19
Modeling randomized computation

> *"Any one who considers arithmetical methods of producing random digits is, of course, in a state of sin."*, John von Neumann, 1951.

So far we have described randomized algorithms in an informal way, assuming that an operation such as "pick a string $x \in \{0,1\}^n$" can be done efficiently. We have neglected to address two questions:

1. How do we actually efficiently obtain random strings in the physical world?

2. What is the mathematical model for randomized computations, and is it more powerful than deterministic computation?

The first question is of both practical and theoretical importance, but for now let's just say that there are various random physical sources. User's mouse movements, (non solid state) hard drive and network latency, thermal noise, and radioactive decay, have all been used as sources for randomness. For example, new Intel chips come with a random number generator built in. One can even build mechanical coin tossing machines (see Fig. 19.1).[1]

In this chapter we focus on the second point - formally modeling probabilistic computation and studying its power. Modeling randomized computation is actually quite easy. We can add the following operations to our NAND, NAND++ and NAND« programming languages:

```
var := RAND
```

where `var` is a variable. The result of applying this operation is that `var` is assigned a random bit in $\{0,1\}$. (Every time the `RAND` operation is invoked it returns a fresh independent random bit.) We

[1] The output of processes such as above can be thought of as a binary string sampled from some distribution μ that might have significant unpredictablity (or *entropy*) but is not necessarily the *uniform* distribution over $\{0,1\}^n$. Indeed, as this paper shows, even (real-world) coin tosses do not have exactly the distribution of a uniformly random string. Therefore, to use the resulting measurements for randomized algorithms, one typically needs to apply a "distillation" or *randomness extraction* process to the raw measurements to transform them to the uniform distribution.

Figure 19.1: A mechanical coin tosser built for Percy Diaconis by Harvard technicians Steve Sansone and Rick Haggerty

call the resulting languages RNAND, RNAND++, and RNAND« respectively.

We can use this to define the notion of a function being computed by a randomized $T(n)$ time algorithm for every nice time bound $T : \mathbb{N} \to \mathbb{N}$, as well as the notion of a finite function being computed by a size S randomized NAND program (or, equivalently, a randomized circuit with S gates that correspond to either NAND or coin-tossing). However, for simplicity we we will not define randomized computation in full generality, but simply focus on the class of functions that are computable by randomized algorithms *running in polynomial time*, which by historical convention is known as **BPP**:

> **Definition 19.1 — BPP.** Let $F :\in \{0,1\}^* \to \{0,1\}$. We say that $F \in$ **BPP** if there exist constants $a, b \in \mathbb{N}$ and an RNAND++ program P such that for every $x \in \{0,1\}^*$, on input x, the program P halts within at most $a|x|^b$ steps and
>
> $$\Pr[P(x) = F(x)] \geq \tfrac{2}{3} \tag{19.1}$$
>
> where this probability is taken over the result of the RAND operations of P. [2]

[2] **BPP** stands for "bounded probability polynomial time", and is used for historical reasons.

The same polynomial-overhead simulation of NAND« programs by NAND++ programs we saw in Theorem 12.4 extends to *randomized* programs as well. Hence the class **BPP** is the same regardless of whether it is defined via RNAND++ or RNAND« programs.

19.0.1 Random coins as an "extra input"

While we presented randomized computation as adding an extra "coin tossing" operation to our programs, we can also model this as being given an additional extra input. That is, we can think of a randomized algorithm A as a *deterministic* algorithm A' that takes *two inputs* x and r where the second input r is chosen at random from $\{0,1\}^m$ for some $m \in \mathbb{N}$. The equivalence to the Definition 19.1 is shown in the following theorem:

Theorem 19.2 — Alternative characterization of BPP. Let $F : \{0,1\}^* \to \{0,1\}$. Then $F \in \textbf{BPP}$ if and only if there exists $a, b \in \mathbb{N}$ and $G : \{0,1\}^* \to \{0,1\}$ such that G is in $\textbf{P}$ and for every $x \in \{0,1\}^*$,

$$\Pr_{r \sim \{0,1\}^{a|x|^b}}[G(xr) = F(x)] \geq \tfrac{2}{3} . \qquad (19.2)$$

Proof Idea: The idea behind the proof is that we can simply replace sampling a random coin with reading a bit from the extra "random input" r and vice versa. Of course to prove this rigorously we need to work through some formal notation, and so this might be one of those proofs that is easier for you to work out on your own than to read. ⋆

Proof of Theorem 19.2. We start by showing the "only if" direction. Let $F \in \textbf{BPP}$ and let P be an RNAND++ program that computes F as per Definition 19.1, and let $a, b \in \mathbb{N}$ be such that on every input of length n, the program P halts within at most an^b steps. We will construct a polynomial-time NAND++ program P' that computes a function G satisfying the conditions of Eq. (19.2). As usual, we will allow ourselves some "syntactic sugar" in constructing this program, as it can always be eliminated with polynomial overhead as in the proof of **??**. The program P' will first copy the bits in positions $n, n + 1, n + 2, \ldots, n + an^b - 1$ of its input into the variables `r_0`, `r_1`, ..., `r_`$\langle an^b - 1\rangle$. We will also assume we have access to an extra index variable `j` which we can increase and decrease (which of course can be simulated via syntactic sugar). The program P' will run the same operations of P except that it will replace a line of the form `foo :=` `RAND` with the two lines `foo := r_j` amd `j    := j + 1`

One can easily verify that **(1)** P' runs in polynomial time and **(2)** if the last an^b bits of the input of P' are chosen at random then its execution when its first n inputs are x is identical to an execution of $P(x)$. By **(2)** we mean that for every $r \in \{0,1\}^{an^b}$ corresponding to the result of the `RAND` operations made by P on its execution on x, the output of P on input x and with random choices r is equal to $P'(xr)$. Hence the distribution of the output $P(x)$ of the randomized

program P on input x is identical to the distribution of $P'(x; r)$ when $r \sim \{0,1\}^{an^b}$.

