Questions Pool

Amnon Ta-Shma and Dean Doron

August 2, 2016

General guidelines

The questions fall into several categories:

K : Make sure you know how to solve. Do not submit.

M : Mandatory questions.

C : Choose one of the questions.

B : Bonus questions.

	Lectures	K	M	C	B
Ex1	Lectures 1 and 2	1,2,15,17	20,7-12	16,23,25,26	3,6,14
Ex2	Lectures 3 and 4	26	27,29,30	32,33	31,28
Ex3	Lectures 5 and 6	34,36,42,44	35,37,38,43,45	39,40	41
Ex4	Lectures 7 and 8	48	46,51,52,53,54	47,49,50	
Ex5	Lectures 9 and 10	61	55,56a-c,57,58,59,60,62		56d
Ex6	Lectures 11 and 12	64,65	63,66-69		70

Statistical distance and min-entropy

1. In class we defined the statistical (also known as variational) distance between two distributions A and B taking values in $\mathcal{U}$ to be $|A-B| = \max_{T \subseteq \mathcal{U}} |\Pr[A \in T] - \Pr[B \in T]|$. $A \times B$ denotes the distribution on (a,b) obtained by independently picking a according to A and b according to B. (A,B) denotes some random variable on $\mathcal{U} \times \mathcal{U}$ with marginal distributions A and B. For $f : \mathcal{U} \to \mathcal{U}'$ and A that is distributed over $\mathcal{U}$ let $f(A)$ denote the distribution on $\mathcal{U}'$ obtained by picking $a \sim A$ and outputting $f(a)$.

   (a) Prove that $|A-B| = \frac{1}{2}\sum_{x \in \mathcal{U}} |A(x) - B(x)|$.

   (b) Prove that $|A \times B - A' \times B'| \le |A - A'| + |B - B'|$.

   (c) Show that it is not always the case that $|(A,B) - (A',B')| \le |A - A'| + |B - B'|$.

   (d) Assume $|(B|A=a) - C| \le \varepsilon$ for every $a \in \mathcal{U}$. Prove that $|(A,B) - A \times C| \le \varepsilon$. Conclude that if $E : \{0,1\}^n \to \{0,1\}^m$ is an extractor for $(n, \frac{n}{2})$ oblivious bit-fixing sources with error ε, and $f(x)$ returns $\frac{n}{2}$ bits of x, then $|E(X) \circ f(X) - U_m \times f(X)| \le \varepsilon$.

(e) Prove that $|f(A) - f(B)| \leq |A - B|$ for every deterministic function f.

(f) Extend the previous item to a probabilistic f.

2. For a distribution X over $\mathcal{U}$, the entropy function is $H(X) = \sum_{a \in \mathcal{U}} X(a) \log \frac{1}{X(a)}$, where we only sum over nonzero $X(a)$. Recall that the min-entropy of X is $H_\infty(X) = \min_{a \in \mathcal{U}} \log \frac{1}{X(a)}$. Also, let $H_0(X) = \log(|\mathrm{Supp}(X)|)$.

   (a) Prove that $H_\infty(X) \leq H(X) \leq H_0(X)$.

   (b) It is a fact that $H(X, Y) \leq H(X) + H(Y)$. Find an example where $H_\infty(X, Y) > H_\infty(X) + H_\infty(Y)$.

   (c) Prove that $H_\infty(X|Y) \geq H_\infty(X, Y) - H_0(Y)$, where $H_\infty(X|Y) = -\log \sum_{y \in \mathcal{U}} \Pr[Y = y] \max_{x \in \mathcal{U}} \Pr[X = x | Y = y]$.

3. (a) Prove that the set of (n, k) sources, treated as vectors in $\mathbb{R}^{2^n}$ is a polytope.

   (b) Prove that the vertices of the polytope are exactly the flat distributions.
   Guidance: Let S be a polytope defined by a set of linear constraints (equalities and inequalities). For a point p, let $n(p)$ denote the number of constraints that are satisfied with equality at p. You can use the fact that a point p of a convex S set is a vertex of the polytope iff $n(p)$ is maximized.

