Introduction to Algorithms (CS 482)	Cornell University
Instructor: Bobby Kleinberg	Lecture Notes, 25 April 2008

The Miller-Rabin Randomized Primality Test

1 Introduction

Primality testing is an important algorithmic problem. In addition to being a fundamental mathematical question, the problem of how to determine whether a given number is prime has tremendous practical importance. Every time someone uses the RSA public-key cryptosystem, they need to generate a private key consisting of two large prime numbers and a public key consisting of their product. To do this, one needs to be able to check rapidly whether a number is prime.

The simplest algorithm to test whether n is prime is trial division: for $k = 2, 3, \ldots, \lfloor\sqrt{n}\rfloor$ test whether $n \equiv 0 \pmod{k}$. This runs in time $O(\sqrt{n}\log^2(n))$, but this running time is *exponential* in the input size since the input represents n as a binary number with $\lceil\log_2(n)\rceil$ digits. (A good public key these days relies on using prime numbers with at least 2^{250} binary digits; testing whether such a number is prime using trial division would require at least 2^{125} operations.)

In 1980, Michael Rabin discovered a *randomized* polynomial-time algorithm to test whether a number is prime. It is called the Miller-Rabin primality test because it is closely related to a deterministic algorithm studied by Gary Miller in 1976. This is still the most practical known primality testing algorithm, and is widely used in software libraries that rely on RSA encryption, e.g. OpenSSL.

2 Randomized algorithms

What does it mean to say that there is a randomized polynomial-time algorithm to solve a problem? Here are some definitions to make this notion precise.

Definition 1 (randomized algorithm, RP, coRP, BPP)**.** A randomized algorithm for a language L is an algorithm $A(x, r)$ which receives an input string x and a random string r, and attempts to output 1 if $x \in L$, 0 if $x \notin L$.

A language L is in **RP** if there exists a randomized algorithm $A(x, r)$ which runs in time polynomial in $|x|$ and satisfies:

- If $x \in L$, $\Pr(A(x, r) = 1) \geq 1/2$, when r is randomly sampled from the uniform distribution on $\{0, 1\}^{|r|}$.
- If $x \notin L$, $A(x, r) = 0$ for every r.

The complexity class **coRP** is defined in the same way except we replace the two conditions with:

- If $x \in L$, $A(x, r) = 1$ for every r.
- If $x \notin L$, $\Pr(A(x, r) = 0) \geq 1/2$.

(Equivalently, we could just say L belongs to **coRP** if its complement belongs to **RP**.) The complexity class **BPP** is defined in the same way except we replace the two conditions with:

- If $x \in L$, $\Pr(A(x, r) = 1) \geq 2/3$.
- If $x \notin L$, $\Pr(A(x, r) = 0) \geq 2/3$.

Theorem 1. *Definition 1 defines the same complexity classes if we change the constant* $1/2$ *to any constant strictly less than* 1*, or if we change the constant* $2/3$ *to any constant strictly between* $1/2$ *and* 1*.*

In particular, this means that if a language is in any of these complexity classes, there is a randomized polynomial-time algorithm $A(x, r)$ such that for *every* input x, $A(x, r)$ outputs the correct answer with probability at least $1 - 2^{-1000}$. So when we discover an efficient randomized algorithm for a problem, it is reasonable to consider that problem to be solved for all practical purposes.

The main theorem in this lecture is:

Theorem 2. PRIMES *is in* **coRP**.

In other words, there is a randomized test which always outputs "prime" if its input is prime, and which outputs "composite" with probability at least $1/2$ if its input is composite. However the algorithm may sometimes output "prime" when its input is actually composite.

In 2002, Agrawal, Kayal, and Saxena discovered a *deterministic* polynomial-time primality test. In other words, they proved PRIMES is in **P**. While this is a great algorithmic discovery, the Miller-Rabin algorithm is still the most widely used primality testing algorithm (and will probably remain so) because its running time is much faster.

