21

Deterministic primality testing

For many years, despite much research in the area, there was no known deterministic, polynomial-time algorithm for testing whether a given integer $n > 1$ is a prime. However, that is no longer the case — the breakthrough algorithm of Agrawal, Kayal, and Saxena, or Algorithm AKS for short, is just such an algorithm. Not only is the result itself remarkable, but the algorithm is striking both in its simplicity, and in the fact that the proof of its running time and correctness are completely elementary (though ingenious).

We should stress at the outset that although this result is an important theoretical result, as of yet, it has no real practical significance: probabilistic tests, such as the Miller–Rabin test discussed in Chapter 10, are *much* more efficient, and a practically minded person should not at all be bothered by the fact that such algorithms may in theory make a mistake with an incredibly small probability.

21.1 The basic idea

The algorithm is based on the following fact:

Theorem 21.1. *Let $n > 1$ be an integer. If n is prime, then for all $a \in \mathbb{Z}_n$, we have the following identity in the ring $\mathbb{Z}_n[X]$:*

$$(X + a)^n = X^n + a. \tag{21.1}$$

Conversely, if n is composite, then for all $a \in \mathbb{Z}_n^*$, the identity (21.1) does not hold.

Proof. Note that

$$(X + a)^n = X^n + a^n + \sum_{i=1}^{n-1} \binom{n}{i} a^i X^{n-i}.$$

If n is prime, then by Fermat's little theorem (Theorem 2.14), we have $a^n = a$, and by Exercise 1.14, all of the binomial coefficients $\binom{n}{i}$, for $i = 1, \ldots, n - 1$, are divisible by n, and hence their images in the ring $\mathbb{Z}_n$ vanish. That proves that the identity (21.1) holds when n is prime.

Conversely, suppose that n is composite and that $a \in \mathbb{Z}_n^*$. Consider any prime factor p of n, and suppose $n = p^k m$, where $p \nmid m$.

We claim that $p^k \nmid \binom{n}{p}$. To prove the claim, one simply observes that

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

and the numerator of this fraction is an integer divisible by p^k, but no higher power of p, and the denominator is divisible by p, but no higher power of p. That proves the claim.

From the claim, and the fact that $a \in \mathbb{Z}_n^*$, it follows that the coefficient of X^{n-p} in $(X + a)^n$ is not zero, and hence the identity (21.1) does not hold. □

Of course, Theorem 21.1 does not immediately give rise to an efficient primality test, since just evaluating the left-hand side of the identity (21.1) takes time $\Omega(n)$ in the worst case. The key observation of Agrawal, Kayal, and Saxena is that if (21.1) holds modulo $X^r - 1$ for a suitably chosen value of r, and for sufficiently many a, then n must be prime. To make this idea work, one must show that a suitable r exists that is bounded by a polynomial in $\text{len}(n)$, and that the number of different values of a that must be tested is also bounded by a polynomial in $\text{len}(n)$.

21.2 The algorithm and its analysis

The algorithm is shown in Fig. 21.1. A few remarks on implementation are in order:

- In step 1, we can use the algorithm for perfect-power testing discussed in Exercise 3.31.
- The search for r in step 2 can just be done by brute-force search; likewise, the determination of the multiplicative order of $[n]_r \in \mathbb{Z}_r^*$ can be done by brute force: after verifying that $\gcd(n, r) = 1$, compute successive powers of n modulo r until we get 1.