   (c) Prove that every (n, k) source is a convex combination of flat distributions with entropy k.

   (d) Conclude that $E : \{0,1\}^n \times \{0,1\}^t \to \{0,1\}^m$ is a (k, ε) extractor iff it is an extractor for all flat distributions with min-entropy k.

4. Let A be a distribution over $\mathcal{U}$. Prove that if A is not ε-close to a k-source then there exists a subset $S \subseteq \mathcal{U}$ of cardinality at most 2^k such that $\Pr_{a \in A}[a \in S] \geq \varepsilon$.

Non-oblivious extractors

5. Let $q > 0$ be a fixed integer and $\alpha = \log_3 2$. Prove that $I_q(\mathrm{IMAJ}_3) = O(\frac{q}{n^\alpha})$, where IMAJ_3 is the iterated majority function defined in class.

6. In the *leader election* problem, n players, q of which are malicious, collectively choose a single leader. The goal is to choose the leader in such a way that the probability of choosing a faulty player as a leader is not too large. Formally, for a leader election protocol P with r rounds, let $I_q(P)$ be the maximal probability of choosing a malicious leader (over the possible strategies of the malicious players).

   Consider the following protocol for leader election, due to Feige, known as the "Lightest-Bin Protocol" (LB):

   - Set $X \leftarrow [n]$.
   - Repeat while $|X| > 1$:
     - Each player in X broadcasts a random bit. Let X_0 denote the set of players who broadcast 0, and X_1 denote the set of players who broadcast 1.
     - If $|X_0| \leq \frac{1}{2}|X|$, set $X \leftarrow X_0$. Otherwise, $X \leftarrow X_1$.
   - Output the single element in X.

(a) Let $\delta > 0$ be such that there are $q = (\frac{1}{2} - \delta)n$ malicious players. Assume that $n - q$ is an exact power of 2 and set $t = \log(n-q)$. Prove that $I_q(\mathrm{LB}) \le 1 - (n-q)^{-t/4}$.

(b) Prove that if there exists a leader election protocol P such that $I_q(P) \le \varepsilon$ then there exists a collective coin flipping protocol f such that $I_q(f) \le \varepsilon$.

Expanders

7. Let G be an undirected, regular graph. Prove:
   - The vector $\mathbf{1}$ is an eigenvector of A with eigenvalue 1.
   - For all $i \in [n]$, $|\lambda_i| \le 1$.

8. Let G be a regular, undirected graph. Prove that the number of connected components of G equals the dimension of its 1-eigenspace. In particular, G is connected iff $\lambda_2 < 1$.

9. Prove that if λ is an eigenvalue of an undirected *bipartite* graph, then do does $-\lambda$. Prove that a D-regular, undirected connected graph G is bipartite iff $\lambda_n = -1$. What is the associated eigenvector v_n ?

10. Let G_n be a cycle on n vertices. Show that $\{\chi_k\}_{k=0}^{n-1}$ is an orthonormal eigenvector basis for G_n with eigenvalues $\cos\left(\frac{2\pi k}{n}\right)$, where $\chi_k(i) = \omega^{ki}$ for ω a primitive n-th root of unity. Is $\{G_n\}_{n\in\mathbb{N}}$ a family of expanders?

11. The diameter of a graph is the maximum minimal distance between two vertices in the graph. Let G be a D regular graph over n vertices.

    (a) Prove that $diam(G) \ge \log_{D-1}(n-1) - 2$.

    (b) Prove that $diam(G) \le 1 + \log_{\frac{1}{\lambda}} n$, where $\lambda = \max\{-\lambda_n, \lambda_2\}$.

12. Prove that in any D-regular, undirected graph with n vertices, if $D \le \frac{n}{2}$ then $\lambda \ge \frac{1}{\sqrt{2D}}$, where $\lambda = \max\{-\lambda_n, \lambda_2\}$.
    Hint: Calculate $\mathrm{Tr}(A^2)$ in two different ways.

