Quantum Proofs, Semester A 2024

Homework # 3 **due: 5pm, February 1st, 2024**

Ground rules:

Homework is due through Moodle. If you are having issues with this, email the instructor (thomas.vidick@weizmann.ac.il) or drop your work in my mailbox at the top of the central stairs in Ziskind. Solutions can be latexed or handwritten. In the latter case, please make sure that your handwriting is legible. Special care should be taken in writing up a precise solution. If I am not able to follow the logic in your argument, if there is a small gap or an uncovered case, you will lose points.

You are encouraged to collaborate with your classmates on homework problems, but each person must write up the final solutions individually. You should note on your homework specifically which problems were a collaborative effort and with whom. You may not search online for solutions, but if you do use research papers or other sources in your solutions, you must cite them.

Late homework will not be accepted or graded.

The first problem asks you to practice with the notion of a semidefinite program, and work through some simple examples where "classical" problems in quantum information can be expressed as the optimum of a semidefinite program. Each question has a very short solution. If you are not already familiar with semidefinite programs, I can recommend checking the lecture notes by Lovasz (see link on course webpage) for a good presentation.
The second problem gives an alternate proof of error amplification for QIP*. It introduces a very useful norm on quantum channels, the diamond norm. This norm is frequently used outside of the study of* QIP*, to compare quantum channels. In that problem, questions (a) and (b) can be accepted on faith at first; the more interesting parts are (c)–(f). This problem requires more algebraic manipulations than the two others.*
*The last problem has a very long description, but it has a short solution! In that problem, I am asking you to take time to understand a new definition for a complexity class (*competing-prover games*) and then apply your understanding to show that this class is contained in a class we know well (the class* QIP*). Once you understand the problem statement, each question (except for the optional (b)) has a quite short solution.*

Any changes since the first posting will be marked in blue.

Problems:

1. (4 points) **Practice with semidefinite programs**

   Recall that a semidefinite program is said in *primal normal form* if it is written as

$$\begin{aligned} \sup \quad & B \bullet X \\ s.t. \quad & A_i \bullet X = a_i \;, \quad \forall\, i = 1, \ldots, m \\ & X \geq 0 \;, \end{aligned}$$

   where we used the shorthand notation $X \bullet Y = \mathrm{Tr}(X^\dagger Y)$, $B, A_1, \ldots, A_m$ are complex Hermitian matrices of the same size as X, and $a_1, \ldots, a_m$ real numbers.

   (a) Suppose given a complex Hermitian matrix $A \in \mathbb{C}^{d \times d}$. Write a semidefinite program, in primal normal form, whose optimum is the largest eigenvalue of A.

   For the next questions, you no longer need to write the semidefinite program in normal form. In practice, this means that you may have multiple matrix variables X, Y, Z, etc., on which you can impose any number of linear or positive-definite inequalities (including such as $Y \leq \mathbb{I}$, etc.).

   (b) Can you do the same with $\|A\|_1$, the sum of the singular values of A?

   (c) Deduce a semidefinite program whose optimum is the trace distance $\|\sigma_0 - \sigma_1\|_{tr} = \frac{1}{2}\|\sigma_0 - \sigma_1\|_1$ between two density matrices σ_0 and σ_1 (given explicitly, as matrices).

   (d) Suppose given an ensemble $\{(p_i, \rho_i) : i \in \mathcal{I}\}$, where: $\mathcal{I}$ is a finite index set; for each i, $p_i \in [0,1]$ such that $\sum_{i \in \mathcal{I}} p_i = 1$; and for each i, ρ_i is a density matrix on n qubits, specified explicitly (in matrix form, as for the previous question). Write the maximum success probability of the adversary in the following game, played against a trusted challenger, as the optimum of a semidefinite program:

      i. The challenger selects $i \in \mathcal{I}$ according to the distribution (p_i). They prepare the quantum state ρ_i and send it to the adversary.

      ii. The adversary performs a measurement and returns to the challenger an index $i' \in \mathcal{I}$.

      iii. The challenger declares that the adversary has won if and only if $i' = i$.

2. **The diamond norm and error amplification**

   In this problem, T denotes a "super-operator," which in general is any linear map $T : L(\mathcal{N}) \to L(\mathcal{M})$. Here, $\mathcal{N}$ and $\mathcal{M}$ are (finite-dimensional) Hilbert spaces and $L(\mathcal{N})$ and $L(\mathcal{M})$ are the space of linear operators on $\mathcal{N}$ and $\mathcal{M}$ respectively. Said in other words, $\mathcal{N} = \mathbb{C}^{d_\mathcal{N}}$ for some integer $d_\mathcal{N}$ and $L(\mathcal{N}) = \mathbb{C}^{d_\mathcal{N} \times d_\mathcal{N}}$, the space of $d_\mathcal{N} \times d_\mathcal{N}$ matrices. So, T is a linear map that sends $d_\mathcal{N} \times d_\mathcal{N}$ matrices to $d_\mathcal{M} \times d_\mathcal{M}$ matrices. (If T is additionally completely positive and trace preserving, then it is a channel; but for the time being we allow general linear T.)

