

# CS265/CME309: Randomized Algorithms and Probabilistic Analysis

## Lecture #1: Computational Models, and the Schwartz-Zippel Randomized Polynomial Identity Test

Gregory Valiant\*

November 17, 2019

### 1 Introduction

Welcome to CS265/CME309!! This course will revolve around two intertwined themes: 1) Analyzing random structures, including a variety of models of natural randomized processes, including random walks, and random graphs/networks, and 2) Designing random structures and algorithms that solve problems we care about. By the end of this course, my hope is that you will be able to think precisely about randomness, and also appreciate that randomness can be a powerful (and somewhat mysterious) tool, if wielded correctly. As we will see, there are a number of basic problems for which extremely simple randomized algorithms outperform (both in theory and in practice) the best deterministic algorithms that we currently know.

### 2 Computational Model

During this course, we will discuss algorithms at a slightly higher level of abstraction, though it will be useful to begin with a formal definition of the model of randomized computation that we will be working with.

**Definition 2.1** A *randomized algorithm* is an algorithm that can be computed by a Turing machine (or random access machine), which has access to an infinite string of random bits. Equivalently, it is an algorithm that can be performed by a Turing machine that has a special instruction "flip-a-coin", which returns the outcome of an independent flip of a fair coin.

---

\*©2019, Gregory Valiant. Not to be sold, published, or distributed without the authors' consent.

We will never describe algorithms at the level of Turing machine instructions, though if you are ever uncertain whether or not an algorithm you have in mind is “allowed”, you can return to this definition. The above definition implies the following two important properties of randomized algorithms:

- Because the computation can depend on the randomness, the output of the algorithm, and the computation path itself, are random variables. As such, it makes sense to ask questions like “What is the probability the output is correct?” or “What is the standard deviation of the runtime of the algorithm”. For these questions, we are thinking of the computation path as being a random variable, defined in terms of the randomness of the random bits that the algorithm uses. In particular, we are thinking of the input to the algorithm as fixed—for example, if we are running a randomized primality testing algorithm to test if the input  $q = 1173$  is prime, we will regard  $q$  as a fixed parameter, and analyze the execution of the algorithm on input  $q$  in terms of the random coin flips. Frequently, we will be able to prove guarantees that hold for *all* inputs; namely, continuing the primality testing example, we will be able to say things like “for every input,  $q$ , with probability at least  $1 - 1/2^{100}$  over the random coin flips of the algorithm, the algorithm will correctly output whether or not  $q$  is prime.
- Because we are working with a model where the algorithm has access to coin-flips (which take one of two possible outcomes), the computation path can be described as a *binary* tree, where the branches correspond to the outcomes of the coin tosses. The topmost branch corresponds to whether the first coin flip landed heads or tails, the next level corresponds to the second coin-flip, etc. Hence if we have a randomized algorithm that terminates after flipping at most  $k$  coins, the probability of every output must be an integer multiple of  $1/2^k$ . [If this isn’t obvious to you, take 3 minutes and convince yourself of it now....I’ll add a pretty picture to these lecture notes at some point : ) ] We won’t be super pedantic about things, and if you want a randomized algorithm that executes command X with probability  $1/3$ , that’s fine (because one can simulate a coin that lands heads with probability  $1/3$  via a *sequence* of flips of a fair coin). But, if you are ever designing a randomized algorithm that uses funny transition probabilities, you can always come back to this model, and double-check whether what you are doing is actually feasible in this model.

There are two broad classes of randomized algorithms, referred to as “Las Vegas” algorithms and “Monte Carlo” algorithms:

- Las Vegas Algorithms: These algorithms always output the correct answer, though the runtime can be unbounded, though we require that it has finite expectation. (I.e. the output is deterministic, but the runtime is a random variable, with finite expectation.)
- Monte Carlo Algorithms: These algorithms have bounded runtime, though they can output the wrong answer. For algorithms that answer yes/no questions (e.g. “Is  $q$  prime?”), we group these into two subclasses, according to whether they have 1-sided or 2-sided error:

