CS59200-PRS: Pseudorandomness Oct 8th, 2025

Lecture 13: Nisan-Zuckerman Generator

Lecturer: Wei Zhan *Scribe: Xiuyu Ye*

1 More Constructions of Extractors

1.1 Extractor from expanders

Recall the stronger form of the expander mixing lemma.

Theorem 1 (Expander Mixing Lemma)**.** *If a d-regular graph $H = (V, E)$ with $|V| = n$ has $(1 - \lambda)$ spectral expansion, then for all $S, S' \subseteq V$,*

$$\left| \frac{e(S, S')}{dn} - \frac{|S|}{n^2} \frac{|S'|}{n} \right| \leq \lambda \cdot \sqrt{\frac{|S|}{n} \cdot \frac{|S'|}{n}}.$$

Theorem 2. *If $H \colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^n$, as a 2^d-regular graph on 2^n vertices, is $(1-\lambda)$ spectral expanding, then H is also a (k, ε)-extractor for $\varepsilon = \lambda \cdot 2^{(n-k)/2}$.*

Proof. Let $H = (V, E)$ be the corresponding 2^d-regular graph where $V = \{0,1\}^n$.

Given a flat k-source X on $\{0,1\}^n$, we think of it as the uniform distribution over $S \subseteq V$ where $|S| = 2^k$. And given a distinguisher $A \colon \{0,1\}^n \to \{0,1\}$, let the subset $S' \subseteq V$ be $S' = \{v \in V \mid A(v) = 1\}$. This way, $A(H(x, r)) = 1$ if and only if the edge from $x \in S$ labeled with $r \in \{0,1\}^d$ lands in S'. Therefore

$$\begin{aligned}
\left| \mathop{\mathbb{E}}_{x \sim X, r \sim \{0,1\}^d} [A(H(x, r))] - \mathop{\mathbb{E}}_{v \sim V} [A(v)] \right| &= \left| \frac{e(S, S')}{|S| \cdot 2^d} - \frac{|S'|}{2^n} \right| \\
&\leq \lambda \cdot \frac{2^n}{|S|} \cdot \sqrt{\frac{|S|}{2^n}} && \text{(Theorem 1)} \\
&= \lambda \cdot 2^{(n-k)/2} && \square
\end{aligned}$$

If we take extractor to be a power of Ramanujan graph, we have $\lambda = O(2^{-d/2})$. Hence the output length $m = n$, which is too good, at the cost of the seed length $d = n - k + 2\log(1/\varepsilon)$ which is exponentially worse than optimal.

1.2 Extractor from expander random walks

Let $H\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^n$ be a 2^d-regular graph $H = (V, E)$. Instead of using the function directly as a extractor, we construct an extractor with random walks on H as follows.

Let $(v_0, r_1, \ldots, r_t)$ be a t-step random walk on H starting from a random vertex $v_0 \sim V$, where $r_i \in \{0,1\}^d$. In each step of the random walk, we let $v_i = H(v_{i-1}, r_i)$. Consider an extractor $\mathrm{Ext}\colon \{0,1\}^{n+dt} \times [t] \to \{0,1\}^n$ that outputs the i-th step vertex in the random walk,

$$\mathrm{Ext}((v_0, r_1, \ldots, r_t), i) = v_i.$$

Theorem 3. *If H is γ-spectral expanding for some constant d and γ, then* Ext *is a $(k = \delta(n+dt), \varepsilon)$-extractor for some constant $\delta < 1$ and $\varepsilon > 0$.*

Proof. For distinguisher $A\colon \{0,1\}^n \to \{0,1\}$, by expander Chernoff bound,

$$\Pr_{v_0, r_1, \ldots, r_t}\left[\left|\frac{1}{t}\sum_{i=1}^{t} A(v_i) - \mathop{\mathbb{E}}_{v\sim V}[A(v)]\right| \geq \varepsilon/2\right] \leq 2 \cdot e^{-\frac{1}{16}\gamma t\varepsilon^2}.$$

That means the number of such random walks (where $\left|\frac{1}{t}\sum A(v_i) - \mathbb{E}[A(v)]\right| \geq \varepsilon/2$) is at most $2e^{-\frac{1}{16}\gamma t\varepsilon^2} \cdot 2^{n+dt}$. If the random walk $(v_0, r_1, \ldots, r_t)$ is sampled from a k-source X instead, then

$$\Pr_{(v_0, r_1, \ldots, r_t)\sim X}\left[\left|\frac{1}{t}\sum_{i=1}^{t} A(v_i) - \mathop{\mathbb{E}}_{v\sim V}[A(v)]\right| \geq \varepsilon/2\right] \leq 2e^{-\frac{1}{16}\gamma t\varepsilon^2} \cdot 2^{n+dt} \cdot 2^{-k}.$$

