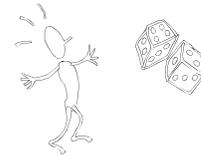


15-251

Great Theoretical Ideas in Computer Science

Randomness and Computation

Lecture 17, October 20, 2009



Super-simple and powerful idea

Drawing balls at random

You have a bucket with n balls
there are $n/100$ green balls (good)
the remaining are red (bad)

What is the probability of drawing a good ball
if you draw a random ball from the bucket?

Now if you draw balls from the bucket at random
(with replacement), how many draws until you
draw a good ball?

Drawing balls at random

You have a bucket with n balls
there are k green balls (good)
the remaining are red (bad)

Probability of getting a good ball
= k/n .

Expected number of draws until a good ball
= n/k .

even simpler idea...

Repeated experiments

Suppose you run a random experiment that fails with probability $\frac{1}{4}$ independent of the past.

What is the probability that you succeed in k steps?

= $1 -$ probability you fail in all k steps

= $1 - (\frac{1}{4})^k$

If probability of failure was at most δ , then probability of success at least once in k steps is at least $1 - \delta^k$

the following (trivial) question

Representing numbers

Question:

Given two numbers a and b , both $\leq n$, how long does it take to add them together?

- a) n
- b) \sqrt{n}
- c) $\log n$
- d) 2^n

Representing the number n takes $\log n$ bits

Representing numbers

Suppose I want to sell you (for \$1M) an algorithm that takes as input a number n , and factors them in $\approx \sqrt{n}$ time, should you accept my offer?

Factoring fast \Rightarrow breaking RSA!

Finally, remember this bit of algebra

The Fundamental theorem of Algebra

A root of a polynomial $p(x)$ is a value r , such that $p(r) = 0$.

If $p(x)$ is a polynomial of degree d , how many roots can it have?

At most d .

How to check your work...

Checking Our Work

Suppose we want to check $p(x)q(x) = r(x)$, where p , q and r are three polynomials.

$$(x-1)(x^3+x^2+x+1) = x^4-1$$

If the polynomials have degree n , requires n^2 mults by elementary school algorithms
-- or can do faster with fancy techniques like the Fast Fourier transform.

Can we check if $p(x)q(x) = r(x)$ more efficiently?

Idea: Evaluate on Random Inputs

Let $f(x) = p(x)q(x) - r(x)$. Is f zero everywhere?

Idea: Evaluate f on a *random* input z .

If we get nonzero $f(z)$, clearly f is not zero.

If we get $f(z) = 0$, this is (weak) evidence that f is zero everywhere.

If $f(x)$ is a degree $2n$ polynomial, it can only have $2n$ roots. We're unlikely to guess one of these by chance!

Equality checking by random evaluation

1. Say $S = \{1, 2, \dots, 4n\}$
2. Select value z uniformly at random from S .
3. Evaluate $f(z) = p(z)q(z) - r(z)$
4. If $f(z) = 0$, output "possibly equal" otherwise output "not equal"

Equality checking by random evaluation

What is the probability the algorithm outputs "not equal" when in fact $f = 0$?

Zero!

If $p(x)q(x) = r(x)$, always correct!

Equality checking by random evaluation

What is the probability the algorithm outputs "maybe equal" when in fact $f \neq 0$?

Let $A = \{z \mid z \text{ is a root of } f\}$.

Recall that $|A| \leq \text{degree of } f \leq 2n$.

Therefore: $P(\text{picked a root}) \leq 2n/4n = 1/2$

Equality checking by random evaluation

By repeating this procedure k times, we are “fooled” by the event

$$f(z_1) = f(z_2) = \dots = f(z_k) = 0$$

when actually $f(x) \neq 0$

with probability no bigger than

$$P(\text{picked root } k \text{ times}) \leq (1/2)^k$$



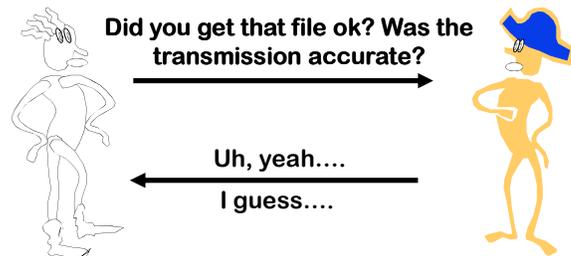
This idea can be used for testing equality of lots of different types of “functions”!