1. 1-sided error: If the true answer is “Yes” then the algorithm will output “Yes”, and if the true answer is “No”, the algorithm will output “No” with probability at least  $p > 0$ . [We could also swap the Yes/No and it would still be called 1-sided error....] For the complexity theorists out there, the class of problems that have polynomial-time 1-sided error Monte Carlo algorithms is known as “RP” (Randomized-Polynomial).
2. 2-sided error: The algorithm is correct with probability at least  $1/2 + \epsilon$  for some  $\epsilon > 0$ . The class of problems that have polynomial-time 2-sided error Monte Carlo algorithms is known as “BPP” (Bounded Error Probabilistic Polynomial).

One reason why we are alright with the above definitions, which only guarantee a slight bias towards being correct, is that we can increase the probability of correctness by running the algorithm multiple times. For example, given a 1-sided error Monte Carlo algorithm, if we run it  $t$  times, and it ever outputs “No” we can safely output “No”, and if it always outputs “Yes”, then we can output “Yes” and the probability we are incorrect is at most  $(1 - p)^t$ . [Make sure you understand why this is true!!]

For the 2-sided error algorithm, intuitively, the right approach is to repeat the algorithm  $t$  times, and output the majority answer. I have included the following calculation, as it is the sort of thing we will not spend much time on in class, though I expect that you will be able to perform such calculations on your own.

**Lemma 2.2** *If an algorithm is correct with probability  $1/2 + \epsilon$ , and we run it  $t$  times and output the majority answer, the probability we are correct is at least  $1 - e^{-2t\epsilon^2}$ .*

*Proof:*

$$\begin{aligned}
 \Pr[\text{majority is wrong}] &\leq \sum_{i=0}^{t/2} \binom{t}{i} \left(\frac{1}{2} + \epsilon\right)^i \left(\frac{1}{2} - \epsilon\right)^{t-i} \\
 &\leq \sum_{i=0}^{t/2} \binom{t}{i} \left(\frac{1}{2} + \epsilon\right)^{t/2} \left(\frac{1}{2} - \epsilon\right)^{t/2} \\
 &= \left(\frac{1 - 4\epsilon^2}{4}\right)^{t/2} \sum_{i=0}^{t/2} \binom{t}{i} \\
 &\leq \left(\frac{1 - 4\epsilon^2}{4}\right)^{t/2} 2^t = (1 - 4\epsilon^2)^{t/2} \leq e^{-2t\epsilon^2},
 \end{aligned}$$

where the last inequality is from the fact that for any  $x$ ,  $1 - x \leq e^{-x}$ . This trick of replacing  $1 - x$  with  $e^{-x}$  is a commonly used trick for getting rid of a tricky minus sign, and replacing it with a nice exponent. For small values of  $x$ , this approximation is quite accurate, because the derivative of  $1 - x$  and  $e^{-x}$  are equal at  $x = 0$ . ■

## 2.1 What Kind of Resource is Randomness?

One of the biggest open questions in complexity theory, is the question of whether randomness *really helps*, and what are the connections between the above classes of randomized algorithm. The following is one example of such a question that you should be able to answer after a bit of thinking.

**Question 1** Which is more powerful, Las Vegas Algorithms or Monte Carlo Algorithms? [Given a Las Vegas algorithm whose expected runtime is polynomial in the input length, can we turn it into a Monte Carlo Algorithm whose runtime is bounded by a polynomial in the input length? Given a Monte Carlo Algorithm with polynomial runtime, can we turn it into a Las Vegas algorithm with polynomial expected runtime?]

While there are many instances of problems for which there seem to be much simpler and faster randomized algorithms, we do not have a rigorous proof of whether or not there problems for which randomness is essential for any efficient algorithm. Phrased succinctly, we can ask, is the set of problems solvable via polynomial time deterministic algorithms equal the set of problems solvable via 2-sided error Monte Carlo algorithms (does  $P = BPP$ ?), does  $P = RP$ ?. One way of approaching these questions is to ask whether we can simply replace the random coin flips via a sequence of deterministic bits that “look kind of random”. For example, what if we run a randomized algorithm, but instead of flipping coins, we just use the bits of the binary representation of  $\pi$ ? Surely that can’t be too bad, no? We could also analyze how our analysis of randomization actually uses the randomness assumption, with an aim towards figuring out whether the algorithm would notice whether or not we fed it less-random bits.... These ideas fall under the broad umbrella of *de-randomization*, and are discussed in the graduate complexity course, taught by Li-Yang Tan or Omer Reingold, two of the world experts on these topics!