We want to prove that Algorithm AKS runs in polynomial time and is correct. To prove that it runs in polynomial time, it clearly suffices to prove that there exists an integer r satisfying the condition in step 2 that is bounded by a polynomial in $\text{len}(n)$, since all other computations can be carried out in time $(r + \text{len}(n))^{O(1)}$. Correctness means that it outputs *true* if and only if n is prime.

```
On input $n$, where $n$ is an integer and $n > 1$, do the following:
1.  if $n$ is of the form $a^b$ for integers $a > 1$ and $b > 1$ then
            return false
2.  find the smallest integer $r > 1$ such that either
            $\gcd(n, r) > 1$
    or
            $\gcd(n, r) = 1$ and
            $[n]_r \in \mathbb{Z}_r^*$ has multiplicative order $> 4\,\mathrm{len}(n)^2$
3.  if $r = n$ then return true
4.  if $\gcd(n, r) > 1$ then return false
5.  for $j \leftarrow 1$ to $2\,\mathrm{len}(n)\lfloor r^{1/2} \rfloor + 1$ do
            if $(X + j)^n \not\equiv X^n + j \pmod{X^r - 1}$ in the ring $\mathbb{Z}_n[X]$ then
                    return false
6.  return true
```

Fig. 21.1. Algorithm AKS

21.2.1 Running time analysis

The question of the running time of Algorithm AKS is settled by the following fact:

Theorem 21.2. *For integers $n > 1$ and $m \geq 1$, the least prime r such that $r \nmid n$ and the multiplicative order of $[n]_r \in \mathbb{Z}_r^*$ is greater than m is $O(m^2 \,\mathrm{len}(n))$.*

Proof. Call a prime r "good" if $r \nmid n$ and the multiplicative order of $[n]_r \in \mathbb{Z}_r^*$ is greater than m, and otherwise call r "bad." If r is bad, then either $r \mid n$ or $r \mid (n^d - 1)$ for some $d = 1, \ldots, m$. Thus, any bad prime r satisfies

$$r \mid n \prod_{d=1}^{m} (n^d - 1).$$

If all primes r up to some given bound $x \geq 2$ are bad, then the product of all primes up to x divides $n \prod_{d=1}^{m} (n^d - 1)$, and so in particular,

$$\prod_{r \leq x} r \leq n \prod_{d=1}^{m} (n^d - 1),$$

where the first product is over all primes r up to x. Taking logarithms, we obtain

$$\sum_{r \le x} \log r \le \log\Big(n \prod_{d=1}^{m} (n^d - 1)\Big) \le (\log n)\Big(1 + \sum_{d=1}^{m} d\Big)$$
$$= (\log n)(1 + m(m+1)/2).$$

But by Theorem 5.7, we have

$$\sum_{r \le x} \log r \ge cx$$

for some constant $c > 0$, from which it follows that

$$x \le c^{-1}(\log n)(1 + m(m+1)/2),$$

and the theorem follows. □

From this theorem, it follows that the value of r found in step 2—which need not be prime—will be $O(\mathrm{len}(n)^5)$. From this, we obtain:

Theorem 21.3. *Algorithm* AKS *can be implemented so that its running time is* $O(\mathrm{len}(n)^{16.5})$.

Proof. As discussed above, the value of r determined in step 2 will be $O(\mathrm{len}(n)^5)$. It is fairly straightforward to see that the running time of the algorithm is dominated by the running time of step 5. Here, we have to perform $O(r^{1/2}\,\mathrm{len}(n))$ exponentiations to the power n in the ring $\mathbb{Z}_n[X]/(X^r - 1)$. Each of these exponentiations takes $O(\mathrm{len}(n))$ operations in $\mathbb{Z}_n[X]/(X^r - 1)$, each of which takes $O(r^2)$ operations in $\mathbb{Z}_n$, each of which takes time $O(\mathrm{len}(n)^2)$. This yields a running time bounded by a constant times

$$r^{1/2}\,\mathrm{len}(n) \times \mathrm{len}(n) \times r^2 \times \mathrm{len}(n)^2 = r^{2.5}\,\mathrm{len}(n)^4.$$

Substituting the bound $O(\mathrm{len}(n)^5)$ for r, we obtain the desired bound. □

21.2.2 Correctness

As for the correctness of Algorithm AKS, we first show:

Theorem 21.4. *If the input to Algorithm* AKS *is prime, then the output is* true.

Proof. Assume that the input n is prime. The test in step 1 will certainly fail. If the algorithm does not return *true* in step 3, then certainly the test in step 4 will fail as well. If the algorithm reaches step 5, then all of the tests in the loop in step 5 will fail—this follows from Theorem 21.1. □

The interesting case is the following:

Theorem 21.5. *If the input to Algorithm* AKS *is composite, then the output is* false.

