Notes for Lecture 25

These notes follow quite literally the notes that I wrote for the PCMI school on computational complexity in Summer 2000. There is some redundancy between these notes and the content of past lectures.

1 Probabilistic Algorithms versus Deterministic Algorithms

A probabilistic algorithm $A(\cdot,\cdot)$ is an algorithm that takes two inputs x and r, where x is an instance of some problem that we want to solve, and r is the output of a *random source*. A random source is an idealized device that outputs a sequence of bits that are uniformly and independently distributed. For example the random source could be a device that tosses coins, observes the outcome, and outputs it. A probabilistic algorithm A is good if it is efficient and if, say, for every x,

$$\mathbf{Pr}_r[A(x,r) = \text{ right answer for } x\ ] \geq \frac{3}{4}$$

We will typically restrict to the case where A solves a decision problem (e.g. it tests whether two read-once branching programs are equivalent). In this case we say that a language L is in **BPP** if there is a polynomial time algorithm $A(\cdot,\cdot)$ (polynomial in the length of the first input) such that for every x

$$\mathbf{Pr}_r[A(x,r) = \chi_L(x)] \geq \frac{3}{4}$$

or, said, another way,

$$x \in L \Rightarrow \mathbf{Pr}_r[A(x,r) = 1] \geq \frac{3}{4}$$

and

$$x \notin L \Rightarrow \mathbf{Pr}_r[A(x,r) = 1] \leq \frac{1}{4}\ .$$

The choice of the constant $3/4$ is clearly quite arbitrary. For any constant $1/2 < p < 1$, if we had defined **BPP** by requiring the probabilistic algorithm to be correct with probability st least p, we would have given an equivalent definition. In fact, for any polynomial p, it would have been equivalent to define **BPP** by asking the algorithm to be correct with probability at least $1/2+1/p(n)$, where n is the size of the input, and it would have also been equivalent if we had asked the algorithm to be correct with probability at least $1-1/2^{p(n)}$. That is, for any two polynomials p and q, if for a decision problem L we have a probabilistic polynomial time A that solves L on every input of length n with probability at least $1/2+1/p(n)$, then there is another probabilistic algorithm A', still running in polynomial time, that solves L on every input of length n with probability at least $1-2^{-q(n)}$.

For quite a few interesting problems, the only known polynomial time algorithms are probabilistic. A well-known example is the problem of testing whether two multivariate low-degree polynomials given in an implicit representation are equivalent. Another example is the problem of extracting "square roots" modulo a prime, i.e. to find solutions, if they exist, to equations of the form $x^2 = a \pmod{p}$ where p and a are given, and p is prime. More generally, there are probabilistic polynomial time algorithms to find roots of polynomials modulo a prime. There is no known deterministic polynomial time algorithm for any of the above problems.

It is not clear whether the existence of such probabilistic algorithms suggests that probabilistic algorithms are inherently more powerful than deterministic ones, or that we have not been able yet to find the best possible deterministic algorithms for these problems. In general, it is quite an interesting question to determine what is the relative power of probabilistic and deterministic computations. This question is the main motivations for the results described in this lecture and the next ones.

1.1 A trivial deterministic simulation

Let A be a probabilistic algorithm that solves a decision problem L. On input x of length n, say that A uses a random string r of length $m = m(n)$ and runs in time $T = T(n)$ (note that $m \leq T$).

It is easy to come up with a deterministic algorithm that solves L in time $2^{m(n)}T(n)$. On input x, compute $A(x, r)$ for every r. The correct answer is the one that comes up the majority of the times, so, in order to solve our problem, we just have to keep track, during the computation of $A(x, r)$ for every r, of the number of strings r for which $A(x, r) = 1$ and the number of strings r for which $A(x, r) = 0$.

Notice that the running time of the simulation depends exponentially on the number of random bits used by A, but only polynomially on the running time of A. In particular, if A uses a logarithmic number of random bits, then the simulation is polynomial. However, typically, a probabilistic algorithm uses a linear, or more, number of random bits, and so this trivial simulation is exponential. As we will see in the next section, it is not easy to obtain more efficient simulations.

1.2 Exponential gaps between randomized and deterministic procedures

For some computational problems (e.g. approximating the size of a convex body) there are probabilistic algorithms that work even if the object on which they operate is exponentially big and given as a black box; in some cases one can prove that deterministic algorithms cannot solve the same problem in the same setting, unless they use exponential time. Let us see a particularly clean (but more artificial) example of this situation.