For the other direction, given a function $G \in \mathbf{P}$ satisfying the condition Eq. (19.2) and a NAND++ program P' that computes G in polynomial time, we will construct an RNAND++ program P that computes F in polynomial time. The idea behind the construction of P is simple: on input a string $x \in \{0,1\}^n$, we will first run for an^b steps and use the RNAND operation to create variables r_0, r_1, ...,r_$\langle an^b - 1 \rangle$ each containing the result of a random coin toss. We will then execute P' on the input x and r_0,...,r_$\langle an^b - 1 \rangle$ (i.e., replacing every reference to the variable x_$\langle n + k \rangle$ with the variable r_$\langle k \rangle$). Once again, it is clear that if P' runs in polynomial time then so will P, and for every input x and $r \in \{0,1\}^{an^b}$, the output of P on input x and where the coin tosses outcome is r is equal to $P'(xr)$. ∎

> **(R)** **Definitions of BPP and NP** The characterization of **BPP** Theorem 19.2 is reminiscent of the characterization of **NP** in Definition 14.1, with the randomness in the case of **BPP** playing the role of the solution in the case of **NP** but there are important differences between the two:
>
> - The definition of **NP** is "one sided": $F(x) = 1$ if *there exists* a solution w such that $G(xw) = 1$ and $F(x) = 0$ if *for every* string w of the appropriate length, $G(xw) = 0$. In contrast, the characterization of **BPP** is symmetric with respect to the cases $F(x) = 0$ and $F(x) = 1$.
>
> - For this reason the relation between **NP** and **BPP** is not immediately clear, and indeed is not known whether **BPP** $\subseteq$ **NP**, **NP** $\subseteq$ **BPP**, or these two classes are incomprable. It is however known (with a non-trivial proof) that if **P** $=$ **NP** then **BPP** $=$ **P** (see Theorem 19.9).
>
> - Most importantly, the definition of **NP** is "ineffective", since it does not yield a way of actually finding whether there exists a solution among the exponentially many possiblities. In contrast, the definition of **BPP** gives us a way to compute the function in practice by simply choosing the second input at random.

"Random tapes" Theorem 19.2 motivates sometimes considering the randomness of an RNAND++ (or RNAND«) program as an extra input, and so if A is a randomized algorithm that on inputs of length n makes at most $p(n)$ coin tosses, we will sometimes use the notation $A(x; r)$ (where $x \in \{0,1\}^n$ and $r \in \{0,1\}^{p(n)}$) to refer to the result of executing x when the coin tosses of A correspond to the coordinates

of r. This second or "auxiliary" input is sometimes referred to as a "random tape", with the terminology coming from the model of randomized Turing machines.

19.0.2 Amplification

The number $2/3$ might seem arbitrary, but as we've seen in Chapter 18 it can be amplified to our liking:

> **Theorem 19.3 — Amplification.** Let P be an RNAND« program, $F \in \{0,1\}^* \rightarrow \{0,1\}$, and $T : \mathbb{N} \rightarrow \mathbb{N}$ be a nice time bound such that for every $x \in \{0,1\}^*$, on input x the program P runs in at most $T(|x|)$ steps and moreover $\Pr[P(x) = F(x)] \geq \frac{1}{2} + \epsilon$ for some $\epsilon > 0$. Then for every k, there is a program P' taking at most $O(k \cdot T(n)/\epsilon^2)$ steps such that on input $x \in \{0,1\}^*$, $\Pr[P'(x) = F(x)] > 1 - 2^{-k}$.

Proof Idea: The proof is the same as we've seen before in the maximum cut and other examples. We use the Chernoff bound to argue that if we run the program $O(k/\epsilon^2)$ times, each time using fresh and independent random coins, then the probability that the majority of the answers will not be correct will be less than 2^{-k}. Amplification can be thought of as a "polling" of the choices for randomness for the algorithm, see Fig. 19.2. ⋆

Proof of Theorem 19.3. We can run P on input x for $t = 10k/\epsilon^2$ times, using fresh randomness each one, to compute outputs $y_0, \dots, y_{t-1}$. We output the value y that appeared the largest number of times. Let X_i be the random variable that is equal to 1 if $y_i = F(x)$ and equal to 0 otherwise. Then all the random variables $X_0, \dots, X_{t-1}$ are i.i.d. and satisfy $\mathbb{E}[X_i] = \Pr[X_i = 1] \geq 1/2 + \epsilon$. Hence by the Chernoff bound (Theorem 17.9) the probability that the majority value is not correct (i.e., that $\sum X_i \leq t/2$) is at most $2e^{-2\epsilon^2 t} < 2^{-k}$ for our choice of t.   ■

There is nothing special about NAND« in Theorem 19.3. The same proof can be used to amplify randomized NAND or NAND++ programs as well.