A natural norm on the space of such linear maps T is the operator norm induced by the 1 norm, i.e.

$$|||T|||_1 := \sup_{X \neq 0} \frac{\|T(X)\|_1}{\|X\|_1} . \qquad (1)$$

Here, $\|X\|_1 = \text{Tr}\sqrt{XX^\dagger}$ is the 1 norm of the matrix X, which is the sum of the singular values. The norm $|||\cdot|||_1$ has the following inconvenient:

(a) Let $T : L(\mathbb{C}^2) \to L(\mathbb{C}^2)$ be defined by $T : |i\rangle\langle j| \mapsto |j\rangle\langle i|$ for all $i, j \in \{0, 1\}$, and extended by linearity to all 2×2 matrices. (So, T is the transpose map!) Show that $|||T|||_1 \leq 1$, but $|||T \otimes \mathbb{I}_2|||_1 \geq 2$, where $\mathbb{I}_2$ is the identity map on 2×2 matrices.

The previous question shows that $|||\cdot|||_1$, when used on super-operators, does not "stabilize". This property is not welcome when discussing quantum channels, as we would not want that the "norm" of a channel tensored with the identity is bigger than the norm of the channel itself. So instead, we define

$$|||T|||_\diamond := \sup_{d \geq 1} \left|\left|\left| T \otimes \mathbb{I}_{L(\mathbb{C}^d)} \right|\right|\right|_1 ,$$

where $|||\cdot|||_1$ is as defined in (1), and $\mathbb{I}_{L(\mathbb{C}^d)}$ denotes the identity super-operator from $L(\mathbb{C}^d)$ to itself.

(b) Show that for any superoperators R, S it holds that $|||RS|||_\diamond \leq |||R|||_\diamond |||S|||_\diamond$. (You may use that the same inequality holds for the norm $\|\cdot\|_1$, without reproving this fact.)

In the remainder of this problem we use the norm $|||\cdot|||_\diamond$ to characterize the maximum acceptance probability of a $\mathsf{QIP}(3)$ verifier, and give an alternate proof of error amplification.

In the following fix a $\mathsf{QIP}(3)$ verifier $V = (V_1, V_2)$ in purified form. Here, V_1 is a unitary that acts on the message $\mathcal{Y}$ received from the prover, and the verifier's private space $\mathcal{Z}$. It produces a message sent back to the prover, which for convenience we assume lies on the same space $\mathcal{Y}$, and a residual memory state. So, V_1 is a unitary on $\mathcal{Z} \otimes \mathcal{Y}$. Similarly, V_2 is the unitary on $\mathcal{Z} \otimes \mathcal{Y}$ applied by the verifier upon receipt of the prover's second message. After V_2 has been applied, the verifier measures using a measurement (Π_{acc}, Π_{rej}) that we assume acts on the entire space $\mathcal{Z} \otimes \mathcal{Y}$. Finally, let Π_{init} denote the projection on the space where all verifier's qubits (the register $\mathcal{Z}$) are initialized to 0.

Let $W_1 = V_1 \Pi_{init}$ and $W_2 = V_2^\dagger \Pi_{acc}$. Let $T : L(\mathcal{Z} \otimes \mathcal{Y}) \to L(\mathcal{Y})$ be the superoperator defined as $T(X) = \text{Tr}_{\mathcal{Z}}(W_1 X W_2^\dagger)$.

(c) Show that $\omega(V) = \max\{|\langle|\phi|W_2^\dagger U W_1|\psi\rangle|^2\}$, where the maximum is taken over all states $|\psi\rangle, |\phi\rangle \in \mathcal{Z} \otimes \mathcal{Y} \otimes \mathcal{W}$ and unitaries U on $\mathcal{Y} \otimes \mathcal{W}$, with $\mathcal{W}$ the prover's private space.

(d) For a fixed space $\mathcal{H}$, show that the maximum of $\|T \otimes \mathbb{I}_{L(\mathcal{H})}(Y)\|_1$ over all Y such that $\|Y\|_1 = 1$ is attained at a Y of the form $Y = |\psi\rangle\langle\phi|$, for normalized vectors $|\psi\rangle, |\phi\rangle$.

(e) Deduce from the previous questions that $\omega(V) = \|\!|T|\!\|_\diamond^2$.

(f) Suppose that V' is another verifier, not necessarily identical to V. Let $V \otimes V'$ denote the verifier that runs V and V' in parallel and accepts if and only if both accept. Use the previous questions to show that $\omega(V \otimes V') \leq \omega(V)\omega(V')$.

3. Competing-prover games

In this problem we consider a variant of the model of quantum interactive proofs studied in class, in which there are *two* provers: the "YES-prover" and the "NO-prover." Formally, a (short) competing-prover game is an interactive game where:

(i) At the first step, the YES-prover sends a quantum message, in register Y_1, to the verifier.

(ii) The verifier applies an arbitrary unitary on Y_1, his private space Z, and another message register X_2. The verifier sends X_2 to the NO-prover.

(iii) The NO-prover performs some action on X_2 and its private register W_2. It returns a message register Y_2 to the verifier.

