CS 880: Pseudorandomness and Derandomization 4/29/2013

Lecture 26: PRGs from shrinkage

Instructors: Holger Dell and Dieter van Melkebeek Scribe: Gautam Prakriya

DRAFT

In the last lecture we showed that with high probability, a pseudorandom restriction ρ picked from an $O(\log n)$ - wise independent distribution shrinks branching programs of size s. Today, we use this result as a black box to construct pseudorandom generators for branching programs. In fact we present a generic construction which works for any non-uniform computational model for which such a shrinkage result holds.

1 Recap

Throughout this lecture f will represent a branching program, n the number of variables in f and s the size of f (The number of nodes in f. With a slight abuse of notation we let $s(.)$ represent the size function. We begin by restating the shrinkage result from the previous lecture.

Lemma 1. *Let R_p be a p -regular $\log n$ - wise independent distribution over $\{1,*\}^n$, then for every branching program f of size s,*

$$\Pr_{\rho\sim R_p}[s(f|_\rho) \geq c\log(1/\epsilon)\cdot p\cdot s] \leq \epsilon].$$

We let s_0 denote $c\log(1/\epsilon)\cdot p\cdot s$. The following corollary is an easy consequence of this lemma.

Corollary 2. *With* $\Pr \geq 1-\epsilon$.

- *$f|_\rho$ depends on at most s_0 variables.*
- *$f|_\rho$ can be described with $O(s_0\log s_0)$ bits.*

2 PRG construction

We first present a construction that doesn't quite work, but will help motivate the actual PRG construction.

2.1 Attempt 1

1. Sample a random restriction $\rho \in R_p$. Corollary 2 tells us that w.h.p the restricted B.P. $f|_\rho$ depends on at most s_0 variables. This takes $O(\log^3 s)$ bits. (We require $O(\log^2 n)$ uniform bits to generate a $O(\log n)$ - wise independent distribution over $\{1,*\}^n$ and to ensure that this distribution is p - regular, we need $O(\log(1/p)\cdot\log^2 n)$ bits.)

2. Set the active variables in ρ by sampling s_0 uniform bits. There is a small technical issue here - $f|_\rho$ could depend on more than s_0 variables, in this case we fill the remaining unset variables with 1's.

Since $s_0 = \sqrt{s}$, the seed length for this construction is $\sqrt{s} + O(\log^3 s) = O(\sqrt{s})$. Let us denote the output distribution of this construction by W. To prove that this is a PRG for branching programs, we need to show that $|\Pr[f(U_n) = 1] - \Pr[f(W) = 1]| = |\Pr[f(U_n) = 1] - \Pr[f|_\rho(U_{c\sqrt{s}}) = 1]|$ is small. This doesn't neccasarily hold. For instance consider a branching program f that computes the Majority function. It is not difficult to see that that $f(W)$ is biased towards 1. One could argue that this counter-example works only because the restrictions are picked from $\{1, *\}^n$. This is true, and it may be possible to find a pseudorandom p-regular distribution over $\{0,1\}^n$ for which this construction does yield a PRG, but we would no longer be using the shrinkage result in a black box fashion.

2.2 Attempt 2

To get around the issue in the above construction, we sample a number of restritions $\rho_1, \dots, \rho_t$, where $t = \log(n/\epsilon)/p$. t is chosen so that w.h.p. every variable is left active by one of the restrictions. Let $A(\rho_i) \subseteq [n]$ be the set of variables left active by ρ_i. Let $W_1, \dots, W_t$ denote the independent distributions obtained as in attempt 1 from $\rho_1, \dots, \rho_t$ respectively, and let $\mathbb{W} = W_1 \oplus \dots \oplus W_t$.