19.1 BPP AND NP COMPLETENESS

Since "noisy processes" abound in nature, randomized algorithms can be realized physically, and so it is reasonable to propose **BPP** rather than **P** as our mathematical model for "feasible" or "tractable" computation. One might wonder if this makes all the previous chapters irrelevant, and in particular does the theory of **NP** completeness still apply to probabilistic algorithms. Fortunately, the answer is *Yes*:

$F(X) = 1$ $F(X) = 0$

Figure 19.2: If $F \in$ **BPP** then there is randomized polynomial-time algorithm P with the following property. In the case $F(x) = 0$ two thirds of the "population" of random choices satisfy $P(x; r) = 0$ and in the case $F(x) = 1$ two thirdsof the population satisfy $P(x; r) = 1$. We can think of amplification as a form of "polling" of the choices of randomness. By the Chernoff bound, if we poll a sample of $O(\frac{\log(1/\delta)}{\epsilon^2})$ random choices r, then with probability at least $1 - \delta$, the fraction of r's in the sample satisfying $P(x; r) = 1$ will give us an estimate of the fraction of the population within an ϵ margin of error. This is the same calculation used by pollsters to determine the needed sample size in their polls.

> **Theorem 19.4 — NP hardness and BPP.** Suppose that F is **NP**-hard and $F \in$ **BPP** then **NP** $\subseteq$ **BPP**.

Before seeing the proof note that Theorem 19.4 in particular implies that if there was a randomized polynomial time algorithm for any **NP**-complete problem such as $3SAT$, $ISET$ etc.. then there will be such an algorithm for *every* problem in **NP**. Thus, regardless of whether our model of computation is deterministic or randomized algorithms, **NP** complete problems retain their status as the "hardest problems in **NP**".

Proof Idea: The idea is to simply run the reduction as usual, and plug it into the randomized algorithm instead of a deterministic one. It would be an excellent exercise, and a way to reinforce the definitions of **NP**-hardness and randomized algorithms, for you to work out the proof for yourself. However for the sake of completeness, we include this proof below. ⋆

Proof of Theorem 19.4. Suppose that F is **NP**-hard and $F \in$ **BPP**. We will now show that this implies that **NP** $\subseteq$ **BPP**. Let $G \in$ **NP**. By the definition of **NP**-hardness, it follows that $G \leq_p F$, or that in other words there exists a polynomial-time computable function $R : \{0,1\}^* \to \{0,1\}^*$ such that $G(x) = F(R(x))$ for every $x \in \{0,1\}^*$. Now if F is in **BPP** then there is a polynomial-time RNAND++ program P

such that

$$\Pr[P(y) = F(y)] \geq 2/3 \qquad (19.3)$$

for *every* $y \in \{0,1\}^*$ (where the probability is taken over the random coin tosses of P). Hence we can get a polynomial-time RNAND++ program P' to compute G by setting $P'(x) = P(R(x))$. By Eq. (19.3) $\Pr[P'(x) = F(R(x))] \geq 2/3$ and since $F(R(x)) = G(x)$ this implies that $\Pr[P'(x) = G(x)] \geq 2/3$, which proves that $G \in$ **BPP**.   ∎

Most of the results we've seen about the **NP** hardness, including the search to decision reduction of Theorem 15.1, the decision to optimization reduction of Theorem 15.2, and the quantifier elimination result of Theorem 15.3, all carry over in the same way if we replace **P** with **BPP** as our model of efficient computation. Thus if **NP** $\subseteq$ **BPP** then we'd get essentially all of the strange and wonderful consequences of **P** $=$ **NP**. Unsurprisingly, we cannot rule out this possiblity. In fact, unlike **P** $=$ **EXP**, which is ruled out by the time hierarchy theorem, we don't even know how to rule out the possiblity that **BPP** $=$ **EXP**! Thus a priori it's possible (though seems highly unlikely) that randomness is a magical tool that allows to speed up arbitrary exponential time computation.[3] Nevertheless, as we discuss below, it is believed that randomization's power is much weaker and **BPP** lies in much more "pedestrian" territory.

[3] At the time of this writing, the largest "natural" complexity class which we can't rule out being contained in **BPP** is the class **NEXP**, which we did not define in this course, but corresponds to non deterministic exponential time. See this paper for a discussion of this question.

19.2 THE POWER OF RANDOMIZATION

A major question is whether randomization can add power to computation. Mathematically, we can phrase this as the following question: does **BPP** $=$ **P**? Given what we've seen so far about the relations of other complexity classes such as **P** and **NP**, or **NP** and **EXP**, one might guess that:

1. We do not know the answer to this question.

2. But we suspect that **BPP** is different than **P**.

One would be correct about the former, but wrong about the latter. As we will see, we do in fact have reasons to believe that **BPP** $=$ **P**. This can be thought of as supporting the *extended Church Turing hypothesis* that deterministic polynomial-time NAND++ program (or, equivalently, polynomial-time Turing machines) capture what can be feasibly computed in the physical world.

We now survey some of the relations that are known between **BPP** and other complexity classes we have encountered, see also Fig. 19.3.

19.2.1 Solving **BPP** in exponential time
It is not hard to see that if F is in **BPP** then it can be computed in *exponential* time.

Figure 19.3: Some possibilities for the relations between **BPP** and other complexity classes. Most researchers believe that **BPP** = **P** and that these classes are *not* powerful enough to solve **NP**-complete problems, let alone all problems in **EXP**. However, we have not even been able yet to rule out the possiblity that randomness is a "silver bullet" that allows exponential speedup on all problems, and hence **BPP** = **EXP**. As we've already seen, we also can't rule out that **P** = **NP**. Interestingly, in the latter case, **P** = **BPP**.