3 Fermat's little theorem, the Fermat test, and Carmichael numbers

Theorem 3 (Fermat's little theorem)**.** *The number* n *is prime if and only if the congruence*

$$x^{n-1} \equiv 1 \pmod{n}$$

is satisfied for every integer x *between* 0 *and* n*.*

We will prove the theorem in a series of steps, beginning with:

Lemma 4. *If p is prime, then every pair of integers a, b satisfies*

$$(a+b)^p \equiv a^p + b^p \pmod{p}.$$

Proof. By the binomial theorem,

$$(a+b)^p = \sum_{k=0}^{p} \binom{p}{k} a^k b^{p-k}.$$

Every term in the sum is divisible by p except the $k = 0$ and $k = p$ terms. □

Proposition 5. *If A is a subset of the integers which is closed under addition and subtraction, then A is equal to $d\mathbb{Z}$, the set of all multiples of d, for some integer d.*

Proof. If $A = \{0\}$ then $d = 0$ and we are done. Otherwise, let d be the absolute value of the smallest non-zero element of A. The set A contains all multiples of d, since it contains $\{\pm d\}$ and is closed under addition and subtraction. Furthermore, A cannot contain any element x which is not divisible by d, since then we could subtract the nearest multiple of d to obtain a non-zero element of A whose absolute value is less than d. □

Proof of Fermat's little theorem. If n is not prime then it has a divisor $d > 1$. The number d^{n-1} is divisible by d so it is not equal to 1 mod n.

If n is prime, let A be the set of integers x which satisfy $x^n \equiv x \pmod{n}$. This set contains $x = 1$, and it is closed under addition and subtraction, by Lemma 4. Hence every integer x belongs to A.

Now let x by any integer not divisible by n. The fact that $x \in A$ means that $n \mid x^n - x = x(x^{n-1} - 1)$. Since n is prime and x is indivisible by n, this implies $n \mid x^{n-1} - 1$, i.e. $x^{n-1} \equiv 1 \pmod{n}$. □

Definition 2. Let n be a composite number. If $n \nmid x$ and $x^{n-1} \not\equiv 1 \pmod{n}$, we say that x is a *Fermat witness* for n. If $x^{n-1} \equiv 1 \pmod{n}$ we say x is a *Fermat liar* for n.

Figure 1 describes a primality testing algorithm based on Fermat's little theorem. The idea of the algorithm is simple: pick a positive integers $x < n$ and checking whether x is a Fermat witness. If so, then output "composite." Otherwise output "prime." To determine whether x is a Fermat witness for n, one needs to compute x^{n-1} mod n; the obvious way of doing this requires $n-2$ iterations of mod-n multiplication. But using the binary expansion of $n-1$ and repeated squaring, we can reduce this to $O(\log n)$ multiplication operations. For example, if $n = 23$ then $n-1 = 22 = 16+4+2$ so

$$x^{22} = x^{16}x^4x^2 = (((x^2)^2)^2)^2 \cdot (x^2)^2 \cdot x^2$$

FermatTest(n)

Choose $x \in \{1, 2, \ldots, n-1\}$ uniformly at random.
If $x^{n-1} \not\equiv 1 \pmod{n}$, return **composite;**
Else return **probably prime.**

Figure 1: The Fermat primality test.

and this describes an efficient algorithm for raising any integer to the 22nd power.

If n is prime, the Fermat primality test will always output "probably prime." But if n is composite, the algorithm will not output "composite" unless it randomly picks a Fermat witness for n. How hard is it to find a Fermat witness? Any proper divisor of n will do, but there may be very few of these. (For example, if $n = pq$ and p, q are distinct primes, the only two proper divisors of n are p and q.) But for most composite numbers, Fermat witnesses are much more prevalent. The next series of lemmas explains why this is so.

