Space-Bounded Computation – Questions Pool

Amnon Ta-Shma and Noam Parzanchevski

January 13, 2019

General guidelines

The questions fall into several categories:

(Know).	Make sure you know how to solve. Do not submit.
(Mandatory).	Mandatory questions.
(Bonus).	Bonus questions.

Put your answers in the dropbox folder. You have to

1. write the solutions yourself,
2. give credit to any source (or any person) you consulted with.

You have to submit solutions to, at least, the mandatory questions.

HW 1

Out: 22.10.2018
Due: 05.11.2018

Unique Perfect Matching

1. (Mandatory). Finish the proof of the probabilistic algorithm we've seen in class: show how to find a Unique Perfect Matching using the Isolation Lemma

k-wise independence

2. (Mandatory). Draw $a \in \{0,1\}^{\log n}$ uniformly and for every $0 \neq i \in \{0,1\}^{\log n}$ let X_i be the random variable $X_i = \langle a, i \rangle \bmod 2$.

   Prove that $X = (X_1, \ldots, X_{n-1})$ is a distribution over $\mathbb{F}_2^{n-1}$ with support size n and is pairwise independent.

3. (Mandatory). You are about to play a game where n coins are laid covered on a table and you uncover and take $\frac{2n}{3}$ coins. You are promised that $k < \frac{n}{3}$ of the coins are pure gold and the rest copper. The catch is that you first have to announce your strategy (be it deterministic or probabilistic) and only then an adversary places the coins on the table. Show that:

   (a) If you use a deterministic strategy, you can guarantee no gold coin.

   (b) If you use n random coins you can almost certainly get $\Omega(k)$ gold coins. What is the failure probability?

   (c) If you use $O(\log n)$ random coins, you can guarantee $\Omega(k)$ gold coins with probability at least $1 - O(\frac{1}{k})$.

4. (Know). Let A, B be two distributions taking values in Λ For $f : \Lambda \to \Lambda'$, $f(A)$ (corr. $f(B)$) denotes the distribution over Λ' obtained by picking $a \sim A$ and outputting $f(a)$. Prove that $\|f(A) - f(B)\|_1 \leq \|A - B\|_1$ for every function f.

PIT

5. (Mandatory). Prove the Schwartz-Zippel lemma.

   If $p : \mathbb{F}^m \to \mathbb{F}$ is a non-zero polynomial of total degree d over a field $\mathbb{F}$ and $\Lambda \subseteq \mathbb{F}$, then $\Pr_{a_1,\ldots,a_m \in \Lambda}[p(a_1, \ldots, a_m) = 0] \leq \frac{d}{|\Lambda|}$.

6. (Mandatory). Give a coRP algorithm for Polynomial Identity Testing (PIT). In the proof work over the finite field $\mathbb{Z}_p$ for an appropriately random prime p, and prove its correctness

Boolean and Arithmetic Circuits

7. (Mandatory). Prove that almost all function $f : \{0,1\}^n \to \{0,1\}$ require circuits of size $\geqslant \frac{2^n}{10n}$, i.e.

$$\Pr_f\left[s(f) \geqslant \frac{2^n}{10n}\right] \xrightarrow[n\to\infty]{} 1$$

8. (Mandatory). Give an algorithm for addition of two integers represented in binary in AC^0
9. (Know). Prove that $\mathsf{NC}^0 \subseteq \mathsf{AC}^0 \subseteq \mathsf{NC}^1 \subseteq \cdots \subseteq \mathsf{NC} = \mathsf{AC} \subseteq \mathsf{P}$
10. (Bonus). Prove that $\mathsf{NC}^k \subseteq \mathsf{Space}(O(\log^k n))$. Note the cost of pointers.

HW 2

Out: 6.11.2018
Due: 19.11.2018

Tail Bounds

1. (Know). Let X_i be i.i.d random variables such that $X_i \sim \mathrm{Ber}(p)$ for some $p \in (0,1)$ and define $X = \sum_i X_i$, then for $0 < q \leqslant p$:

$$\Pr\left[X < qn\right] \leqslant e^{-\mathrm{KL}(q||p)n}$$

Maximal Independent Set

2. (Mandatory). In class, we derandomized the RNC algorithm for MIS using simplifying (and unjustified) assumptions. You are now asked to removed the unjustified assumptions.

   Show an NC algorithm for MIS that works for the general case. I assume that as part of the construction you will use a distribution $X = X_1, \ldots, X_n$ with small support size. Write precisely:

   - What properties you need from X,
   - How you construct X,
   - Why X has the desired properties, and,
   - Why these properties suffice for solving MIS in NC,

3. (Mandatory). Fix a finite field $\mathbb{F} = \mathbb{F}_q$. Show how to "naturally extend" the 2UFOHF family we've seen in class for larger independence. I.e.,

   - Define k-wise independence, and,
   - Show an explicit distribution over $(X_1, \ldots, X_q)$ that is k-wise independent, each X_i is distributed over $\mathbb{F}_q$, and has support size q^k.