Therefore we can bound the total variation distance as

$$\begin{aligned}\left|\mathop{\mathbb{E}}_{(v_0, r_1, \ldots, r_t)\sim X, i\sim[t]}[A(v_i)] - \mathop{\mathbb{E}}_{v\sim V}[A(v)]\right| &\leq \mathop{\mathbb{E}}_{(v_0, r_1, \ldots, r_t)\sim X}\left[\left|\frac{1}{t}\sum_{i=1}^{t} A(v_i) - \mathop{\mathbb{E}}_{v\sim V}[A(v)]\right|\right] \\ &\leq \frac{\varepsilon}{2} + 2e^{-\frac{1}{16}\gamma t\varepsilon^2} \cdot 2^{n+dt} \cdot 2^{-k}.\end{aligned}$$

When $k = \delta(n+dt)$ for some δ close enough to 1, it suffices to have $t = O(n/\varepsilon^2)$ to bound it by $\varepsilon/2 + 2^{-n}$, which is at most ε as long as $n \geq \log(1/\varepsilon) + 1$. □

The above extractor has output length $m = n$, but the actual input length is $n' = n + dt$. Since $t = O(n/\varepsilon^2)$, for constant ε we have $m = O(n')$, which is asymptotically optimal. And the amount of extra randomness is $\log t = \log n + 2\log(1/\varepsilon) + O(1)$, which is also close to optimal.

The drawback is that the input source has to be good: the min-entropy must be at least $\delta n'$ for some δ close to 1. The following extractor, which we are not going to describe in details, allows freer choices of the parameters while achieving the same asymptotics:

Theorem 4 (Guruswami-Umans-Vadhan)**.** *For every constant $\delta < 1$ and $\varepsilon > 0$, there exists an explicit (k, ε)-extractor* $\mathrm{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *with output length $m = \delta k$ and seed length $d = O(\log n + \log(1/\varepsilon))$.*

2 Nisan-Zuckerman Generator

Back to the initial intuition that extractors are another abstraction of the notion of recycling random bits, we now think of how to construct PRGs from extractors. An immediate idea is to do something similar to the INW generator: We recursively construct, from a given PRG $G_n : \{0,1\}^n \to \{0,1\}^{n'}$, a new PRG

$$G'(s, r) = (G_n(s), G_m(\mathrm{Ext}(s, r)))$$

using the extractor $\mathrm{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$. However, for good extractors with small seed length, the output length m will be much smaller than the input length n. As a result, the output length of the PRG will increase by merely an $o(1)$ factor instead of doubling in each recursion, making the number of recursion steps and thus the amount of total extra randomness much larger than desired.

The Nisan-Zuckerman PRG takes the other route by sequentially applying the extractor instead of recursively. The basic Nisan-Zuckerman construction for a PRG $G\colon \{0,1\}^{n+dt} \to \{0,1\}^{tm}$ is

$$G(s, r_1, \ldots, r_t) = (\mathrm{Ext}(s, r_1), \mathrm{Ext}(s, r_2), \ldots, \mathrm{Ext}(s, r_t))$$

Lemma 5. *If* $\mathrm{Ext}\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *is a* (k, ε)*-extractor, then generator* G δ*-fools every width-*w*, length-*tm *ROBP for* $\delta = t \cdot (\varepsilon + w \cdot 2^{k-n})$.

Proof. Consider the following hybrid. In each step i, define

$$G_i(s, r_1, \ldots, r_t) \;:=\; (\mathrm{Ext}(s, r_1), \ldots, \mathrm{Ext}(s, r_i), x_{i+1}, \ldots, x_t)$$

where $x_{i+1}, \ldots, x_t$ are random bits.

Given a width-w, length-tm ROBP B, let $F_{i,r_1,\ldots,r_i}(s) \in [w]$ be the state of the ROBP at the im-th layer, reached by following the partial input $(\mathrm{Ext}(s, r_1), \ldots, \mathrm{Ext}(s, r_i))$. Then

$$\begin{aligned}
|\mathbb{E}[B(G_i)] - \mathbb{E}[B(G_{i+1})]| &\le d_{\mathrm{TV}}\left((F_{i,r_1,\ldots,r_i}(s), \mathrm{Ext}(s, r_{i+1})), (F_{i,r_1,\ldots,r_i}(s), x_{i+1})\right) \\
&\le \mathop{\mathbb{E}}_{r_1,\ldots,r_i}\left[d_{\mathrm{TV}}\left((F_{i,r_1,\ldots,r_i}(s), \mathrm{Ext}(s, r_{i+1})), (F_{i,r_1,\ldots,r_i}(s), x_{i+1})\right)\right] \\
&\le \varepsilon + w \cdot 2^{k-n},
\end{aligned}$$

where the last inequality follows from the fact that extractor implies recycling. Hence,

$$|\mathbb{E}[B(G)] - \mathbb{E}[B(x)]| \le t \cdot (\varepsilon + w \cdot 2^{k-n}). \qquad \square$$

Notice that this PRG does not have very good stretch: it maps a seed of length $(n + dt)$ to an output of length tm. Since $m \le n + d$, it is easy to see that $n + dt \ge \sqrt{tm}$. So the stretch is polynomial instead of exponential, which is not desirable for typical logspace derandomization.

