CS59200-PRS: Pseudorandomness Sep 17, 2025

Lecture 6,7: Nisan's RPG

Lecturer: Wei Zhan *Scribe: Arvind Ramaswami*

1 Nisan's PRG

The starting point of Nisan's PRG is the idea of recursive construction. Suppose that we already have a PRG $G\colon \{0,1\}^\ell \to \{0,1\}^{n/2}$ that ε-fools all length-$n/2$, width-w ROBPs. How can we double the output length to obtain a new PRG that fools all length-n, width-w ROBPs?

The natural idea is to use the PRG G twice, so that the new PRG is $(G(s), G(s'))$ for some $s, s' \in \{0,1\}^\ell$. If s is uniformly random so that the first half of the ROBP is fooled, what should s' be?

- If s' is uniform and independent from s, we can indeed show that $(G(s), G(s'))$ is pseudorandom by a hybrid argument:

  $$(G(s), G(s')) \approx (G(s), x_2) \approx (x_1, x_2)$$

  for $x_1, x_2 \sim \{0,1\}^{n/2}$. However, while the output length of the PRG is doubled, the seed length also doubles, so in the end we cannot really save any randomness.

- If we went to the other extreme and let $s' = s$, then $G(s') = G(s)$ and they do not look like random even to a width-2 ROBP (which for instance could check the first bits of $G(s)$ and $G(s')$ being the same).

- Therefore, although we want s' to be deterministically computed from s so that we don't have to introduce extra randomness, we want the function $h\colon s \mapsto s'$ be as complicated as possible, so that an ROBP with bounded width cannot tell the relation between s and s'. A (pseudo)random function is with high probability complicated, and thus is perfectly suitable for the job.

Lemma 1. *If $h : \{0,1\}^\ell \to \{0,1\}^l$ is drawn from a pairwise uniform hash function family, then for all $A : \{0,1\}^{2\ell} \to [0,1]$, w.p. at least $1 - \frac{1}{2^\ell \varepsilon^2}$ over h, we have*

$$\left| \mathop{\mathbb{E}}_{s,s' \sim \{0,1\}^\ell} [A(s,s')] - \mathop{\mathbb{E}}_{s \sim \{0,1\}^\ell} [A(s, h(s))] \right| \le \varepsilon.$$

Proof. This is a direct implied by the pairwise independent Chebyshev's inequality, since the random variables $A(s,h(s))$ for $s \in \{0,1\}^\ell$ (with randomness from h) are pairwise independent and each has expectation $\mathop{\mathbb{E}}_h[A(s,h(s))] = \mathop{\mathbb{E}}_{s'}[A(s,s')]$. □

Now for the branching program $B : \{0,1\}^n \to \{0,1\}$, we can take $A(s, s') = B(G(s), G(s'))$ and use the hybrid argument to easily show that:

Corollary 2. *If $G : \{0,1\}^\ell \to \{0,1\}^{n/2}$ ε-fools all length-$\frac{n}{2}$, width w ROBPs, then for every length-n, width-w ROBP B, w.p. at least $1 - \frac{1}{2^\ell \delta^2}$ over h from a pairwise uniform hash family, the PRG $G_h \colon \{0,1\}^\ell \to \{0,1\}^n$ defined as*

$$G_h(s) = (G(s), G(h(s)))$$

$(2\varepsilon + \delta)$-fools B.

This motivates us to recursively construct Nisan's PRG, doubly the output length in each step. Initially let $G \colon \{0,1\}^\ell \to \{0,1\}^\ell$ be $G(s) = s$, so that G is truly random. And in the i-th step, let

$$G_{h_1,\ldots,h_k}(s) = (G_{h_1,\ldots,h_{k-1}}(s), G_{h_1,\ldots,h_{k-1}}(h_k(s)))$$

for a new hash function h_k drawn from the pairwise uniform family. For instance

- $G_{h_1}(s) = (s, h_1(s))$,
- $G_{h_1,h_2}(s) = (s, h_1(s), h_2(s), h_1(h_2(s)))$,
- $G_{h_1,\ldots,h_k}(s) = (s, h_1(s), h_2(s), h_1(h_2(s)), h_3(s), \ldots, h_1(h_2(\ldots h_k(s)\ldots)))$.

