0368-4283: Space-bounded computation 12/06/18

Take-Home Exam

Amnon Ta-Shma and Dean Doron

General instructions:

1. The deadline for the exam is 12/07/18.
2. Submit your (typed) solution by mail to amnon@tau.ac.il and deandoron@mail.tau.ac.il.
3. Work must be done alone.
4. If you had to use an electronic source, state it explicitly within the relevant question.

Question 1

ZPL is the probabilistic, logspace class with zero-sided error (for every input, with probability at least half the answer is the correct one, and otherwise it is *quit*). Let $\mathsf{ZPL}^\star$ be the variant of ZPL where the machine has *multiple* access to the random tape.

Prove that $\mathsf{BPL} \subseteq \mathsf{ZPL}^\star$.

Question 2

We want to find an explicit, polynomial-length UTS for *consistently labeled* undirected d-regular graphs over n vertices with $\bar{\lambda} < \frac{3}{4}$. Call this family of graphs $\mathcal{G}_n$.

Definition 1. *$\sigma \in [d]^*$ is a UTS for $\mathcal{G}_n$, if for every $G = (V, E) \in \mathcal{G}_n$, $v \in V$, walking on G from v according to σ reaches all the vertices in the graph. σ is explicit, if there exists an algorithm running in* $\mathrm{poly}(n)$ *time outputting it.*

Definition 2. *We say $S \subseteq [d]^\star$ is good for (G, v_0) ($G \in \mathcal{G}_n$) if there exists $\sigma \in S$ such that for every vertex v, a walk over G according to σ starting from v will visit v_0. We say S is explicit if $|S|$ is polynomial, and there exists a* $\mathrm{poly}(n)$ *algorithm that given $i \in [|S|]$ outputs the i-th string in S.*

1. Show that given an explicit S that is good for (G, v_0) for every $G \in \mathcal{G}$ and v_0, one can construct an explicit UTS for $\mathcal{G}$.
2. Construct a sequence of sets $S_i \subseteq [d]^\star$ such that

   $$S_0 = \{\varepsilon\}\,,\ S_{i+1} = \{utu : u \in S_i, t \in [d]\}\,.$$

   Fix $G = (V, E) \in \mathcal{G}$ and v_0. Define

   $$R^{-1}_{G,v_0}(\sigma) = \{v \in V \mid \text{The walk on } G \text{ from } v \text{ according to } \sigma \text{ vists } v_0\}\,.$$

   Define

   $$r_i = \max_{\sigma \in S_i} |R^{-1}_{G,v_0}(\sigma)|.$$

   Prove that $r_0 = 1$ and $r_{i+1} > r_i + \delta$, where $\delta = \frac{r_i(n-r_i)}{4n}$.

   Guidance:

   - Choose $\sigma \in S_i$ that attains r_i. Show that walks according to σ from $V \setminus R^{-1}_{G,v_0}(\sigma)$ end in exactly $n - r_i$ vertices. Denote these vertices by B. Show that there exist enough edges between $R^{-1}_{G,v_0}(\sigma)$ and B.
3. Show that $r_k = n$ for $k = O(\log n)$.
4. Conclude an explicit UTS. What is its cardinality as a function of n and d?

Question 3

In this question we will construct PRGs fooling "almost-balanced" halfspaces. Given $w \in \mathbb{R}^n$ and $\theta \in \mathbb{R}$, the halfspace $H_{w,\theta}\colon \{1,-1\}^n \to \{1,-1\}$ is the function $H_{w,\theta}(x) = \mathsf{sign}(\langle w, x\rangle - \theta)$.

Definition 3. *We say $G\colon \{0,1\}^\ell \to \{1,-1\}^n$ ε-fools $H_{w,\theta}$ if*

$$\left|\mathbb{E}[H_{w,\theta}(U_{\{1,-1\}^n})] - \mathbb{E}[H_{w,\theta}(G(U_\ell))]\right| \le \varepsilon.$$

We measure distance between real-valued distributions P, Q by the CDF distance,

$$d(P,Q) = \|\mathrm{CDF}(P) - \mathrm{CDF}(Q)\|_\infty = \max_{t\in\mathbb{R}} \left| \Pr_{x\sim P}[x < t] - \Pr_{x\sim Q}[x < t] \right|.$$

1. Fix $w \in \mathbb{R}^n$ and $\theta \in \mathbb{R}$. Prove that if $G\colon \{0,1\}^\ell \to \{1,-1\}^n$ is such that

   $$d(\langle w, G(U_\ell)\rangle, \langle w, U_{\{1,-1\}^n}\rangle) \le \varepsilon$$

   then G 2ε-fools $H_{w,\theta}$.