Suppose that there is some function $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ that is given as an oracle; we want to devise an algorithm that on input x finds an approximation (say, to within an additive factor 1/10) to the value $\mathbf{Pr}_y[f(x, y) = 1]$. A probabilistic algorithm would pick $O(1)$ points $y_1, \ldots, y_t$ at random, evaluate $f(x, y_i)$, and then output the fraction of i such that $f(x, y_i) = 1$. This will be an approximation to within 1/10 with good probability. However a deterministic subexponential algorithm, given x, can only look at a negligible fraction of the values $f(x, y)$. Suppose that f is zero everywhere. Now consider the function $g(x, y)$ that is equal to f on all the points that our algorithm queries, and is 1 elsewhere (note that, by this definition, the queries of the algorithm on input x will be the same for f and g). If the algorithm takes sub-exponential time, g is almost everywhere one, yet the algorithm will give the same answer as when accessing f, which is everywhere zero. If our algorithm makes less than 2^{n-1} oracle queries, it cannot solve the problem with the required accuracy.

2 De-randomization Under Complexity Assumptions

It is still not known how to improve, in the general case, the deterministic simulation of Section 1.1, and the observation of Section 1.2 shows one of the difficulties in achieving an improvement. If we want to come up with a general way of transforming probabilistic procedures into deterministic sub-exponential procedures, the transformation cannot be described and analyzed by modeling in a "black box" way the probabilistic procedure.[1] If we want to deterministically and sub-exponentially simulate **BPP** algorithms, we have to exploit the fact that a **BPP** algorithm $A(\cdot, \cdot)$ is not an arbitrary function, but an efficiently computable one, and this is difficult because we still have a very poor understanding of the nature of efficient computations.

The results described in these notes show that it is indeed possible to deterministically simulate probabilistic algorithms in sub-exponential (or even polynomial) time, provided that certain complexity-theoretic assumptions are true. It is quite usual in complexity theory that, using reductions, one can show that the answer to some open question is implied by (or even equivalent to) the answer to some other question, however the nature of the results of these notes is somewhat unusual. Typically a reduction from a computational problem A to a problem B shows that if B has an efficient algorithm then A has also an efficient algorithm, and, by counterpositive, if A is intractable then B is also intractable. In general, using reductions one shows that algorithmic assumptions imply algorithmic consequences, and intractability assumptions imply intractability consequences. In these notes we will see instead that the existence of efficient derandomized algorithms is implied by the *intractability* of some other problem, so that a hardness condition implies an algorithm consequence.

In the next section we will introduce some notation about computational problems and complexity measures, and then we will state some results about conditional de-randomization.

2.1 Formal Definitions of Complexity Measures and Complexity Classes

For a decision problem L and an integer n we denote by L_n the restriction of L to inputs of length n. It will be convenient to think of L_n as a Boolean function $L_n : \{0,1\}^n \to \{0,1\}$ (with the convention that $x \in L_n$ if and only if $L_n(x) = 1$).

For a function $f : \{0,1\}^n \to \{0,1\}$, consider the size of the smallest circuit that solves f; denote this number by $CC(f)$. By definition, we have that if C is a circuit with n inputs of size less than $CC(f)$ then there exists an $x \in \{0,1\}^n$ such that $C(x) \neq f(x)$.

[1] More precisely, it is impossible to have a sub-exponential time deterministic "universal derandomization procedure" that given x and oracle access to an arbitrary function $A(\cdot, \cdot)$ outputs 1 when $\mathbf{Pr}_r[A(x,r) = 1] \geq 3/4$ and outputs 0 when $\mathbf{Pr}_r[A(x,r) = 1] \leq 1/4$. In fact, more generally, it is impossible to give sub-exponential time algorithms for all **BPP** problems by using "relativizing" techniques. It is beyond the scope of these notes to explain what this means, and why it is more general. "Relativizations" are discussed in [Pap94], where it is possible to find pointers to the relevant literature.

Consider now, for every n, what is the largest s such that for every circuit C of size $\leq s$ we have $\mathbf{Pr}_{x \in \{0,1\}^n}[C(x) = f(x)] \leq 1/2 + 1/s$; denote this number by $H(f)$.

Recall that $DTIME(T(n))$ is the class of decision problems that can be solved by deterministic algorithms running in time at most $T(n)$ on inputs of length n. We have the classes $\mathbf{E} = DTIME(2^{O(n)})$ and $\mathbf{EXP} = DTIME(2^{n^{O(1)}})$.