Theorem 19.5 — Simulating randomized algorithm in exponential time.
BPP ⊆ **EXP**

> P   The proof of Theorem 19.5 readily follows by enu-
> merating over all the (exponentially many) choices
> for the random coins. We omit the formal proof, as
> doing it by yourself is an excellent way to get comfort
> with Definition 19.1

19.2.2 Simulating randomized algorithms by circuits or straightline programs.

We have seen in Theorem 12.7 that if F is in **P**, then there is a polynomial $p : \mathbb{N} \to \mathbb{N}$ such that for every n, the restriction F_n of F to inputs $\{0,1\}^n$ is in $SIZE(p(n))$. A priori it is not at all clear that the same holds for a function in **BPP**, but this does turn out to be the case.

Theorem 19.6 — Randomness does not help for non uniform computation: BPP ⊆ **P**$_{/\textbf{poly}}$. For every $F \in$ **BPP**, there exist some $a, b \in \mathbb{N}$ such that for every $n > 0$, $F_n \in SIZE(an^b)$ where F_n is the restriction of F to inputs in $\{0,1\}^n$.

Proof Idea: The idea behind the proof is that we can first amplify by repetition the probability of success from $2/3$ to $1 - 0.1 \cdot 2^{-n}$. This will allow us to show that there exists a single fixed choice of "favorable coins" that would cause the algorithm to output the right answer on *all* of the possible 2^n inputs. We can then use the standard "unravelling the loop" technique to transform an RNAND++ program to an RNAND program, and "hardwire" the favorable choice of random coins to transform the RNAND program into a plain-old deterministic NAND program. ⋆
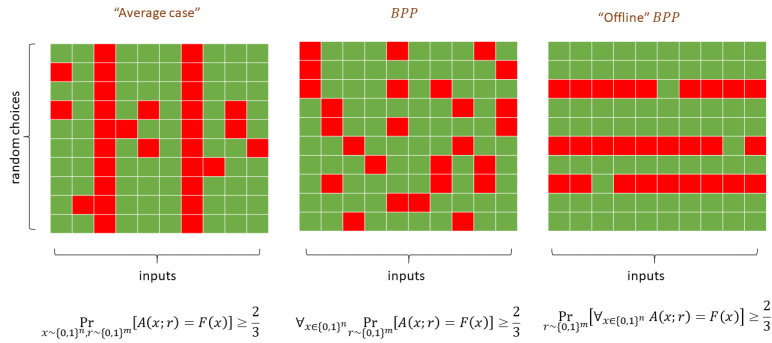
$$\Pr_{x\sim\{0,1\}^n, r\sim\{0,1\}^m}[A(x;r) = F(x)] \geq \frac{2}{3} \qquad \forall_{x\in\{0,1\}^n} \Pr_{r\sim\{0,1\}^m}[A(x;r) = F(x)] \geq \frac{2}{3} \qquad \Pr_{r\sim\{0,1\}^m}[\forall_{x\in\{0,1\}^n} A(x;r) = F(x)] \geq \frac{2}{3}$$

Figure 19.4: The possible guarantees for a randomized algorithm A computing some function F. In the tables above the columns correspond to different inputs, the rows to different choices of the random tape. A cell at position r, x is colored green if $A(x;r) = F(x)$ (i.e., the algorithm outputs the correct answer) and red otherwise. The standard **BPP** guarantee corresponds to the middle figure, where for every input x, at least two thirds of the choices r for a random tape will result in A computing the correct value. That is, every column is colored green in at least two thirds of its coordinates. In the left figure we have an "average case" guarantee where the algorithm is only guaranteed to output the correct answer with probabilty two thirds over a *random* input (i.e., at most one third of the total entries of the table are colored red, but there could be an all red column). The right figure corresponds to the "offline **BPP**" case, with probability at least two thirds over the random choice r, r will be good for *every* input. That is, at least two thirds of the rows are all green. Theorem 19.6 (**BPP** $\subseteq$ **P**$_{/\text{poly}}$) is proven by amplifying the success of a **BPP** algorithm until we have the "offline **BPP**" guarantee, and then hardwiring the choice of the randomness r to obtain a nonuniform deterministic algorithm.

Proof of Theorem 19.6. Suppose that $F \in$ **BPP** and let P be a polynomial-time RNAND++ program that computes F as per Definition 19.1. Using Theorem 19.3 we can *amplify* the success probability of P to obtain an RNAND++ program P' that is at most a factor of $O(n)$ slower (and hence still polynomial time) such that for every $x \in \{0,1\}^n$

$$\Pr_{r\sim\{0,1\}^m}[P'(x;r) = F(x)] \geq 1 - 0.1 \cdot 2^{-n}, \qquad (19.4)$$

where m is the number of coin tosses that P' uses on inputs of length n, and we use the notation $P'(x;r)$ to denote the execution of P' on input x and when the result of the coin tosses corresponds to the string r.

