CS 593/MA 592 - Intro to Quantum Computation
Homework 2

Due Monday, January 29 at 8pm (upload to Brightspace)

1. Do the following exercises from Nielsen and Chuang: 2.57, 2.58, 2.59, 2.60, 2.61, 2.66.

2. (a) A vector $|\psi\rangle$ in a tensor product Hilbert space $V \otimes W$ is called *separable* (or *unentangled*) if there exist vectors $|v\rangle \in V$ and $|w\rangle \in W$ such that $|\psi\rangle = |v\rangle \otimes |w\rangle$. Give an exampe of a state $|\psi\rangle \in (\mathbb{C}^2)^{\otimes 2}$ on two qubits that is not separable (in other words, it is entangled). Justify your answer.

   (b) Show that $V \otimes W$ has no entangled states if and only if V or W is 0 or 1 dimensional.

3. Let's work through the details of quantum state tomography via repeated measurements in the computational basis.

   Let
   $$|\psi\rangle = \sum_{b=0}^{2^n-1} z_b |b\rangle \in (\mathbb{C}^2)^{\otimes n}$$
   be some unknown state on n qubits, which we will assume is normalized. The goal of quantum state tomography is to determine what the amplitudes z_b are—up to a given error, with high confidence. We don't yet have the tools to do things at this level of precision quite yet, but we can at least ask about trying to determine, say, $|z_0|^2$ up to some given accuracy.

   Since measurement collapses the state, we will assume that we are able to prepare copies of this state for free. On each copy, we will perform projective measurement in the computational basis. The outcomes will be independent and identically distributed. If we do this k times, we get a sequence of outcomes $(i_1, \dots, i_k)$ where each $i_j \in \{0, \dots, 2^n - 1\}$. From this, we may compute an empirical probability distribution $\tilde{p}_k$ on the set $\{0, \dots, 2^{n-1}\}$ simply by counting the different outcomes and dividing by k
   $$\tilde{p}_k(i) := \frac{\#\{j \mid i_j = i\}}{k}.$$
   Of course, the *true* distribution of outcomes is given by the Born rule:
   $$p(i) = p(i \mid |\psi\rangle) = |z_i|^2 = z_i z_i^*.$$

   Let $\epsilon > 0$. We would like to know how many rounds of our experiment we need to perform—that is, how large k needs to be—in order for us to be able to *confidently* say that our empirical estimate $\tilde{p}_k(0)$ is within ϵ of the true value $p(0)$. This requires a little bit of explaining, basically having to do with the fact that $\tilde{p}_k(0)$ is itself a random variable (on the set $\{0, 1/k, 2/k, \dots, k/k = 1\}$, but don't think too hard about this).

   Let us say that we are δ-*confident* that our observed $\tilde{p}_k(0)$ is within ϵ if we pick k large enough so that
   $$\text{Prob}(|\tilde{p}_k(0) - p(0)| \geq \epsilon) \leq \delta.$$

Our goal is to find a lower bound on k (as a function of ϵ, but independent of everything else) that makes this inequality true.

To do so, we can use Chebyshev's inequality (see Appendix 1 in Nielsen and Chuang). This problem will walk you through this. The idea is exactly the same as trying to get a good estimate of the bias of an unfair coin with high confidence.

(a) Let Y be the random variable on the set $\{0, 1\}$ with $p(0) = 1 - |z_0|^2$ and $p(1) = |z_0|^2$.[1] Show that $\mathbb{E}(Y) = \mathbb{E}(Y^2) = |z_0|^2$. Use this to show the variance $\mathrm{var}(Y) = |z_0|^2 - |z_0|^4 = |z_0|^2(1 - |z_0|^2)$.

(b) Show that $\max_{0 \le p \le 1} p(1-p) = 1/4$. Conclude that $\mathrm{var}(Y) \le 1/4$.

(c) Now let $Y_1, \ldots, Y_k$ be k i.i.d variables all having the same distribution as Y.[2] Let X_k be the sample mean

$$\frac{1}{k}\sum_{i=1}^{k} Y_i.$$

Show that X_k is exactly the same thing as $\tilde{p}_k(0)$. (This should be very easy.)

(d) Use the fact that expectation values are linear to show $\mathbb{E}(X_k) = \mathbb{E}(\tilde{p}_k(0)) = p(0)$. (In the language of probability theory, this shows that $\tilde{p}_k(0)$ is an "unbiased estimator" of the true probability $p(0)$.)

(e) Since the Y_i are independent, the variance of their sum is the sum of their variances. Use this to show $\mathrm{var}(X) = \frac{1}{k}\,\mathrm{var}(Y)$.

(f) Now use Chebyshev's inequality to argue that we should take $k \ge \frac{1}{4\epsilon^2\delta}$.

(g) How big should k be if we want to be 95% confident that our estimate of $|z_0|^2$ is correct up to b bits?

Let me conclude by noting that there are better ways to do quantum state tomography!

[1]So, we should interpret outcome 0 for Y as "after measuring $|\psi\rangle$ once in the computational basis, we did not see outcome 0." Similarly, we should interpret outcome 1 as "after measuring $|\psi\rangle$ once in the computational basis, we DID see outcome 0."

[2]Think of these as the different measurements we perform on k copies of $|\psi\rangle$.