CSE 291 / Math 277 - Quantum Complexity Theory (Fall 2024)
Homework 2
Due Thursday, October 17, 3:30pm

Note: It is highly recommended (though not required) that you type your answers. It is your responsibility to make any handwriting clear and legible for grading. You may work with 1-2 other collaborators, but you must write the solutions separately and clearly mark the names of each person you worked with.

Problems:

1. **Trace distance and the max distinguishing probability of quantum states**

   How do we determine how close two n-qubit quantum states with density matrices ρ and σ are to each other? One popular metric is called the *trace distance*:

   $$\|\rho - \sigma\|_{\mathrm{tr}} = \frac{1}{2} \max_U \mathrm{tr} \left| U\rho U^\dagger - U\sigma U^\dagger \right|,$$

   where the maximum is over all unitary matrices U, and $\mathrm{tr}\,|\cdot|$ is the sum of the absolute values of the elements along the diagonal. (Note that because the unitary matrices form a compact group, the maximum in the definition is equivalent to a supremum.)

   (a) The trace distance is particularly useful because it is related to the maximum probability of distinguishing two quantum states: suppose you are given state ρ with probability 1/2 and state σ with probability 1/2. Prove that the highest probability with which you can determine the state given to you is

   $$\frac{1 + \|\rho - \sigma\|_{\mathrm{tr}}}{2}.$$

   You may assume that the strategy consists of applying some unitary matrix U followed by a computational basis measurement where the measurement outcomes are partitioned into two sets. If the measurement falls in one set, you conclude that the state you were given was ρ; otherwise, it falls in the other set, and you conclude that the state was σ. (Hint: write out the probability of success for any strategy, and then take the maximum.)

   (b) Show that the trace distance satisfies the triangle inequality: for all density matrices ξ, $\|\rho - \sigma\|_{\mathrm{tr}} \leq \|\rho - \xi\|_{\mathrm{tr}} + \|\sigma - \xi\|_{\mathrm{tr}}$.

2. **Quantum errors only accumulate linearly**

   Suppose we wish to apply the n-qubit quantum operation $U = U_t U_{t-1} \cdots U_1$, where unitary U is the product of t individual gates. Unfortunately, for each gate, there is some noise in our implementation so that when we try to apply U_i, we actually apply V_i. Thankfully, for any state ρ and any gate U_i, we have the guarantee that $\|V_i \rho V_i^\dagger - U_i \rho U_i^\dagger\|_{\mathrm{tr}} \leq \epsilon$. Prove that these errors only add up linearly as we apply the entire sequence of gates that computes U—that is, setting $V = V_t V_{t-1} \cdots V_1$, show that $\|V\rho V^\dagger - U\rho U^\dagger\|_{\mathrm{tr}} \leq t\epsilon$.

3. **Quantum computation does not require complex numbers**

   Suppose we have an n-qubit circuit constructed from a sequence of 1- and 2-qubit gates $g_1, \dots, g_m$. Show that there is another $(n+1)$-qubit circuit constructed from gates $g'_1, \dots, g'_m$ such that

   - Each gate g'_i can be efficiently constructed from g_i and only has *real* entries.
   - The probability distribution resulting from measuring the first qubit is the same for each circuit.

   (Hint: each gate g'_i may act on more qubits than g_i.)