“Random Fingerprinting”

Find a small random “fingerprint” of a large object: e.g., the value $f(z)$ of a polynomial at a point z .

This fingerprint captures the essential information about the larger object: if two large objects are different, their fingerprints are usually different!

Earth has huge file X that she transferred to Moon. Moon gets Y .



Did you get that file ok? Was the transmission accurate?

Uh, yeah...
I guess....

How do we quickly check for accuracy? More soon...

How do you pick a random 1000-bit prime?

Picking A Random Prime

“Pick a random 1000-bit prime.”

Strategy:

- 1) Generate random 1000-bit number
- 2) Test for primality
[more on this later in the lecture]
- 3) Repeat until you find a prime.

How many retries until we succeed?

Recall the balls-from-bucket experiment?

If n = number of 1000-bit numbers = 2^{1000}

and k = number of primes in $0 \dots 2^{1000}-1$

then $E[\text{number of rounds}] = n/k$.

Question:

How many primes are there between 1 and n ?

(approximately...)



Legendre

Let $\pi(n)$ be the number of primes between 1 and n .

I wonder how fast $\pi(n)$ grows?

Conjecture [1790s]:

$$\lim_{n \rightarrow \infty} \frac{\pi(n)}{n/\ln n} = 1$$



Gauss

Their estimates

x	$\pi(x)$	Gauss' Li	Legendre	$x/(\log x - 1)$
1000	168	178	172	169
10000	1229	1246	1231	1218
100000	9592	9630	9588	9512
1000000	78498	78628	78534	78030
10000000	664579	664918	665138	661459
100000000	5761455	5762209	5769341	5740304
1000000000	50847534	50849235	50917519	50701542
10000000000	455052511	455055614	455743004	454011971



De la Vallée Poussin

Two independent proofs of the Prime Density Theorem [1896]:

$$\lim_{n \rightarrow \infty} \frac{\pi(n)}{n/\ln n} = 1$$



J-S Hadamard

The Prime Density Theorem

This theorem remains one of the celebrated achievements of number theory.

In fact, an even sharper conjecture remains one of the great open problems of mathematics!

The Riemann Hypothesis [1859]:



$$\lim_{n \rightarrow \infty} \frac{\pi(n) - n / \ln n}{\sqrt{n}} = 0$$

still unproven!

The Prime Density Theorem

$$\lim_{n \rightarrow \infty} \frac{\pi(n)}{n / \ln n} = 1$$



Slightly easier to show
 $\pi(n)/n \geq 1/(2 \log n)$.

In other words, at least (1/2B) of all B-bit numbers are prime

So, for this algo...

“Pick a random 1000-bit prime.”

Strategy:

- 1) Generate random 1000-bit number
- 2) Test for primality
[more on this later in the lecture]
- 3) Repeat until you find a prime.

These are the facts:

If we're picking 1000-bit numbers,

number of numbers is $n = 2^{1000}$

number of primes is $k \geq n/(2 \log n)$

Hence, expected number of trials before we get a prime number = $n/k \leq 2 \log n = 2000$.

Moral of the story

Picking a random B-bit prime is “almost as easy as”^{*} picking B random B-bit numbers.

Need to try at most 2 B times, in expectation.

(*Provided we can check for primality. More on this later.)

Earth has huge file X that she transferred to Moon. Moon gets Y.



Earth: X

Did you get that file ok? Was the transmission accurate?

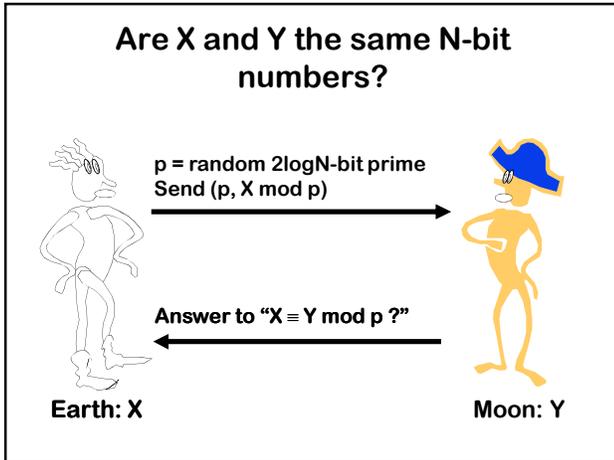
→



Moon: Y

←

Uh, yeah.



Why is this any good?

Easy case:
If $X = Y$, then $X \equiv Y \pmod{p}$

Why is this any good?

