

15-251
Great Ideas in
Theoretical Computer Science

Lecture 23:
Randomized Algorithms I

November 14th, 2017




Randomness is an essential tool in
modeling and analyzing nature.

It also plays a key role in **computer science.**

Randomness and Computer Science

Statistics via Sampling



Population: 300m **Random sample size:** 2000


Theorem:

Randomized Algorithms

Dimer Problem:

Given a region, in how many different ways can you tile it with 2x1 rectangles (dominoes)?

e.g.




→ 1024 tilings

Captures thermodynamic properties of matter.

- Fast randomized algs can approximately count.
- No fast deterministic alg known.

Distributed Computing



Nash Equilibria in Games

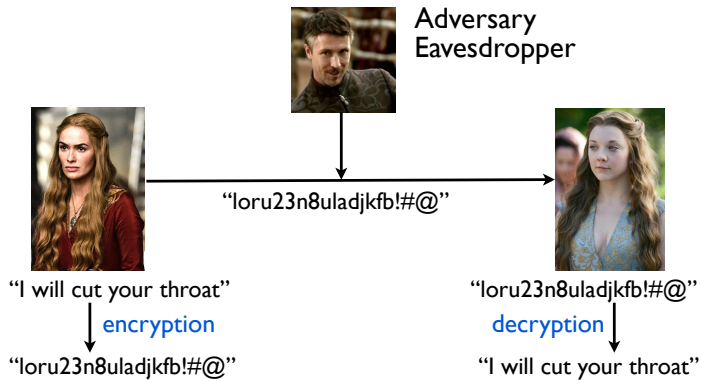
The Chicken Game



	Swerve	Straight
Swerve	1 1	0 2
Straight	2 0	-3 -3


Theorem (Nash):

Cryptography



Shannon:

Error-Correcting Codes



Alice


"bit.ly/vrxUBN"
noisy channel



Bob

Each symbol can be corrupted with a certain probability.
How can Alice still get the message across?

Communication Complexity



Want to check if the contents of two databases are exactly the same.

How many bits need to be communicated?

Interactive Proofs

Verifier



*poly-time
skeptical*


Prover



*omniscient
untrustworthy*

Can I convince you that I have proved $P \neq NP$ without revealing any information about the proof?

Quantum Computing



Probability Theory: The CS Approach

The Big Picture


The Non-CS Approach



The Big Picture

Real World \longrightarrow Mathematical Model

Flip a coin.



Ω = "sample space"
= set of all possible outcomes


$\Pr : \Omega \rightarrow [0, 1]$ prob. distribution

$$\sum_{\ell \in \Omega} \Pr[\ell] = 1$$

The Big Picture

Real World \longrightarrow Mathematical Model


Flip two coins.



The Big Picture


Real World \longrightarrow Mathematical Model

Flip a coin.
If it is Heads, throw
a 3-sided die.
If it is Tails, throw a
4-sided die.



The Big Picture

The CS Approach



The Big Picture

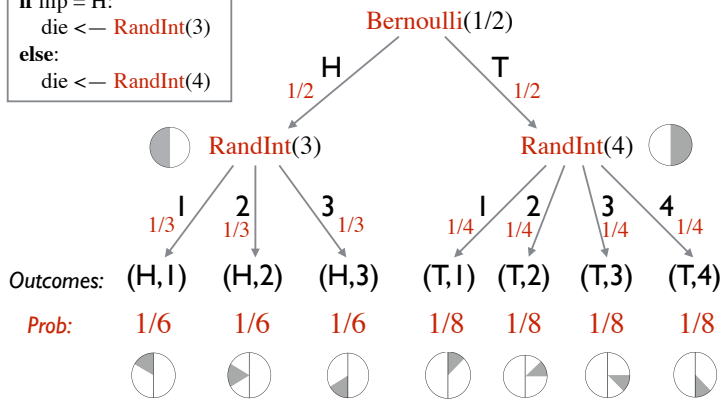
Real World → Code → Probability Tree

Flip a coin.
If it is Heads, throw
a 3-sided die.
If it is Tails, throw a
4-sided die.

```
flip ← Bernoulli(1/2)
if flip = 1: # i.e. Heads
  die ← RandInt(3)
else:
  die ← RandInt(4)
```

Probability Tree

```
flip ← Bernoulli(1/2)
if flip = H:
  die ← RandInt(3)
else:
  die ← RandInt(4)
```



What is a Random Variable?

A **random variable** is a variable in some randomized code (more accurately, the variable's value at the end of the execution) of type 'real number'.

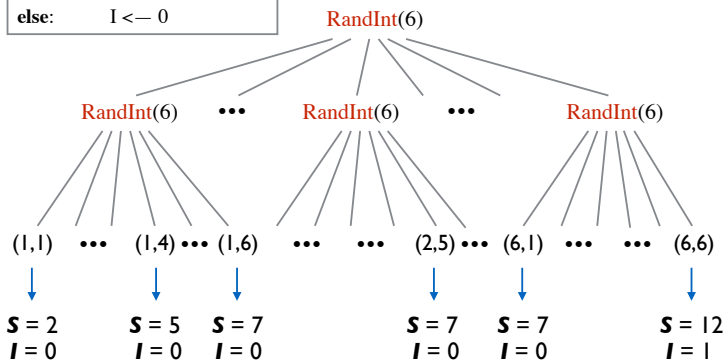
Example:

```
S ← RandInt(6) + RandInt(6)
if S = 12: I ← 1
else:     I ← 0
```

Random variables:

What is a Random Variable?

```
S ← RandInt(6) + RandInt(6)
if S = 12: I ← -1
else:     I ← 0
```



Markov's Inequality

A non-negative random variable X is rarely much bigger than its expectation $E[X]$.



