Nisan's Pseudorandom Generator for Space-Bounded Computation

Ian Mertz

University of Toronto

November 8, 2021

Pseudorandom Generators

Pseudorandom generator $G : \{0,1\}^\ell \to \{0,1\}^n$ ϵ-fools the class $\mathcal{C}$ if for all $C \in \mathcal{C}$,

$$\left| \Pr_{x \sim \mathcal{U}_n}[C(x) = 1] - \Pr_{s \sim \mathcal{U}_\ell}[C(G(s)) = 1] \right| \leq \epsilon$$

small technical note: C only gets one pass over the input

Nisan's pseudorandom generator

Main result [Nisan'90]: for any S, there is a PRG G which 2^{-S}-fools $SPACE(S)$, where the seed length ℓ is $O(S \log n)$.

Additionally, G can be computed in space $O(\ell)$.

(note that $S = o(\log n)$ isn't very interesting)

Nisan's pseudorandom generator

Main result [Nisan'90]: for any S, there is a PRG G which 2^{-S}-fools $SPACE(S)$, where the seed length ℓ is $O(S \log n)$.

This means we can fool logspace machines using $O(\log^2 n)$ random bits, aka $BPL \subseteq L^2$.

(BPL: two-sided error)

Nisan's pseudorandom generator

Main result [Nisan'90]: for any S, there is a PRG G which 2^{-S}-fools $SPACE(S)$, where the seed length ℓ is $O(S \log n)$.

Savich's theorem: $SPACE(S^2) \supseteq NSPACE(S)$

($\supseteq RSPACE(S)$: one-sided error)

Advantages:
1) we get $BPSPACE(S)$
2) gives us a black-box strategy.

Prelims: universal hash family

Hash family $\mathcal{H} \subseteq \{h : \{0,1\}^m \to \{0,1\}^m\}$.

One nice property: for all $x, y \in \{0,1\}^m$, $\Pr_{h\sim\mathcal{H}}[h(x) = y] = 2^{-m}$

Pairwise independence: for all $x_1 \neq x_2, y_1, y_2 \in \{0,1\}^m$,
$\Pr_{h\sim\mathcal{H}}[h(x_1) = y_1 \wedge h(x_2) = y_2] = 2^{-2m}$

Prelims: universal hash family

side note: can actually get a pairwise independent hash family with description length $2m$!

If we just wanted nice property, could pick XOR mask: $h_a(x) = (x \oplus a)$.

Can instead pick $h_{a,b}(x) = (a * x) \oplus b$ (where $(a * x)_j = \sum_i a_{i+j \mod m} \cdot x_i \mod 2$ is the convolution operation).

Prelims: easier view of $SPACE(S)$

Distinguisher Q: move from space S to a DFA with 2^S states (one for each setting of the work tape).

Our PRG will output n blocks of m bits, so we'll let each state of Q have 2^m transitions (aka we'll let it read the whole block at once).

Our goal will be to approximate M^n, where M is the transition matrix defined by $M[i,j] = \Pr_{x \sim \{0,1\}^m}[i \to_x j]$ for each $i, j \in [2^S]$.

Taking two steps

Main idea: can approximate two steps, aka $\Pr_{x_1,x_2}[i \to_{x_1,x_2} j]$, by using $h(x_1)$ in place of x_2.

Define M_h by $M_h[i,j] = \Pr_x[i \to_{x,h(x)} j]$, will show $M_h \approx_\epsilon M^2$.

After that it's just going to be a matter of iterating $\log n$ times.

Taking two steps

Fix $i, j \in [2^S]$, and for $p \in [2^S]$ let $A_{ip} := \{x \in \{0,1\}^m : i \rightarrow_x p\}$ and $B_{pj} := \{x \in \{0,1\}^m : p \rightarrow_x j\}$.

$$
\begin{aligned}
|M_h[i,j] - M^2[i,j]| &= |\Pr_x[i \rightarrow_{x,h(x)} j] - \Pr_{x_1,x_2}[i \rightarrow_{x_1,x_2} j]| \\
&= \sum_p |\Pr_x[i \rightarrow_x p \wedge p \rightarrow_{h(x)} j] - \\
&\quad \Pr_{x_1,x_2}[i \rightarrow_{x_1} p \wedge p \rightarrow_{x_2} j]| \\
&= \sum_p |\Pr_x[x \in A_{ip} \wedge h(x) \in B_{pj}] - \frac{|A_{ip}|}{2^m} \cdot \frac{|B_{pj}|}{2^m}|
\end{aligned}
$$

Mixing Lemma

Let $A, B \subseteq \{0,1\}^m$. We say h is δ-independent for (A, B) if $|\Pr_x[x \in A \wedge h(x) \in B] - \frac{|A|}{2^m} \cdot \frac{|B|}{2^m}| \leq \delta$.

