mathcal{P}$ to $\mathbb{F}_p[x]$, as claimed. □

Claim 6. *We have* $|\overline{\mathcal{P}}| \geq \binom{|\overline{\mathcal{E}}| + \lfloor \sqrt{\phi(r)} \log n \rfloor}{|\overline{\mathcal{E}}| - 1}$.

Proof of claim. Consider polynomials $f(x), g(x) \in \mathcal{P} \subset \mathbb{Z}[x]$ that map to the same element of $\overline{\mathcal{P}} \subset \overline{\mathbb{F}}_p$ and whose degrees are less than $|\overline{\mathcal{E}}|$. We claim that $f = g$. To prove this, by Claim 5 it is enough to prove that f and g map to the same element of $\mathbb{F}_p[x]$. Assume otherwise. Their difference $f(x) - g(x)$ is thus a polynomial of degree less than $|\overline{\mathcal{E}}|$ that does not map to the zero polynomial in $\mathbb{F}_p[x]$. It follows that $f(x) - g(x)$ has strictly fewer than $|\overline{\mathcal{E}}|$ roots in $\overline{\mathbb{F}}_p$. Since f and g map to the same element of $\overline{\mathcal{P}}$, we have $f(\mu) \equiv g(\mu)$, so μ is one of these roots. But now for all $e \in \mathcal{E}$ we have

$$f(\mu^e) \equiv f(\mu)^e \equiv g(\mu)^e \equiv g(\mu^e),$$

so in fact μ^e is a root for all $e \in \mathcal{E}$. Since μ is a primitive r-th root of unity and $\overline{\mathcal{E}}$ is the mod-r reduction of $\mathcal{E}$, this gives $|\overline{\mathcal{E}}|$ roots, a contradiction.

We deduce that $|\overline{\mathcal{P}}|$ is greater than or equal to the number of elements of $\mathcal{P}$ whose degrees are at most $|\overline{\mathcal{E}}|-1$. Such elements are in bijection with collections of at most $|\overline{\mathcal{E}}|-1$ elements of the set

$$\left\{x-a \mid 0 \le a \le \lfloor\sqrt{\phi(r)}\log n\rfloor\right\}$$

of multiplicative generators for $\mathcal{P}$. These collections are allowed to have repeated elements. Since we allow $a=0$ in the above set, it has $\lfloor\sqrt{\phi(r)}\log n\rfloor+1$ elements. The number of collections of at most $|\overline{\mathcal{E}}|-1$ elements (potentially with repetitions) from a set of size $\lfloor\sqrt{\phi(r)}\log n\rfloor+1$ is precisely

$$\binom{|\overline{\mathcal{E}}|+\lfloor\sqrt{\phi(r)}\log n\rfloor}{|\overline{\mathcal{E}}|-1}.$$

The claim follows. □

Claim 7. *We have* $|\overline{\mathcal{P}}| > n^{\sqrt{|\overline{\mathcal{E}}|}}$.

Proof of claim. We start with the estimate

$$|\overline{\mathcal{P}}| \ge \binom{|\overline{\mathcal{E}}|+\lfloor\sqrt{\phi(r)}\log n\rfloor}{|\overline{\mathcal{E}}|-1} \tag{6.4}$$

given by Claim 6. Claim 4 says that $|\overline{\mathcal{E}}| \le \sqrt{\phi(r)}$, so the right hand side of (6.4) is at least

$$\binom{|\overline{\mathcal{E}}|+\lfloor\sqrt{|\overline{\mathcal{E}}|}\log n\rfloor}{|\overline{\mathcal{E}}|-1} \tag{6.5}$$

Claim 3 says that $|\overline{\mathcal{E}}| > \log^2 n$, so $\sqrt{|\overline{\mathcal{E}}|} > \log n$ and hence $|\overline{\mathcal{E}}| \ge \lfloor\sqrt{|\overline{\mathcal{E}}|}\log n\rfloor+1$. The quantity (6.5) is thus at least[1]

$$\binom{2\lfloor\sqrt{|\overline{\mathcal{E}}|}\log n\rfloor+1}{\lfloor\sqrt{|\overline{\mathcal{E}}|}\log n\rfloor}. \tag{6.6}$$

We now use Claim 3 again to see that

$$\lfloor\sqrt{|\overline{\mathcal{E}}|}\log n\rfloor \ge \lfloor\log^2 n\rfloor \ge \lfloor\log^2 2\rfloor = 1,$$

so (6.6) is strictly greater than[2]

$$2^{\lfloor\sqrt{|\overline{\mathcal{E}}|}\log n\rfloor+1} \ge 2^{\sqrt{|\overline{\mathcal{E}}|}\log n} = (2^{\log n})^{\sqrt{|\overline{\mathcal{E}}|}} = n^{\sqrt{|\overline{\mathcal{E}}|}}.$$

The claim follows. □

As was noted above, Claims 2 and 7 contradict each other, so our assumption (6.1) must be wrong, i.e. n must be a power of p, as desired. □

[1] Here we are using the inequality $\binom{a+c}{b+c} \ge \binom{a}{b}$, which holds for all $a,b,c \ge 0$. To see this, observe that there is an injection from the b-element subsets of $\{1,\dots,a\}$ to the $(b+c)$-element subsets of $\{1,\dots,a+c\}$ taking I to $I \cup \{a+1,\dots,a+c\}$.

[2] Here we are using the strict inequality $\binom{2a+1}{a} > 2^{a+1}$, which holds for all $a \ge 1$. To see this, observe that there is a non-injective surjection from the a-element subsets of $\{1,\dots,2a+1\}$ to the set of all subsets of $\{1,\dots,a+1\}$ taking I to $I \cap \{1,\dots,a+1\}$.

References

[1] M. Agrawal, N. Kayal and N. Saxena, PRIMES is in P, Ann. of Math. (2) 160 (2004), no. 2, 781–793.

[2] M. Agrawal, N. Kayal and N. Saxena, Errata: PRIMES is in P, Ann. of Math. (2) 189 (2019), no. 1, 317–318.

[3] A. Granville, It is easy to determine whether a given integer is prime, Bull. Amer. Math. Soc. (N.S.) 42 (2005), no. 1, 3–38.

[4] M. Nair, On Chebyshev-type inequalities for primes, Amer. Math. Monthly 89 (1982), no. 2, 126–129.

Andrew Putman
Department of Mathematics
University of Notre Dame
255 Hurley Hall
Notre Dame, IN 46556
andyp@nd.edu