The proof of this theorem is rather long, and is the subject of the remainder of this section.

Suppose the input n is composite. If n is a prime power, then this will be detected in step 1, so we may assume that n is not a prime power. Assume that the algorithm has found a suitable value of r in step 2. Clearly, the test in 3 will fail. If the test in step 4 passes, we are done, so we may assume that this test fails; that is, we may assume that all prime factors of n are greater than r. Our goal now is to show that one of the tests in the loop in step 5 must pass. The proof will be by contradiction: we shall assume that none of the tests pass, and derive a contradiction.

The assumption that none of the tests in step 5 fail means that in the ring $\mathbb{Z}_n[X]$, the following congruences hold:

$$(X + j)^n \equiv X^n + j \pmod{X^r - 1} \quad (j = 1, \dots, 2\,\mathrm{len}(n)\lfloor r^{1/2} \rfloor + 1). \tag{21.2}$$

For the rest of the proof, we fix a particular prime divisor p of n—the choice of p does not matter. Since $p \mid n$, we have a natural ring homomorphism from $\mathbb{Z}_n[X]$ to $\mathbb{Z}_p[X]$ (see Examples 7.52 and 7.46), which implies that the congruences (21.2) hold in the ring of polynomials over $\mathbb{Z}_p$ as well. *From now on, we shall work exclusively with polynomials over* $\mathbb{Z}_p$.

Let us state in somewhat more abstract terms the precise assumptions we are making in order to derive our contradiction:

(A0) $n > 1$, $r > 1$, and $\ell \geq 1$ are integers, p is a prime dividing n, and $\gcd(n, r) = 1$;

(A1) n is not a prime power;

(A2) $p > r$;

(A3) the congruences

$$(X + j)^n \equiv X^n + j \pmod{X^r - 1} \quad (j = 1, \dots, \ell)$$

hold in the ring $\mathbb{Z}_p[X]$;

(A4) the multiplicative order of $[n]_r \in \mathbb{Z}_r^*$ is greater than $4\,\mathrm{len}(n)^2$;

(A5) $\ell > 2\,\mathrm{len}(n)\lfloor r^{1/2} \rfloor$.

The rest of the proof will rely only on these assumptions, and not on any other details of Algorithm AKS. From now on, only assumption (A0) will be implicitly in force. The other assumptions will be explicitly invoked as necessary. Our goal is to show that assumptions (A1), (A2), (A3), (A4), and (A5) cannot all be true simultaneously.

Define the $\mathbb{Z}_p$-algebra $E := \mathbb{Z}_p[X]/(X^r - 1)$, and let $\xi := [X]_{X^r-1} \in E$, so that $E = \mathbb{Z}_p[\xi]$. Every element of E can be expressed uniquely as $g(\xi) = [g]_{X^r-1}$, for $g \in \mathbb{Z}_p[X]$ of degree less than r, and for an arbitrary polynomial $g \in \mathbb{Z}_p[X]$, we have $g(\xi) = 0$ if and only if $(X^r - 1) \mid g$. Note that $\xi \in E^*$ and has multiplicative order r: indeed, $\xi^r = 1$, and $\xi^s - 1$ cannot be zero for $s < r$, since $X^s - 1$ has degree less than r.

Assumption (A3) implies that we have a number of interesting identities in the $\mathbb{Z}_p$-algebra E:

$$(\xi + j)^n = \xi^n + j \quad (j = 1, \ldots, \ell).$$

For the polynomials $g_j := X + j \in \mathbb{Z}_p[X]$, with j in the given range, these identities say that $g_j(\xi)^n = g_j(\xi^n)$.

