

Lecture 18: Time-Bounded Derandomization

Instructor: Dieter van Melkebeek

Scribe: Ozcan ILIKHANI

DRAFT

In the last lecture we completed our discussion of Space-Bounded Derandomization and start considering Time-Bounded Derandomization.

The results we are aiming for are conditional results as we mentioned last time in the space bounded setting. In a time bounded setting we do not have any unconditional results but conditional ones. However, condition seemed to be reasonable. One result is the following:

Theorem 1. *If there exists a language $L \in \mathbf{E}$ such that $H_L(m) \geq m^{\omega(1)}$, then we can build a G_r with $\ell(r) = r^O(1)$, so $BPP \subseteq SUBEXP = \bigcap_{\epsilon > 0} DTIME(2^{n^\epsilon})$*

Definition 1. *For a language L , the average-case hardness of L at input length m , denoted $H_L(m)$, is the largest s such that no circuit of size at most s can compute L correctly on at least a $\frac{1}{2} + \frac{1}{s}$ fraction of the inputs of length m .*

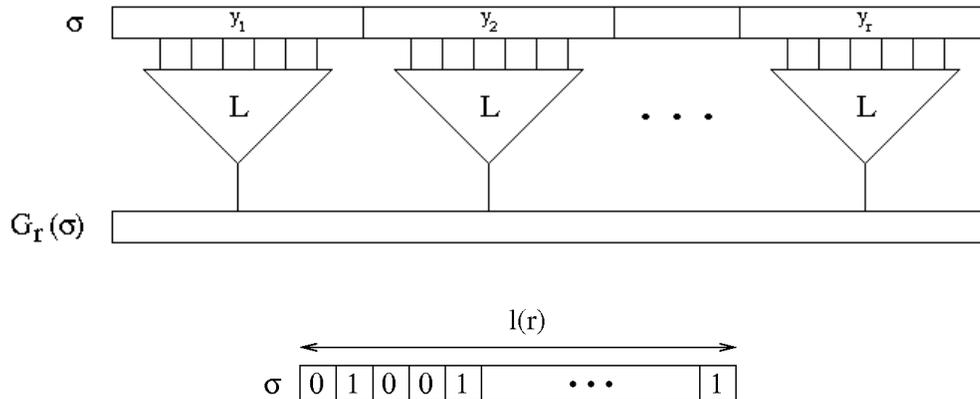
Note that computing L correctly on at least a $1/2$ fraction of the inputs at a given length is trivial—either constant 0 or constant 1 will do the job. An average-case hard function is one that is not only hard to compute exactly, but also hard to compute correctly on noticeably more than half the inputs.

One might wonder why the above definition uses s to refer to both the size of the circuits under consideration and the degree of hardness. The main reason is to keep the number of parameters small, so that our analysis works out cleanly. We combine the constraints of having the chance of being correct be very close to $1/2$ and the size of the circuit be very large. One might also wonder why we are measuring hardness against *nonuniform* circuits. There is a very good reason for this—our arguments will crucially use this nonuniformity.

Let us gain some intuition about Definition 1. As mentioned in a previous lecture, every predicate on m bits can be computed exactly by a circuit of size at most 2^m , so $H_L(m) < 2^m$ for all L . As s gets smaller, the condition of Definition 1 gets easier to satisfy: the circuits under consideration become more computationally restricted, *and* they're required to compute L on more inputs. Thus $H_L(m)$ serves as a measure of the average-case hardness of L at input length m .

0.1 PRG Construction

Suppose we have a language L such that $H_L(m) \geq r/\epsilon$. Then no circuit of size r/ϵ , and in particular no circuit of size r , can compute L with probability at least $\frac{1}{2} + \frac{\epsilon}{r}$ over inputs of length m chosen uniformly at random. This suggests the following approach for constructing a PRG: given $y_1, \dots, y_r \in \{0, 1\}^m$ chosen independently and uniformly at random, apply L (viewed as a function producing a bit) to each y_i , yielding an output string of length r . Intuitively, if some bit of the output distribution of this PRG were predictable, then since the samples y_i are chosen independently, such a predictor would have an advantage in computing L . Thus the output distribution of this PRG would be unpredictable and hence indistinguishable from the uniform distribution, as desired.



We will not proceed to formalize this very vague idea because it is seriously flawed. One issue that naturally arises is that computing the output of such a function would require computing L , which is assumed to be hard to compute. However, the length m at which we would be computing L would ideally be much less than the output length r , so the complexity of computing L might not be prohibitive. A much more critical problem is that this construction takes a seed of length mr but only outputs r bits! We want our seed length to be much smaller than r , and certainly not larger. It is trivial to build a PRG when the seed length is at least as large as the output length — we can just output some of the bits of the seed, yielding a uniform output distribution. The reason the above construction requires such a long seed is that all y_i 's are to be chosen independently. We would like to show that by sacrificing some independence of the y_i 's, we can drastically reduce the seed length without the quality of the output distribution deteriorating by too much. Say we let the y_i 's overlap a little bit thus reducing the incoming number of random bits. If they overlap in a logarithmic number of positions, they will still be random enough.