Harder case:
What if $X \neq Y$? We mess up if $p \mid (X-Y)$.

Define $Z = (X-Y)$. To mess up, p must divide Z .

Z is an N -bit number.
 $\Rightarrow Z$ is at most 2^N .

But each prime ≥ 2 .
Hence Z has at most N prime divisors.

Almost there...

$Z = (X-Y)$ has at most N prime divisors.

How many $2\log N$ -bit primes?

A random B -bit number has at least a $1/2B$ chance of being prime.

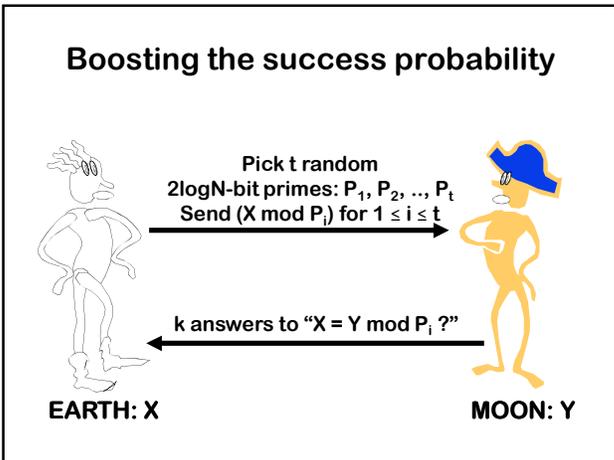
at least $2^{2\log N} / (2 * 2\log N) = N^2 / (4\log N) \gg 2N$ primes.

Only (at most) half of them divide Z .

Theorem:
Let X and Y be distinct N -bit numbers. Let p be a random $2\log N$ -bit prime.

Then
Prob $[X = Y \bmod p] < 1/2$

Earth-Moon protocol makes mistake with probability at most $1/2!$



Exponentially smaller error probability

If $X=Y$, always accept.

If $X \neq Y$,

Prob [$X = Y \bmod P_i$ for all i] $\leq (1/2)^t$

Picking A Random Prime

“Pick a random B-bit prime.”

Strategy:

- 1) Generate random B-bit numbers
- 2) Test each one for primality

How do we test if a number n is prime?

Primality Testing: Trial Division On Input n

Trial division up to \sqrt{n}

for $k = 2$ to \sqrt{n} do
if $k | n$ then
return “ n is not prime”
otherwise return “ n is prime”

about \sqrt{n} divisions

Trial division performs \sqrt{n} divisions
on input n .

Is that efficient?

For a 1000-bit number, this will take
about 2^{500} operations.

That's not very efficient at all!!!



More on efficiency and run-times
in a future lecture...

But so many cryptosystems,
like RSA and PGP, use fast
primality testing as part of
their subroutine to generate
a random n -bit prime!

What is the fast primality
testing algorithm that they
use?



There are fast *randomized*
algorithms to do primality
testing.



Miller-Rabin test



Solovay-Strassen test

If n is composite, how would you show it?

Give a non-trivial factor of n .

But, we don't know how to factor numbers fast.

We will use a *different* certificate of compositeness that does not require factoring.



simple idea #1

Recall that for prime p , $a \neq 0 \pmod p$:
Fermat Little Thm: $a^{p-1} = 1 \pmod p$.

Hence, $a^{(p-1)/2} = \pm 1$.

So if we could find some $a \neq 0 \pmod p$ such that $a^{(p-1)/2} \neq \pm 1$

$\Rightarrow p$ must not be prime.

$\text{Good}_n = \{ a \in \mathbb{Z}_n^* \mid a^{(n-1)/2} \neq \pm 1 \}$
(these prove that n is not prime)

$\text{Useless}_n = \{ a \in \mathbb{Z}_n^* \mid a^{(n-1)/2} = \pm 1 \}$
(these don't prove anything)

Theorem:
if Good_n is not empty, then Good_n contains at least half of \mathbb{Z}_n^* .



simple idea #2

Remember Lagrange's theorem:

If G is a group, and U is a subgroup then $|U|$ divides $|G|$.

In particular, if $U \neq G$ then $|U| \leq |G|/2$.

Proof

$\text{Good}_n = \{ a \in \mathbb{Z}_n^* \mid a^{(n-1)/2} \neq \pm 1 \}$
 $\text{Useless}_n = \{ a \in \mathbb{Z}_n^* \mid a^{(n-1)/2} = \pm 1 \}$

Fact 1: Useless_n is a subgroup of \mathbb{Z}_n^*

Fact 2: If H is a subgroup of G then $|H|$ divides $|G|$.