Lemma 6. *Let a, b be any two integers and let $d = \gcd(a, b)$. The set $a\mathbb{Z} + b\mathbb{Z} = \{ar + bs : r, s \in \mathbb{Z}\}$ is equal to $d\mathbb{Z}$ where $d = \gcd(a, b)$.*

Proof. The set $a\mathbb{Z} + b\mathbb{Z}$ is closed under addition and subtraction, so $a\mathbb{Z} + b\mathbb{Z} = c\mathbb{Z}$ for some integer c. If $d = \gcd(a, b)$ then every element of $a\mathbb{Z} + b\mathbb{Z}$ is divisible by d, so $d \mid c$. But a and b are both elements of $c\mathbb{Z}$, i.e. they are both divisible by c. This means c is a common divisor of a and b, so $c \mid d$. It follows that $c = d$. □

Lemma 7. *If $\gcd(a, n) = 1$ then there is an integer a^{-1} such that $a \cdot a^{-1} \equiv 1 \pmod{n}$.*

Proof. By Lemma 6, the set $a\mathbb{Z} + n\mathbb{Z}$ is equal to $\mathbb{Z}$, the set of all integers. In particular, this means there are integers r, s such that $ar + ns = 1$. This implies that $a \cdot r \equiv 1 \pmod{n}$, as desired. □

Lemma 8. *If b, c, n are positive integers such that $\gcd(c, n) = 1$ and the congruence $x^b \equiv c \pmod{n}$ has $k > 0$ solutions, then the congruence $x^b \equiv 1 \pmod{n}$ also has k solutions.*

Proof. Let x_0 be a solution of $x^b \equiv c \pmod{n}$. We must have $\gcd(x_0, n) = 1$, since otherwise $\gcd(x_0^b, n) = \gcd(c, n)$ would be greater than 1, contradicting our hypothesis. Lemma 7 now says that there is a number x_0^{-1} such that $x_0 \cdot x_0^{-1} \equiv 1 \pmod{n}$. A one-to-one correspondence between the solution sets of $x^b \equiv c \pmod{n}$ and of $x^b \equiv 1 \pmod{n}$ is given by the mapping $y \mapsto y \cdot x_0^{-1}$. □

Corollary 9. *If a composite number n has at least one Fermat witness x such that $\gcd(x, n) = 1$, then at least half of the elements of $1, 2, \ldots, n-1$ are Fermat witnesses for n.*

Proof. If $\gcd(x, n) = 1$ and x is a Fermat witness for n, then $x^{n-1} \equiv c \pmod{n}$ for some $c \neq 1$ satisfying $\gcd(c, n) = 1$. Now we can use Lemma 8 to show that there are at least as many Fermat witnesses as Fermat liars. □

Definition 3. An odd composite number n is a *Carmichael number* if every x satisfying $\gcd(x, n) = 1$ is a Fermat liar for n.

So far we have established that the Fermat test FermatTest(n) always outputs "prime" when n is prime, and that it outputs "composite" with probability at least $1/2$ when n is an odd composite number but not a Carmichael number. Obviously, it is easy to test whether an even number is prime. But we still don't have a good algorithm for distinguishing Carmichael numbers from prime numbers. The Miller-Rabin test is a more sophisticated version of the Fermat test which accomplishes this.

4 The Miller-Rabin test

So far, we know of two ways to prove that a number n is composite:

1. Exhibit a factorization $n = ab$, where $a, b > 1$.
2. Exhibit a Fermat witness for n, i.e. a number x satisfying $x^{n-1} \not\equiv 1 \pmod{n}$.

The Miller-Rabin test is based on a third way to prove that a number is composite.

3. Exhibit a "fake square root of 1 mod n," i.e. a number x satisfying $x^2 \equiv 1 \pmod{n}$ but $x \not\equiv \pm 1 \pmod{n}$.

The following lemma explains why this is a satisfactory proof of compositeness.

Lemma 10. *If x, n are positive integers such that $x^2 \equiv 1 \pmod{n}$ but $x \not\equiv \pm 1 \pmod{n}$, then n is composite.*

Proof. The hypotheses of the lemma imply that n is a divisor of $x^2 - 1 = (x+1)(x-1)$, but n divides neither $x + 1$ nor $x - 1$. This is impossible when n is prime. □

Later on, we will need the following generalization of Lemma 10.