Main lemma: for any A, B, $\Pr_h[h$ is not δ-independent for $(A, B)] < \frac{1}{2^m \delta^2}$.

If true, then for a random h, $|M_h[i,j] - M^2[i,j]| \leq \frac{\epsilon}{2^{2S}}$ except with probability $\frac{2^{6S}}{2^m \epsilon^2}$ (set $\delta := \epsilon / 2^{3S}$, sum over all $p \in [2^S]$).

...which implies M_h and M^2 are ϵ-close in total distance (sum over all $i, j \in [2^S]$) except w.p. $\frac{2^{6S}}{2^m \epsilon^2}$.

Proof of mixing lemma

Main lemma: $\Pr_h[h$ **is not** δ**-independent for** $(A, B)] < \frac{1}{2^m \delta^2}$.

Define $C := \frac{1}{2^m} |\{x \in A : h(x) \in B\}|$. Then

$$\begin{aligned}
\mathbb{E}_h[C] &= \frac{1}{2^m} \sum_{x \in A} \Pr_h[h(x) \in B] \\
&= \frac{1}{2^m} \sum_{x \in A} \frac{|B|}{2^m} \\
&= \frac{|A|}{2^m} \cdot \frac{|B|}{2^m} < 1
\end{aligned}$$

fact: $Var_h[2^m \cdot C] < \mathbb{E}_h[2^m C]$, and so $Var_h[C] < \frac{1}{2^m} \mathbb{E}_h[C] < \frac{1}{2^m}$

Chebyshev: $\Pr[|C - \frac{|A|}{2^m} \cdot \frac{|B|}{2^m}| > \delta] < \frac{Var_h[C]}{\delta^2} < \frac{1}{2^m \delta^2}$

Iterating

Recap: M_h and M^2 are ϵ-close in total distance except with probability $\frac{2^{6S}}{2^{m}\epsilon^2}$ over the choice of h_i.

Taking two steps: $(x, h(x))$ was almost as good as (x_1, x_2).

Taking four steps: $((x, h_1(x)), (h_2(x), h_2(h_1(x))))$ should be almost as good as (x_1, x_2, x_3, x_4).

Proof omitted, but the point is that for each new h_i we double our length and only (roughly) double our ϵ price in closeness, plus an additive $\frac{2^{6S}}{2^{m}\epsilon^2}$ in the potential error of the new h_i.

Defining Nisan's PRG

Seed will be $x \in \{0,1\}^m, h_1, h_2 \dots h_{\log n}$, length is $m + 2m \cdot \log n = O(m \log n)$.

$$G_0(x) := x$$

$$G_k(x, h_1 \dots h_k) := G_{k-1}(x, h_1 \dots h_{k-1}) \circ G_{k-1}(h_k(x), h_1 \dots h_{k-1})$$

Constraints:
- $|M^n - M_{h_1 \dots h_{\log n}}| \leq \epsilon \cdot (n-1) \leq 2^{-S}$
- $\frac{2^{6S} \cdot \log n}{2^m \cdot \epsilon^2} \leq 2^{-S}$

Fix $\epsilon := \frac{2^{-S}}{(n-1)}$, end up with $m := 9S + 2\log(n-1) = O(S)$.

More passes?

Note that we only allow one pass over the random tape (most reasonable definition for space-bounded complexity classes).

$RL[k], BPL[k]$: allow k passes (R^*L, BP^*L: unlimited)

Need to be careful, BP^*L can equal $PSPACE$ if we don't restrict the runtime, and is not known to be in P even if we do...

More error buys two passes

Claim [David-Papakonstantinou-Sidiropoulos'10]: any PRG G which ϵ-fools $SPACE(2S)$ also $\epsilon \cdot 2^{2S}$-fools $SPACE(S)$ with two passes.

Note that we could've picked $\epsilon \geq 2^{-CS}$ for no real cost

We could even pick 2^{-CS^k} if we let $m = (C+1)S^k$, so if we are ok with seed length $O(\log^{O(1)} n)$, we can fool $O(\log^{O(1)} n)$ space even with $O(\log^{O(1)} n)$ passes (iterate $O(\log \log n)$ times).

More error buys two passes

Claim [David-Papakonstantinou-Sidiropoulos'10]: any PRG G which ϵ-fools $SPACE(2S)$ also $\epsilon \cdot 2^{2S}$-fools $SPACE(S)$ with two passes.