Oblivious extractors

13. Let n be an integer, and $\varepsilon \ge 0$. For every $q \le n - \log n - 2\log\frac{1}{\varepsilon} - O(1)$ and $m \le n - q - 2\log\frac{1}{\varepsilon} - O(1)$ there exists an ε-error extractor $E : \{0,1\}^n \to \{0,1\}^m$ for (n,q) oblivious bit-fixing sources.

14. (a) Let M be an $m \times n$ matrix over $\mathbb{F}_2$ for $n \ge m$. Define the oblivious bit fixing extractor $f : \{0,1\}^n \to \{0,1\}^m$ by $f(x) = Mx$. Prove that if f is a zero error extractor for (n,q) oblivious bit fixing sources, then $C_f = \left\{M^\dagger a \mid a \in \mathbb{F}_2^m\right\}$ is an error correcting code with distance at least $q+1$.

    (b) Conclude that if there exists a linear zero-error extractors for $(n, \frac{2n}{3}+1)$ oblivious bit-fixing sources that extract two bits then there exists a binary linear code with dimension 2 and distance $\frac{2n}{3}+2$. The Plotkin bound says there is no such code, and therefore we may conclude that there is no such extractor.

Deterministic extractors

15. Show that there is no 1-source extractor. Namely, for every $E : \{0,1\}^n \to \{0,1\}$ there exists an $(n, k = n-1)$ source X such that $f(X)$ is fixed.

16. Let n be an integer, and $\varepsilon \geq 0$. For every $m \leq 2k - \log n - 2\log\frac{1}{\varepsilon} - O(1)$ there exists a two-source extractor $E : (\{0,1\}^n)^2 \to \{0,1\}^m$ with ε-error for independent sources with min-entropy k.

17. Let H be the Hadamard matrix of dimension 2^n, where $H_{i,j} = (-1)^{\langle x,y\rangle}$ (the indices are interpreted as binary vectors of length n). Find $\| H \|_2$.

Ramsey graphs

18. Prove that there is no graph on n vertices that is $\left(\frac{1}{2}\log n - 2\right)$-Ramsey.
Hint: Let $n = R(a,b)$ be the smallest integer such that for every graph on n vertices, there is always a clique of size a or an independent set of size b. Prove that $R(a,b) \leq R(a-1,b) + R(a,b-1)$.

19. Prove that there exists a graph on n vertices that is $2\log n$-Ramsey. By the same method, prove that there exists a *bipartite* graph on n vertices that is $2\log n$-Ramsey.

20. Show how to explicitly translate a $(k,0)$ 2-source disperser $E : (\{0,1\}^n)^2 \to \{0,1\}$ into a 2^{k+1}-Ramsey graph on 2^n vertices.

Affine extractors

21. Prove that there exists an affine extractor $E : \mathbb{F}_2^n \to \{0,1\}^m$ for min-entropy $k \geq 2\log n + O(1)$, where $m \leq k - O(1)$.

22. Find an explicit affine extractor $E : \mathbb{F}_2^n \to \{0,1\}$ for min-entropy larger than $\frac{n}{2}$. Hint: You may want to use the Chor-Goldreich two-source extractor that we saw in class.

Extractors and dispersers

23. Prove that for every integers $n \geq k$ and $\varepsilon > 0$ there exists a (K,ε) disperser $E : [N] \times [D] \to [M]$ with $m = k + d - \log\log\frac{1}{\varepsilon} - O(1)$ and $d = \log(n-k) + \log\frac{1}{\varepsilon} + O(1)$.

24. Let $E : [N] \times [D] \to [M]$ be a (K,ε) disperser such that $D \leq \frac{(1-\varepsilon)M}{2}$ and $\varepsilon \leq \frac{1}{2}$. Then, $D = \Omega\left(\frac{1}{\varepsilon}\log\frac{N}{K}\right)$. Deduce that every (K,ε) disperser $E : [N] \times [D] \to [M]$ has an entropy loss of at least $\log\log\frac{1}{\varepsilon} - O(1)$.