Lemma 11. *If p is prime, then for any $k > 0$ the number of $x \in \{1, 2, \ldots, p-1\}$ satisfying $x^k \equiv 1 \pmod{p}$ is at most k.*

Proof. We will prove, more generally, that for any nonzero polynomial

$$P(x) = a_0 + a_1 x + \ldots + a_k x^k,$$

the number of $x \in \{1, 2, \ldots, p-1\}$ satisfying $P(x) \equiv 0 \pmod{p}$ is at most k. The proof is by induction on k, the base case $k = 0$ being trivial. Otherwise, suppose a

MillerRabin(n)

If $n > 2$ and n is even, return **composite.**
/* Factor $n-1$ as $2^s t$ where t is odd. */
$s \leftarrow 0$
$t \leftarrow n-1$
while t is even
    $s \leftarrow s+1$
    $t \leftarrow t/2$
end /* Done. $n-1 = 2^s t$. */
Choose $x \in \{1, 2, \ldots, n-1\}$ uniformly at random.
Compute each of the numbers $x^t, x^{2t}, x^{4t}, \ldots, x^{2^s t} = x^{n-1} \bmod n$.
If $x^{n-1} \not\equiv 1 \pmod{n}$, return **composite.**
for $i = 1, 2, \ldots, s$
    If $x^{2^i t} \equiv 1 \pmod{n}$ and $x^{2^{i-1} t} \not\equiv \pm 1 \pmod{n}$, return **composite.**
end /* Done checking for fake square roots. */
Return **probably prime.**

Figure 2: The Miller-Rabin primality test.

satisfies $P(a) \equiv 0 \pmod{p}$. We may write $P(x) = (x-a)Q(x) + c$, where $Q(x)$ is a polynomial of degree $k-1$ with integer coefficients. The congruence $P(a) \equiv 0 \pmod{p}$ implies that c is divisible by p. If b satisfies $P(b) \equiv 0 \pmod{p}$ but $Q(b) \not\equiv 0 \pmod{p}$ then p is a divisor of $(b-a)Q(b)$ but not of $Q(b)$, hence $b \equiv a \pmod{p}$. It follows that every $b \in \{1, 2, \ldots, p-1\}$ satisfying $P(b) \equiv 0 \pmod{p}$ satisfies either $b = a$ or $Q(b) \equiv 0 \pmod{p}$. By the induction hypothesis, at most $k-1$ elements of $\{1, 2, \ldots, p-1\}$ satisfy the second congruence. □

The Miller-Rabin test is shown in Figure 2. The idea of the test is to pick a random number x in $\{1, 2, \ldots, n-1\}$ and use it to try finding either a Fermat witness or a fake square root of 1 mod n.

Why does the Miller-Rabin test work? We have seen that when n is prime, the test always outputs "probably prime." When n is composite but is not a Carmichael number, we have seen that it outputs "composite" with probability at least $1/2$.

What if n is a Carmichael number? The workings of the Miller-Rabin test in this case can best be understood in terms of the *Chinese Remainder Theorem.*

Theorem 12 (Chinese Remainder Theorem)**.** *Let $n_1, n_2, \ldots, n_k$ be numbers, no two of which have a common factor. For any numbers $a_1, a_2, \ldots, a_k$, the system of con-*

gruences

$$\begin{aligned} x &\equiv a_1 \pmod{n_1} \\ x &\equiv a_2 \pmod{n_2} \\ &\vdots \\ x &\equiv a_k \pmod{n_k} \end{aligned}$$

has a solution. Any two solutions x_1, x_2 are congruent mod $n_1 n_2 \dots n_k$.

Proof. Let m_i denote the product of all elements of the set $\{n_1, n_2, \dots, n_k\}$ other than n_i. Note that $\gcd(m_i, n_i) = 1$ so Lemma 7 implies that there is a number r_i such that $m_i r_i \equiv 1 \pmod{n_i}$. Now let $x = \sum_{i=1}^{k} a_i m_i r_i$ and check that x satisfies the given system of congruences. If x_1, x_2 both satisfy the given system of congruences, then $x_1 - x_2$ is divisible by each of $n_1, n_2, \dots, n_k$. As these numbers have no common factors, we may conclude that $x_1 - x_2$ is divisible by $n_1 n_2 \dots n_k$. □