There is also a large body of work understanding how to actually obtain random bits, and what can be done with imperfect randomness. For example, if the coins aren’t quite unbiased or independent, are there ways of *distilling* or *extracting* the randomness? Given a large number of kind-of-sort-of-maybe random coins (that have some randomness, but maybe are correlated in unknown ways), can we produce a smaller number of almost perfectly random coins?

We won’t discuss these topics much more in this course. Instead, our focus will be on designing and analyzing algorithms, assuming access to perfect random coins.

## 3 Randomized Polynomial Identity Testing

The first randomized algorithm that we will see is an algorithm for deciding whether a multivariate polynomial is actually 0: Given  $P(x_1, x_2, \dots, x_n)$ , we want to know whether  $P = 0$ . This problem isn’t as simple as it might seem at first. For example, consider

$$P(x_1, x_2, x_3) = -x_3^2(x_1+x_2)(x_1-x_2) + ((x_1^2 + (1+x_2)(1-x_2))x_3)^2 - x_3^2 - x_3^2(x_1+x_2)(x_1-x_2))^2.$$

It turns out that, if you were to expand out this polynomial, all the terms cancel. A degree  $d$  polynomial in  $n$  variables could end up with  $\gg \binom{n}{d} \approx n^d$  monomials—so the naive algorithm of simply expanding out all the monomials and then seeing what cancels, could, in the worst case, take time exponential in the length of the input representation of  $P$ .

Before describing the Schwartz-Zippel efficient randomized algorithm, it is worth pointing out that the task of identifying whether a polynomial is 0 or not is, in fact, an *incredibly* useful primitive, largely because polynomials occur everywhere. Below are a few common use-cases:

- Graph analysis: Graphs can be represented via adjacency matrices, and the determinants of these matrices (or variant of them) convey information about the structure of these graphs (especially the question of whether the determinant is 0 or not). Determinants are just polynomials!! [One can use this approach to efficiently test whether a graph has a perfect matching....]
- Coding theory and pattern matching....
- Primality Checking: As we will see in a week or two, the polynomial  $P(z) := (1 + z)^n - 1 - z^n \pmod n$  is 0 if, and only if,  $n$  is prime. Working modulo  $n$  is a bit a twist that we won't discuss this lecture....

### 3.1 Schwartz-Zippel Polynomial Identity Test

This testing algorithm was published independently by Jack Schwartz and Richard Zippel in 1979 and should also be attributed to Richard DeMillo and Richard Lipton who published a version of it in 1978. Based on the fact that it was independently discovered at the same time by so many people, you might guess that it is fairly simple...

#### Algorithm 1

##### POLYNOMIAL IDENTITY TEST

Given a polynomial  $P(x_1, \dots, x_n)$ :

1. Fix a finite set,  $S$ .
2. Choose  $n$  values,  $r_1, \dots, r_n$  independently and uniformly at random from  $S$ .
3. Evaluate  $P(r_1, \dots, r_n)$ . If  $P(r_1, \dots, r_n)$  output “ $P = 0$ ” otherwise return “ $P \neq 0$ ”.