25. Consider the *privacy amplification problem.* Alice holds a secret $x \in \{0,1\}^n$. Eve knows t bits of information about x, $w = f(x) \in \{0,1\}^t$. To overcome this problem, Alice would like to hash x to a shorter string, $h(x) \in \{0,1\}^m$, that looks truly uniform to Eve. Formally, Alice wants h such that $|h(X) \circ f(X) - U_m \circ f(X)| \leq \varepsilon$. What can Alice do?

Limited independence

26. Prove that the k-wise sample space of size n^k we saw in class is indeed k-wise independent.

27. For every $a \in \{0,1\}^{\log n}, 0 \neq i \in \{0,1\}^{\log n}$, consider the construction $X_i(a) = \langle a, i \rangle \bmod 2$. Prove that $X_1, \ldots, X_{n-1}$ forms a pairwise independent sample space.

28. Prove: If $X = X_1, \ldots, X_n$ is k-wise independent and each X_i is boolean then $|\mathrm{Supp}(X)| \geq B(k/2, n)$, where $B(r,n)$ is the number of words of weight at most r in the n-dimensional boolean cube.

29. Let $V = \{0,1\}^m$ and $\mathcal{H} \subseteq V \to V$ a two universal family of hash functions (see definition in Lecture 2). 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 $A, B \subseteq V$, $\varepsilon > 0$,

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

30. For a set C, let U_C denote the uniform distribution over C.

Let $\mathcal{H} \subseteq \Lambda \to \Gamma$ be a two universal family of hash functions (see definition in Lecture 2). For a distribution D over Λ let $(H, H(D))$ denote the distribution over $\mathcal{H} \times \Gamma$ obtained by picking d according to D, picking h uniformly from $\mathcal{H}$ and outputting $(h, h(d))$.

- Prove that for any distribution X over C, $\parallel X - U_C \parallel_2^2 = \parallel X \parallel_2^2 - \parallel U_C \parallel_2^2$.
- Prove that for any distribution X, $\parallel X \parallel_2^2 = \Pr_{x_1, x_2 \in X}[x_1 = x_2]$.
- Prove that $\parallel (H, H(D)) \parallel_2^2 \leq \parallel U_H \parallel_2^2 \cdot [\parallel U_\Gamma \parallel_2^2 + \parallel D \parallel_2^2]$.
- Conclude that $\parallel (H, H(D)) - U_H \times U_\Gamma \parallel_2 \leq \parallel U_H \parallel_2 \cdot \parallel D \parallel_2$.
- Prove that $\parallel (H, H(D)) - U_H \times U_\Gamma \parallel_1 \leq \sqrt{|\Gamma|} \cdot \parallel D \parallel_2$.

31. Among known NP-complete problems, reductions can be found that preserve the number of solutions (or witnesses). A characteristic of such NP-complete problems is that their instances have widely varying numbers of solutions. It is then natural to ask whether it is inherent, and Valiant and Vazirani shows that it is not.

We say that $L_1 \leq_r L_2$ if there exists a randomized polynomial-time Turing machine M and a polynomial p such that if $x \in L_1$ then $\Pr[M(x) \in L_2] \geq \frac{1}{p(|x|)}$ and if $x \notin L_1$ then $\Pr[M(x) \in L_2] = 0$.

(a) Find a two universal family of hash functions (see definition in Lecture 2) $\mathcal{H} \subseteq \{0,1\}^n \to \{0,1\}^k$ for which you can express the condition $h(x_1, \ldots, x_n) = 0$ as a CNF with the variables $x_1, \ldots, x_n$.

(b) Let UniqSAT be the language of satisfiable CNFs with only one satisfying assignments. Prove that SAT $\leq_r$ UniqSAT.
Hint: Pick k at random and combine (a) with your original CNF formula.