For every $x \in \{0,1\}^n$, define the "bad" event B_x to hold if $P'(x) \neq F(x)$, where the sample space for this event consists of the coins of P'. Then by Eq. (19.4), $\Pr[B_x] \leq 0.1 \cdot 2^{-n}$ for every $x \in \{0,1\}^n$. Since there are 2^n many such x's, by the union bound we see that the probability that the *union* of the events $\{B_x\}_{x\in\{0,1\}^n}$ is at most 0.1. This means that if we choose $r \sim \{0,1\}^m$, then with probability at least 0.9 it will be the case that for *every* $x \in \{0,1\}^n$, $F(x) = P'(x;r)$. (Indeed, otherwise the event B_x would hold for some x.) In particular, because of the mere fact that the the probability of $\cup_{x\in\{0,1\}^n} B_x$ is smaller than 1, this

means that *there exists* a particular $r^* \in \{0,1\}^m$ such that

$$P'(x; r^*) = F(x) \qquad (19.5)$$

for every $x \in \{0,1\}^n$. Now let us use the standard "unravelling the loop" the technique and transform P' into a NAND program Q of polynomial in n size, such that $Q(xr) = P'(x;r)$ for every $x \in \{0,1\}^n$ and $r \in \{0,1\}^m$. Then by "hardwiring" the values $r_0^*, \dots, r_{m-1}^*$ in place of the last m inputs of Q, we obtain a new NAND program Q_{r^*} that satisfies by Eq. (19.5) that $Q_{r^*}(x) = F(x)$ for every $x \in \{0,1\}^n$. This demonstrates that F_n has a polynomial sized NAND program, hence completing the proof of Theorem 19.6   ∎

> (R) **Randomness and non uniformity** The proof of
> Theorem 19.6 actually yields more than its state-
> ment. We can use the same "unrolling the loop"
> arguments we've used before to show that the
> restriction to $\{0,1\}^n$ of every function in **BPP** is
> also computable by a polynomial-size RNAND
> program (i.e., NAND program with the RAND op-
> eration), but like in the **P** vs $SIZE(poly(n))$ case,
> there are functions outside **BPP** whose restrictions
> can be computed by polynomial-size RNAND pro-
> grams. Nevertheless the proof of Theorem 19.6
> shows that even such functions can be computed
> by polynomial sized NAND programs without us-
> ing the rand operations. This can be phrased as
> saying that $BPSIZE(T(n)) \subseteq SIZE(O(nT(n)))$
> (where $BPSIZE$ is defined in the natural way us-
> ing RNAND progams). Theorem 19.6 can also be
> phrased as saying that **BPP** $\subseteq$ **P**$_{/\text{poly}}$, and the stronger
> result can be phrased as **BPP**$_{/\text{poly}} = $ **P**$_{/\text{poly}}$.

19.3 DERANDOMIZATION

The proof of Theorem 19.6 can be summarized as follows: we can replace a $poly(n)$-time algorithm that tosses coins as it runs, with an algorithm that uses a single set of coin tosses $r^* \in \{0,1\}^{poly(n)}$ which will be good enough for all inputs of size n. Another way to say it is that for the purposes of computing functions, we do not need "online" access to random coins and can generate a set of coins "offline" ahead of time, before we see the actual input.

But this does not really help us with answering the question of whether **BPP** equals **P**, since we still need to find a way to generate these "offline" coins in the first place. To derandomize an RNAND++ program we will need to come up with a *single* deterministic algo-rithm that will work for *all input lengths*. That is, unlike in the case of

RNAND programs, we cannot choose for every input length n some string $r^* \in \{0,1\}^{poly(n)}$ to use as our random coins.

Can we derandomize randomized algorithms, or does randomness add an inherent extra power for computation? This is a fundamentally interesting question but is also of practical significance. Ever since people started to use randomized algorithms during the Manhattan project, they have been trying to remove the need for randomness and replace it with numbers that are selected through some deterministic process. Throughout the years this approach has often been used successfully, though there have been a number of failures as well.[4]

A common approach people used over the years was to replace the random coins of the algorithm by a "randomish looking" string that they generated through some arithmetic progress. For example, one can use the digits of π for the random tape. Using these type of methods corresponds to what von Neumann referred to as a "state of sin". (Though this is a sin that he himself frequently committed, as generating true randomness in sufficient quantity was and still is often too expensive.) The reason that this is considered a "sin" is that such a procedure will not work in general. For example, it is easy to modify any probabilistic algorithm A such as the ones we have seen in Chapter 18, to an algorithm A' that is *guaranteed to fail* if the random tape happens to equal the digits of π. This means that the procedure "replace the random tape by the digits of π" does not yield a *general* way to transform a probabilistic algorithm to a deterministic one that will solve the same problem. Of course, this procedure does not *always* fail, but we have no good way to determine when it fails and when it succeeds. This reasoning is not specific to π and holds for every deterministically produced string, whether it obtained by π, e, the Fibonacci series, or anything else.

An algorithm that checks if its random tape is equal to π and then fails seems to be quite silly, but this is but the "tip of the iceberg" for a very serious issue. Time and again people have learned the hard way that one needs to be very careful about producing random bits using deterministic means. As we will see when we discuss cryptography, many spectacular security failures and break-ins were the result of using "insufficiently random" coins.

19.3.1 Pseudorandom generators

So, we can't use any *single* string to "derandomize" a probabilistic algorithm. It turns out however, that we can use a *collection* of strings to do so. Another way to think about it is that rather than trying to *eliminate* the need for randomness, we start by focusing on *reducing* the amount of randomness needed. (Though we will see that if we reduce the randomness sufficiently, we can eventually get rid of it altogether.)

[4] One amusing anecdote is a recent case where scammers managed to predict the imperfect "pseudorandom generator" used by slot machines to cheat casinos. Unfortunately we don't know the details of how they did it, since the case was sealed.