We will accomplish this by taking a seed σ of length $\ell(r) > m$ and selecting r subsets S_i ($i = 1, \dots, r$) of the bit positions of the seed, and letting $y_i = \sigma|_{S_i}$ be the bits of the seed indexed by S_i . For example, if $S_1 = \{1, 3, \ell(r)\}$ and the seed σ is as illustrated below, then $y_1 = 001$.

The desired subset construction is formalized in the following definition.

Definition 2. An (m, a) -design of size r over $[\ell] = \{1, \dots, \ell\}$ is a sequence of subsets $S_1, \dots, S_r \subseteq [\ell]$ such that $|S_i| = m$ for all i , and $|S_i \cap S_j| \leq a$ for all $i \neq j$.

We want our PRG output length r to be large, but at the same time we want the pairwise intersections of the S_i 's to be small so that the y_i 's are “as independent as possible.” These two goals are at odds with each other, but the following lemma shows that, in fact, not only do there exist such designs with large r , but the subsets can be efficiently computed.

Lemma 1. For all r and $m \geq \log r$, there exists an efficiently computable $(m, \log r)$ -design of size r over $[\ell]$ where $\ell = O(m^2)$.

Proof. Assume m is a prime power so that $\text{GF}(m)$ is a field. If it's not, we can pad to the next prime power to make it a field. Let the S_i 's be the graph of a univariate polynomial of degree at most n over $\text{GF}(m)$. The graph here mean the pairs $(x, f(x))$. $S_i = \{(x, q_i(x)) : x \in \text{GF}(m)\}$

where q_i is the i th univariate polynomial. So the universe is $GF(m) \times GF(m)$, that is, with x in the first is $GF(m)$ and $f(x)$ in the second. $|S_i| = m$ for all i . Then $|S_i \cap S_j| \leq n$ follows from the fact that the polynomials of degree less than n represent distinct functions and can have at most n points in common. We still need enough polynomials, so we need m large enough. That is, we need $m^n \geq r$. This however is implied by $m \geq \log r$. To do this efficiently, we can assume that m is prime, not just a prime power, to make the calculations easier. The calculations can then be done in $2^{O(\ell(n))}$, which is fast enough. □

Note that under a different construction, we can make $\ell = O(\log r)$ instead of $\ell = O(\log r^2)$.

The condition that $m \geq \log r$ is no problem for us since we will want m to be such that $H_L(m) \geq \frac{r}{\epsilon} \geq r$, and since $H_L(m) \leq 2^m$, we get the constraint $m \geq \log r$ anyway.

We are now in a position to fully specify our PRG. For a given r and m , we can set $\ell(r) = O(m^2)$ and obtain an $(m, \log r)$ -design S_1, \dots, S_r via Lemma 1. Then our PRG G_r will be

$$G_r(\sigma) = L(\sigma|_{S_1})L(\sigma|_{S_2}) \cdots L(\sigma|_{S_r}).$$

It is straightforward to verify that if L is computable in linear exponential time, then this PRG is quick. All that remains is to show that this construction fools circuits of size r if L is sufficiently hard.

Theorem 2. *If $H_L(m) \geq \frac{r}{\epsilon}$ and $\epsilon \leq \frac{1}{r}$, then the above construction is an ϵ -PRG for circuits of size r .*

Proof. We will prove the theorem by contradiction. Suppose that for some circuit C of size r , we have

$$\left| Pr_{\sigma \in \{0,1\}^{\ell(r)}} [C(G_r(\sigma)) = 1] - Pr_{\rho \in \{0,1\}^r} [C(\rho) = 1] \right| \geq \epsilon.$$

We will show that then $H_L(m) < \frac{r}{\epsilon}$ by exhibiting a circuit of size at most $\frac{r}{\epsilon}$ that solves L at length m on at least a $\frac{1}{2} + \frac{\epsilon}{r}$ fraction of the inputs, thus contradicting the assumed hardness of L .