\Rightarrow If Good is not empty, then $|\text{Useless}| \leq |\mathbb{Z}_n^*|/2$

$\Rightarrow |\text{Good}| \geq |\mathbb{Z}_n^*|/2$

Randomized Primality Test

Let's suppose that $\text{Good}_n = \{ a \in \mathbb{Z}_n^* \mid a^{(n-1)/2} \neq \pm 1 \}$ contains at least half the elements of \mathbb{Z}_n^* .

Randomized Test:

For $i = 1$ to k :

 Pick random $a_i \in [2 \dots n-1]$;

 If $\text{GCD}(a_i, n) \neq 1$, Halt with "Composite";

 If $a_i^{(n-1)/2} \neq \pm 1$, Halt with "Composite";

Halt with "I think n is prime. I am only wrong $(1/2)^k$ fraction of times I think that n is prime."

Is Good_n non-empty for all primes n ?

Recall: $\text{Good}_n = \{ a \in \mathbb{Z}_n^* \mid a^{(n-1)/2} \neq \pm 1 \}$

Good_n may be empty even if n is not a prime.

A Carmichael number is a number n such that $a^{(n-1)/2} \equiv 1 \pmod{n}$ for all numbers a with $\text{gcd}(a,n)=1$.

Example: $n = 561 = 3 \cdot 11 \cdot 17$ (the smallest Carmichael number)

$$1105 = 5 \cdot 13 \cdot 17$$

$$1729 = 7 \cdot 13 \cdot 19$$

And there are many of them. For sufficiently large m , there are at least $m^{2/7}$ Carmichael numbers between 1 and m .

The saving grace

The randomized test fails only for Carmichael numbers.

But, there is an efficient way to test for Carmichael numbers.

Which gives an efficient algorithm for primality.

Randomized Primality Test

Let's suppose that Good_n contains at least half the elements of \mathbb{Z}_n^* .

Randomized Test:

If n is Carmichael, Halt with "Composite"

For $i = 1$ to k :

 Pick random $a_i \in [2 \dots n-1]$;

 If $\text{GCD}(a_i, n) \neq 1$, Halt with "Composite";

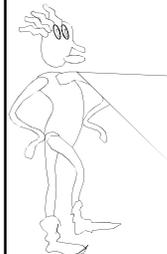
 If $a_i^{(n-1)/2} \equiv \pm 1 \pmod{n}$, Halt with "Composite";

Halt with "I think n is prime. I am only wrong $(\frac{1}{2})^k$ fraction of times I think that n is prime."

Primality Versus Factoring

Primality has a fast randomized algorithm.

Factoring is not known to have a fast algorithm. The fastest randomized algorithm currently known takes $\exp(O(n \log n \log n)^{1/3})$ operations on n -bit numbers.



number	digits	prize	factored
RSA-100	100		Apr. 1991
RSA-110	110		Apr. 1992
RSA-120	120		Jun. 1993
RSA-129	129	\$100	Apr. 1994
RSA-130	130		Apr. 10, 1996
RSA-140	140		Feb. 2, 1999
RSA-150	150		Apr. 16, 2004
RSA-155	155		Aug. 22, 1999
RSA-160	160		Apr. 1, 2003
RSA-200	200		May 9, 2005
RSA-576	174	\$10,000	Dec. 3, 2003
RSA-640	193	\$20,000	Nov 2, 2005
RSA-704	212	\$30,000	open
RSA-768	232	\$50,000	open
RSA-896	270	\$75,000	open
RSA-1024	309	\$100,000	open
RSA-1536	463	\$150,000	open
RSA-2048	617	\$200,000	open

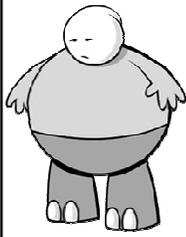
Google: RSA Challenge Numbers (the challenge is no longer active)

The techniques we've been discussing today are sometimes called "fingerprinting."

The idea is that a large object such as a string (or document, or function, or data structure...) is represented by a much smaller "fingerprint" using randomness.



If two objects have identical sets of fingerprints, they're likely the same object.



**Here's What
You Need to
Know...**

Primes

Prime number theorem
How to pick random primes

Fingerprinting

How to check if a polynomial
of degree d is zero
How to check if two n -bit strings
are identical

Primality

Fermat's Little Theorem
Algorithm for testing primality