4. (Mandatory). Let $V = \{0,1\}^m$ and $\mathcal{H} \subseteq \{h : V \to V\}$ a two universal family of hash functions. Fix two sets $A, B \subseteq V$. Call a hash function $h \in \mathcal{H}$ *ε-good for A, B* if

$$\left| \Pr_{x \in V}[x \in A \ \cap \ h(x) \in B] - \rho(A)\rho(B) \right| \ \leq \ \varepsilon,$$

   where $\rho(C) = \frac{|C|}{|V|}$.

   Prove that for any $A, B \subseteq V$, $\varepsilon > 0$,

$$\Pr_{h \in \mathcal{H}}\left[h \text{ is not } \varepsilon\text{-good for } A, B\right] \ \leq \ \frac{\rho(A)\rho(B)(1-\rho(B))}{\varepsilon^2 \cdot |V|} \ \leq \ \frac{1}{\varepsilon^2 |V|}.$$

Error Correcting Codes

5. (Mandatory). Let $\mathcal{C}$ be a linear $[n, k, d]_q$ code.

   - Show that $d = \min_{\bar{0} \neq x \in \mathcal{C}} \operatorname{wt}(x)$ Where $\operatorname{wt}(x) = |\{i : x_i \neq 0\}|$.
   - Prove the Singleton bound: $d \leqslant n - k + 1$.
   - Recall the parity code we've seen in class $\operatorname{Par} : \{0,1\}^k \to \{0,1\}^{k+1}$ where $\operatorname{Par}(x) = x \circ \bigoplus_i x_i$. Show that Par is a linear code and describe its generator matrix $G \in \mathbb{F}_2^{k+1 \times k}$.

6. (Mandatory). In this exercise we will prove two simple (non-matching) upper and lower bounds on the number of codewords in the best code with distance d. Let Σ be some alphabet and denote $q = |\Sigma|$. Let $d \leq n$ be integers. We also let $B_q(r) = \{x \in [q]^n \mid wt(x) \leq r\}$, the ball around zero of radius $r - 1$ in the Hamming weight distance.

   - (The Gilbert-Varshamov bound, A lower bound on the number of codewords) Prove that there exists a distance d code $\mathcal{C} \subseteq \Sigma^n$ with

   $$|\mathcal{C}| \geq \frac{|\Sigma|^n}{|B_q(d-1)|}.$$

   - (The Hamming bound) An upper bound on the number of codewords) Prove that for any distance d code $\mathcal{C} \subseteq \Sigma^n$,

   $$|\mathcal{C}| \leq \frac{|\Sigma|^n}{|B_q(\frac{d-1}{2})|}.$$

   - Use the asymptotic estimates given in class for $q = 2$ to show that for any $0 < \delta < 1/2$ there exists a code $\mathcal{C} \subseteq \{0,1\}^n$ with relative distance δ and relative rate $r \geqslant 1 - H(\delta) - o(1)$, and every code $\mathcal{C} \subseteq \{0,1\}^n$ with relative distance δ has relative rate $r \leqslant 1 - H(\delta/2) + o(1)$.

A reminder. If $\mathcal{C} \subseteq \Sigma^n$ has distance d, then $\mathcal{C}$ has *relative distance* $\delta = \frac{d}{n}$ and *relative rate* $r = \frac{\log |\mathcal{C}|}{n \log |\Sigma|}$. You may use $|B_2(\lambda n)| \leq 2^{H(\lambda)n}$ and $\lim_{n\to\infty} \frac{\log(|B_2(\lambda n)|)}{n} = H(\lambda)$ for any $0 < \lambda < 1/2$.

HW 3

Out: 27.11.2018
Due: 17.12.2018

Error Correcting Codes

1. (Mandatory). Recall that the Justesen code we've constructed in class gave us an $\left[n, \frac{r}{2}, \frac{1-r}{21}\right]_{\{0,1\}}$ linear code for any $0 < r < 1$.

   Show that we can pick an r such that in every non-zero codeword the number of zero and one symbols is somewhat balanced in the sense that there exists some constant $0 < c < \frac{1}{2}$ such that for any non-zero codeword $x \in \text{JUS}$, it holds that $\frac{1}{2} - c \leqslant \frac{\mathrm{wt}(x)}{n} \leqslant \frac{1}{2} + c$.

AKS Primality Testing

2. (Mandatory).