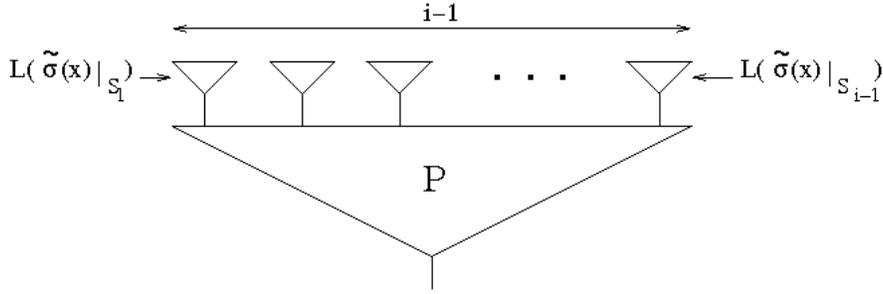
There exists an $i \in \{1, \dots, r\}$ and a circuit P of size at most r such that

$$Pr [P((G_r(\sigma))_1, \dots, (G_r(\sigma))_{i-1}) = (G_r(\sigma))_i] \geq \frac{1}{2} + \frac{\epsilon}{r}.$$

We would like to use P to construct a small circuit that will approximate L well at length m . Intuitively, P seems to be approximating L on input $\sigma|_{S_i}$, and in fact, by another averaging argument we can fix some setting to the bits of σ other than those indexed by S_i such that the predictor P maintains its ϵ/r advantage. Here we are again critically use the fact that we are working with nonuniform circuits. Renaming $\sigma|_{S_i}$ to x and letting $\tilde{\sigma}(x)$ denote the $\ell(r)$ bits where x fills the positions indexed by S_i and the rest of the positions are fixed as above, we have

$$Pr [P((G_r(\tilde{\sigma}(x)))_1, \dots, (G_r(\tilde{\sigma}(x)))_{i-1}) = L(x)] \geq \frac{1}{2} + \frac{\epsilon}{r}$$

where the probability is over x chosen uniformly at random from $\{0, 1\}^m$. This is exactly the sort of behavior we would like, but we need to construct a circuit that takes input x , whereas P takes the first $i - 1$ bits of G_r 's output. We cannot just attach a circuit computing G_r to P , since computing G_r involves computing L exactly. But now we come to the most critical observation of the entire



argument: each input to P depends only on at most $\log r$ bits of x , since $|S_i \cap S_j| \leq \log r$ for $j \neq i$, and can thus be computed from x by a circuit of size at most $2^{\log r} = r$. Having the limited pairwise intersections of the subsets in the design is the key to avoiding the inherent complexity of computing L , allowing us to get a contradiction. To the j th input of P , we can attach a circuit of size at most r computing $(G_r(\tilde{\sigma}(x)))_j = L(\tilde{\sigma}(x)|_{S_j})$ from x , as illustrated below.

Since $i < r$, we have thus obtained a circuit of size at most $r^2 \leq r/\epsilon$ that succeeds in computing $L(x)$ with probability at least $\frac{1}{2} + \frac{\epsilon}{r}$ over the choice of x . Since $|x| = m$, we conclude that $H_L(m) < \frac{r}{\epsilon}$. We have our contradiction, and the theorem is proved. \square

To summarize the above proof, we assumed that our PRG's output distribution was distinguishable from the uniform distribution by a small circuit, obtained a small predictor circuit for some bit of the pseudorandom distribution, and connected some additional small circuitry to this predictor to convert it into a small circuit that approximated L well, thus showing that L cannot be too hard. We can eliminate the ϵ parameter by setting $\epsilon = 1/r$ to obtain the following clean corollary.

Corollary 1. *If $L \in \mathbf{E}$, m , and r are such that $H_L(m) \geq r^2$, then there exists a quick $\frac{1}{r}$ -PRG for circuits of size r with seed length $O(m^2)$.*

We remark that no languages in \mathbf{E} are known to satisfy the hardness condition $H_L(m) \geq r^2$ on arbitrary circuits for interesting values of m , say $m = r^{o(1)}$. We list some results with reasonable assumptions.

If $L \in \mathbf{E}$, m , and r are such that $H_L(m) \geq m^{\omega(1)}$, a superpolynomial, then $\ell(r) = r^{o(1)}$. We need $H_L(m) \geq r^2$. with $\ell = m^2$. So r^ϵ is enough. This gives us subexponential time.

If $L \in \mathbf{E}$, m , and r are such that $H_L(m) \geq 2^{m^{\Omega(1)}}$ then $\ell(r) = \log r^{O(1)}$. Quasi-polynomial time

The most we can hope for is that $L \in \mathbf{E}$, m , and r are such that $H_L(m) \geq 2^{\Omega(m)}$. In which case $\ell(r) = \log r^2$ which while doesn't give us the polynomial time, we can do the other setup to get $\ell(r) = \log r$

Again we don't know if any such L actually exist.

However, we will see in the next lecture that if there exist languages in \mathbf{E} requiring large circuits in the worst case, which is conjectured to be true, then such average-case hard languages do exist. We will use error correction codes. Notice also that the results give us a relationship between hardness and randomness. Intuitively, if there exists a hard problem, then randomness is useless.

We can do a similar construction with branching problems. Derandomization gives BPP lower bounds.

Acknowledgements

In writing the notes for this lecture, I perused the notes by Tom Watson and Andrew Bolanowski for lecture 17 from the Spring 2010 offering of CS 710.