In order to exploit these identities, we study more generally functions σ_k, for various integer values k, that send $g(\xi) \in E$ to $g(\xi^k)$, for arbitrary $g \in \mathbb{Z}_p[X]$, and we investigate the implications of the assumption that such functions behave like the k-power map on certain inputs. To this end, let $\mathbb{Z}^{(r)}$ denote the set of all positive integers k such that $\gcd(r, k) = 1$. Note that the set $\mathbb{Z}^{(r)}$ is **multiplicative**, by which we mean $1 \in \mathbb{Z}^{(r)}$, and $kk' \in \mathbb{Z}^{(r)}$ for all $k, k' \in \mathbb{Z}^{(r)}$. Also note that because of our assumption (A0), both n and p are in $\mathbb{Z}^{(r)}$. For $k \in \mathbb{Z}^{(r)}$, let $\hat{\sigma}_k : \mathbb{Z}_p[X] \to E$ be the polynomial evaluation map that sends $g \in \mathbb{Z}_p[X]$ to $g(\xi^k)$. This is of course a $\mathbb{Z}_p$-algebra homomorphism, and we have:

Lemma 21.6. *For all $k \in \mathbb{Z}^{(r)}$, the kernel of $\hat{\sigma}_k$ is $(X^r - 1)$, and the image of $\hat{\sigma}_k$ is E.*

Proof. Let $J := \operatorname{Ker} \hat{\sigma}_k$, which is an ideal of $\mathbb{Z}_p[X]$. Let k' be a positive integer such that $kk' \equiv 1 \pmod{r}$, which exists because $\gcd(r, k) = 1$.

To show that $J = (X^r - 1)$, we first observe that

$$\hat{\sigma}_k(X^r - 1) = (\xi^k)^r - 1 = (\xi^r)^k - 1 = 1^k - 1 = 0,$$

and hence $(X^r - 1) \subseteq J$.

Next, we show that $J \subseteq (X^r - 1)$. Let $g \in J$. We want to show that $(X^r - 1) \mid g$. Now, $g \in J$ means that $g(\xi^k) = 0$. If we set $h := g(X^k)$, this implies that $h(\xi) = 0$, which means that $(X^r - 1) \mid h$. So let us write $h = (X^r - 1)f$, for some $f \in \mathbb{Z}_p[X]$. Then

$$g(\xi) = g(\xi^{kk'}) = h(\xi^{k'}) = (\xi^{k'r} - 1)f(\xi^{k'}) = 0,$$

which implies that $(X^r - 1) \mid g$.

That finishes the proof that $J = (X^r - 1)$.

Finally, to show that $\hat{\sigma}_k$ is surjective, suppose we are given an arbitrary element of E, which we can express as $g(\xi)$ for some $g \in \mathbb{Z}_p[X]$. Now set $h := g(X^{k'})$, and observe that

$$\hat{\sigma}_k(h) = h(\xi^k) = g(\xi^{kk'}) = g(\xi). \quad \square$$

Because of Lemma 21.6, then by Theorem 7.26, the map $\sigma_k : E \to E$ that sends $g(\xi) \in E$ to $g(\xi^k)$, for $g \in \mathbb{Z}_p[X]$, is well defined, and is a ring automorphism—indeed, a $\mathbb{Z}_p$-*algebra* automorphism—on E. Note that for all $k, k' \in \mathbb{Z}^{(r)}$, we have

- $\sigma_k = \sigma_{k'}$ if and only if $\xi^k = \xi^{k'}$ if and only if $k \equiv k' \pmod{r}$, and
- $\sigma_k \circ \sigma_{k'} = \sigma_{k'} \circ \sigma_k = \sigma_{kk'}$.

So in fact, the set $\{\sigma_k : k \in \mathbb{Z}^{(r)}\}$ under composition forms an abelian group that is isomorphic to $\mathbb{Z}_r^*$.