32. A universe $\mathcal{U}$ has an (unknown) set $T \subseteq \mathcal{U}$ of *good* elements of size n and an (unknown) set $S \subseteq \mathcal{U}$ of *bad* elements of size n. The elements $\mathcal{U} \setminus (S \cup T)$ are *neutral*. Our goal is to choose a subset $A \subseteq \mathcal{U}$ that contains at least one good element and no bad elements. We call such set an *appropriate* set.

    (a) Suppose we choose A so that every element in $\mathcal{U}$ is in A with probability p. The choices are independent. Prove that if $p = \frac{1}{n}$ then the probability that A is appropriate is lower-bounded by some constant that is independent of n.

    (b) Choose p such that if instead of independently, we choose in a pairwise-independent manner, we can still get a constant probability of choosing an appropriate A.

33. 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 $1 - O(\frac{1}{k})$.

Fourier analysis

34. Prove Claim 2 from Lecture 3.

35. 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 $G \cong \widehat{G}$.

36. 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 $C(H, S)$.

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

37. We will use Fourier analysis to prove the result of Question 28. Let k be a constant.

    (a) Prove that D is a k-wise independent distribution if and only if $\hat{D}(S) = 0$ for every S with $0 < |S| \leq k$.

    (b) Let $d = \frac{k}{2}$ and denote $t = \sum_{i=0}^{d-1} \binom{n}{i}$. Let $X \subseteq \{0,1\}^n$ with $|X| < t$. Show that there is a function $f : \{0,1\}^n \to \mathbb{R}$ which satisfies:

- f is not identically zero.
- $\hat{f}(S) = 0$ for every S with $|S| > d$.
- $f(x) = 0$ for every $x \in X$.

(c) Conclude that any k-wise independent distribution over $\{0,1\}^n$ has support size at least $\Omega(n^{k/2})$.

ε-bias and almost independence

38. Fill in the details of the construction of Subsection 2.3 from Lecture 3. Specifically, prove that the concatenated code meets these parameters and that it is balanced.

39. 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.

    (a) What are the eigenvectors of A (you might want to look at previous questions).

    (b) Prove that G has a spectral gap of at least $1 - 2\varepsilon$.

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

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

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

    (b) Prove that for any ε-biased distribution D over $\{0,1\}^n$, $\| D \|_2^2 \leq \varepsilon^2 + \frac{1}{2^n}$.

    (c) 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 $|\mathrm{Supp}(D^{(t)})| \leq \binom{|\mathrm{Supp}(D)|+t}{t}$.

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

Finite fields

42. We represent $\mathbb{F}_{16}$ as $\mathbb{F}_2[X] (\mathrm{mod}\ X^4 + X + 1)$. X is a generating element for $\mathbb{F}_{16}^\star$. Below you can find a table relating the vector space representation to powers of X. Find how many elements generate $\mathbb{F}_{16}^\star$. How many elements generate $\mathbb{F}_q^\star$ for an arbitrary prime power q?

Power	Element	Vector space representation
0	0	$(0,0,0,0)$
$X^0 = X^{15} = 1$	1	$(0,0,0,1)$
X	X	$(0,0,1,0)$
X^2	X^2	$(0,1,0,0)$
X^3	X^3	$(1,0,0,0)$
X^4	$1+X$	$(0,0,1,1)$
X^5	$X+X^2$	$(0,1,1,0)$
X^6	X^2+X^3	$(1,1,0,0)$
X^7	$1+X+X^3$	$(1,0,1,1)$
X^8	$1+X^2$	$(0,1,0,1)$
X^9	$X+X^3$	$(1,0,0,1)$
X^{10}	$1+X+X^2$	$(0,1,1,1)$
X^{11}	$X+X^2+X^3$	$(1,1,1,0)$
X^{12}	$1+X+X^2+X^3$	$(1,1,1,1)$
X^{13}	$1+X^2+X^3$	$(1,1,0,1)$
X^{14}	$1+X^3$	$(1,0,0,1)$

43. Give an efficient algorithm (polynomial in the input length) that given a degree m polynomial $E(X)$ that is irreducible over $\mathbb{F}_p$, and a non-zero element $x \in \mathbb{F}_q = \mathbb{F}_p[X] (\text{mod } E)$, finds x^{-1}.

