Problem Set 2

Submit your problem set by April 24 by emailing it to Fermi (fermima1@gmail.com) and James (bartusek.james@gmail.com). Full credit will be given as long as you make an honest attempt to solve each problem.

1. Give an interactive protocol between a prover and verifier that satisfies the following properties:

   - Under a (plausibly post-quantum) hardness assumption, *any* efficient classical prover convinces the verifier with probability at most $\mathsf{negl}(\lambda)$.
   - There exists an efficient quantum prover that makes the verifier accept with probability $1 - \mathsf{negl}(\lambda)$.

   You should not need to design a new protocol from scratch; instead, your protocol should be based on a protocol covered in the lectures. You may use any claims stated in Additional Resources.

2. The four-message proof of quantumness protocol we saw in the class can be seen as "enforcing a qubit" — for any efficient prover that convinces the verifier to accept with overwhelming probability, the measurement outcomes the verifier gets must be consistent with measuring some one-qubit state $|\psi\rangle$ in the standard/Hadamard bases. In this problem, we'll formally prove this.

   (a) First, we'll need a formal definition of an "approximate qubit." Consider the following two definitions:

   - (Approximate qubit, Definition 1) An ε-approximate-qubit is a triple $(|\psi\rangle, X, Z)$ such that $|\psi\rangle \in \mathcal{H}$,[1] and X and Z are any binary observables (see Additional Resources for a definition) satisfying

   $$\langle\psi| (XZ + ZX)^2 |\psi\rangle \leq \varepsilon.$$

   - (Approximate qubit, Definition 2) An ε-approximate-qubit is a triple $(|\psi\rangle, X, Z)$ such that there exists a Hilbert space $\mathcal{H}'$ and an isometry $V : \mathcal{H} \to \mathbb{C}^2 \otimes \mathcal{H}'$ (see Additional Resources for a definition) such that

   $$VX |\psi\rangle \approx_\varepsilon (\sigma_X \otimes \mathsf{Id}) V |\psi\rangle \quad \text{and} \quad VZ |\psi\rangle \approx_\varepsilon (\sigma_Z \otimes \mathsf{Id}) V |\psi\rangle,$$

   where $|\psi\rangle \approx_\varepsilon |\phi\rangle$ means $\||\psi\rangle - |\phi\rangle\| \leq \varepsilon$ and σ_X, σ_Z denote the Pauli operators

   $$\sigma_X \coloneqq \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_Z \coloneqq \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

[1] Here, $\mathcal{H}$ denotes a finite-dimensional Hilbert space, i.e., a space isomorphic to $\mathbb{C}^d$ for some positive integer d.

Prove that these definitions are equivalent, up to a constant multiplicative factor in ε. While not required, you may find it helpful to prove the $\varepsilon = 0$ case first (this corresponds to two definitions of an *exact* qubit).[2]

(b) Consider any QPT prover that is accepted in Protocol 2 with probability 1.[3] Suppose we run the first two messages of the protocol to obtain (pk, y). Show how to use the description of the prover's current state and remaining algorithms to define a triple $(|\psi\rangle, X, Z)$ such that, with probability $1 - \mathsf{negl}(\lambda)$ over (pk, y), $(|\psi\rangle, X, Z)$ is a $\mathsf{negl}(\lambda)$-approximate qubit.

3. (Quantum Goldreich-Levin) Let $s \in \{0,1\}^n$ be an unknown string. Suppose you are given access to a classical oracle $f_s : \{0,1\}^n \to \{0,1\}$ such that

$$\Pr_{r \leftarrow \{0,1\}^n}[f_s(r) = s \cdot r] = \frac{1}{2} + \varepsilon.$$

Show how to output s with probability $\mathsf{poly}(\varepsilon)$ given only one (quantum) query to this oracle.