Remark. It is perhaps helpful (but not necessary for the proof) to examine the behavior of the map σ_k in a bit more detail. Let $\alpha \in E$, and let

$$\alpha = \sum_{i=0}^{r-1} a_i \xi^i$$

be the canonical representation of α. Since $\gcd(r, k) = 1$, the map $\pi : \{0, \ldots, r-1\} \to \{0, \ldots, r-1\}$ that sends i to $ki \bmod r$ is a permutation whose inverse is the permutation π' that sends i to $k'i \bmod r$, where k' is a multiplicative inverse of k modulo r. Then we have

$$\sigma_k(\alpha) = \sum_{i=0}^{r-1} a_i \xi^{ki} = \sum_{i=0}^{r-1} a_i \xi^{\pi(i)} = \sum_{i=0}^{r-1} a_{\pi'(i)} \xi^i.$$

Thus, the action of σ_k is to permute the coordinate vector $(a_0, \ldots, a_{r-1})$ of α, sending α to the element in E whose coordinate vector is $(a_{\pi'(0)}, \ldots, a_{\pi'(r-1)})$. So we see that although we defined the maps σ_k in a rather "highbrow" algebraic fashion, their behavior in concrete terms is actually quite simple.

Recall that the p-power map on E is a $\mathbb{Z}_p$-algebra homomorphism (see Theorem 19.7), and so for all $\alpha \in E$, if $\alpha = g(\xi)$ for $g \in \mathbb{Z}_p[X]$, then (by Theorem 16.7) we have

$$\alpha^p = g(\xi)^p = g(\xi^p) = \sigma_p(\alpha).$$

Thus, σ_p acts just like the p-power map on all elements of E.

We can restate assumption (A3) as follows:

$$\sigma_n(\xi + j) = (\xi + j)^n \quad (j = 1, \ldots, \ell).$$

That is to say, the map σ_n acts just like the n-power map on the elements $\xi + j$ for $j = 1, \ldots, \ell$.

Now, although the σ_p map must act like the p-power map on all of E, there is no good reason why the σ_n map should act like the n-power map on any particular element of E, and so the fact that it does so on all the elements $\xi + j$ for $j = 1, \ldots, \ell$ looks decidedly suspicious. To turn our suspicions into a contradiction, let us start by defining some notation. For $\alpha \in E$, let us define

$$C(\alpha) := \{k \in \mathbb{Z}^{(r)} : \sigma_k(\alpha) = \alpha^k\},$$

and for $k \in \mathbb{Z}^{(r)}$, let us define

$$D(k) := \{\alpha \in E : \sigma_k(\alpha) = \alpha^k\}.$$

In words: $C(\alpha)$ is the set of all k for which σ_k acts like the k-power map on α, and $D(k)$ is the set of all α for which σ_k acts like the k-power map on α. From the discussion above, we have $p \in C(\alpha)$ for all $\alpha \in E$, and it is also clear that $1 \in C(\alpha)$ for all $\alpha \in E$. Also, it is clear that $\alpha \in D(p)$ for all $\alpha \in E$, and $1_E \in D(k)$ for all $k \in \mathbb{Z}^{(r)}$.

The following two simple lemmas say that the sets $C(\alpha)$ and $D(k)$ are multiplicative.

Lemma 21.7. *For every $\alpha \in E$, if $k \in C(\alpha)$ and $k' \in C(\alpha)$, then $kk' \in C(\alpha)$.*

Proof. If $\sigma_k(\alpha) = \alpha^k$ and $\sigma_{k'}(\alpha) = \alpha^{k'}$, then

$$\sigma_{kk'}(\alpha) = \sigma_k(\sigma_{k'}(\alpha)) = \sigma_k(\alpha^{k'}) = (\sigma_k(\alpha))^{k'} = (\alpha^k)^{k'} = \alpha^{kk'},$$

where we have made use of the homomorphic property of σ_k. □

Lemma 21.8. *For every $k \in \mathbb{Z}^{(r)}$, if $\alpha \in D(k)$ and $\beta \in D(k)$, then $\alpha\beta \in D(k)$.*

Proof. If $\sigma_k(\alpha) = \alpha^k$ and $\sigma_k(\beta) = \beta^k$, then

$$\sigma_k(\alpha\beta) = \sigma_k(\alpha)\sigma_k(\beta) = \alpha^k\beta^k = (\alpha\beta)^k,$$

where again, we have made use of the homomorphic property of σ_k. □

Let us define

- s to be the multiplicative order of $[p]_r \in \mathbb{Z}_r^*$, and
- t to be the order of the subgroup of $\mathbb{Z}_r^*$ generated by $[p]_r$ and $[n]_r$.