**Theorem 3.1** *If  $P = 0$ , then the algorithm will always output “ $P = 0$ ”. If  $P \neq 0$ , then the algorithm will be correct with probability at least  $1 - \frac{\text{degree}(P)}{|S|}$ , where  $\text{degree}(P)$  is the maximum sum of the degrees of terms in any monomial (e.g.  $\text{degree}(x_1^3 x_2^2) = 5$ .)*

*Proof:* The first part of the proof is trivial, since if  $P$  is 0, then it will evaluate to 0. Hence, for the remainder, it suffices to consider the case that  $P(x_1, \dots, x_n)$  is a degree  $d$  polynomial. The proof leverages the following basic fact about univariate polynomials that we all learned at some point: a degree  $d$  polynomial in one variable can have at most  $d$  roots. Given this fact, we will prove the theorem by induction on the number of variables,  $n$ :

Base case,  $n = 1$ : The only way the algorithm will output the wrong answer is if  $P(r_1) = 0$ , in which case, by definition,  $r_1$  is a root of  $P$ . Since  $P$  has degree  $d$ , by assumption, from our fact about degree  $d$  univariate polynomials,  $\Pr[P(r_1) = 0] \leq \frac{d}{|S|}$ , since at most  $d$  of the  $|S|$  elements of our set can be roots of  $P$ .

Induction step: Assuming the theorem holds for polynomials in  $\leq n - 1$  variables, we will consider  $P(x_1, \dots, x_n)$ , and try to reduce this case to an argument about the univariate polynomial  $P(x_1, r_2, \dots, r_n)$ , obtained after plugging in the values for  $x_2, \dots, x_n$ . To make this argument, we will leverage our induction hypothesis. Without loss of generality, assume that  $P(x_1, \dots, x_n) = x_1^k Q(x_2, \dots, x_n) + T(x_1, \dots, x_n)$  for some polynomials  $Q$  and  $T$ , such that  $Q \neq 0$ , and the maximum degree of  $x_1$  in  $T$  is  $< k$ . Namely, we have factored out the monomials in  $P$  for which  $x_1$  has the highest degree. [If  $x_1$  does not actually appear in  $P$ , then we would do this argument for a variable  $x_i$  that does appear....since  $P$  is nonzero such a variable exists.]

Consider the event that  $Q(r_2, \dots, r_n) = 0$ . By our inductive hypothesis,  $\Pr[Q(r_2, \dots, r_n) = 0] \leq \frac{d-k}{|S|}$ . If  $Q(r_2, \dots, r_n) \neq 0$ , then what happens to  $P$  after we plug in  $r_2, \dots, r_n$ ? Well, then  $P(x_1, r_2, \dots, r_n)$  is a degree  $k$  nonzero polynomial of  $x_1$ . Applying our fact about univariate polynomials, the probability that this evaluates to 0 when we plug in  $r_1$  for  $x_1$  is at most  $\frac{k}{|S|}$ . Hence the probability the algorithm successfully evaluates to something other than 0 is at least

$$\left(1 - \frac{d-k}{|S|}\right) \left(1 - \frac{k}{|S|}\right) = 1 - \frac{d}{|S|} + \frac{d(d-k)}{|S|^2} \geq 1 - \frac{d}{|S|}.$$

■

### 3.2 Discussion

Given that the Schwartz-Zippel algorithm is the simplest thing imaginable, you might be tempted to wonder whether there is a simple deterministic algorithm that also works. Sadly, we do not know of any sub-exponential time deterministic algorithm! For certain special families of polynomials, deterministic algorithms are known (see e.g. the work of Ran Raz and Amir Shpilka [2]). In fact, Kabanets and Impagliazzo [1] have shown that if you find a deterministic polynomial time polynomial identity testing algorithm, this would imply some extremely strong complexity-theoretic results (such as  $NEXP \not\subseteq polyCircuits$ ). We believe that these complexity theoretic results are true, and believe that a deterministic algorithm exists, but both seem far beyond what we can achieve with current approaches/techniques.

One reason to begin this class with polynomial identity testing is because the story is fairly faithful to our current understanding of randomized algorithms in general: a very simple/clean randomized algorithm exists, whose analysis is clever but fairly easy. We believe

that an efficient deterministic algorithm exists, but that it is almost certainly extremely complicated.

## References

- [1] Valentine Kabanets and Russell Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Computational Complexity*, 13(1-2):1–46, 2004.
- [2] Ran Raz and Amir Shpilka. Deterministic polynomial identity testing in non-commutative models. *Computational Complexity*, 14(1):1–19, 2005.