Lemma 3. *$\mathbb{W}$ is ϵ close to the uniform distribution U_n.*

Proof. Note that $\mathbb{W}$ is the uniform distribution $\cup A(\rho_i) = [n]$. So we only need to bound the probability of the event $(\exists j : j \notin \cup A(\rho_i))$.

$$\Pr[\exists j : j \notin \cup A(\rho_i)] \leq \sum_j \prod_i \Pr[j \notin \cup A(\rho_i)] = \sum_j (1-p)^t \leq n \cdot e^{-pt} = \epsilon.$$

We showed that $d(\mathbb{W}, U_n) \leq \epsilon$. But generating $\mathbb{W}$ takes more than n random bits. So we need a more conservative way of setting the active variables in the restrictions. The idea is to use an extractor with a fixed source X and t different seeds, $Y_1, \dots Y_t$. Below we describe the PRG:

1. Sample $\rho_1, \dots, \rho_t$ independently from R_p. Sample a source X of the extractor. X will be a uniformly random binary string of length $O(s_0 \log s_0)$.

2. Sample independent strings $Y_1, \dots, Y_t$ and for each i, let V_i be the distribution obtained by setting the active variables in ρ_i using the string $E(X, Y_i)$. As earlier if the length of the output of the extractor is less than the number of active variables in ρ_i, set the remaining variables to 1.

3. Output $\mathbb{V} = V_1 \oplus \dots \oplus V_t$.

Lemma 4. *$d(f(\mathbb{V}), U_n) \leq 5 \cdot \epsilon \cdot t$.*

Proof. The proof is by a hybrid argument. Define for $1 \leq i \leq t+1$, the hybrid distributions $Z_i = W_1 \oplus \dots \oplus W_{i-1} \oplus V_i \oplus \dots V_t$. Note that $Z_{t+1} = \mathbb{W}$ and $Z_1 = \mathbb{V}$. Lemma 3 tells us that $d(f(Z_{t+1}), f(U)) \leq \epsilon$, it is therefore sufficient to show that $d(f(Z_i), f(Z_{i+1})) \leq 5 \cdot \epsilon$. For notational convenience, let us define $Z := W_1 \oplus \dots \oplus W_{i-1} \oplus V_{i+1} \oplus \dots V_t$ also Define $f_Z(x) := f(x \oplus Z)$. Notice that using this notation, $Z_i = Z \oplus V_i$ and $Z_{i+1} = Z \oplus W_i$. We need to prove that $d(f_Z(V_i), f_Z(W_i)) \leq 5 \cdot \epsilon$. The intuition behind the argument is as follows, By Corollary 2, we know that w.h.p. $f_Z|_{\rho_i}$ can be described with $O(s_0 \log s_0)$ bits. Let us call this event G. We can then argue that conditioned on G, X has min-entropy at least $H_\infty(X) - O(s_0 \log s_0)$. So if $H_\infty(X)$ is sufficiently large, $E(X, Y_i)$ is close to the uniform distribution.

To show that conditioned on G, X has min-entropy at least $H_\infty(X) - O(s_0 \log s_0)$, we use an argument we saw in an earlier lecture. Let h denote the random function $f_Z|_{\rho_i}$. G denotes the event $(s(h) \leq s_0)$. Let $\mathcal{F} = \{g : \Pr_{Z,\rho_i}[h = g] \geq \epsilon / s_0^{cs_0}\}$. Using the fact that there are at most $s^{O(s)}$ branching programs of size s, it is not difficut to see $\Pr[\not(G \wedge h \in \mathcal{F})] \leq 2 \cdot \epsilon$. Now let $g \in \mathcal{F}$,

$$H_\infty(X|h = g) \geq H_\infty(X) - \log(1/\epsilon) - cs_0 \log s_0.$$

Fact 5. *There exists an explicit functions* $E : 0, 1^N \times 0, 1^d 0, 1^m$ *that is a* $(N/2, \epsilon)$*-extractor with* $m = N/4$ *and* $d = O(log(N/\epsilon))$.

Since we start with a truly random source, $H_\infty(X) = N$. If $N \geq 2 \log(1/\epsilon) + 2cs_0 \log s_0$, then $E(X, Y_i)$ is ϵ - close to the uniform distribution. This implies that $d(f_Z(V_i), f_Z(W_i)) \leq 5 \cdot \epsilon$. We skip the details of this argument. □

We leave it as an exercise to show that the seed length of this construction is $s^{1/2+o(s)}$.