(iv) The verifier now has (Y_1, Z, Y_2). It performs a measurement to make its decision.

Initially, the registers (Y_1, W_1) (for the YES-prover), (Z, X_2) (for the verifier), and W_2 (for the NO-prover) are all initialized to $|0\rangle$. Note that the two provers do not interact directly, only with the verifier. (They also do not share any entanglement.)

We say that a promise problem $L = (L_y, L_n)$ is in the class $\mathsf{CPG}_{c,s}$ if there is a polynomial-time mapping from instances $z \in \{0,1\}^*$ to descriptions of verifiers V_z for short competing-prover games of the form above such that

- If $z \in L_y$ then there is a YES-prover such that for any action of the NO-prover, V_z accepts with probability at least c; and
- If $z \in L_n$ then there is a NO-prover such that for any action of the YES-prover, V_z accepts with probability at most s.

The goal of this problem is to show that $\mathsf{CPG}_{c,s}$ contains QIP, for some choice of c and s which you will determine at the end of the problem. Our starting point is the QIP-complete problem "close images" (CI) seen in class. To recall, an instance of CI is a pair of (circuit implementations of) quantum channels Φ_0 and Φ_1, each mapping n to m qubits, such that either there exists ρ_0, ρ_1 such that $\Phi_0(\rho_0) = \Phi_1(\rho_1)$ (yes-instance), or for any ρ_0, ρ_1, $F(\Phi_0(\rho_0), \Phi_1(\rho_1))^2 \leq \varepsilon$ (no-instance). Here, ε is a small parameter which we consider fixed; in practice you can assume that ε is any small enough constant.

We suggest a CPG protocol for the instance $z = (\Phi_0, \Phi_1)$ of CI as follows:

- The verifier receives two n-qubit quantum registers, R_0 and R_1, from the YES-prover. (So, what is called Y_1 above is now $Y_1 = R_0R_1$.)
- The verifier chooses $i \in \{0,1\}$ uniformly at random and applies Φ_i to the state in register R_i. Let the output be in the m-qubit register Z. The verifier sends Z to the prover.
- The prover responds with a bit $b \in \{0,1\}$. The verifier accepts iff $b \neq i$.

(a) Describe a strategy for the YES-prover in the case that (Φ_0, Φ_1) is a yes-instance of CI. What is the maximum probability that the verifier accepts (over all actions of the NO-prover)? (You do not need to show that the strategy you describe is optimal. But it should be intuitively clear that this strategy is in the YES-prover's best interest, assuming they want the verifier to accept with the highest possible probability.)

To analyze no-instances of CI, first define two convex sets $\mathcal{A}_i = \{\Phi_i(\rho) : \rho \in D(\mathbb{C}^{2^n})\}$, for $i \in \{0,1\}$. Here $D(\mathbb{C}^{2^n})$ is the convex set of all density matrices on n qubits (which we identify with the n input qubits to Φ_i). Define the distance between $\mathcal{A}_0$ and $\mathcal{A}_1$ as

$$d_{\mathcal{A}_0,\mathcal{A}_1} := \inf_{\rho_0 \in \mathcal{A}_0, \rho_1 \in \mathcal{A}_1} \|\rho_0 - \rho_1\|_1 \ .$$

(b) *(This question is optional. You may assume its outcome without showing it.)* Show that for any two convex subsets $\mathcal{B}_0$, $\mathcal{B}_1$ of the space of density matrices $D(\mathcal{H})$ on some finite-dimensional Hilbert space $\mathcal{H}$, there is a (not necessarily projective) measurement $\{P_0, P_1\}$ on $\mathcal{H}$ such that, for any $\sigma_0 \in \mathcal{B}_0$ and $\sigma_1 \in \mathcal{B}_1$,

$$\frac{1}{2}\mathrm{Tr}(P_0\sigma_0) + \frac{1}{2}\mathrm{Tr}(P_1\sigma_1) \geq \frac{1}{2} + \frac{1}{4}d_{\mathcal{B}_0,\mathcal{B}_1} \ .$$

This shows that if $\rho \in \{\rho_0, \rho_1\}$ is chosen uniformly at random, measuring it with $\{P_0, P_1\}$ will correctly identify ρ with good probability. *[The content of the question is in the fact that $\{P_0, P_1\}$ depends on the sets $\mathcal{B}_0$ and $\mathcal{B}_1$, but then it works for any σ_0 and σ_1 chosen from those sets. For the proof, use that disjoint convex sets have a separating hyperplane.]*

(c) Show that if (Φ_0, Φ_1) is a no-instance of CI then $d_{\mathcal{A}_0,\mathcal{A}_1} \geq 2(1 - \sqrt{\varepsilon})$. *[Hint: you may use the relation $F(\rho,\sigma) \geq 1 - \frac{1}{2}\|\rho - \sigma\|_{tr}$, which is valid for any density matrices ρ, σ.]*

(d) Use (b) to deduce that in this case, there is a strategy for the NO-prover such that the verifier accepts with probability at most $\sqrt{\varepsilon}/2$.