Corollary 13. *Let $n = p_1^{b_1} p_2^{b_2} \dots p_k^{b_k}$. Numbers $x \in \{1, 2, \dots, n-1\}$ which are relatively prime to n (i.e. satisfy $\gcd(x, n) = 1$) are in one-to-one correspondence with ordered k-tuples $(x_1, x_2, \dots, x_k)$ such that $1 \le x_i \le p_i^{b_i} - 1$ and $p_i \nmid x_i$.*

Now let's take an example of a Carmichael number and see what happens when we run the Miller-Rabin primality test. The smallest Carmichael number is $n = 561 = 3 \cdot 11 \cdot 17$. We have $n - 1 = 560 = 2^4 \cdot 35$. Let's take a random number, e.g. $x = 245$, and see what happens when we run the algorithm.

Number	mod 3	mod 11	mod 17
x	-1	3	7
x^{35}	-1	1	3
x^{70}	1	1	9
x^{140}	1	1	-4
x^{280}	1	1	-1
x^{560}	1	1	1

As we read down the numbers in each column, if any entry is equal to 1, then all subsequent entries in that column are also equal to 1. Define the "crossover row" for column j to be the row in which the number 1 first appears. In order for the Miller-Rabin test to output "probably prime", the crossovers must be *synchronized*, i.e. the crossover row must be the same in each column. We will see that this is an improbable event.

To begin with, we need the following description of Carmichael numbers.

Lemma 14. *If n is a Carmichael number, then n has at least three distinct prime factor and is not divisible by the square of any prime.*

Proof. First suppose that p is prime and $p^2 \mid n$. Write $n = p^k q$ where $k > 1$ and $p \nmid q$. Using the Chinese Remainder Theorem we may find a number x such that $x \equiv p+1 \pmod{p^k}$ and $x \equiv 1 \pmod{q}$. We claim that $(x, n) = 1$ and x is a Fermat witness for n. The fact that $(x, n) = 1$ is obvious: by construction, x has no common factors with p^k or with q, and $n = p^k q$. To see that $x^{n-1} \not\equiv 1 \pmod{n}$, it suffices to prove that $x^{n-1} \not\equiv 1 \pmod{p^2}$, or equivalently $x^n \not\equiv x \pmod{p^2}$. The formula $(p+1)^p = \sum_{k=0}^{p} \binom{p}{k} p^k$ implies $(p+1)^p \equiv 1 \pmod{p^2}$, which in turn implies $x^n \equiv 1 \pmod{p^2}$ and, which establishes that $x^n \not\equiv x \pmod{p^2}$ as claimed.

It remains to prove that when p, q are distinct odd primes, their product $n = pq$ is not a Carmichael number. Assume without loss of generality that $p < q$. Lemma 11 ensures that there we can choose $x \in \{1, 2, \ldots, q-1\}$ such that congruence $x^{p-1} \not\equiv 1 \pmod{q}$. We claim x is a Fermat witness for n. To prove this, observe that

$$x^{n-1} = x^{pq-1} = x^{pq-p} x^{p-1} = (x^p)^{q-1} x^{p-1} \equiv x^{p-1} \not\equiv 1 \pmod{q}.$$

Since q is a divisor of n, it follows that $x^{n-1} \not\equiv 1 \pmod{n}$ as claimed. □

Theorem 15. *If n is a Carmichael number, then* MillerRabin*(n) outputs "composite" with probability at least* $3/4$.

Proof. Suppose n is a Carmichael number. By analogy with the definitions of "Fermat witness" and "Fermat liar", let us call a number x a "Miller-Rabin witness" (or MR-witness) for n if the algorithm MillerRabin(n) outputs "composite" when x is the element of $\{1, 2, \ldots, n-1\}$ randomly chosen by the algorithm; otherwise we call x a "Miller-Rabin liar" (or MR-liar) for n. We must prove that there are at most $\frac{n-1}{4}$ MR-liars.

Write $n-1 = 2^s t$ where t is odd. Partition $\{1, 2, \ldots, n-1\}$ into sets $X, Y, Z_1, Z_2, \ldots, Z_s$ defined as follows:

- $x \in X$ if x shares a common divisor with n;
- $x \in Y$ if $x^t \equiv 1 \pmod{n}$.
- $x \in Z_j$ $(1 \le j \le s)$ if $x^{2^j t} \equiv 1 \pmod{n}$ but $x^{2^{j-1} t} \not\equiv 1 \pmod{n}$.