Notice that the i-th block of $G_{h_1,\ldots,h_k}$ can be easily computed from the binary representation of i.

We should be able to argue that with high probability over the choices of $h_1, \ldots, h_k$, $G_{h_1,\ldots,h_k}$ fools all length-$2^k\ell$, width-w ROBPs. However, Corollary 2 is not sufficient, as its assumption too strong for its conclusion to match (the assumption is that the PRG from the previous step universally fools all ROBP, but we could only conclude the PRG after the current step fools a single ROBP). We need some weaker assumption, that instead of G fooling all ROBPs, it only needs to fool every ROBP with the same transition functions as the half of B (but the initial states and final states could vary).

1.1 Matrix Formulation

Recall the in the lecture from last week we showed that the transition functions of the ROBP B have a succinct matrix formulation. Assuming each layer of B has the same transition function, we let $M_0, M_1 \in \{0,1\}^{w \times w}$ be that $M_0[i,j] = 1$ (resp. $M_1[i,j] = 1$) if and only if state j goes to state i in the next layer with an edge labeled with 0 (resp. 1). Then the execution of B on input $x \in \{0,1\}^n$ can be represented by $M_x = M_{x_n} \cdots M_{x_1}$, and for $M = (M_0 + M_1)/2$,

$$\mathop{\mathbb{E}}_{x \sim \{0,1\}^n}[M_x] = M^n.$$

The matrix M is a *stochastic matrix* in the sense that every column of M is a distribution. To fool B, we want to approximately compute the n-th power of the stochastic matrix M, and for convenience we give the following definition.

Definition 1. *The PRG $G : \{0,1\}^\ell \to \{0,1\}^n$ ε-fools M^n, for $M = (M_0 + M_1)/2$ corresponding to a width-w ROBP, if*

$$\left\| \mathop{\mathbb{E}}_{s \sim \{0,1\}^\ell} [M_{G(s)}] - M^n \right\|_1 \leq \varepsilon.$$

Remark. Here $\|\cdot\|_1$ is the operator norm induced by the ℓ_1-norm on vectors. In other words, for $M \in \mathbb{R}^{n \times n}$,

$$\|M\|_1 = \max_{|v|_1 = 1} |Mv|_1 = \max_i \sum_j |M[j,i]| .$$

We will use two properties of the norm: First, $\|M\|_1 = 1$ for every stochastic matrix M. Second, the norm is submultiplcaitive, i.e. $\left\|MM'\right\|_1 \leq \|M\|_1 \left\|M'\right\|_1$.

Lemma 3. *If $h : \{0,1\}^\ell \to \{0,1\}^\ell$ is drawn from a pairwise uniform hash function family, then for every $M \in \mathbb{R}^{w \times w}$ corresponding to a width-w ROBP, and every function $G : \{0,1\}^\ell \to \{0,1\}^{n/2}$, w.p. at least $1 - \frac{w^4}{2^\ell \varepsilon^2}$ over h we have*

$$\left\| \mathop{\mathbb{E}}_{s,s' \sim \{0,1\}^\ell} [M_{G(s),G(s')}] - \mathop{\mathbb{E}}_{s \sim \{0,1\}^\ell} [M_{G(s),G(h(s))}] \right\|_1 \leq \varepsilon.$$

Proof. In Lemma 1, taking $A(s,s') = M_{G(s),G(s')}[i,j]$, we have that the difference in each entry of the matrix

$$\left| \mathop{\mathbb{E}}_{s,s'} [M_{G(s),G(s')}[i,j]] - \mathop{\mathbb{E}}_{s} [M_{G(s),G(h(s))}[i,j]] \right| \leq \frac{\epsilon}{w}$$

w.p. at least $1 - \frac{w^2}{2^\ell \varepsilon^2}$ over h. The claim of the theorem is implied by using the union bound over $i,j \in [w]$, and the fact that the difference in matrix 1-norm is at most the sum of w differences in a single column. □

We can now apply the hybrid argument to prove a more useful version of Corollary 2 as follows.