Since $r \mid (p^s - 1)$, if we take any extension field F of degree s over $\mathbb{Z}_p$ (which we know exists by Theorem 19.12), then since F^* is cyclic (Theorem 7.29) and has order $p^s - 1$, we know that there exists an element $\zeta \in F^*$ of multiplicative order r (Theorem 6.32). Let us define the polynomial evaluation map $\hat{\tau} : \mathbb{Z}_p[X] \to F$ that sends $g \in \mathbb{Z}_p[X]$ to $g(\zeta) \in F$. Since $X^r - 1$ is clearly in the kernel of $\hat{\tau}$, then by Theorem 7.27, the map $\tau : E \to F$ that sends $g(\xi)$ to $g(\zeta)$, for $g \in \mathbb{Z}_p[X]$, is a well-defined ring homomorphism, and actually, it is a $\mathbb{Z}_p$-algebra homomorphism.

For concreteness, one could think of F as $\mathbb{Z}_p[X]/(f)$, where f is an irreducible factor of $X^r - 1$ of degree s. In this case, we could simply take ζ to be $[X]_f$ (see

Example 19.1), and the map $\hat{\tau}$ above would be just the natural map from $\mathbb{Z}_p[X]$ to $\mathbb{Z}_p[X]/(f)$.

The key to deriving our contradiction is to examine the set $S := \tau(D(n))$, that is, the image under τ of the set $D(n)$ of all elements $\alpha \in E$ for which σ_n acts like the n-power map.

Lemma 21.9. *Under assumption (A1), we have*

$$|S| \le n^{2\lfloor t^{1/2} \rfloor}.$$

Proof. Consider the set of integers

$$I := \{n^u p^v : u, v = 0, \ldots, \lfloor t^{1/2} \rfloor\}.$$

We first claim that $|I| > t$. To prove this, we first show that each distinct pair (u, v) gives rise to a distinct value $n^u p^v$. To this end, we make use of our assumption (A1) that n is not a prime power, and so is divisible by some prime q other than p. Thus, if $(u', v') \neq (u, v)$, then either

- $u \neq u'$, in which case the power of q in the prime factorization of $n^u p^v$ is different from that in $n^{u'} p^{v'}$, or
- $u = u'$ and $v \neq v'$, in which case the power of p in the prime factorization of $n^u p^v$ is different from that in $n^{u'} p^{v'}$.

The claim now follows from the fact that both u and v range over a set of size $\lfloor t^{1/2} \rfloor + 1 > t^{1/2}$, and so there are strictly more than t such pairs (u, v).

Next, recall that t was defined to be the order of the subgroup of $\mathbb{Z}_r^*$ generated by $[n]_r$ and $[p]_r$; equivalently, t is the number of distinct residue classes of the form $[n^u p^v]_r$, where u and v range over all non-negative integers. Since each element of I is of the form $n^u p^v$, and $|I| > t$, we may conclude that there must be two distinct elements of I, call them k and k', that are congruent modulo r. Furthermore, any element of I is a product of two positive integers each of which is at most $n^{\lfloor t^{1/2} \rfloor}$, and so both k and k' lie in the range $1, \ldots, n^{2\lfloor t^{1/2} \rfloor}$.

Now, let $\alpha \in D(n)$. This is equivalent to saying $n \in C(\alpha)$. We always have $1 \in C(\alpha)$ and $p \in C(\alpha)$, and so by Lemma 21.7, we have $n^u p^v \in C(\alpha)$ for all non-negative integers u, v, and so in particular, $k, k' \in C(\alpha)$.