4. We saw in class that given access to a classical algorithm $\mathcal{O}_{\mathsf{LWE},A}$ that on input y outputs the shortest vector in $\mathsf{span}(A^\top, y)$, there exists an algorithm for solving the Short Integer Solution (SIS) problem for A with respect to some bound β.[4] That is, with probability $1 - \mathsf{negl}(\lambda)$, the algorithm outputs a non-zero x such that $Ax = 0$ and $\|x\|_\infty \leq \beta$.

(a) Suppose instead that $\mathcal{O}_{\mathsf{LWE},A}(y)$ only outputs the shortest vector in $\mathsf{span}(A^\top, y)$ with probability $1/\mathsf{poly}(\lambda)$. That is,

$$\Pr_{s,e}[\mathcal{O}_{\mathsf{LWE},A}(As + e) = e] = 1/\mathsf{poly}(\lambda).$$

Show how to solve the SIS problem for A with probability $1/\mathsf{poly}(\lambda)$ given access to this algorithm.

(b) Suppose instead that we have oracle access to a *quantum* algorithm that outputs the shortest vector in $\mathsf{span}(A^\top, y)$ with probability $1/\mathsf{poly}(\lambda)$. That is, we have access to a unitary $U_{\mathsf{LWE},A}$ such that measuring the first register of $U_{\mathsf{LWE},A}|As + e\rangle|0\rangle$ yields e with probability $1/\mathsf{poly}(\lambda)$. Show how to solve SIS with probability $1/\mathsf{poly}(\lambda)$ given access to this unitary and its inverse.

5. In class, we proved that the YZ problem (Definition 8) is hard for any classical adversary that (a) makes $\mathsf{poly}(\lambda)$-many non-adaptive queries to the random oracle and (b) outputs $c = (c_1, \ldots, c_n)$ with the property that each c_i was previously queried to the random oracle at

[2] This problem is inspired by Exercise 2.1 in Thomas Vidick's notes [Vid20]. Feel free to reference the notes if you get stuck.

[3] Note that this is technically unreasonable, since even the honest strategy only has some $1 - \mathsf{negl}(\lambda)$ of being accepted on the equation test. However, this will (slightly) simplify the proof, and it is not too difficult to extend what we will do to provers that are accepted with probability $1 - \mathsf{negl}(\lambda)$.

[4] You will not need to worry about the exact parameters (i.e. dimension of A, field size, value of β, etc.) in order to answer this problem.

some point, assuming that the code C is "list-recoverable." Show that if C is "list-recoverable with errors" (Definition 9) the same result holds for an arbitrary classical adversary, i.e., one that makes $\mathsf{poly}(\lambda)$-many adaptive queries to the random oracle.

6. The quantum YZ solver that we saw in class generates a *uniformly random* solution to the YZ problem. Assuming that the Aaronson-Ambainis conjecture is true (Conjecture 1), prove that there is no efficient quantum algorithm for the YZ problem that outputs a *deterministic solution*. In other words, show that any quantum algorithm outputs a solution with non-zero entropy (showing even negligible entropy will suffice for this problem).

7. (a) For $(x, \theta) \in \{0,1\}^2$, let $\left|x^\theta\right\rangle := H^\theta \left|x\right\rangle$ (where $H = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ denotes the Hadamard gate). Consider the following challenger-adversary game:

   i. The challenger samples uniformly random $(x_1, \ldots, x_n), (\theta_1, \ldots, \theta_n)$ and sends $|x_1^{\theta_1}\rangle \otimes \cdots \otimes |x_n^{\theta_n}\rangle$ to the adversary.

   ii. The adversary outputs two n-bit strings r, s, and wins if for all indices i where $\theta_i = 0$, $r_i = x_i$, and for all indices i where $\theta_i = 1$, $s_i = x_i$.

   Prove that no (even unbounded) adversary can win this game with probability better than $\mathsf{negl}(n)$.

   (b) Construct a *quantum-communication* bit commitment scheme that satisfies (a) classical-style binding but is *not* (b) collapse-binding.

Additional Resources

You may use (without proof) any of the claims in this section.

