Derandomizing Space-Bounded Computation

Winter 2025

Course Summary & Review

Instructor: William Hoza

The University of Chicago

The complexity class BPL

- Let $f : \{0,1\}^* \to \{0,1\}$
- By definition, $f \in \mathrm{BPL}$ if there exists a Turing machine M such that:
  - There is a read-only input tape 
  - There is a read/write work tape of size $O(\log n)$ 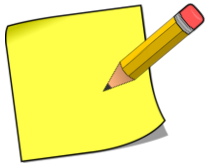
  - There is a read-once random tape 
  - For every $x \in \{0,1\}^*$, we have $\Pr[M(x) = f(x)] \geq 2/3$
  - M halts for every input and every setting of the random tape

Undirected s-t connectivity

- **Theorem** [AKLLR 1979]**:** The undirected s-t connectivity problem is in BPL
- Algorithm: Take a polynomial-length random walk from s, and accept if you ever visit t
- We analyzed this algorithm using the spectral expansion parameter

Spectral expansion parameter

- Let H be a directed regular multigraph
- Identify H with its transition probability matrix. Definition:

$$\lambda(H) = \max_{\pi} \frac{\|\pi H - u\|_2}{\|\pi - u\|_2},$$

where π is a probability vector and u is the uniform probability vector

Derandomization

- AKLLR 1979: Does $\mathrm{L} = \mathrm{BPL}$? (*actually they asked about RL)
- **Conjecture:** $\mathrm{L} = \mathrm{BPL}$
- $\mathrm{L} = \mathrm{BPL}$ would mean that randomness is never necessary for space-efficient computation
- Intensely studied since AKLLR 1979 paper… with considerable success!

Read-once branching programs (ROBPs)

- To prove $\mathrm{L} = \mathrm{BPL}$, it suffices to design a deterministic log-space algorithm for the following problem:

- **Input:** The description of a standard-order ROBP f
- **Output:** A number μ such that $|\mathbb{E}[f] - \mu| \leq 0.1$

Read-once branching programs (ROBPs)