   - Show that $x^b - 1 \mid x^a - 1$ iff $b \mid a$:
     - In $\mathbb{Z}$
     - In $\mathbb{F}_p[x]$ for a prime p
   - Using the notation in class, show that for every $f \in A$, $f \in P_{\frac{n}{p}}$.
   - Using the notation in class, show that $\frac{n}{p} \in G$.

3. For the following question, recall that $\Phi_r(x) = \prod_{i:(i,r)=1}(x - \omega_r^i)$ where ω_r is an r-primitive root of unity.

   - (Mandatory). Show that if $(m, r) = 1$ then $\Phi_r \mid \Phi_r(x^m)$
   - (Bonus). Say we replace the line in the algorithm

     $$\forall 1 \leqslant j \leqslant \ell : (x + j)^n = x^n + j \bmod x^r - 1$$

     with

     $$\forall 1 \leqslant j \leqslant r : (x + j)^n = x^n + j \bmod \Phi_r.$$

     Will the algorithm still work? Prove your answer.

Expanders

4. (Mandatory). For the following, let $G = (V, E)$ be a D-regular undirected graph over n vertices, let A be the adjacency matrix of G and let $\lambda_1 \geqslant \lambda_2 \geqslant \cdots \geqslant \lambda_n$ be the spectrum of A. Prove the following:

   - $\lambda_1 = D$.
   - $\lambda_2 < D$ iff G is connected.
   - $\lambda_n \geq -D$. Furthermore, if G is connected, $\lambda_n = -D$ iff G is bipartite.

- If G has no self-loops then there exists a negative eigenvalue $\lambda_i < 0$

5. (Mandatory). The second item from the question above implies that the multiplicity of the eigenvalue D is 1 iff G is connected. In this question we will show that if we do not require G to be finite but only locally-finite then this does not necessarily hold.

   To that end, show that there exists a regular, undirected, connected graph $G = (V, E)$ of finite, (even constant) degree D where $|V| = \infty$ such that there exists a non-constant vector $f \in \mathbb{Z}^{|V|}$, $f \notin \text{span}\{1^{|V|}\}$, such that $Af = Df$.

   Hint: Consider the 2-regular graph G where $V = \mathbb{Z}$ where for any $j \in \mathbb{Z}$ we have the edges $(j, j-1)$ and $(j, j+1)$

Deterministic Amplification

6. (Mandatory). Let A be some probabilistic algorithm using n coins with success probability $\frac{1+\alpha}{2}$. For a sequence of coin tosses $x_1, \ldots, x_t$ where $x_i \in \{0,1\}^n$ define

$$\text{MAJ}(A, x_1, \ldots, x_t) = \begin{cases} 1, & \text{if } \sum_i A(x_i) > t/2 \\ 0, & \text{otherwise} \end{cases}.$$

   We consider two amplification protocols:

   (a) One where $x_1, \ldots, x_t$ are chosen uniformly and independently from $x_i \in \{0,1\}^n$.

   (b) One where $x_1, \ldots, x_t$ are chosen from a pairwise independent distribution over $\{0,1\}^n$.

   What is the required t to amplify the success probability $\frac{1+\alpha}{2}$ of $A(x)$ to success probability $1 - \delta$ of $\text{MAJ}(A, x_1, \ldots, x_t)$ in each of these two cases?

   In particular does pairwise independent work well when:

   - when $\alpha = \frac{1}{n}$ and we want δ to be a small constant,
   - when α is a constant and we want δ to be exponentially small.

HW 4

Out: 24.12.2018
Due: 7.1.2019

Bias Amplification

1. (Mandatory). Let G be a graph over $W = [n]$ vertices with a normalized adjacency matrix A and let $f : W \to \mathbb{F}_2$ be an ε-biased function. As in class, we let Π be the $n \times n$ diagonal matrix $\Pi_{w,w} = (-1)^{f(w)}$ and let $A_0 = \{w \in [n] : f(w) = 0\}$ (define A_1 likewise).

   Prove that the probability a random walk of length t over G visits A_1 an even number of times minus the probability it visits A_1 an odd number of times is:

$$\left| \mathbb{1}^{\dagger} \Pi \cdot (A \cdot \Pi)^t \cdot \frac{1}{n} \mathbb{1} \right|,$$

   where $\mathbb{1}$ is the all one vector.

Fourier analysis

2. (Mandatory). Let G be a finite Abelian group.

   (a) Prove that there are exactly $|G|$ characters of G.

   (b) Let $\widehat{G}$ be the set of characters. Prove that $\widehat{G}$ is a group.

   (c) Prove that $G \cong \widehat{G}$.

3. (Mandatory). Let H be a group and S a set of generators. The Caylely graph $C(H, S)$ is defined as follows: The vertices are labeled with elements of H, and (a, b) is an edge iff $a = bs^{-1}$ for some $s \in S$.