44. Let $\mathbb{F}$ be a finite field. What is the expected number of roots of a random univariate polynomial of degree k over $\mathbb{F}$?

45. Let $\mathbb{F}_q$ be a finite field of odd characteristic. An element $x \in \mathbb{F}_q^\star$ is a *quadratic residue* if there exists $y \in \mathbb{F}_q$ such that $x = y^2$. What is the number of quadratic residues in $\mathbb{F}_q$? Prove that the set of quadratic residues is a multiplicative group.

AC circuits

46. Prove: For every constant $d \geq 2$, there are circuits of size $2^{O(n^{1/(d-1)})}$ that compute the parity of n bits.

47. Prove Lemma 9 from Lecture 4 using a two-universal family of hash functions.

48. Prove: The worst-case, $\frac{1}{3}$-approximation of the Parity function requires degree n, and this is tight.

49. Prove that the Majority function is not in AC^0.

50. Prove that multiplication of two n-bit inputs is not in AC^0.

51. Prove that Parity is hard on average for AC^0.

52. Let $\mathsf{AC}^0(3)$ be the class of languages decidable by AC^0 where we allow MOD3 gates in addition to the usual operations $\{\wedge, \vee, \neg\}$. A MOD3 gate on inputs $x_1, \ldots, x_n$ returns 0 if $(\sum_i x_i) \bmod 3 = 0$ and 1 otherwise. We say $p : \mathrm{GF}(3)^n \to \mathrm{GF}(3)$ solves parity on average with ε error, if $\Pr_{x \in \{0,1\}^n}[p(x) = \mathrm{Parity}(x)] \geq 1 - \varepsilon$.

Prove that any $\mathsf{AC}^0(3)$ circuit of size s and degree d can be approximated by a $(O(\log \frac{s}{\varepsilon} \log s))^d$-degree polynomial over GF(3) with average case error (with respect to the uniform distribution) at most ε.

53. In this question we prove parity cannot be solved on average (with respect to the uniform distribution) by low-degree polynomials over GF(3).

    (a) Prove that every function $f : \mathrm{GF}(3)^n \to \mathrm{GF}(3)$ can be uniquely expressed as a polynomial $p : \mathrm{GF}(3)^n \to \mathrm{GF}(3)$ of local degree at most 2.

    (b) Let $f : \{0,1\}^n \to \{0,1\}$ be a boolean function. Let $\phi : \{0,1\} \to \mathrm{GF}(3)$ be such that $\phi(1) = -1$ and $\phi(0) = 1$. Define $f' : \{-1,1\} \to \{-1,1\}$ such that $f'(\phi(x_1), \dots, \phi(x_n)) = \phi(f(x_1, \dots, x_n))$. What is parity$'$?

    (c) Prove that Parity is not in $\mathsf{AC}^0(3)$.

    (d) Prove that Parity is hard on average for $\mathsf{AC}^0(3)$.

54. Let RAC^0 be the class of languages such that $L \in \mathsf{RAC}^0$ if there exists an AC^0 circuit $C_L : \{0,1\}^n \times \{0,1\}^{\mathrm{poly}(n)} \to \{0,1\}$ such that:

    - If $x \in L$ then $\Pr_y[C_L(x,y) = 1] \geq \frac{2}{3}$.
    - If $x \notin L$ then $\Pr_y[C_L(x,y) = 1] \leq \frac{1}{3}$.

    Prove that $\mathsf{RAC}^0 \subseteq \bigcup_c \mathsf{DSPACE}(\log^c n)$.

Designs and Trevisan's extractor

55. Prove that for every $\ell, a \geq 1$, there exists an (ℓ, a)-design $S_1, \dots, S_m \subseteq [t]$ with $t = O(\frac{\ell^2}{a})$ and $m = 2^{\Omega(a)}$.