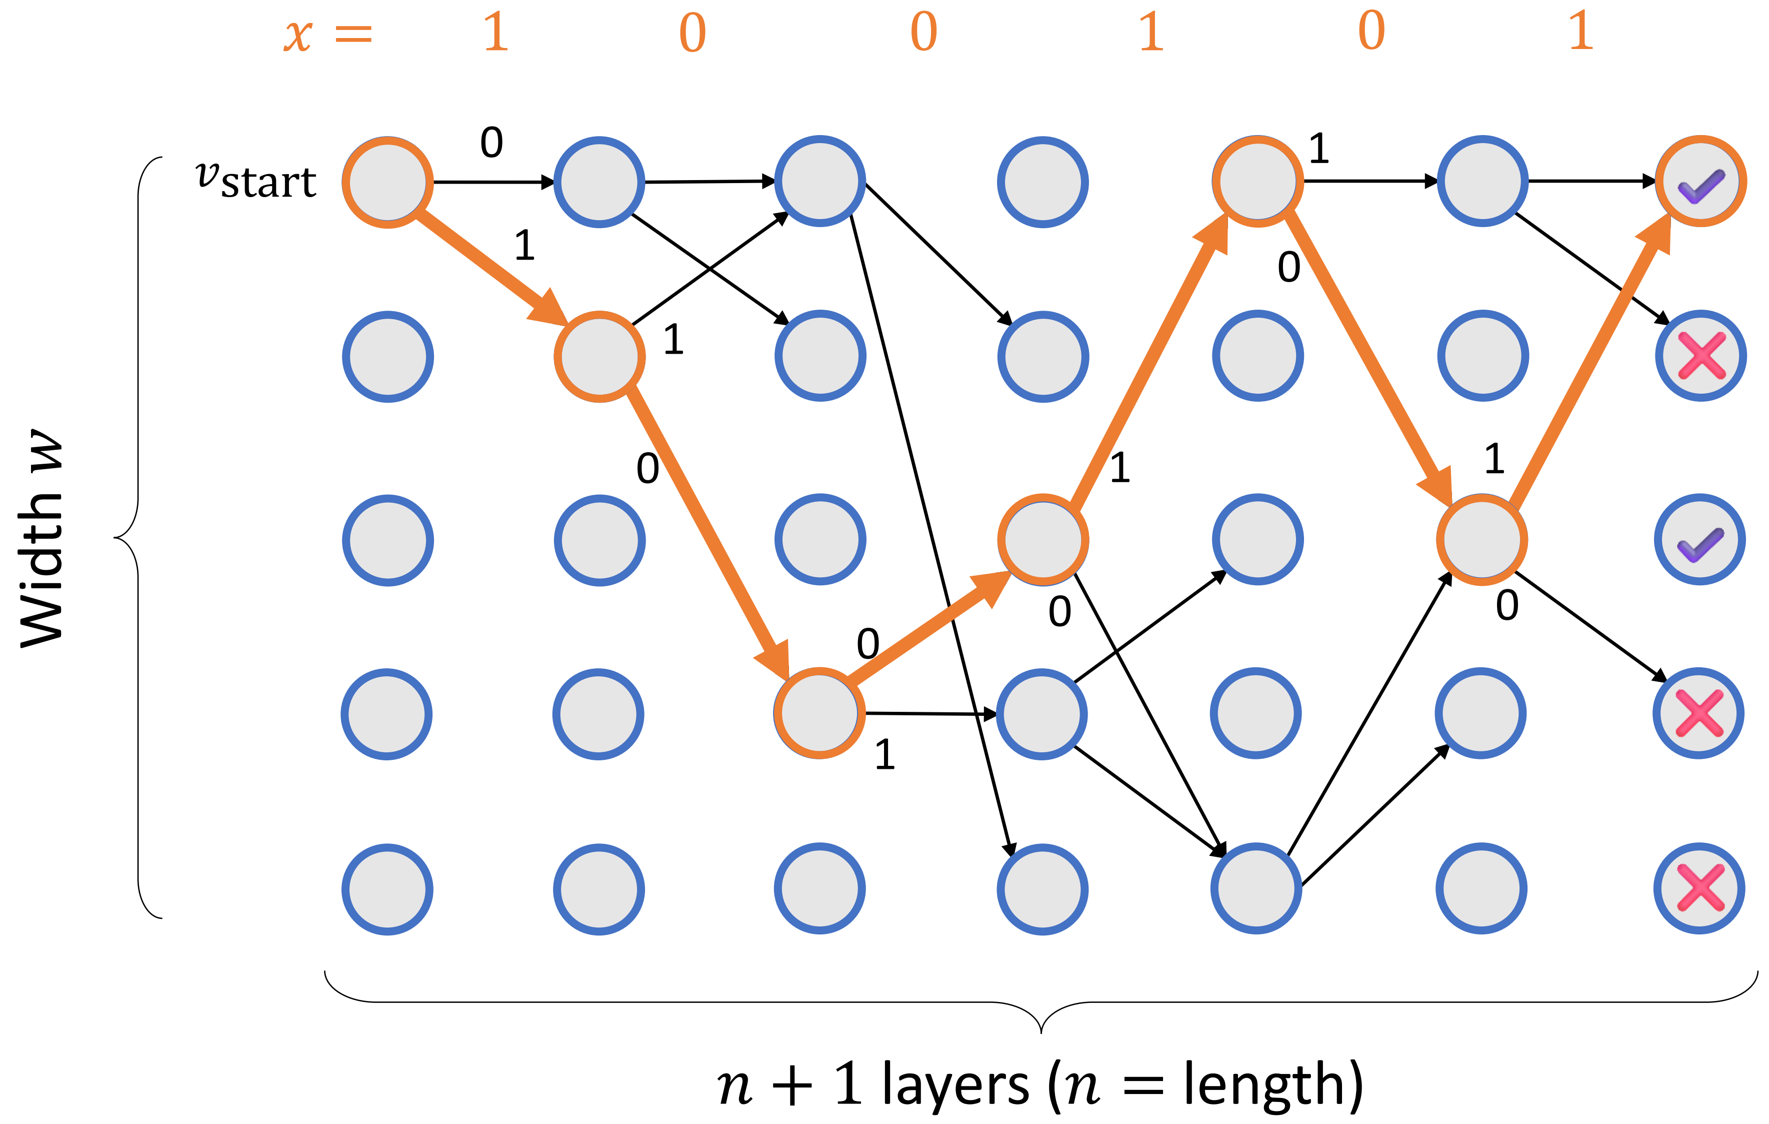

Computes
$f\colon \{0,1\}^n \to \{0,1\}$

Four approaches

- In this course, we studied four approaches to derandomizing $\mathbf{BPL}$:

  1. The INW Approach
  2. The Iterated Restrictions Approach
  3. The Nisan Approach
  4. The Inverse Laplacian Approach

1. The INW Approach

Pseudorandom generators

- A pseudorandom generator (PRG) is a function $G: \{0,1\}^s \to \{0,1\}^n$
- The PRG fools $f: \{0,1\}^n \to \{0,1\}$ with error ε if

$$\left|\mathbb{E}[f] - \mathbb{E}[f(G(U_s))]\right| \leq \varepsilon$$

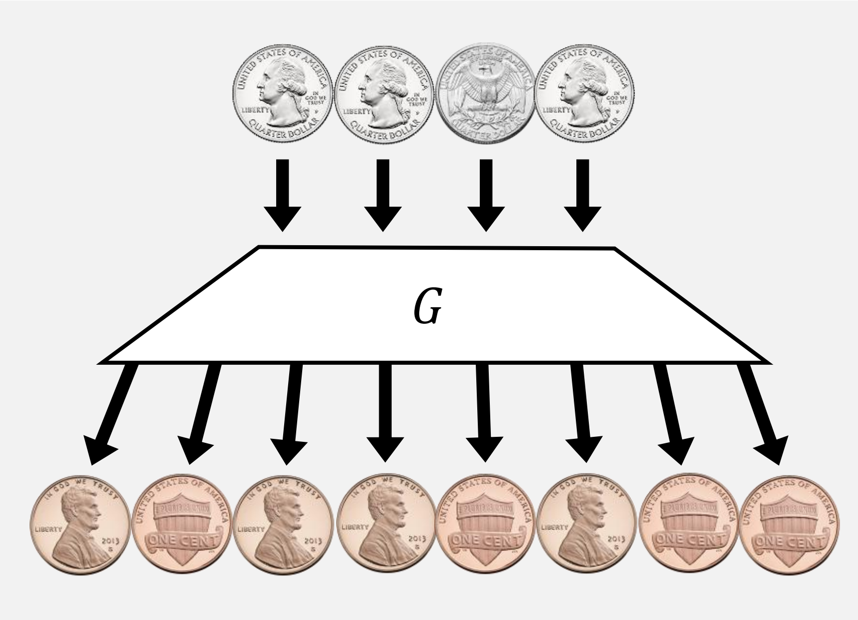

The INW PRG

- **Theorem** [Nisan 1992]**:** For every w, n, ε, there is an explicit PRG that fools width-w length-n standard-order ROBPs with error ε and seed length $O(\log(wn/\varepsilon) \cdot \log n)$
- One example of such a PRG: The INW PRG [Impagliazzo, Nisan, Wigderson 1994]
- Base case: $G_0(x) = x$
- Recursive step: $G_{i+1}(x, y) = \left(G_i(x), G_i(H_{i+1}[x, y])\right)$ for some expander graph H_{i+1}

Expander graphs

- Let H be a regular undirected multigraph
- Informally, we say that H is an expander if H has low degree, and yet $\lambda(H)$ is small
- **Fact:** For every $n \in \mathbb{N}$ and $\lambda \in (0,1)$, there exists an explicit expander on n vertices with $\lambda(H) \leq \lambda$ and $\deg(H) \leq \mathrm{poly}(1/\lambda)$

Analysis of the INW PRG

- Assume by induction that G_i fools width-w programs with error ε_i
- Expander Mixing Lemma $\Rightarrow G_{i+1}$ fools width-w programs with error

$$2 \cdot \varepsilon_i + \lambda(H_{i+1}) \cdot w$$

- Consequently, if $\lambda(H_i) \leq \lambda$ for every i, then $G_{\log n}$ fools width-w programs with error $\lambda \cdot w \cdot n$
- Choose $\lambda = \frac{\varepsilon}{wn}$ ✔

Regular ROBPs

- An ROBP is regular if every vertex has two incoming edges (except the vertices in layer 0)
- **Theorem** [Lee, Pyne, Vadhan 2023]: If $f: \{0,1\}^n \to \{0,1\}$ can be computed by a standard-order ROBP of width w, then f can also be computed by a standard-order regular ROBP of width $O(wn)$

Fooling regular ROBPs

- **Theorem** [Braverman, Rao, Raz, Yehudayoff 2014]**:** If $\lambda(H_i) \leq \lambda$ for every i, then the INW generator $G_{\log n}$ fools width-w standard-order regular ROBPs with error $\lambda \cdot \mathrm{poly}(w) \cdot \log n$
  - Proof based on analyzing the weight of a regular ROBP
- **Corollary:** Can fool such programs with seed length $\tilde{O}(\log(w/\varepsilon) \cdot \log n)$