2.2 Hardness versus Randomness

From our previous arguments, we have $\mathbf{BPP} \subseteq \mathbf{EXP}$. Since there are settings where probabilistic procedures require exponential time to be simulated, one would conjecture that $\mathbf{BPP} \not\subseteq 2^{n^{o(1)}}$; on the other hand, **BPP** seems to still represent a class of feasible computations, and it would be *very* surprising if $\mathbf{BPP} = \mathbf{EXP}$. As we will see in a moment, something is wrong with the above intuition. Either $\mathbf{BPP} = \mathbf{EXP}$, which sounds really impossible, or else it must be the case that **BPP** has sub-exponential time deterministic algorithms (that will work well only on average, but that would be quite remarkable enough).

Theorem 1 ([IW98]) *Suppose* $\mathbf{BPP} \neq \mathbf{EXP}$*; then for every* **BPP** *language* L *and every* $\epsilon > 0$ *there is a deterministic algorithm* A *that works in time* 2^{n^ϵ} *and, for infinitely many* n*, solves* L *on a fraction* $1 - 1/n$ *of the inputs of length* n.

This gives a non-trivial simulation of **BPP** under an uncontroversial assumption. We can also get an optimal simulation of **BPP** under an assumption that is much stronger, but quite believable.

Theorem 2 ([IW97]) *Suppose there is a problem* L *in* **E** *and a fixed* $\delta > 0$ *such that, for all sufficiently large* n*,* $CC(L_n) \geq 2^{\delta n}$*; then* $\mathbf{P} = \mathbf{BPP}$.

We will call the statement "there is a problem L in **E** and a fixed $\delta > 0$ such that, for all sufficiently large n, $CC(L_n) \geq 2^{\delta n}$" the "IW assumption." Note that if the IW assumption is true, then it is true in the case where

$$L = \{(M, x, 1^k) : \text{ machine } M \text{ halts within } 2^k \text{ steps on input } x \ \}$$

Notice also that L cannot be solved by algorithms running in time $2^{o(n)}$, and so it would be a little bit surprising if it could be solvable by circuits of size $2^{o(n)}$, because it would mean that, for general exponential time computations, non-uniformity buys more than a polynomial speed-up. In fact it would be very surprising if circuits of size $2^{.99n}$ existed for L.

The two theorems that we just stated are the extremes of a continuum of results showing that by making assumptions on the hardness of problems in **E** and **EXP** it is possible to devise efficient deterministic algorithms for all **BPP** problems. The stronger the assumption, the more efficient the simulation.

Notice that the assumption in Theorem 2 is stronger than the assumption in Theorem 1 in two ways, and that, similarly, the conclusion of Theorem 2 is stronger than the conclusion in Theorem 1 in two ways. On the one hand, the assumption in Theorem 2 refers to circuit size, that is, to a non-uniform measure of complexity, whereas the assumption in Theorem 1 uses a uniform measure of complexity (running time of probabilistic algorithms). This difference accounts for the fact that the conclusion of Theorem 2 gives an algorithm that works for all inputs, while the conclusion of Theorem 1 gives an algorithm that works only for most inputs. The other difference is that Theorem 2 assumes exponential hardness, while Theorem 2 assumes only super-polynomial hardness. This is reflected in the running time of the consequent deterministic simulations (respectively, polynomial and sub-exponential).

When one makes the non-uniform assumption that there is a problem in **E** that requires circuits of size $s(n)$, then the consequence is a deterministic simulation of **BPP** in time roughly $2^{s^{-1}(n^{O(1)})}$ [ISW99, SU01, Uma02]. So if one assumes that **E** requires super-polynomial circuits, **BPP** can be simulated in time $2^{n^{o(1)}}$, if one assumes that **E** requires circuits of size $2^{\Omega(n)}$ then the simulation runs in time $n^{O(1)}$, if one assumes that **E** requires circuits of size $n^{\log n}$ then the simulation runs in time $2^{2^{O(\sqrt{\log n})}}$, and so on. The result of [IW98] does not scale up so well when one is willing to make stronger uniform assumptions. In particular, the following is an open question.