56. Two norm-one vectors $v_1, v_2 \in \mathbb{R}^n$ are almost orthogonal if $|\langle v_1, v_2 \rangle| \leq \varepsilon$.

    (a) Show how to convert an (ℓ, a)-design $S_1, \dots, S_m \subseteq [t]$ into:
        - A set of m nearly orthogonal norm-one vectors.
        - A binary error-correcting code of length t with m codewords and large distance.

    (b) How many norm-one orthogonal vectors can one put into $\mathbb{R}^d$?

    (c) How many norm-one ε-almost orthogonal vectors can one put into $\mathbb{R}^d$? Give a lower bound.

    (d) How many norm-one ε-almost orthogonal vectors can one put into $\mathbb{R}^d$? Give an upper bound. Can you reach tight estimations?

57. Read the definition of weak design given in the lecture notes. It is known that there exist explicit weak $(\ell, \rho = 1)$ design $S_1, \dots, S_m \subseteq [t]$ with $t = O(\ell^2 \log m)$. Notice that $\rho = 1$. Use this weak design, and the NW construction to construct a strong (k, ε) extractor $E : \{0,1\}^n \times \{0,1\}^t \to \{0,1\}^m$ with $t = O(\log^2(\frac{n}{\varepsilon}) \log k)$ and $m = k - O(t)$.

58. In previous lectures we saw how to construct an extractor from pair-wise independence, namely $E(x,h) = h(x)$ where h is sampled from a two-universal family of hash functions. Use almost pair-wise (and almost t-wise) independence to construct a (k, ε) strong extractor $E : \{0,1\}^n \times \{0,1\}^t \to \{0,1\}^m$ with $t = O(m + \log \frac{n}{\varepsilon})$ and almost optimal entropy loss $2 \log \frac{1}{\varepsilon} + O(1)$.

59. Let $E : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ be a (k, ε) strong extractor. Define $E' : \{0,1\}^n \times (\{0,1\}^d)^2 \to \{0,1\}^{2m}$ such that $E'(x; y_1, y_2) = E(x; y_1) \circ E(x; y_2)$. Prove that E' is a $(k + m + \log \frac{1}{\varepsilon}, O(\varepsilon))$ strong extractor.

60. Show how to combine the extractors in the previous questions, to obtain an explicit extractor with almost minimal entropy loss $2\log(\frac{1}{\varepsilon}) + O(1)$ and poly-logarithmic seed length.

PRGs

61. (Distinguishably implies predictability) Let $f : \{0,1\}^n \to \{0,1\}$ and suppose $C : \{0,1\}^n \times \{0,1\} \to \{0,1\}$ is a circuit such that

$$\Pr_{x \sim U_n}[C(x, f(x)) = 1] - \Pr_{x \sim U_n, b \sim U_1}[C(x, b) = 1] > \delta.$$

Prove that there exists another circuit $C' : \{0,1\}^n \to \{0,1\}$ such that

$$\Pr_{x \sim U_n}[C'(x) = f(x)] > \frac{1}{2} + \delta.$$

62. Prove: If there exists a uniform function $G : \{0,1\}^\ell \to \{0,1\}^{\ell+1}$ running in $poly(2^\ell)$ time, which is pseudo-random against circuits of size $s(\ell^c)$ for some $c > 0$, and if G is also one-to-one, then there exists a uniform language in EXP that cannot be approximated by circuits of size $s(\ell^d)$ for some $d > 0$.

Can you prove the same without the assumption that G is one-to-one. What is the problem? (Remark: It is known how to remove the one-to-one requirement using things we have not learnt).