Reingold's theorem

- **Theorem** [Reingold 2005]**:** Undirected s-t connectivity is in L
- Algorithm idea [Rozenman, Vadhan 2005]:
  - Try all seeds for the INW generator, with suitable $\lambda(H_i)$ values
  - Accept if there is a seed that brings us from s to t
- Analysis based on the derandomized square operation

$$\lambda(G \circledS H) \leq \left(1 - \lambda(H)\right) \cdot \lambda(G)^2 + \lambda(H)$$

2. The Iterated Restrictions Approach

The Forbes-Kelley PRG

- Let D, T, U be independent random variables, each distributed over $\{0,1\}^n$
- Assume U is uniform random, D is $(2k)$-wise uniform, T is k-wise uniform
- **Theorem** [Forbes, Kelley 2018]: $D + (T \wedge U)$ fools width-w length-n ROBPs with error $w \cdot n \cdot 2^{-k/2}$
- Proof uses Fourier analysis

Iterated restrictions

- Define $X \in \{0, 1, \star\}^n$ by

$$X_i = \begin{cases} D_i, & T_i = 0 \\ \star, & T_i = 1 \end{cases}$$

- $D + (T \wedge U)$ fools f, so $\mathbb{E}[f] \approx \mathbb{E}_{X,U}[f|_X(U)]$
- One round ⇒ Assign values to half the variables. Cost $O(\log(wn/\varepsilon) \cdot \log n)$
- Repeat for $O(\log(n/\varepsilon))$ rounds
- ⇒ PRG fooling ROBPs with seed length $O(\log(wn/\varepsilon) \cdot \log(n/\varepsilon) \cdot \log n)$

Arbitrary-order ROBPs

- The Forbes-Kelley seed length is a bit worse than the INW seed length
- However, FK fools arbitrary-order ROBPs!
- That is, if we let $G_\pi(x) = \left(G(x)_{\pi(1)}, \ldots, G(x)_{\pi(n)}\right)$, then G_π fools ROBPs for any permutation $\pi: [n] \to [n]$
- Reason: D_π is still $(2k)$-wise uniform and T_π is still k-wise uniform

The constant-width case

- **Theorem** [Forbes, Kelley 2018]: Using only $\tilde{O}(\log(n/\varepsilon))$ truly random bits, it is possible to assign values to ≈half the variables of a constant-width ROBP while preserving its expectation to within error ε
  - Construction based on small-bias generators

Iterated restrictions with early termination

- Let $\mathcal{F}$ be a subclass of constant-width ROBPs, e.g., read-once CNFs
- Strategy for fooling $\mathcal{F}$ with seed length $\tilde{O}(\log(n/\varepsilon))$:
  1. Do $\mathrm{poly}(\log\log(n/\varepsilon))$ rounds of Forbes-Kelley restrictions
  2. Prove that w.h.p., f simplifies under the restrictions
  3. Use some other approach to fool the simplified f with a short seed

3. The Nisan Approach

Nisan's PRG

- Let $\mathcal{H}$ be a pairwise uniform family of hash functions $h\colon \{0,1\}^k \to \{0,1\}^k$
  - $k = O(\log(wn/\varepsilon))$
- Nisan's PRG:

$$G_{h_1,\ldots,h_{\log n}}(x) = \left(G_{h_1,\ldots,h_{\log n-1}}(x), G_{h_1,\ldots,h_{\log n-1}}\left(h_{\log n}(x)\right)\right)$$

- Pairwise Uniformity Mixing Lemma $\Rightarrow$ Can generate n bits that fool w-state automata with error ε and seed length $O(\log(wn/\varepsilon)) \cdot \log n)$

Good hash functions

- The seed length of Nisan's PRG is not any better than that of the INW PRG
- However, Nisan's PRG has some useful extra structure
- With high probability, h_i is "good" relative to the automaton M and the previous hash functions $h_1, \ldots, h_{i-1}$
  - I.e., M doesn't distinguish $G_{h_1,\ldots,h_i}$ from two copies of $G_{h_1,\ldots,h_{i-1}}$

$\mathrm{BPL} \subseteq \mathrm{SC}$

- **Theorem** [Nisan 1994]: Every problem in BPL can be solved by a deterministic algorithm that simultaneously uses $O(\log^2 n)$ bits of space and $\mathrm{poly}(n)$ time
- **Proof idea:** Exhaustively search for a good h_1, then exhaustively search for a good h_2, then a good h_3, etc.