Theorem:

New Topic:

Randomized Algorithms

Randomness and algorithms

How can randomness be used in computation?

Given some algorithm that solves a problem:

- (i) the input can be chosen randomly
average-case analysis
- (ii) the algorithm can make random choices
randomized algorithm

Which one will we focus on?

Randomness and algorithms

What is a randomized algorithm?

A *randomized algorithm* is an algorithm that is allowed to “**flip a coin**” (i.e., has access to random bits).

In 15-251:

A randomized algorithm is an algorithm that is allowed to call:

- `RandInt(n)` (we'll assume these take $O(1)$ time)
- `Bernoulli(p)`

Deterministic vs Randomized

Deterministic

```
def A(x):  
    y = 1  
    if(y == 0):  
        while(x > 0):  
            x = x - 1  
    return x+y
```

Randomized

```
def A(x):  
    y = Bernoulli(0.5)  
    if(y == 0):  
        while(x > 0):  
            x = x - 1  
    return x+y
```

For any fixed input (e.g. $x = 3$):

- | | |
|---------------------------|---------------------------|
| - the output | - the output |
| - the running time | - the running time |

Deterministic vs Randomized

A **deterministic algorithm** A computes $f : \Sigma^* \rightarrow \Sigma^*$ in time $T(n)$ means:

- **correctness:** $\forall x \in \Sigma^*, A(x) = f(x)$.
- **running time:** $\forall x \in \Sigma^*, \# \text{ steps } A(x) \text{ takes is } \leq T(|x|)$.

Note: we require **worst-case** guarantees for **correctness** and **run-time**.

Deterministic vs Randomized

A Try

A **randomized algorithm** A computes $f : \Sigma^* \rightarrow \Sigma^*$ in time $T(n)$ means:

- **correctness:** $\forall x \in \Sigma^*, \boxed{A(x)} = f(x)$.
- **running time:** $\forall x \in \Sigma^*, \boxed{\# \text{ steps } A(x) \text{ takes is}} \leq T(|x|)$.

these are random

Deterministic vs Randomized

A Try

A **randomized algorithm** A computes $f : \Sigma^* \rightarrow \Sigma^*$ in time $T(n)$ means:

- **correctness:** $\forall x \in \Sigma^*, \Pr[A(x) = f(x)] = 1$.
- **running time:** $\forall x \in \Sigma^*, \Pr[\# \text{ steps } A(x) \text{ takes is } \leq T(|x|)] = 1$

Is this interesting? No.

A randomized algorithm should gamble with either **correctness** or **run-time**.

Formal Definitions

Formal Definition: Deterministic

Let $f : \Sigma^* \rightarrow \Sigma^*$ be a computational problem.

We say that deterministic algorithm A computes f in time $T(n)$ if:

$$\forall x \in \Sigma^*, \quad A(x) = f(x)$$

$$\forall x \in \Sigma^*, \quad \# \text{ steps } A(x) \text{ takes is } \leq T(|x|).$$

Picture:



Deterministic:

Each input x induces a deterministic path.

Formal Definition: Monte Carlo


Let $f : \Sigma^* \rightarrow \Sigma^*$ be a computational problem.

We say that randomized algorithm A is a $T(n)$ -time **Monte Carlo algorithm** for f with ϵ error probability if:

$$\forall x \in \Sigma^*,$$

$$\forall x \in \Sigma^*,$$

Picture:



Monte Carlo:

Each input x induces a probability tree.

Formal Definition: Las Vegas


Let $f : \Sigma^* \rightarrow \Sigma^*$ be a computational problem.

We say that randomized algorithm A is a $T(n)$ -time **Las Vegas algorithm** for f if:

$$\forall x \in \Sigma^*,$$

$$\forall x \in \Sigma^*,$$

Picture:



Las Vegas:

Each input x induces a probability tree.

Examples

3 IMPORTANT PROBLEMS

Integer Factorization

Input: integer N
Output: a prime factor of N

isPrime

Input: integer N
Output: True if N is prime.

Generating a random n -bit prime

Input: integer n
Output: a random n -bit prime

Most crypto systems start like:

- pick two random n-bit primes P and Q.
- let $N = PQ$. (N is some kind of a "key")
- (more steps...)

We should be able to do **efficiently** the following:

- check if a given number is prime.
- generate a random prime.

We should **not** be able to do **efficiently** the following:

- given N, find P and Q. (the system is broken if we can do this!!!)

isPrime

```
def isPrime(N):  
    if (N < 2): return False  
    maxFactor = round(N**0.5)  
    for factor in range(2, maxFactor+1):  
        if (N % factor == 0): return False  
    return True
```

Problems:

isPrime

Amazing result from 2002:

There is a poly-time algorithm for isPrime.




Agrawal, Kayal, Saxena

However, best known implementation is $\sim O(n^6)$ time.
Not feasible when $n = 2048$.

isPrime

So that's **not** what we use in practice.

Everyone uses the **Miller-Rabin** algorithm (1975).



The running time is $\sim O(n^2)$.

Why is the previous result a breakthrough?

Generating a random prime

```
repeat:  
  let N be a random n-bit number  
  if isPrime(N): return N
```

Prime Number Theorem (informal):

About $1/n$ fraction of n-bit numbers are prime.

\implies expected run-time of the above algorithm:

No poly-time deterministic algorithm is known to generate an n-bit prime!!!