Definition 1 (Trapdoor claw-free function with adaptive hard-core bit)**.** *The family of functions* $\mathcal{F} = \{f_{\mathsf{pk}} : \{0,1\} \times \{0,1\}^{n(\lambda)} \to \{0,1\}^{n(\lambda)}\}_{\lambda\in\mathbb{N},\mathsf{pk}\in\{0,1\}^{k(\lambda)}}$ *is a* trapdoor claw-free function with adaptive hard-code bit *if it satisfies the following properties.*

- ***Two-to-one.*** *For each* $y \in \{0,1\}^{n(\lambda)}$*, the set* $f_{\mathsf{pk}}^{-1}(y)$ *contains exactly two elements* $(0, x_0), (1, x_1)$*.*
- ***Adaptive hard-core bit.*** *For any QPT adversary* $\{\mathcal{A}_\lambda\}_{\lambda\in\mathbb{N}}$*,*

$$\Pr\left[d \neq 0^{n+1} \wedge d \cdot ((0,x_0) \oplus (1,x_1)) = 0 : \begin{array}{r} \mathsf{pk} \leftarrow \{0,1\}^k \\ (b, x_b, d) \leftarrow \mathcal{A}_\lambda(\mathsf{pk}) \\ ((0,x_0),(1,x_1)) \coloneqq f_{\mathsf{pk}}^{-1}(f_{\mathsf{pk}}(b,x_b)) \end{array}\right] \leq \frac{1}{2} + \mathsf{negl}(\lambda).$$

- ***Trapdoor.*** *There exist algorithms* $(\mathsf{pk}, \mathsf{td}) \leftarrow \mathsf{Gen}(1^\lambda)$ *and* $(x_0, x_1) \leftarrow \mathsf{Invert}(\mathsf{td}, y)$ *such that the distribution of* pk *when sampled from* $\mathsf{Gen}(1^\lambda)$ *is uniform, and* $(x_0, x_1) \leftarrow \mathsf{Invert}(\mathsf{td}, y)$ *is such that* $f_{\mathsf{pk}}((0, x_0)) = f_{\mathsf{pk}}((1, x_1)) = y$*.*

Protocol 2. *(Four-message proof of quantumness) Let* $\mathcal{F} = \{f_{\mathsf{pk}} : \{0,1\} \times \{0,1\}^{n(\lambda)} \to \{0,1\}^{n(\lambda)}\}_{\lambda\in\mathbb{N},\mathsf{pk}\in\{0,1\}^{k(\lambda)}}$ *be a* trapdoor claw-free function with adaptive hard-core bit *(see Definition 1). Recall the following four-message protocol sketched in class based on* $\mathcal{F}$*.*

- *The verifier samples* $(\mathsf{pk}, \mathsf{td}) \leftarrow \mathsf{Gen}(1^\lambda)$ *and sends* pk *to the prover.*
- *The prover returns a string* $y \in \{0,1\}^n$*.*
- *The verifier samples a uniformly random challenge* $c \leftarrow \{0,1\}$ *and sends* c *to the prover.*
  - *In the case* $c = 0$ *(a preimage test), the prover returns a string* $(b, x_b) \in \{0,1\}^{n+1}$*, and the verifier accepts if and only if* $f_{\mathsf{pk}}((b, x_b)) = y$*.*
  - *In the case* $c = 1$ *(an equation test), the prover returns a string* $d \in \{0,1\}^{n+1}$*, the verifier runs* $(x_0, x_1) \leftarrow \mathsf{Invert}(\mathsf{td}, y)$*, and accepts if and only if* $d \cdot ((0, x_0) \oplus (1, x_1)) = 0$*.*

Definition 3 (Binary Observable)**.** *A binary observable is any Hermitian (i.e.,* $O^\dagger = O$*) operator* O *that satisfies* $O^2 = \mathbb{I}$*.*

Since Hermitian operators must have real eigenvalues, a binary observable has at most two eigenvalues: $+1$ and -1. Consequently, a binary observable can always be written in the form $O = \Pi_{+1} - \Pi_{-1}$ where Π_{+1}, Π_{-1} are projections onto the $+1, -1$ eigenspaces and $\Pi_{+1} + \Pi_{-1} = \mathbb{I}$.