Conjecture 1 *Suppose* $\mathbf{E} \not\subseteq \bigcap_{\delta>0} BPTIME(2^{\delta n})$*; then for every* **BPP** *language* L *there is a deterministic polynomial time algorithm* A *that, for infinitely many* n*, solves* L *on a fraction* $1-1/n$ *of the inputs of length* n*.*

3 Pseudorandom Generators

We say that a function $G : \{0,1\}^t \to \{0,1\}^m$ is a (s,ϵ)-pseudorandom generator if for every circuit D of size $\leq s$ we have

$$|\mathbf{Pr}_r[D(r) = 1] - \mathbf{Pr}_z[D(G(z)) = 1]| \leq \epsilon$$

Suppose that we have a probabilistic algorithm A such that for inputs x of length n the computation $A(x,\cdot)$ can be performed by a circuit of size $s(n)$; suppose that for every x we have $\mathbf{Pr}_r[A(x,r) = \text{ right answer }] \geq 3/4$, and suppose that we have a $(s,1/8)$ pseudorandom generator $G : \{0,1\}^{t(n)} \to \{0,1\}^{m(n)}$. Then we can define a new probabilistic algorithm A' such that $A'(x,z) = A(x,G(z))$. It is easy to observe that for every x we have

$$\mathbf{Pr}_z[A'(x,z) = \text{ right answer }] \geq 5/8$$

and that, using the trivial derandomization we can get a deterministic algorithm A'' that always works correctly and whose running time is 2^t times the sum of the running time of A plus the running time of G.

If t is logarithmic in m and s, and if G is computable in poly(m,s) time, then the whole simulation runs in deterministic polynomial time. Notice also that if we have a (s,ϵ)-pseudorandom generator $G : \{0,1\}^t \to \{0,1\}^m$, then for every $m' \leq m$ we also have, for a stronger reason, a (s,ϵ) pseudorandom generator $G' : \{0,1\}^t \to \{0,1\}^{m'}$ (G' just computes G and omits the last $m-m'$ bits of the output). So there will be no loss in generality if we consider only generators for the special case where, say, $s = 2m$. (This is not really necessary, but it will help reduce the number of parameters in the statements of theorems.) We have the following easy theorem.

Theorem 3 *Suppose there is a family of generators* $G_m : \{0,1\}^{O(\log m)} \to \{0,1\}^m$ *that are computable in poly*(m) *time and that are* $(2m,1/8)$*-pseudorandom; then* $\mathbf{P} = \mathbf{BPP}$*.*

Of course this is only a sufficient condition. There could be other approaches to proving (conditionally) $\mathbf{P} = \mathbf{BPP}$, without passing through the construction of such strong generators. Unfortunately we hardly know of any other approach, and anyway the (arguably) most interesting results are proved using pseudorandom generators.[2]

4 The two main theorems

4.1 The Nisan-Wigderson Theorem

Theorem 4 (Special case of [NW94]) *Suppose there is $L \in \mathbf{E}$ and $\delta > 0$ such that, for all sufficiently large n, $H(L_n) \geq 2^{\delta n}$; then there is a family of generators $G_m : \{0,1\}^{O(\log m)} \to \{0,1\}^m$ that are computable in poly(m) time and that are $(2m, 1/8)$-pseudorandom (in particular, $\mathbf{P} = \mathbf{BPP}$).*

Notice the strength of the assumption. For almost every input length n, our problem has to be so hard that even circuits of size $2^{\delta n}$ have to be unable to solve the problem correctly on more than a fraction $1/2 + 2^{-\delta n}$ of the inputs. A circuit of size 1 can certainly solve the problem on a fraction at least $1/2$ of the inputs (either by always outputting 0 or by always outputting 1). Furthermore, a circuit of size 2^n always exist that solves the problem on every input. A circuit of size $2^{\delta n}$ can contain, for example, the right solution to our problem for every input whose first $(1-\delta)n$ bits are 0; the circuit can give the right answer on these $2^{\delta n}$ inputs, and answer always 0 or always 1 (whichever is better) on the other inputs. This way the circuit is good on about a fraction $1/2 + 2^{-(1-\delta)n}$ of the inputs. So, in particular, for every problem, there is a circuit of size $2^{n/2}$ that solves the problem on a fraction $1/2 + 2^{-n/2}$ of the inputs. It is somewhat more tricky to show that there is in fact even a circuit of size $2^{(1/3+o(1))n}$ that solves the problem on a fraction $1/2 + 2^{-(1/3+o(1))n}$ of the inputs, and this is about best possible for general problems [ACR97].

4.2 Worst-case to Average-case Reduction

Theorem 5 ([BFNW93, Imp95, IW97]) *Suppose there is $L \in \mathbf{E}$ and $\delta > 0$ such that, for all sufficiently large n, $CC(L_n) \geq 2^{\delta n}$; Then there is $L' \in \mathbf{E}$ and $\delta' > 0$ such that, for all sufficiently large n, $H(L'_n) \geq 2^{\delta' n}$.*

This is quite encouraging: the (believable) IW assumption implies the (a priori less believable) NW assumption. Notice how Theorem 2 follows from Theorems 4 and 5.