We will assume $\|w\| = 1$, as this can be done without loss of generality. We say a unit-norm $w \in \mathbb{R}^n$ is η-balanced if $\|w\|_\infty \le \eta$. We say a halfspace $H_{w,\theta}$ is η-balanced if w is.

Fooling monotone BPs and a first try at fooling halfspaces

Throughout, all of our BPs are read-once.

Definition 4. *Given a $[W,T]_\Sigma$ BP M and a vertex v at layer i, we let $Acc_M(v)$ be the set of all $z \in \Sigma^{T-i}$ such that starting from v, M accepts z. We say M is monotone if for every layer i there exists an ordering $v_{i_1}, \ldots, v_{i_W}$ of the vertices in layer i such that if $j < k$ then $A_M(v_j) \subseteq A_M(v_k)$.*

2. Prove that for any $w \in \mathbb{R}^n$ and $\theta \in \mathbb{R}$ there exists a monotone $[W,n]_{\Sigma=\{-1,1\}}$ BP M that solves $H_{w,\theta}$ for some W.

Definition 5. *Given a $[W,T]_\Sigma$ BP M, we say that the BPs M_{down} and M_{up} ε-sandwich M if:*

- *For every $z \in \Sigma^T$, $M_{down}(z) \le M(z) \le M_{up}(z)$.*
- $\Pr[M_{up}(U) = 1] - \Pr[M_{down}(U) = 1] \le \varepsilon$.

3. Prove that for any monotone $[W,T]_\Sigma$ BP M there exists M_{down} and M_{up} that ε-sandwich M with width $\frac{2T}{\varepsilon}$ (note that it is independent of the original width W).

4. Prove that if $G\colon \{0,1\}^\ell \to (\Sigma)^T$ δ-fools monotone $[\frac{2T}{\varepsilon}, T]_\Sigma$ BPs then it also $(\varepsilon+\delta)$-fools monotone $[W,T]_\Sigma$ BPs for arbitrary W.

5. Prove that if $G\colon \{0,1\}^\ell \to (\Sigma)^T$ δ-fools monotone $[\frac{2T}{\varepsilon}, T]_\Sigma$ BPs then

   $$d(\langle w, G(U_\ell)\rangle, \langle w, U_{\{1,-1\}^n}\rangle) \le O(\varepsilon + \delta).$$

   and conclude that there exists an explicit PRG that ε-fools any halfspace on n variables with seed-length $O(\log^2 \frac{n}{\varepsilon})$.

A second try at fooling balanced halfspaces

In this section you may use the following theorem, which is a corollary of the Berry-Esseen theorem.

Theorem 6. *Let $Y_1, \ldots, Y_t$ be independent random variables with $\mathbb{E}[Y_i] = 0$, and denote $\sum_i \mathbb{E}[Y_i^2] = \sigma^2$ and $\sum_i \mathbb{E}[Y_i^4] = \rho$. Let $S_n = \frac{1}{\sigma}\sum_i Y_i$. Then,*

$$d(S_n, \mathcal{N}(0,1)) \;\leq\; \frac{\sqrt{\rho}}{\sigma^2},$$

where $\mathcal{N}(0,1)$ is the normal distribution with mean 0 and variance 1.

To define the construction, given n and η we set $t = \frac{1}{\eta^2}$ and require the following ingredients:

- A 2-UFOHFs $\mathcal{H} \subseteq \{h\colon [n] \to [t]\}$, such that $\forall_{i\in[t]},\ |h^{-1}(i)| = |\{x \in [n] \mid h(x) = i\}| = \frac{n}{t}$.
- A 4-UFOHFs $\mathcal{F} \subseteq \left\{f\colon [\frac{n}{t}] \to \{1,-1\}\right\}$, i.e., if we let Y_i be the random variable with value $f(i)$ where f is uniform over F, then for every distinct $i_1, i_2, i_3, i_4 \in [\frac{n}{t}]$, $(Y_{i_1}, Y_{i_2}, Y_{i_3}, Y_{i_4}) = U_{\{1,-1\}^4}$. For $f \in F$ we let $string(f) \in \{-1,1\}^{n/t}$ be the string that at location i has value $f(i)$.