Corollary 4. *Let $M \in \mathbb{R}^{w \times w}$ correspond to a width-w ROBP. If $G : \{0,1\}^\ell \to \{0,1\}^{n/2}$ ε-fools $M^{n/2}$, then w.p. $\geq 1 - \frac{w^4}{2^\ell \delta^2}$ over h from a pairwise uniform hash function family,*

$$G_h(s) = (G(s), G(h(s)))$$

$(2\varepsilon + \delta)$-fools M^n.

Proof. By triangle inequality we have

$$\begin{aligned}\left\| \mathop{\mathbb{E}}_{s\sim\{0,1\}^\ell}[M_{G_h(s)}] - M^n \right\|_1 \le & \left\| \mathop{\mathbb{E}}_{s\sim\{0,1\}^\ell}[M_{G(s),G(h(s))}] - \mathop{\mathbb{E}}_{s,s'\sim\{0,1\}^\ell}[M_{G(s),G(s')}] \right\|_1 \\ & + \left\| \mathop{\mathbb{E}}_{s,s'\sim\{0,1\}^\ell}[M_{G(s),G(s')}] - \mathop{\mathbb{E}}_{\substack{s\sim\{0,1\}^\ell \\ x_2\sim\{0,1\}^{n/2}}}[M_{G(s),x_2}] \right\|_1 \\ & + \left\| \mathop{\mathbb{E}}_{\substack{s\sim\{0,1\}^\ell \\ x_2\sim\{0,1\}^{n/2}}}[M_{G(s),x_2}] - \mathop{\mathbb{E}}_{x_1,x_2\sim\{0,1\}^{n/2}}[M_{x_1,x_2}] \right\|_1\end{aligned}$$

The first term is at most δ w.p. at least $1 - \frac{w^4}{2^\ell\delta^2}$ over h, by Lemma 3. The second term can be bounded using submultiplicativity since

$$\begin{aligned}\left\| \mathop{\mathbb{E}}_{s,s'}[M_{G(s),G(s')}] - \mathop{\mathbb{E}}_{s,x_2}[M_{G(s),x_2}] \right\|_1 &= \left\| \mathop{\mathbb{E}}_{s}[M_{G(s)}] \left(\mathop{\mathbb{E}}_{s'}[M_{G(s')}] - \mathop{\mathbb{E}}_{x_2}[M_{x_2}] \right) \right\|_1 \\ &\le \left\| \mathop{\mathbb{E}}_{s}[M_{G(s)}] \right\|_1 \left\| \mathop{\mathbb{E}}_{s'}[M_{G(s')}] - \mathop{\mathbb{E}}_{x_2}[M_{x_2}] \right\|_1 \\ &= \left\| \mathop{\mathbb{E}}_{s}[M_{G(s)}] \right\|_1 \left\| \mathop{\mathbb{E}}_{s'}[M_{G(s')}] - M^{n/2} \right\|_1 \le \varepsilon.\end{aligned}$$

This is because $\mathop{\mathbb{E}}_{s}[M_{G(s)}]$ is stochastic, and G ε-fools $M^{n/2}$. Similarly, the third term is also at most ε, and thus the total error is bounded by $2\varepsilon + \delta$. □

To examine the pseudorandomness of Nisan's PRG, notice that initially $G(s) = s$ perfectly fools M^ℓ, and applying Corollary 4 iteratively gives us

- $G_{h_1}(s)$ δ-fools $M^{2\ell}$;
- $G_{h_1,h_2}(s)$ 3δ-fools $M^{4\ell}$;
- ...
- $G_{h_1,\ldots,h_k}(s)$ $(2^k-1)\delta$-fools $M^{2^k\ell}$

with high probability over the choices of $h_1, \ldots, h_k$ by a union bound. We thus conclude the following theorem.