We make the following definition:

> **Definition 19.7 — Pseudorandom generator.** A function $G : \{0,1\}^\ell \to \{0,1\}^m$ is a (T, ϵ)-*pseudorandom generator* if for every NAND program P with m inputs and one output of at most T lines,
>
> $$\left| \Pr_{s \sim \{0,1\}^\ell}[P(G(s)) = 1] - \Pr_{r \sim \{0,1\}^m}[P(r) = 1] \right| < \epsilon \qquad (19.6)$$


Figure 19.5: A pseudorandom generator G maps a short string $s \in \{0,1\}^\ell$ into a long string $r \in \{0,1\}^m$ such that an small program P cannot distinguish between the case that it is provided a random input $r \sim \{0,1\}^m$ and the case that it is provided a "pseudorandom" input of the form $r = G(s)$ where $s \sim \{0,1\}^\ell$. The short string s is sometimes called the *seed* of the pseudorandom generator, as it is a small object that can be thought as yielding a large "tree of randomness".

> **P**   This is a definition that's worth reading more than once, and spending some time to digest it. Note that it takes several parameters:
>
> - T is the limit on the number of lines of the program P that the generator needs to "fool". The larger T is, the stronger the generator.
> - ϵ is how close is the output of the pseudorandom generator to the true uniform distribution over $\{0,1\}^m$. The smaller ϵ is, the stronger the generator.
> - ℓ is the input length and m is the output length. If $\ell \geq m$ then it is trivial to come up with such

a generator: on input $s \in \{0,1\}^\ell$, we can output $s_0, \dots, s_{m-1}$. In this case $\Pr_{s \sim \{0,1\}^\ell}[P(G(s)) = 1]$ will simply equal $\Pr_{r \in \{0,1\}^m}[P(r) = 1]$, no matter how many lines P has. So, the smaller ℓ is and the larger m is, the stronger the generator, and to get anything non-trivial, we need $m > \ell$.

Furthermore note that although our eventual goal is to fool probabilistic randomized algorithms that take an unbounded number of inputs, Definition 19.7 refers to *finite* and *deterministic* NAND programs.

We can think of a pseudorandom generator as a "randomness amplifier". It takes an input s of ℓ bits chosen at random and expands these ℓ bits into an output r of $m > \ell$ *pseudorandom* bits. If ϵ is small enough then the pseudorandom bits will "look random" to any NAND program that is not too big. Still, there are two questions we haven't answered:

- *What reason do we have to believe that pseudorandom generators with non-trivial parameters exist?*

- *Even if they do exist, why would such generators be useful to derandomize randomized algorithms?* After all, Definition 19.7 does not involve RNAND++ or RNAND«programs but deterministic NAND programs with no randomness and no loops.

We will now (partially) answer both questions.

For the first question, let us come clean and confess we do not know how to *prove* that interesting pseudorandom generators exist. By *interesting* we mean pseudorandom generators that satisfy that ϵ is some small constant (say $\epsilon < 1/3$), $m > \ell$, and the function G itself can be computed in $poly(m)$ time. Nevertheless, Lemma 19.10 (whose statement and proof is deferred to the end of this chapter) shows that if we only drop the last condition (polynomial-time computability), then there do in fact exist pseudorandom generators where m is *exponentially larger* than ℓ.

P At this point you might want to skip ahead and look at the *statement* of Lemma 19.10. However, since its *proof* is somewhat subtle, I recommend you defer reading it until you've finished reading the rest of this chap.

19.3.2 From existence to constructivity

The fact that there *exists* a pseudorandom generator does not mean that there is one that can be efficiently computed. However, it turns

out that we can turn complexity "on its head" and used the assumed *non existence* of fast algorithms for problems such as 3SAT to obtain pseudorandom generators that can then be used to transform randomized algorithms into deterministic ones. This is known as the *Hardness vs Randomness* paradigm. A number of results along those lines, most of whom are ourside the scope of this course, have led researchers to believe the following conjecture:

> **Optimal PRG conjecture:** There is a polynomial-time computable function $PRG : \{0,1\}^* \rightarrow \{0,1\}$ that yields an *exponentially secure pseudorandom generator*. Specifically, there exists a constant $\delta > 0$ such that for every ℓ and $m < 2^{\delta\ell}$, if we define $G : \{0,1\}^\ell \rightarrow \{0,1\}^m$ as $G(s)_i = PRG(s,i)$ for every $s \in \{0,1\}^\ell$ and $i \in [m]$, then G is a $(2^{\delta\ell}, 2^{-\delta\ell})$ pseudorandom generator.

> **P**   The "optimal PRG conjecture" is worth while reading more than once. What it posits is that we can obtain (T, ϵ) pseudorandom generator G such that every output bit of G can be computed in time polynomial in the length ℓ of the input, where T is exponentially large in ℓ and ϵ is exponentially small in ℓ. (Note that we could not hope for the entire output to be computable in ℓ, as just writing the output down will take too long.)
>
> To understand why we call such a pseudorandom generator "optimal", it is a great exercise to convince yourself that there exists no (T, ϵ) pseudorandom generator unless T is smaller than (say) $2^{2\ell}$ and ϵ is at least (say) $2^{-2\ell}$. For the former case note that if we allow a NAND program with much more than 2^ℓ lines then this NAND program could "hardwire" inside it all the outputs of G on all its 2^ℓ inputs, and use that to distinguish between a string of the form $G(s)$ and a uniformly chosen string in $\{0,1\}^m$. For the latter case note that by trying to "guess" the input s, we can achieve a $2^{-\ell}$ advantage in distinguishing a pseudorandom and uniform input. But working out these details is a highly recommended exercise.