The results on polynomial reconstruction from Lecture 11 imply Theorem 5, although this is not the way it was originally proved. Our proof followed [STV01].

[2] Some exceptions are discussed below. Andreev et al. [ACR98] show that in order to deterministically simulate probabilistic algorithms it is enough to construct *hitting set generators*, a seemingly weaker primitive than a pseudorandom generator. The complicated proof of [ACR98] was simplified in subsequent work [ACRT99, BF99, GW99]. Andreev et al. [ACR99] also show how to construct hitting set generators, but only under very strong complexity assumptions. Miltersen and Vinodchandran [MV99] give a very elegant construction of hitting set generators, but it also requires a stronger hardness assumption than in [IW97]. On the other hand, [MV99] also gets a stronger conclusion, and, in particular, it is not known how to prove the main result of [MV99] (about the "derandomization" of two-rounds interactive proofs) using pseudorandom generators.

References

[ACR97] A.E. Andreev, A.E.F. Clementi, and J.D.P. Rolim. Optimal bounds for the approximation of boolean functions and some applications. *Theoretical Computer Science*, 180:243–268, 1997.

[ACR98] A.E. Andreev, A.E.F. Clementi, and J.D.P. Rolim. A new general derandomization method. *Journal of the ACM*, 45(1):179–213, 1998.

[ACR99] A.E. Andreev, A.E.F. Clementi, and J.D.P. Rolim. Worst-case hardness suffices for derandomization: A new method for hardness vs randomness trade-offs. *Theoretical Computer Science*, 221:3–18, 1999.

[ACRT99] A.E. Andreev, A.E.F. Clementi, J.D.P. Rolim, and L. Trevisan. Weak random sources, hitting sets, and BPP simulations. *SIAM Journal on Computing*, 28(6):2103–2116, 1999. Preliminary version in *Proc of FOCS'97*.

[BF99] H. Buhrman and L. Fortnow. One-sided versus two-sided error in probabilistic computation. In *STACS'99*, pages 100–109, 1999.

[BFNW93] L. Babai, L. Fortnow, N. Nisan, and A. Wigderson. BPP has subexponential time simulations unless EXPTIME has publishable proofs. *Computational Complexity*, 3(4):307–318, 1993.

[GW99] O. Goldreich and A. Wigderson. Improved derandomization of BPP using a hitting set generator. In *RANDOM'99*, pages 131–137, 1999.

[Imp95] R. Impagliazzo. Hard-core distributions for somewhat hard problems. In *Proceedings of the 36th IEEE Symposium on Foundations of Computer Science*, pages 538–545, 1995.

[ISW99] R. Impagliazzo, R. Shaltiel, and A. Wigderson. Near-optimal conversion of hardness into pseudo-randomness. In *Proceedings of the 40th IEEE Symposium on Foundations of Computer Science*, 1999.

[IW97] R. Impagliazzo and A. Wigderson. $P = BPP$ unless E has sub-exponential circuits. In *Proceedings of the 29th ACM Symposium on Theory of Computing*, pages 220–229, 1997.

[IW98] R. Impagliazzo and A. Wigderson. Randomness versus time: De-randomization under a uniform assumption. In *Proceedings of the 39th IEEE Symposium on Foundations of Computer Science*, pages 734–743, 1998.

[MV99] P.B. Miltersen and N.V. Vinodchandran. Derandomizing Arthur-Merlin games using hitting sets. In *Proceedings of the 40th IEEE Symposium on Foundations of Computer Science*, pages 71–80, 1999.

[NW94] N. Nisan and A. Wigderson. Hardness vs randomness. *Journal of Computer and System Sciences*, 49:149–167, 1994. Preliminary version in *Proc. of FOCS'88*.

[Pap94] C.H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

[STV01] M. Sudan, L. Trevisan, and S. Vadhan. Pseudorandom generators without the XOR lemma. *Journal of Computer and System Sciences*, 62(2):236–266, 2001.

[SU01] R. Shaltiel and C. Umans. Simple extractors for all min-entropies and a new pseudo-random generator. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*, pages 648–657, 2001.

[Uma02] C. Umans. Pseudo-random generators for all hardnesses. In *Proceedings of the 34th ACM Symposium on Theory of Computing*, pages 627–634, 2002.