It is clear that X does not contain any MR-liars. We claim that $|Y| \le \frac{n-1}{8}$ and that each set Z_j contains at most $|Z_j|/7$ MR-liars. Assuming this claim, the total number of MR-liars is at most

$$|Y| + \frac{n-1-|Y|}{7} = \frac{6|Y|}{7} + \frac{n-1}{7} \le \frac{6(n-1)}{7 \cdot 8} + \frac{n-1}{7} = \frac{n-1}{4},$$

which will establish the theorem.

Suppose $n = p_1 p_2 \ldots p_k$ is the prime factorization of n; recall from Lemma 14 that $k \ge 3$ and that the primes $p_1, \ldots, p_k$ are all distinct. According to the Chinese Remainder Theorem, choosing a number $x \in \{1, 2, \ldots, n-1\} \setminus X$ is equivalent to

choosing numbers $x_i \in \{1, 2, \ldots, p_i - 1\}$ for each $i = 1, 2, \ldots, k$. The number x belongs to Y if and only if $x_i^t \equiv 1 \pmod{p_i}$ for every i. We claim that at most half of the elements $y \in \{1, 2, \ldots, p_i - 1\}$ satisfy $y^t \equiv 1 \pmod{p_i}$. Let A_i be the set of all integers u such that every $y \in \{1, 2, \ldots, p_i - 1\}$ satisfies $y^u \equiv 1 \pmod{p_i}$. This set is closed under addition and subtraction, hence it is equal to $d\mathbb{Z}$ for some integer d. Moreover, Fermat's Little Theorem ensures that $p_i - 1 \in A_i$ and Lemma 11 ensures that A_i does not contain any number between 0 and $p_i - 1$. This implies that A_i is the set of all multiples of $p_i - 1$; in particular, every element of A_i is even and this means that $t \notin A_i$. Thus there is at least one $y \in \{1, 2, \ldots, p_i - 1\}$ satisfying $y^t \not\equiv 1 \pmod{p_i}$, and by Lemma 8 this means that at most half of the elements of $\{1, 2, \ldots, p_i - 1\}$ satisfy $y^t \equiv 1 \pmod{p_i}$. When we pick elements x_i in $\{1, 2, \ldots, p_i - 1\}$ uniformly at random, for each $i = 1, 2, \ldots, k$, the probability that each of these numbers x_i satisfies $x_i^t \equiv 1 \pmod{p_i}$ is at most $(1/2)^k$, which is less than or equal to $1/8$ since $k \geq 3$. Thus the probability that a random $x \in \{1, 2, \ldots, n-1\} \setminus X$ belongs to Y is at most $1/8$, which establishes that $|Y| \leq \frac{n-1}{8}$.

We turn now to showing that each set Z_j contains at most $|Z_j|/7$ MR-liars. As before, an element $x \in Z_j$ may be represented by a k-tuple of numbers $x_i \in \{1, 2, \ldots, p_i - 1\}$. The fact that $x \in Z_j$ means that $x_i^{2^j t} \equiv 1 \pmod{p_i}$ for each value of i. Note that this implies $x_i^{2^{j-1} t} \equiv \pm 1 \pmod{p_i}$ for each i, so we can associate to each $x \in Z_j$ a sequence of k +/- signs, with the i-th sign being + or - according to whether $x_i^{2^{j-1} t} \equiv +1$ or $-1 \pmod{p_i}$. (The sign sequence is never $(+, +, \ldots, +)$, since this would imply that $x \in Z_\ell$ for some $\ell < j$.) An element $x \in Z_j$ is a MR-liar if and only if its sign sequence is $(-, -, \ldots, -)$. Let us assume there is at least one MR-liar in Z_j; otherwise there is nothing to prove. If Z_j contains a MR-liar, this implies that each of the congruences $y^{2^{j-1} t} \equiv -1 \pmod{p_i}$ $(1 \leq i \leq k)$ has a solution y_i. Let w_i be an element of $\{1, 2, \ldots, n-1\}$ which satisfies $w_i \equiv y_i \pmod{p_i}$ and $w_i \equiv 1 \pmod{p_{i'}}$ for every $i' \neq i$. If σ, σ' are two sign sequences which differ only by flipping the i-th sign in the sequence, then multiplication by w_i defines a one-to-one correspondence between the set of elements of Z_j with sign sequence σ and those with sign sequence σ'. It follows that each of the $2^k - 1$ possible sign sequences defines a subset of Z_j whose cardinality is exactly $|Z_j|/(2^k - 1)$. In particular, the set of MR-liars is defined by the sign sequence $(-, -, \ldots, -)$ and consequently has cardinality $|Z_j|/(2^k - 1) \leq |Z_j|/7$. □