We emphasize again that the optimal PRG conjecture is, as its name implies, a *conjecture*, and we still do not know how to *prove* it. In particular, it is stronger than the conjecture that **P** $\neq$ **NP**. But we do have some evidence for its truth. There is a spectrum of different types of pseudorandom generators, and there are weaker assumption than the optimal PRG conjecture that suffice to prove that **BPP** $=$ **P**. In particular this is known to hold under the assumption that there exists a

function $F \in \mathbb{TIME}(2^{O(n)})$ and $\epsilon > 0$ such that for every sufficiently large n, F_n is not in $SIZE(2^{\epsilon n})$. The name "Optimal PRG conjecture" is non standard. This conjecture is sometimes known in the literature as the existence of exponentially strong pseudorandom functions.[5]

19.3.3 Usefulness of pseudorandom generators

We now show that optimal pseudorandom generators are indeed very useful, by proving the following theorem:

> **Theorem 19.8 — Derandomization of BPP.** Suppose that the optimal PRG conjecture is true. Then **BPP = P**.

Proof Idea: The optimal PRG conjecture tells us that we can achieve *exponential expansion* of ℓ truly random coins into as many as $2^{\delta\ell}$ "pseudorandom coins". Looked at from the other direction, it allows us to reduce the need for randomness by taking an algorithm that uses m coins and converting it into an algorithm that only uses $O(\log m)$ coins. Now an algorithm of the latter type by can be made fully deterministic by enumerating over all the $2^{O(\log m)}$ (which is polynomial in m) possiblities for its random choices. We now proceed with the proof details. ⋆

Proof of Theorem 19.8. Let $F \in$ **BPP** and let P be a NAND++ program and a, b, c, d constants such that for every $x \in \{0,1\}^n$, $P(x)$ runs in at most $c\dot{n}^d$ steps and $\Pr_{r\sim\{0,1\}^m}[P(x;r) = F(x)] \geq 2/3$. By "unrolling the loop" and hardwiring the input x, we can obtain for every input $x \in \{0,1\}^n$ a NAND program Q_x of at most, say, $T = 10c\dot{n}^d$ lines, that takes m bits of input and such that $Q(r) = P(x;r)$.

Now suppose that $G : \{0,1\}^\ell \to \{0,1\}$ is a $(T, 0.1)$ pseudorandom generator. Then we could deterministically estimate the probability $p(x) = \Pr_{r\sim\{0,1\}^m}[Q_x(r) = 1]$ up to 0.1 accuracy in time $O(T \cdot 2^\ell \cdot m \cdot cost(G))$ where $cost(G)$ is the time that it takes to compute a single output bit of G. The reason is that we know that $\tilde{p}(x) = \Pr_{s\sim\{0,1\}^\ell}[Q_x(G(s)) = 1]$ will give us such an estimate for $p(x)$, and we can compute the probability $\tilde{p}(x)$ by simply trying all 2^ℓ possibillites for s. Now, under the optimal PRG conjecture we can set $T = 2^{\delta\ell}$ or equivalently $\ell = \frac{1}{\delta}\log T$, and our total computation time is polynomial in $2^\ell = T^{1/\delta}$, and since $T \leq 10c\dot{n}^d$, this running time will be polynomial in n. This completes the proof, since we are guaranteed that $\Pr_{r\sim\{0,1\}^m}[Q_x(r) = F(x)] \geq 2/3$, and hence estimating the probability $p(x)$ to within 0.1 accuracy is sufficient to compute $F(x)$. ∎

19.4 P $=$ NP AND BPP VS P

Two computational complexity questions that we cannot settle are:

- Is **P** $=$ **NP**? Where we believe the answer is *negative*.

- If **BPP** $=$ **P**? Where we believe the answer is *positive*.

However we can say that the "conventional wisdom" is correct on at least one of these questions. Namely, if we're wrong on the first count, then we'll be right on the second one:

> **Theorem 19.9 — Sipser–Gács Theorem.** If **P** $=$ **NP** then **BPP** $=$ **P**.

Proof Idea: The construction follows the "quantifier elimination" idea which we have seen in Theorem 15.3. We will show that for every $F \in$ **BPP**, we can reduce the question of some input x satisfies $F(x) = 1$ to the question of whether a formula of the form $\exists_{u \in \{0,1\}^m} \forall_{v \in \{0,1\}^k} P(x, y)$ is true where m, k are polynomial in the length of x and P is polynomial-time computable. By Theorem 15.3, if **P** $=$ **NP** then we can decide in polynomial time whether such a formula is true or false. $\star$

Proof of Theorem 19.9. TO BE COMPLETED   ∎

19.5 NON-CONSTRUCTIVE EXISTENCE OF PSEUDORANDOM GENERATORS

We now show that, if we don't insist on *constructivity* of pseudorandom generators, then we can show that there exists pseudorandom generators with output that *exponentially larger* in the input length.

Lemma 19.10 — Existence of inefficient pseudorandom generators. There is some absolute constant C such that for every ϵ, T, if $\ell > C(\log T + \log(1/\epsilon))$ and $m \leq T$, then there is an (T, ϵ) pseudorandom generator $G : \{0,1\}^\ell \to \{0,1\}^m$.

Proof Idea: The proof uses an extremely useful technique known as the "probabilistic method" which is not too hard mathematically but can be confusing at first.[6] The idea is to give a "non constructive" proof of existence of the pseudorandom generator G by showing that if G was chosen at random, then the probability that it would be a valid (T, ϵ) pseudorandom generator is positive. In particular this means that there *exists* a single G that is a valid (T, ϵ) pseudorandom generator. The probabilistic method is just a *proof technique* to demonstrate the existence of such a function. Ultimately, our goal is to show the existence of a *deterministic* function G that satisfies $\star$

[6] There is a whole (highly recommended) book by Alon and Spencer devoted to this method.