Since both k and k' are in $C(\alpha)$, we have

$$\sigma_k(\alpha) = \alpha^k \text{ and } \sigma_{k'}(\alpha) = \alpha^{k'}.$$

Since $k \equiv k' \pmod{r}$, we have $\sigma_k = \sigma_{k'}$, and hence

$$\alpha^k = \alpha^{k'}.$$

Now apply the homomorphism τ, obtaining

$$\tau(\alpha)^k = \tau(\alpha)^{k'}.$$

Since this holds for all $\alpha \in D(n)$, we conclude that all elements of S are roots of the polynomial $X^k - X^{k'}$. Since $k \neq k'$, we see that $X^k - X^{k'}$ is a non-zero polynomial of degree at most $\max\{k, k'\} \leq n^{2\lfloor t^{1/2} \rfloor}$, and hence can have at most $n^{2\lfloor t^{1/2} \rfloor}$ roots in the field F (Theorem 7.14). □

Lemma 21.10. *Under assumptions (A2) and (A3), we have*

$$|S| \geq 2^{\min(t,\ell)} - 1.$$

Proof. Let $m := \min(t, \ell)$. Under assumption (A3), we have $\xi + j \in D(n)$ for $j = 1, \ldots, m$. Under assumption (A2), we have $p > r > t \geq m$, and hence the integers $j = 1, \ldots, m$ are distinct modulo p. Define

$$P := \Big\{ \prod_{j=1}^{m} (X + j)^{e_j} \in \mathbb{Z}_p[X] : e_j \in \{0, 1\} \text{ for } j = 1, \ldots, m, \text{ and } \sum_{j=1}^{m} e_j < m \Big\}.$$

That is, we form P by taking products over all subsets $S \subsetneq \{X + j : j = 1, \ldots, m\}$. Clearly, $|P| = 2^m - 1$.

Define $P(\xi) := \{f(\xi) \in E : f \in P\}$ and $P(\zeta) := \{f(\zeta) \in F : f \in P\}$. Note that $\tau(P(\xi)) = P(\zeta)$, and that by Lemma 21.8, $P(\xi) \subseteq D(n)$.

Therefore, to prove the lemma, it suffices to show that $|P(\zeta)| = 2^m - 1$. Suppose that this is not the case. This would give rise to distinct polynomials $g, h \in \mathbb{Z}_p[X]$, both of degree at most $t - 1$, such that

$$g(\xi) \in D(n),\ h(\xi) \in D(n), \text{ and } \tau(g(\xi)) = \tau(h(\xi)).$$

So we have $n \in C(g(\xi))$ and (as always) $1, p \in C(g(\xi))$. Likewise, we have $1, n, p \in C(h(\xi))$. By Lemma 21.7, for all integers k of the form $n^u p^v$, where u and v range over all non-negative integers, we have

$$k \in C(g(\xi)) \text{ and } k \in C(h(\xi)).$$

For each such k, since $\tau(g(\xi)) = \tau(h(\xi))$, we have $\tau(g(\xi))^k = \tau(h(\xi))^k$, and hence

$$\begin{aligned} 0 &= \tau(g(\xi))^k - \tau(h(\xi))^k \\ &= \tau(g(\xi)^k) - \tau(h(\xi)^k) \quad (\tau \text{ is a homomorphism}) \\ &= \tau(g(\xi^k)) - \tau(h(\xi^k)) \quad (k \in C(g(\xi)) \text{ and } k \in C(h(\xi))) \\ &= g(\zeta^k) - h(\zeta^k) \quad (\text{definition of } \tau). \end{aligned}$$

Thus, the polynomial $f := g - h \in \mathbb{Z}_p[X]$ is a non-zero polynomial of degree at most $t - 1$, having roots ζ^k in the field F for all k of the form $n^u p^v$. Now, t is by definition the number of distinct residue classes of the form $[n^u p^v]_r \in \mathbb{Z}_r^*$. Also, since ζ has multiplicative order r, for all integers k, k', we have $\zeta^k = \zeta^{k'}$ if and only if $k \equiv k' \pmod{r}$. Therefore, as k ranges over all integers of the form $n^u p^v$,