   (a) What is $C(\mathbb{Z}_n, \{1, -1\})$? What is $C(\mathbb{Z}_2^n, \{e_1, , e_n\})$ (where e_i has 1 in the i-th coordinate and 0 otherwise)?

   (b) Prove that if H is Abelian then the characters of H form an orthonormal basis for $C(H, S)$. (Please do not cite the result from class, but do the calculation from scratch).

   (c) Calculate the eigenvalues and the spectral gap of $C(\mathbb{Z}_2^n, \{e_1, , e_n\})$.

   (d) Let S be an ε-biased set. Define the graph $G = C(\mathbb{Z}_2^n, S)$ and let A be its normalized adjacency matrix. What are the eigenvectors of A? Prove that G has a spectral gap of at least $1 - \varepsilon$.

4. (Mandatory). The variational distance between two distributions P, Q over Λ is $\frac{1}{2}\sum_{x \in \Lambda} |P(x) - Q(x)|$. A distribution D on $\{0, 1\}^n$ is (k, ε)-wise independent, if for every $S \subseteq [n]$ of cardinality at most k, the marginal distribution of D on S is ε close to uniform in the variational distance.

   Prove that if a distribution D is (k, ε)-wise independent then there exists a distribution X that is k-wise independent and $|X - D| \leq 2n^k \varepsilon$.

   Hint: Consider the bias of linear tests of size at most k. Construct X explicitly. If you wish you can see the easy proof at [1].

5. (Mandatory). (due to Swastik Kopparty) Let $f, g : \{0,1\}^n \to \mathbb{C}$. We define their *convolution* $h = f \star g$ to be

$$h(x) = \mathop{\mathbb{E}}_{y} f(x \oplus y) g(y).$$

Note that if D_1 and D_2 are distributions, the distribution $D_1 \star D_2$ corresponds to the distribution of $d_1 \oplus d_2$ where $d_1 \sim D_1$ and $d_2 \sim D_2$ are picked independently.

Also, if D is a distribution over Λ we let $Col(D) = \Pr_{x,x' \in D}[x = x'] = \sum_{x \in \Lambda} (D(x))^2$.

   (a) Prove that for every $S \subseteq [n]$, $\hat{h}(S) = \hat{f}(S) \cdot \hat{g}(S)$.

   (b) Let D be an ε-biased distribution, and let $D^{(t)} = D \star \ldots \star D$ (t times). Prove that $D^{(t)}$ is ε^t-biased and that $|\mathrm{Supp}(D^{(t)})| \leq \binom{|\mathrm{Supp}(D)|+t}{t}$.

   (c) Prove that for any ε-biased distribution D over $\{0,1\}^n$, $Col(D) \leq \varepsilon^2 + 2^{-n}$.

   (d) Let D be an ε-biased distribution over $\{0,1\}^n$. Use the previous items to prove that $|\mathrm{Supp}(D)| \geq \Omega\left(\frac{n}{\varepsilon^2 \log \frac{1}{\varepsilon}}\right)$.

   Hint: Use (c) to derive a lower bound on $|\mathrm{Supp}(D)|$ and then choose t accordingly.

6. (Mandatory). A distribution D over Λ has k min-entropy if the largest probability mass given to any element in Λ is 2^{-k} (i.e., for all $a \in \Lambda$, $\Pr(D = a) \leq 2^{-k}$, and for some a it is 2^{-k}). We denote $H_\infty(D) = k$.

For two distributions X_1, X_2 over $\{0,1\}^n$ let $X_1 + X_2$ denote the distribution over $\{0,1\}^n$ obtained by sampling $x_i \in X_i$ and outputting $x_1 + x_2$, where the sum is addition mod 2 coordinate wise.

   - Say X_1, X_2 are two independent distributions over $\{0,1\}^n$. Prove that $H_\infty(X_1 + X_2) \geq H_\infty(X_1)$.
   - Let $X_1, \ldots, X_t$ be independent distributions over $\{0,1\}^n$ such that each X_i is ε-close (in the variational distance) to a distribution having min-entropy k. Prove that $X = \sum_{i=1}^t X_i$ is ε^t-close to having min-entropy $k - \log \frac{1}{1-\varepsilon}$.
   - (Dori Medini) Find a distribution X that is $\varepsilon = 1/2$ close to uniform and $X_1 + \ldots + X_1$, where $X_1, \ldots, X_t$ are i.i.d with marginal distribution X, is also ε close to uniform.

References

[1] Noga Alon, Oded Goldreich, and Yishay Mansour. Almost k-wise independence versus k-wise independence. *Information Processing Letters*, 88(3):107–110, 2003.