Theorem 5. *Let $h_1, \ldots, h_k : \{0,1\}^\ell \to \{0,1\}^\ell$ be drawn from a pairwise uniform hash family. For every width-w, length-n ROBP ($n = 2^k\ell$) B, w.p. at least $1 - \frac{w^4k}{2^\ell\delta^2}$ over $h_1, \ldots, h_k$, $G_{h_1,\ldots,h_k}(s)$ ε-fools B with $\varepsilon = 2^k\delta$.*

Looking at the parameters, we can choose $k = O(\log n)$ and $\delta = O(1/n)$ to achieve a small constant ε. Therefore we have $\ell = O(\log(w^4 k/\delta^2)) = O(\log(nw))$ for $h_1, \ldots, h_k$ to exist. However, in this case $G_{h_1,\ldots,h_k}$ is not explicit as we don't know which choices of $h_1, \ldots, h_k$ are good. In fact, if we don't care about explicitness then the probabilistic argument in our very first lecture already gives a PRG of seed length $O(\log(nw))$, since there are at most $w^{2nw} \cdot 2^w$ many length-n, width-w ROBPs.

Therefore, we need to think of the randomness used to sample $h_1, \ldots, h_k$ as part of the seed, resulting in the actual seed length $O(k\ell) = O(\log n \cdot \log(nw))$. The good thing is that now $G_{h_1,\ldots,h_k}(s)$ is explicit and can be efficiently computed from the seed, and the PRG fools every width-w, length-n ROBP with error $\varepsilon + \frac{w^4 k}{2^\ell \delta^2}$ which is still small.

2 $\mathsf{BPL} \subseteq \mathsf{SC}$

For logspace computation, Nisan's PRG has seed length $O(\log^2 n)$, so if we naively use it to derandomize BPL we could only get $\mathsf{BPL} \subseteq \mathsf{L}^2$, which is no better than the simple recursive matrix powering algorithm. However, Nisan's PRG has the good property that its seed can be divided into $O(\log n)$ parts and each part "works" with high probability, allowing Nisan to show that $\mathsf{BPL} \subseteq \mathsf{SC}$. Here

$$\mathsf{SC} = \bigcup_{c>0} \mathsf{TISP}(n^c, \log^c n)$$

consists of all languages decidable by an algorithm running simultaneously in polynomial time and poly-logarithmic space (we had two matrix powering algorithms that each separately runs in polynomial time and poly-log space). The name SC means Steve's Class, as a tribute to Stephen Cook.

Here we give a very informal proof of the fact. The observation is that, each h_i in Nisan's PRG is good, if in the corresponding step of the recursive construction, the claim in Lemma 3 holds. If we can compute $G = G_{h_1,\ldots,h_{i-1}}$ efficiently, then we can check if each h_i from the pairwise uniform family is good by computing every entry of the two matrices

$$\mathop{\mathbb{E}}_{s,s' \sim \{0,1\}^\ell}[M_{G(s),G(s')}], \qquad \mathop{\mathbb{E}}_{s \sim \{0,1\}^\ell}[M_{G(s),G(h(s))}]$$

in $n \cdot w^2 \cdot 2^{O(\ell)} = \mathrm{poly}(n)$ time. Once we found the good h_i we can store the length-$O(\ell)$ seed for it, and the bottleneck of space usage is for storing all k good seeds, thus $O(\log^2 n)$ space. Finally, $G_{h_1,\ldots,h_{i-1}}$ can indeed be computed efficiently if we have already stored the seeds for $h_1, \ldots, h_{i-1}$. In particular, we can use the construction $h(x) = s_1 x + s_0$ for $s_0, s_1 \in \mathbb{F}_{2^\ell}$, so that each part of G, which is just different h iteratively applied on s, can be computed in $\mathrm{poly}(\ell)$ time and $O(\ell)$ space.