CS 258: Quantum Cryptography (Fall 2025)
Homework 5 (100 points)

1 Problem 1 (30 points)

Consider a distribution over quantum states, where $|\psi_i\rangle$ is sampled with probability p_i. Let $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$ be the resulting density matrix.

- **Part (a). 10 points.** Let U be a unitary, and consider computing $|\phi_i\rangle = U|\psi_i\rangle$. Taking the probability over i, this gives a new mixed state described by density matrix ρ_a. Show that $\rho_a = U\rho U^\dagger$.
- **Part (b). 10 points.** For any ρ, consider measuring in the computational basis. Show that the probability of a measurement outcome x is given by $\langle x|\rho|x\rangle$.
- **Part (c). 10 points.** Measuring in the computational basis gives x with some probability (as computed in Part (b)), and the post-measurement state is then $|x\rangle$. This gives a new probability distribution over quantum states, which is described by a density matrix ρ_c. Show that ρ_c is a diagonal matrix obtained from ρ by erasing all the off-diagonal entries.

2 Problem 2 (30 points)

For two classical probability distributions D_0, D_1 their distance is captured by the *total variational distance* $\Delta(D_0, D_1) = \frac{1}{2}\sum_x |\Pr[x \leftarrow D_0] - \Pr[x \leftarrow D_1]|$.

- **Part (a), 20 points.** Prove the following: suppose we choose a random bit b, and then sample $x \leftarrow D_b$ and apply some procedure P to make a guess b'. Define $\epsilon(P)$ such that $\Pr[b' = b] = \frac{1+\epsilon}{2}$. Prove that $\Delta(D_0, D_1)$ is the maximum over all possible (potentially inefficient) procedures P of $|\epsilon(P)|$. This contains two parts: (1) show that any procedure has $|\epsilon(P)| \leq \Delta(D_0, D_1)$, and show that (2) there exists some potentially inefficient procedure such that $\epsilon(P) = \Delta(D_0, D_1)$. For simplicity, you may assume the procedures are deterministic.

Thus, $\Delta(D_0, D_1) = 0$ means that no algorithm can do better than random guessing, while $\Delta(D_0, D_1) = 1$ means it is possible to perfectly distinguish the two distributions.

The way to quantify the distance between two mixed states represented by density matrices ρ_0, ρ_1 is through the trace distance. The trace distance has different notations throughout the literature, but is often denoted $\|\rho_0 - \rho_1\|_1$. It is defined as follows: Let $\lambda_1, \cdots, \lambda_n$ be the eigenvalues of $\rho_0 - \rho_1$. Then $\|\rho_0 - \rho_1\|_1 = \sum_i |\lambda_i|$.

- **Part (b), 10 points.** Consider some process for distinguishing ρ_0 from ρ_1. For simplicity, assume the process simply applies a unitary U, and then measures to get a string x. Call the resulting distributions over x D_0 and D_1, respectively. Show that there exists a unitary U such that $\Delta(D_0, D_1) = \|\rho_0 - \rho_1\|_1/2$, where D_0, D_1 are the probabilities obtained from applying U and then measuring. *[Hint: Think about diagonalization.]*

It turns out that for *any* unitary U, we have $\Delta(D_0, D_1) \leq \|\rho_0 - \rho_1\|_1/2$ (though you do not need to show this). Thus, trace distance is the direct quantum analog (up to a factor of two) of total variational distance, in that it exactly captures the ability to distinguish two quantum states.

3 Problem 3 (40 points)

A pseudorandom state (PRS) is a collection of 2^λ states $\{|\psi_k\rangle\}_{k\in\{0,1\}^\lambda}$. Let q be the number of qubits of the $|\psi_k\rangle$. The goal of a PRS is for $q > \lambda$, but for $|\psi_k\rangle$ for a random choice of k to look like a truly random state. Note that the density matrix for a truly random state on q qubits is $\frac{1}{2^q}\mathbf{I}$, where $\mathbf{I}$ is the identity matrix of dimension 2^q.

- **Part (a). 20 points.** Show that for $q > \lambda$, there is an inefficient quantum attack which distinguishes $|\psi_k\rangle$ for a random k from truly random. To do so, consider the density matrix ρ for $|\psi_k\rangle$, and consider the possible eigenvalues of ρ. How many are non-zero? What does this tell you about the trace distance from $\frac{1}{2^q}\mathbf{I}$? Thus, PRS's require computational assumptions

- **Part (b). 10 points** Consider the following commitment scheme built from a PRS. To commit to 0, construct the superposition $\frac{1}{\sqrt{2^q}}\sum_{x\in\{0,1\}^q}|x\rangle|x\rangle$, and give the second register to Bob, keeping the first register for ourselves. To commit to 1, construct the superposition $\frac{1}{\sqrt{2^\lambda}}\sum_{k\in\{0,1\}^\lambda}|k\rangle|\psi_k\rangle$, and give the second register to Bob, keeping the first register for ourselves.

  Show that the scheme is computationally hiding, assuming the PRS is secure.

- **Part (c). 10 points.** Suppose Alice has committed to 0. Explain why there is no unitary she can apply to her state that allows her to transform the joint state into a commitment to 1. This is not a full proof of statistical binding, but gives the idea.