For example, the Pauli operator

$$\sigma_Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} = |0\rangle\langle 0| - |1\rangle\langle 1|$$

is a binary observable with $+1$-eigenspace $|0\rangle\langle 0|$ and -1 eigenspace $|1\rangle\langle 1|$. The Pauli operator

$$\sigma_X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} = |+\rangle\langle +| - |-\rangle\langle -|$$

is a binary observable with $+1$-eigenspace $|+\rangle\langle +|$ and -1 eigenspace $|-\rangle\langle -|$.

Definition 4 (Isometry). *An isometry is a linear map V from a Hilbert space $\mathcal{H}$ to a (potentially larger) Hilbert space $\mathcal{H}'$ such that $V^\dagger V = \mathsf{Id}$. Note that for any orthonormal set $\{v_i\}_i$ in $\mathcal{H}$ and $\{w_i\}_i$ in $\mathcal{H}'$, the operator $\sum_i |w_i\rangle\langle v_i|$ is an isometry.*

Lemma 5 (Jordan's Lemma). *Let P, Q be projectors on a finite-dimensional Hilbert space $\mathcal{H}$. Then there exists an orthogonal decomposition $\mathcal{H} = \oplus_i \mathcal{S}_i$ such that each $\mathcal{S}_i$ is a 1- or 2-dimensional subspace that is invariant under P and Q. Furthermore, whenever $\mathcal{S}_i$ is 2-dimensional, there is a basis for it in which P and Q take the form*

$$P = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad Q = \begin{pmatrix} \cos^2(\theta) & \cos(\theta)\sin(\theta) \\ \cos(\theta)\sin(\theta) & \sin^2(\theta) \end{pmatrix}$$

for some $\theta \in [0, \pi/2)$.[5]

Fact 6. *For any two binary observables X, Z, it holds that*

$$\frac{1}{4}(XZ + ZX)^2 = XZ_0XZ_0 + Z_1XZ_1X,$$

where

$$Z_b \coloneqq \frac{1}{2}\left(\mathsf{Id} + (-1)^b Z\right).$$

That is, Z_b is the projection onto the $(-1)^b$ eigenspace of Z (e.g. when $Z = \sigma_Z$, $Z_0 = |0\rangle\langle 0|$ and $Z_1 = |1\rangle\langle 1|$).

Remark 7 (Quantum access to a classical oracle). *Given a deterministic classical functionality $f : \{0,1\}^n \to \{0,1\}^m$, quantum oracle access to f means that we have access to the map $|x\rangle|y\rangle \to |x\rangle|y \oplus f(x)\rangle$. This is equivalent to having access to the map $|x\rangle \to (-1)^{f(x)}|x\rangle$ (in the case where f has 1-bit output).*

Definition 8 (Yamakawa-Zhandry Problem). *Let C be a subspace of $\mathbb{Z}_q^n$ and let $H : [n] \otimes \mathbb{Z}_q \to \{0,1\}$ be a random oracle. Find $(c_1, \ldots, c_n) \in C$ such that $H(i, c_i) = 0$ for all $i \in [n]$.*

Definition 9 (List-Recoverable With Errors). *A code $C \subset \Sigma^n$ is* list recoverable with errors[6] *if there exists a constant $\varepsilon > 0$ such that for any subsets $S_i \subseteq \Sigma$ such that $|S_i| \leq \mathsf{poly}(n)$ for $i \in [n]$, it holds that*

$$|\{(c_1, \ldots, c_n) \in C : |\{i \in [n] : c_i \in S_i\}| \geq (1-\varepsilon)n\}| \leq \mathsf{poly}(n).$$

[5] That is, P is the projection onto $|0\rangle$, and Q is a projection onto the vector $\cos(\theta)|0\rangle + \sin(\theta)|1\rangle$ (i.e., the result of rotating $|0\rangle$ counterclockwise by θ).