Non-malleable extractors

63. Consider the following two definition for a t-non-malleable extractor:

Def 1: A function $E : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^1$ is a (t, k, ε) non-malleable extractor if it satisfies the following property: If X is an (n, k) source and Y is uniform on $\{0,1\}^d$, and $f_1, \dots, f_t$ are arbitrary functions from d bits to d bits with no fixed points then

$$(E(X,Y), E(X, f_1(Y)), \dots, E(X, f_t(Y)), Y) \approx_\varepsilon (U_1, E(X, f_1(Y)), \dots, E(X, f_t(Y)), Y).$$

Def 2: A function $E : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^1$ is a (t, k, ε) non-malleable-second-def extractor if it satisfies the following property: For every X that is an (n, k) source, there exists a set $G \subset \{0,1\}^d$ of size $(1 - \sqrt{\varepsilon})2^d$ such that $\{E(X, y)\}_{y \in G}$ is $(t, 10t\sqrt{\varepsilon})$ wise independent.

Prove that if E is (t, k, ε) non-malleable then it is $(t + 1, k, \varepsilon)$ non-malleable-second-def.

64. Let X, Y, Z be random variables such that for any $y \in \mathrm{Supp}(Y)$, the random variables $(X|Y = y)$ and $(Z|Y = y)$ are independent. Assume that X is supported on $\{0,1\}^n$. Prove:

$$|(X, Y, Z) - (U_n, Y, Z)| = |(X, Y) - (U_n, Y)|.$$

65. Let $(A \approx_\varepsilon U|B)$ denote $|(A, B) - (U, B)| \le \varepsilon$. Let f be an arbitrary (deterministic or probabilistic) function. Prove that if $(A \approx_\varepsilon U|B)$ then $(A \approx_\varepsilon U|B, f(B))$.

66. Give a full proof to the following claim we proved in class:

    Let X, X' be two (possibly correlated) sources over $\{0,1\}^n$, and Y, Y' two sources over $\{0,1\}^\ell$ and assume (as usual) that

    - (X, X') are independent of (Y, Y'),
    - X is a $k + m + O(\log \frac{1}{\varepsilon})$ source, and
    - Y is uniform.

    Now, we make the further (strong) assumption that Y is independent of Y'. Let $E : \{0,1\}^n \times \{0,1\}^\ell \to \{0,1\}^m$ be an arbitrary (k, ε) extractor. Then $(E(X,Y), E(X',Y'))$ is $O(\varepsilon)$-close to $U_m \times E(X',Y')$.

67. Suppose that for any $k \geq k_0$ there is an explicit $(k, \varepsilon(n))$-extractor $E_k : \{0,1\}^n \times \{0,1\}^{t(n)} \mapsto \{0,1\}^{\frac{k}{f(n)}}$. Then, for any k, there is an explicit $(k, \log(n)(\varepsilon + 2^{-t(n)}))$-extractor $E : \{0,1\}^n \times \{0,1\}^{O(f(n)\log(n)t(n))} \mapsto \{0,1\}^{k-\bar{k}}$.

68. Let X, X' be two (possibly correlated) sources over $\{0,1\}^n$, and Y, Y' two sources over $\{0,1\}^\ell$ and assume (as usual) that

    - (X, X') are independent of (Y, Y'),
    - For every x', $H_\infty(X|X' = x') \geq k$, and
    - Y is uniform.

    Prove that if E is a strong (k, ε) extractor then $(E(X,Y), E(X',Y'))$ is ε-close to $U_m \times E(X',Y')$.

69. (from C16) Let $Cond : [N] \times [D] \to [M]$ be a $k \to_\varepsilon k'$ condenser (if you don't know what a condenser is, look for the definition in the internet). Let X be an (n,k)-source and let S be an independent random variable that is uniformly distributed over d-bit strings. Then, for any $\delta > 0$, with probability $1 - \delta$ over $s \sim S$ it holds that $Cond(X, s)$ is $\frac{2\varepsilon}{\delta}$-close to having min-entropy $k' - d - \log(\frac{2}{\delta})$.

70. (from BIW06, C16) Let $X_1, \ldots, X_t$ be independent n-bit random variables such that each X_i is ε-close to having min-entropy k. Then, $X = \oplus_{i=1}^t X_i$ is ε^t-close to having min-entropy k.

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.