The above discussion might be rather abstract at this point, but would become clearer after seeing the proof.

Proof of Lemma 19.10. Let ϵ, T, ℓ, m be as in the lemma's statement. We need to show that there exists a function $G : \{0,1\}^{\ell} \to \{0,1\}^{m}$ that "fools" every T line program P in the sense of Eq. (19.6). We will show that this follows from the following claim:

Claim I: For every fixed NAND program P, if we pick $G : \{0,1\}^{\ell} \to \{0,1\}^{m}$ *at random* then the probability that Eq. (19.6) is violated is at most 2^{-T^2}.

Before proving Claim I, let us see why it implies Lemma 19.10. We can identify a function $G : \{0,1\}^{\ell} \to \{0,1\}^{m}$ with its "truth table" or simply the list of evaluations on all its possible 2^{ℓ} inputs. Since each output is an m bit string, we can also think of G as a string in $\{0,1\}^{m \cdot 2^{\ell}}$. We define $\mathcal{G}_{\ell}^{m}$ to be the set of all functions from $\{0,1\}^{\ell}$ to $\{0,1\}^{\ell}$. As discussed above we can identify $\mathcal{F}_{\ell}^{m}$ with $\{0,1\}^{m \cdot 2^{\ell}}$ and choosing a random function $G \sim \mathcal{F}_{\ell}^{m}$ corresponds to choosing a random $m \cdot 2^{\ell}$-long bit string.

For every NAND program P let B_P be the event that, if we choose G at random from $\mathcal{F}_{\ell}^{m}$ then Eq. (19.6) is violated with respect to the program P. It is important to understand what is the sample space that the event B_P is defined over, namely this event depends on the choice of G and so B_P is a subset of $\mathcal{F}_{\ell}^{m}$. An equivalent way to define the event B_P is that it is the subset of all functions mapping $\{0,1\}^{\ell}$ to $\{0,1\}^{m}$ that violate Eq. (19.6), or in other words:

$$B_P = \left\{ G \in \mathcal{F}_{\ell}^{m} \ \Big| \ \left| \tfrac{1}{2^{\ell}} \sum_{s \in \{0,1\}^{\ell}} P(G(s)) - \tfrac{1}{2^m} \sum_{r \in \{0,1\}^m} P(r) \right| > \epsilon \right\} .$$
(19.7)

(We've replaced here the probability statements in Eq. (19.6) with the equivalent sums so as to reduce confusion as to what is the sample space that B_P is defined over.)

To understand this proof it is crucial that you pause here and see how the definition of B_P above corresponds to Eq. (19.7). This may well take re-reading the above text once or twice, but it is a good exercise at parsing probabilistic statements and learning how to identify the *sample space* that these statements correspond to.

Now, we've shown in Theorem 5.5 that up to renaming variables (which makes no difference to program's functionality) there are $2^{O(T \log T)}$ NAND programs of at most T lines. Since $T \log T < T^2$ for sufficiently large T, this means that if the Claim I is true, then by the union bound it holds that the probability of the union of B_P over *all* NAND programs of at most T lines is at most $2^{O(T \log T)} 2^{-T^2} < 0.1$ for sufficiently large T. What is important for us about the number 0.1

is that it is smaller than 1. In particular this means that there *exists* a single $G^* \in \mathcal{F}_\ell^m$ such that G^* *does not* violate Eq. (19.6) with respect to any NAND program of at most T lines, but that precisely means that G^* is a (T, ϵ) pseudorandom generator.

Hence conclude the proof of Lemma 19.10, it suffices to prove Claim I. Choosing a random $G : \{0,1\}^\ell \to \{0,1\}^m$ amounts to choosing $L = 2^\ell$ random strings $y_0, \dots, y_{L-1} \in \{0,1\}^m$ and letting $G(x) = y_x$ (identifying $\{0,1\}^\ell$ and $[L]$ via the binary representation). Hence the claim amounts to showing that for every fixed function $P : \{0,1\}^m \to \{0,1\}$, if $L > 2^{C(\log T + \log \epsilon)}$ (which by setting $C > 4$, we can ensure is larger than $10T^2/\epsilon^2$) then the probability that

$$\left| \frac{1}{L} \sum_{i=0}^{L-1} P(y_s) - \Pr_{s \sim \{0,1\}^m}[P(s) = 1] \right| > \epsilon \tag{19.8}$$

is at most 2^{-T^2}. Eq. (19.8) follows directly from the Chernoff bound. If we let for every $i \in [L]$ the random variable X_i denote $P(y_i)$, then since $y_0, \dots, y_{L-1}$ is chosen independently at random, these are independently and identically distributed random variables with mean $\mathbb{E}_{y \sim \{0,1\}^m}[P(y)] = \Pr_{y \sim \{0,1\}^m}[P(y) = 1]$ and hence the probability that they deviate from their expectation by ϵ is at most $2 \cdot 2^{-\epsilon^2 L/2}$. ∎

19.6 LECTURE SUMMARY

19.7 EXERCISES

> **R** **Disclaimer** Most of the exercises have been written in the summer of 2018 and haven't yet been fully debugged. While I would prefer people do not post online solutions to the exercises, I would greatly appreciate if you let me know of any bugs. You can do so by posting a GitHub issue about the exercise, and optionally complement this with an email to me with more details about the attempted solution.

19.8 BIBLIOGRAPHICAL NOTES

19.9 FURTHER EXPLORATIONS

Some topics related to this chapter that might be accessible to advanced students include: (to be completed)

19.10 ACKNOWLEDGEMENTS

V
ADVANCED TOPICS