Assume otherwise, so FSM Q with 2^S states has $|\Pr_{s \sim \mathcal{U}_\ell}[Q(G^2(s)) = 1] - \Pr_{x \sim \mathcal{U}_n}[Q(x^2) = 1]| > \epsilon \cdot 2^{2S}$.

Define $p_{i,j} = \Pr_s[1 \to_{G(s)} i \wedge i \to_{G(s)} j]$ and $q_{i,j} = \Pr_x[1 \to_x i \wedge i \to_x j]$.

$$\sum_{i,j} |p_{i,j} - q_{i,j}| \geq |\Pr_{s \sim \mathcal{U}_\ell}[Q(G^2(s)) = 1] - \Pr_{x \sim \mathcal{U}_n}[Q(x^2) = 1]| > \epsilon \cdot 2^{2S}$$

and so there exist $i^*, j^* \in [2^S]$ such that $|p_{i^*,j^*} - q_{i^*,j^*}| > \epsilon$.

More error buys two passes

Claim [David-Papakonstantinou-Sidiropoulos'10]: any PRG G which ϵ-fools $SPACE(2S)$ also $\epsilon \cdot 2^{2S}$-fools $SPACE(S)$ with two passes.

New machine Q' to break G in a single pass with probability at least ϵ:

2^{2S} states (i, j) such that $(i, j) \rightarrow_x (i', j')$ iff $i \rightarrow_x i' \wedge j \rightarrow_x j'$.

Start state $(1, i^*)$, accept state (i^*, j^*).

$$\left| \Pr_{s \sim \mathcal{U}_\ell}[Q'(G(s)) = 1] - \Pr_{x \sim \mathcal{U}_n}[Q'(x) = 1] \right| = |p_{i^*,j^*} - q_{i^*,j^*}| > \epsilon$$

Even more passes

Claim [David-Papakonstantinou-Sidiropoulos'10]: for $S = \log n$, Nisan's PRG can be broken in logspace if given $n^{O(1)}$ passes, even for $m = 2^{O(\sqrt{\log n})}$.

No longer true of every PRG, but I believe it is true of every known PRG against logspace (since they're all modifications of Nisan's PRG).

In fact, only need that h_1 is affine (might not be hard to guess how we break it now...)

They make a claim in their paper that if you could fool Q with an arbitrary number of passes, then $L \subsetneq NP$, but we couldn't figure out why that's true.

Even more passes

Claim [David-Papakonstantinou-Sidiropoulos'10]: for $S = \log n$, Nisan's PRG can be broken in logspace if given $n^{O(1)}$ passes, even for $m = 2^{O(\sqrt{\log n})}$.

Treat the blocks as $(y_1 \dots y_{n/2}), (z_1 \dots z_{n/2})$, where either all z_is are uniform or each z_i is $h_1(y_i)$.

$h_1(x) = f_1(x) + b_1$ where f_1 is a linear function (no constant terms). Thus if $y_{i_1} \dots y_{i_t}$ are linearly dependent and t is even,

$$\sum_j h_1(y_{i_j}) = \sum_j f_1(y_{i_j}) + t \cdot b_1 = f_1(\sum_j y_{i_j}) = f_1(0) = 0$$

In other words, if we knew dependent $y_{i_1} \dots y_{i_t}$, we can simply test if $\sum_j z_{i_j} = 0$.

Even more passes

Claim [David-Papakonstantinou-Sidiropoulos'10]: for $S = \log n$, Nisan's PRG can be broken in logspace if given $n^{O(1)}$ passes, even for $m = 2^{O(\sqrt{\log n})}$.

[Mulmuley'87]: finding a set of linearly dependent m-dimensional vectors can be done in $NC^2 \subseteq SPACE(\log^2 m) \cap TIME(m^{O(1)})$.

$m \leq 2^{O(\sqrt{\log n})} < n/2 - 1$, and so a dependency exists and finding it only takes space S. The time to find the dependency and add up all the corresponding z_is is at most $m^{O(1)} < n$.

Even more passes

Claim [David-Papakonstantinou-Sidiropoulos'10]: for $S = \log n$, Nisan's PRG can be broken in logspace if given $n^{O(1)}$ passes, even for $m = 2^{O(\sqrt{\log n})}$.

Technicalities:

- ensuring the connection has even size: find a dependency in $(y_1 \dots y_{n/4})$ and a dependency in $(y_{n/4+1} \dots y_{n/2})$, if either is an even size collection then test that one, otherwise test the union.

- none of the vectors $y_1 \dots y_{n/2}$ are the all-zeroes vector with exponentially large probability

- in the case of random $z_1 \dots z_{n/2}$, $\sum_j z_{i_j} \neq 0$ with exponentially large probability

Open problems

Logarithmic seed length!

Resistance to more passes!