ζ^k ranges over precisely t distinct values in F. But since all of these values are roots of the polynomial f, which is non-zero and of degree at most $t-1$, this is impossible (Theorem 7.14). □

We are now (finally!) in a position to complete the proof of Theorem 21.5. Under assumptions (A1), (A2), and (A3), Lemmas 21.9 and 21.10 imply that

$$2^{\min(t,\ell)} - 1 \leq |S| \leq n^{2\lfloor t^{1/2} \rfloor}. \tag{21.3}$$

The contradiction is provided by the following:

Lemma 21.11. *Under assumptions (A4) and (A5), we have*

$$2^{\min(t,\ell)} - 1 > n^{2\lfloor t^{1/2} \rfloor}.$$

Proof. Observe that $\log_2 n \leq \text{len}(n)$, and so it suffices to show that

$$2^{\min(t,\ell)} - 1 > 2^{2\,\text{len}(n)\lfloor t^{1/2} \rfloor},$$

and for this, it suffices to show that

$$\min(t,\ell) > 2\,\text{len}(n)\lfloor t^{1/2} \rfloor,$$

since for all integers a, b with $a > b \geq 1$, we have $2^a > 2^b + 1$.

To show that $t > 2\,\text{len}(n)\lfloor t^{1/2} \rfloor$, it suffices to show that $t > 2\,\text{len}(n) t^{1/2}$, or equivalently, that $t > 4\,\text{len}(n)^2$. But observe that by definition, t is the order of the subgroup of $\mathbb{Z}_r^*$ generated by $[n]_r$ and $[p]_r$, which is at least as large as the multiplicative order of $[n]_r$ in $\mathbb{Z}_r^*$, and by assumption (A4), this is larger than $4\,\text{len}(n)^2$.

Finally, directly by assumption (A5), we have $\ell > 2\,\text{len}(n)\lfloor t^{1/2} \rfloor$. □

That concludes the proof of Theorem 21.5.

EXERCISE 21.1. Show that if Conjecture 5.24 is true, then the value of r discovered in step 2 of Algorithm AKS satisfies $r = O(\text{len}(n)^2)$.

21.3 Notes

The algorithm presented here is due to Agrawal, Kayal, and Saxena [6].

If fast algorithms for integer and polynomial arithmetic are used, then using the analysis presented here, it is easy to see that the algorithm runs in time $O(\text{len}(n)^{10.5+o(1)})$. More generally, it is easy to see that the algorithm runs in time $O(r^{1.5+o(1)}\,\text{len}(n)^{3+o(1)})$, where r is the value determined in step 2 of the algorithm. In our analysis of the algorithm, we were able to obtain the bound $r = O(\text{len}(n)^5)$, leading to the running-time bound $O(\text{len}(n)^{10.5+o(1)})$. Using a result of Fouvry [37], one can show that $r = O(\text{len}(n)^3)$, leading to a running-time bound of $O(\text{len}(n)^{7.5+o(1)})$. Moreover, if Conjecture 5.24 on the density of Sophie Germain primes were true, then one could show that $r = O(\text{len}(n)^2)$ (see Exercise 21.1), which would lead to a running-time bound of $O(\text{len}(n)^{6+o(1)})$. This running-time bound can be achieved rigorously by a different algorithm, due to Lenstra and Pomerance [62].

Prior to this algorithm, the fastest deterministic, rigorously proved primality test was one introduced by Adleman, Pomerance, and Rumely [5], called the **Jacobi sum test**, which runs in time

$$O(\text{len}(n)^{c\,\text{len}(\text{len}(\text{len}(n)))})$$

for some constant c. Note that for numbers n with less than 2^{256} bits, the value of $\text{len}(\text{len}(\text{len}(n)))$ is at most 8, and so this algorithm runs in time $O(\text{len}(n)^{8c})$ for any n that one could ever actually write down.

We also mention the earlier work of Adleman and Huang [3], who gave a probabilistic algorithm whose output is always correct, and which runs in expected polynomial time (i.e., a *Las Vegas* algorithm, in the parlance of §9.7).