Remark 16. Above, we proved that the error probability of MillerRabin(n) is at most $1/4$ when n is a Carmichael number. In fact, the error probability is bounded above by $1/4$ even when n is not a Carmichael number, though the material in Section 3 established only the weaker bound $1/2$.

Remark 17. The running time of MillerRabin(n) is $O(\log^3 n)$. To see this, recall our earlier observation that it is possible to compute $x^t \pmod{n}$ using $O(\log n)$ mod-n multiplication operations. (Each mod-n multiplication takes time $O(\log^2 n)$ using the naive algorithms for integer multiplication and division.) Once we have computed

$x^t \pmod n$, the remaining numbers $x^{2t}, x^{4t}, \ldots, x^{2^s t} \pmod n$ may be obtained by $s \leq \log_2(n)$ iterations of repeated squaring mod n, which again entails only $O(\log n)$ mod-n multiplication operations. All the remaining operations in the Miller-Rabin algorithm require much less running time.

Remark 18. Miller proved that if one assumes the Extended Riemann Hypothesis (a number-theoretic conjecture generally believed to be true), then for every composite number n the set $\{1, 2, \ldots, 2\ln^2(n)\}$ contains a MR-witness for n. Thus, assuming the Extended Riemann Hypothesis, there is a deterministic algorithm to test primality in time $O(\log^5 n)$.

Without assuming the Extended Riemann Hypothesis, a deterministic primality test with running time $O((\log n)^{O(\log\log\log n)})$ was discovered in 1983 by Adleman, Pomerance, and Rumely.

5 Epilogue: The AKS primality test

For almost 20 years no one made significant further progress toward proving PRIMES is in **P**. Then in 2002, Agrawal, Kayal, and Saxena proved that PRIMES is in **P** by discovering a completely different way to prove a number n is composite without using Fermat witnesses or fake square roots of 1. Their algorithm is based on the following theorem.

Theorem 19. *Let $n \geq 2$ and $a \geq 0$ be integers. If n is prime then the polynomials $P(x) = (x-a)^n$ and $Q(x) = x^n - a$ are congruent mod n. If n is composite and $\gcd(a, n) = 1$, then $P(x)$ and $Q(x)$ are not congruent mod n.*

In general, comparing $P(x)$ and $Q(x)$ mod n requires exponential time because it requires enumerating all n coefficients of the polynomials, and n is exponential in the input size. Here, the AKS algorithm uses a clever trick to reduce the work of enumerating coefficients: instead of comparing $P(x)$ and $Q(x)$ mod n, they divide both polynomials by $x^r - 1$ (for some value of r which is only polynomial in $\log(n)$) and compare the remainders mod n. The remainders have only r coefficients, and can be efficiently computed without enumerating all the coefficients of $P(x)$ and $Q(x)$. Agrawal, Kayal, and Saxena proved that if n is composite, then one can always find a proof of compositeness by searching exhaustively through pairs of numbers a, r such that both a and r are bounded by some polynomial in $\log(n)$. Having found the proper pair a, r, one verifies that the polynomials $(x-a)^n \pmod{x^r - 1}$ and $x^n - a \pmod{x^r - 1}$ are not congruent mod n. Subsequent improvements to the AKS algorithm have brought the running time down to $O(\log^6 n)$, still much slower than the $O(\log^3 n)$ randomized test presented in this lecture.