However, there is a scenario where the above PRG excels, that is when the length of the ROBP is small (or equivalently when the width is large). For instance, if the length of the ROBP is $O(\log^{1.5} w)$ (instead of $\mathrm{poly}(w)$), we can get a RPG of seed length $O(\log w)$ which is still none trivial. In fact, this is true every ROBP of polylog$w$ length by recursively applying the above PRG.

2.1 Full Nisan-Zuckerman generator for short and wide ROBPs

Suppose B is a width-w, length-$\log^c w$ ROBP for some constant c. Recall that both Nisan's PRG and INW generator has seed length $O(\log(nw) \cdot \log w)$ which will be $O(\log^2 w)$ in this scenario.

Let us fix some parameters in Lemma 5 to construct the initial PRG against B. Let $\delta > 0$ be a small enough constant, and let

$$\ell_0 = \log^c w,\ m = \log w,\ t = \log^{c-1} w,\ n = 4\log w,\ k = 2\log w,\ \varepsilon = \frac{1}{2}\delta t^{-1}.$$

Instantiate Lemma 5 with the GUV extractor gives $d = O(\log n + \log(1/\varepsilon))$, and we get a PRG $G\colon \{0,1\}^{\ell_1} \to \{0,1\}^{\ell_0}$ that δ-fools B with

$$\begin{aligned}\ell_1 &= n + dt \\ &= O(\log w + (\log\log w + \log t) \cdot \log^{c-1} w) \\ &= O(\log w + \log^{c-1} w \log\log w).\end{aligned}$$

We need another good property of the GUV extractor that it can be computed in space linear in its input size, i.e. $O(n+d)$ space. As a result, the PRG G can also be computed in $O(n+d) = O(\log w)$ space. We now think of the composite function $B \circ G : \{0,1\}^{\ell_1} \to \{0,1\}$, which simulates the behaviour of B since

$$\left| \mathop{\mathbb{E}}_{s\sim\{0,1\}^{\ell_1}}[B(G(s))] - \mathop{\mathbb{E}}_{x\sim\{0,1\}^{\ell_0}}[B(x)] \right| \le \delta,$$

while $B \circ G$ can be computed by a ROBP of length ℓ_1 and width $w' = w \cdot 2^{O(n+d)} = \mathrm{poly}(w)$. Applying Lemma 5 on $B \circ G$ with the parameters

$$m' = \log w',\ t' = \frac{l_1}{m} = O(\log^{c-2} w' \log\log w),\ n' = 4\log w',\ k' = 2\log w',\ \varepsilon' = \frac{1}{2}\delta t'^{-1}$$

gives another PRG $G\colon \{0,1\}^{\ell_2} \to \{0,1\}^{\ell_1}$ that δ-fools $B \circ G$ with

$$\ell_2 = O(\log w + \log^{c-2} w \log\log w).$$

By repeating the above process a constant $\lceil c \rceil$ rounds of recursion, we get a series of PRG $G^{(i)}\colon \{0,1\}^{\ell_i} \to \{0,1\}^{\ell_{i-1}}$ that δ-fools $B \circ G \circ G' \circ \ldots \circ G^{(i-1)}$, with

$$\ell_i = O(\log w + \log^{c-i} w \log\log w).$$

Notice that during the process the parameters m, n and k will always be $O(\log w)$, even though the width becomes polynomially larger in every round of recursion.

Our final generator will be $G \circ G' \circ \ldots \circ G^{(c)}$, which has seed length $\ell_c = O(\log w)$ and can also be computed in space $O(\log w)$. It $c\delta$-fools B since

$$\left| \mathbb{E}[B(x)] - \mathbb{E}[B \circ G \circ \ldots \circ G^{(c)}(s)] \right|$$

$$\leq \sum_{i=0}^{c-1} \left| \mathbb{E}[B \circ G \circ \ldots \circ G^{(i)}(s_i)] - \mathbb{E}[B \circ G \circ \ldots \circ G^{(i+1)}(s_{i+1})] \right| \leq c\delta.$$

One immediate implication of the Nisan-Zuckerman generator is that, randomized logspace algorithms that only uses polylog(n) many random bits can be fully derandomized into deterministic logspace algorithms.