$\mathrm{BPL} \subseteq \mathrm{L}^{1.5}$

- **Theorem** [Saks, Zhou 1995]: $\mathrm{BPL} \subseteq \mathrm{DSPACE}\left(\log^{3/2} n\right)$
- **Proof idea:** Sample only $\sqrt{\log n}$ hash functions $\vec{h} = \left(h_1, \dots, h_{\sqrt{\log n}}\right)$
- Repeatedly use Nisan's PRG $G_{\vec{h}}$ to approximate $M^{2^{\sqrt{\log n}}}$ (same $\vec{h}$)
- After each application of $G_{\vec{h}}$, perturb and round the entries of the transition probability matrix, to break the correlations with $\vec{h}$

4. The Inverse Laplacian Approach

Inverse Laplacian of an ROBP

- Let f be an ROBP on N vertices
- Let $M \in [0,1]^{N \times N}$ be the transition probability matrix
- Let L be the Laplacian matrix: $L = I - M$
- Then $L^{-1} = M^0 + \cdots + M^n$
- L^{-1} is the matrix of expectations of all subprograms $f_{u \to v}$

Non-black-box error reduction

- **Theorem** [Ahmadinejad, Kelner, Murtagh, Peebles, Sidford, Vadhan 2020]**:** Given the description of a width-n length-n ROBP f, it is possible to deterministically compute μ such that $|\mu - \mathbb{E}[f]| \leq \varepsilon$ using space $O\left(\log^{3/2} n + \log n \cdot \log\log(1/\varepsilon)\right)$
- Proof is based on Richardson iteration: If $A \approx L^{-1}$, then $A \cdot \sum_{i=0}^{m}(I - LA)^i$ is a better approximation for L^{-1}

Weighted PRGs

- A WPRG is a pair (G, ρ) where $G: \{0,1\}^s \to \{0,1\}^n$ and $\rho: \{0,1\}^s \to \mathbb{R}$
- We say that the WPRG fools $f: \{0,1\}^n \to \{0,1\}$ with error ε if

$$\left|\mathbb{E}[f] - \mathbb{E}_{x\in\{0,1\}^s}\left[f\big(G(x)\big) \cdot \rho(x)\right]\right| \leq \varepsilon$$

Low-error WPRGs

- **Theorem** [Braverman, Cohen, Garg 2018]: For every w, n, ε, there is an explicit WPRG that fools width-w length-n standard-order ROBPs with error ε and seed length $\tilde{O}(\log(wn) \cdot \log n + \log(1/\varepsilon))$
- **Proof idea** [Cohen, Doron, Renard, Sberlo, Ta-Shma 2021; Pyne, Vadhan 2021]:
  1. Reverse-engineer Richardson iteration
  2. Use the INW generator to sample a sequence of correlated seeds for A^i term

Hitting sets

- Let $H \subseteq \{0,1\}^n$ and let $\mathcal{F}$ be a class of $f: \{0,1\}^n \to \{0,1\}$
- H is an ε-hitting set for $\mathcal{F}$ if, for every $f \in \mathcal{F}$ such that $\mathbb{E}[f] > \varepsilon$, there is some $x \in H$ such that $f(x) = 1$
- PRG ⇒ WPRG ⇒ Hitting Set

Using hitting sets to derandomize BPL

- **Theorem** [Cheng, H 2020]: Assume $\exists\ O(\log n)$-space-computable 0.5-hitting set for width-n length-n standard-order ROBPs. Then $\mathrm{L} = \mathrm{BPL}$
- Proof idea: Each $x \in H$ is the truth table of a candidate PRG $G^{(x)}: \{0,1\}^{O(\log n)} \to \{0,1\}^n$
- Each candidate PRG $G^{(x)}$ induces a candidate approximation $A^{(x)}$ for L^{-1}
- To judge whether $G^{(x)}$ is a good PRG, check whether $LA^{(x)} \approx I$

Conclusions

- To me, L vs. BPL is the most exciting topic in modern complexity theory
- It is an extremely fundamental topic, like P vs. NP, L vs. P, etc.
- L vs. BPL is special because we can feel optimistic about resolving it!
- We already have many powerful and interesting techniques
- Maybe you have what it takes to prove $\mathrm{L} = \mathrm{BPL}$!

Thank you!

- Being your instructor has been a privilege
- Please fill out the Graduate Course Feedback Form using My.UChicago (deadline is Sunday, March 16)