For $x \in \Sigma^n$ and $S \subseteq [n]$, we let $x_{|S}$ denote the substring of x keeping only locations indexed by S (so $x_{|S} \in \Sigma^{|S|}$). Define $G\colon \mathcal{H} \times \mathcal{F}^t \to \{1,-1\}^n$ by

$$G(h; f_1, \ldots, f_t) \;=\; x,$$

where $x_{|h^{-1}(i)} = string(f_i)$ for every $i \in [t]$.

6. Show the construction can be efficiently implemented with seed length $O(\frac{\log n}{\eta^2})$.
7. Prove that when w is η-balanced, $\langle w, U_{\{1,-1\}^n}\rangle$ is η-close to $\mathcal{N}(0,1)$ in the CDF distance.
8. Fix $h \in \mathcal{H}$ and let the probability space be choosing $f_1, \ldots, f_t \in_R \mathcal{F}$. Prove that $\langle w, G(U))\rangle$ is $\sqrt{\zeta_h}$-close to $\mathcal{N}(0,1)$ in the CDF distance, where $\zeta_h = O\left(\sum_{i=1}^t \left(\sum_{j\in h^{-1}(i)} |w_j|^2\right)^2\right)$.
9. A typical $h \in \mathcal{H}$ spreads its weights almost evenly among the buckets. Prove if w is η-balanced, then $\mathbb{E}_{h\in\mathcal{H}}[\zeta_h] = O(\eta^2)$.
10. Prove that G $\varepsilon = O(\eta)$ fools halfspaces on n variables with η-balanced w-s.

A third try at fooling balanced halfspaces

11. Derandomize the selection of $f_1, \ldots, f_t$ to achieve a shorter seed-length, using ideas similar to what we did above. Conclude that there exists an explicit PRG that $O(\eta)$-fools halfspace on n variables having η-balanced w-s with seed-length $O(\log n + \log^2 \frac{1}{\eta})$.

We remark that in the paper they also remove the balance requirement, and achieve error ε with seed length $O(\log n + \log^2(\frac{1}{\varepsilon}))$, and even that was improved in a later paper to $\tilde{O}(\log \frac{n}{\varepsilon})$.

Question 4

Fix parameters T, W, ε. Our goal is to construct an ε-PRG against $[W = 2^s, T]$ BPs. The construction gives better results than Nisan's PRG when $T \leq 2^{s^{1-\alpha}}$ and generalizes the Nisan-Zuckerman PRG. From now on we assume $T \leq 2^{s^{1-\alpha}}$.

1. Show that you can fix

$$\begin{aligned} \Sigma &= \text{poly}\left(\frac{T \cdot \log W}{\varepsilon}\right), \\ \Gamma &= \text{poly}\left(\frac{TW}{\varepsilon}\right), \end{aligned}$$

   such that:

   - $\Gamma = \Sigma^{2\ell}$ for some integer ℓ,
   - There exists a $(\frac{3\ell}{2} \log \Sigma, \frac{\varepsilon}{T})$ extractor $E \colon \Gamma \times \Sigma \to \Sigma^{\ell}$.

2. Let n be an arbitrary integer. Prove that $G_1 \colon \Gamma \times \Sigma^n \to \Sigma^{\ell \cdot n}$ defined by

$$G_1(x; y_1, \ldots, y_n) = E(x, y_1) \circ \ldots \circ E(x, y_n),$$

   is a PRG against $[W, \ell n]_\Sigma$ BPs.

3. Set $n_i = \ell^{i+1}$. define $G_2 : \Gamma^2 \times \Sigma^{n_0} \to \Sigma^{n_2}$ by

$$G_2(x_1, x_2; \overline{y} = (y_1, \ldots, y_{n_0})) \quad = \quad G_1(x_1; G_1(x_2, \overline{y})),$$

   where we view $G_1(x_2, \overline{y})$ as an element of Σ^{n_1}. Prove that G_2 is a PRG against $[W, n_2]_\Sigma$ BPs.

4. Define G_r and prove that G_r is a PRG against $[W, n_r]_\Sigma$ BPs.

5. Conclude that when $T \leq 2^{s^{1-\alpha}}$ for a constant $\alpha > 0$, there exists an explicit construction of a PRG against $[W = 2^s, T = 2^{s^{1-\alpha}}]_\Sigma$ BPs with error $\varepsilon = \frac{1}{T}$ and seed length $O(\log T \cdot \frac{s}{\log s})$.