[6] The notion of "list-recoverable"/"rectangle-evasive" that we saw in class corresponds to the $\varepsilon = 0$ case.

Conjecture 1 (Aaronson-Ambainis). *Let $\varepsilon, \delta > 0$ be real numbers. Given any quantum algorithm A that makes Q quantum queries to a random oracle $H : \{0,1\}^n \to \{0,1\}$, there exists a deterministic classical algorithm B that makes* $\mathsf{poly}(Q, 1/\varepsilon, 1/\delta)$ *classical queries to H and satisfies*

$$\Pr_H \left[\left| \Pr\left[A^H \to 1\right] - B^H \right| \leq \varepsilon \right] \geq 1 - \delta.$$

Hints

2. For part (a), use Jordan's lemma (Lemma 5). For part (b), without loss of generality, we can model the prover as follows.

   - After the second message, it has an internal state $|\psi\rangle$ which may depend on pk and y. We will write this state on two registers (M, A), where M holds $n+1$ qubits and A holds an arbitrary (polynomial) number of qubits.
   - Upon receiving the challenge bit $c \leftarrow \{0,1\}$, it first applies some unitary U_c to $|\psi_{\mathsf{pk},y}\rangle$. Then, if $c = 0$, it measures register M in the standard basis to obtain (b, x_b) and if $c = 1$, it measures register M in the Hadamard basis to obtain d.
   - Moreover, we can assume that the prover applies U_0 before receiving c, and in the case that $c = 0$ immediately measures register M in the standard basis, while in the case of $c = 1$ applies $U := U_1 U_0^\dagger$ and then measures register M the Hadamard basis.

   Now consider the following family of observables associated with the prover's actions, where each is defined based on (pk, y).

$$Z = \sum_{(b,x)\in f_{\mathsf{pk}}^{-1}(y)} (-1)^b \left|(b,x)\right\rangle\!\left\langle(b,x)\right|,$$

$$X = \sum_{d\in\{0,1\}^{n+1}} (-1)^{g(\mathsf{pk},y,d)} U^\dagger (H_M^{\otimes n+1} \oplus \mathsf{Id}_A)(|d\rangle\langle d|_M \otimes \mathsf{Id}_A)(H_M^{\otimes n+1} \oplus \mathsf{Id}_A) U,$$

   for $g(\mathsf{pk}, y, d) \to \{0,1\}$ that is defined to output 0 iff $d \neq 0^{n+1}$ and $d \cdot ((0,x_0),(1,x_1)) = 0$, where $((0,x_0),(1,x_1)) := f_{\mathsf{pk}}^{-1}(f_{\mathsf{pk}}(y))$.

   Finally, it may be helpful to use Fact 6.

3. Run the Bernstein-Vazirani algorithm. It is easy to show that when $\varepsilon = 1/2$, the algorithm outputs s with probability 1. Now, show that for any ε, the state of the algorithm right after its oracle query has inner product $\mathsf{poly}(\varepsilon)$ with the state that would have resulted from querying the "perfect" oracle. Use this fact to show that the output of the algorithm is s with probability $\mathsf{poly}(\varepsilon)$.

4. For part (a), use a similar argument as the previous question. For part (b), consider the state of the algorithm

$$|\psi\rangle := \sum_{s,e} \exp\left(\frac{-|x|^2}{t^2}\right) |e\rangle_X \left|A^\top s + e\right\rangle_Y |0\rangle_Z$$

   right before applying the LWE solver. Let (Π_0, Π_1) be the binary projective measurement that corresponds to running $U_{\mathsf{LWE},A}$ on registers (Y, Z) and measuring whether the output register is equal to the vector stored in register X. Show that the inner product between $|\psi\rangle$ and $\Pi_0 |\psi\rangle / \|\Pi_0 |\psi\rangle\|$ is at least $1/\mathsf{poly}(\lambda)$, and use this to derive your algorithm and proof.