CSE 190 / Math 152 - Introduction to Quantum Computing
Homework 4
Due **Tuesday, April 30th, 1:30pm**

Instructions: **You may work individually or in a team of 2 people.** You may switch teams for different assignments. Please ensure your name(s) and PID(s) are clearly visible on the first page of your submission, and then upload the PDF to Gradescope. If working in a group, submit only one submission per group: one partner uploads the submission through their Gradescope account and then adds the other group member to the Gradescope submission by selecting their name in the "Add Group Members" dialog box. You will need to re-add your group member every time you resubmit a new version of your assignment.

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. A LaTeX template for the homework is provided on Canvas.

We will only be grading some of the problems below for correctness. However, because all of the concepts are important, we will not reveal which problems are being graded for correctness until after the assignment has been submitted. The remaining problems will be graded for completeness (i.e., does it look like there was a good-faith effort to solve the problem?).

Problems:

1. **Constructing states with partial measurement**

   In the last homework assignment, you built circuits for various states by giving an explicit sequence of unitary gates to construct them. Sometimes, however, it will be helpful to use partial measurements as another tool. In this problem, we will explore the use of partial measurements to construct a special class of states. First, for a bit string $x \in \{0,1\}^n$, we define the *Hamming weight* to be the number of 1's in x. For integers k, n such that $k \leq n$, we define the state $|D_k^n\rangle$ to be the uniform superposition of all n-qubit classical basis states that have Hamming weight k.

   For example, we have the state

   $$|D_1^2\rangle = \frac{|01\rangle + |10\rangle}{\sqrt{2}}$$

   which is superposition of states with Hamming weight 1. While this is a fairly simple state to construct without partial measurement, let's still look at it as a toy example:

   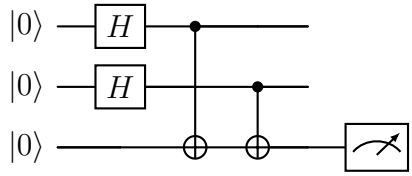

   Suppose the measurement on the third qubit has outcome $b \in \{0,1\}$.

   (a) Show that the first two qubits are in the state $|D_1^2\rangle$ if and only if $b = 1$.

   What you just showed above is that state contstruction based on partial measurement can have some probability of failure. This can be an acceptable trade-off for many quantum applications—if you fail to create the state you wanted to, just start over and try again.

   (b) Using the same idea for the construction of the $|D_1^2\rangle$ state, construct the $|D_2^3\rangle$ state with partial measurement. Your circuit should have 5 total qubits—the top three for the state and the bottom two for the measurement. Furthermore, you can only use Hadmard, CNOT, and Toffoli gates.

   (c) With what probability does your construction succeed? That is, with what probability are the first three qubits in the state $|D_2^3\rangle$ after measurement.

   Recall that you also constructed $|D_2^3\rangle$ in the previous homework using a sequence of unitary gates. Notice in this case that we were able to construct the same state with a simpler set of gates. This comes at the cost of having some probability of failure.

   Sometimes, we can boost our probability of success by post-processing the quantum state based on the measurement result. Consider the circuit:

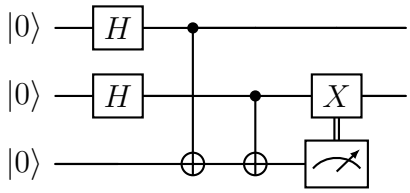

Here, the vertical double wires indicates that application of the X gate depends on the outcome of the measurement. Once again, let $b \in \{0, 1\}$ be the outcome of the measurement on the third qubit.

(d) Suppose we apply the X gate only when $b = 0$. Show that the above circuit always constructs the state $|D_1^2\rangle$. In other words, by applying an extra gate depending on the outcome of the partial measurement, we can boost our probability of success all the way to 1.

Such a strategy can be useful more generally:

(e) Show that you can double the probability of success of your construction in part (b) by applying a layer of single-qubit gates that depend on the measurement result.

Let's now generalize the ideas above:

(f) For all values of n and k, describe a quantum circuit that constructs the state $|D_k^n\rangle$. What is the probability of success of your construction? It will turn out that there are constructions of this state that succeed with high probability. However, for this problem, it's okay if your construction succeeds with very small probability (as long its non-zero).

In your construction, you are allowed to use a unitary that computes Hamming weights. That is, for any bit string $x \in \{0, 1\}^n$, you can apply a unitary which maps $|x\rangle |0 \cdots 0\rangle$ to $|x\rangle |\mathrm{hw}(x)\rangle$, where $\mathrm{hw}(x)$ is the Hamming weight of x. Here, it is assumed that the "$0 \cdots 0$" part of the input register has at least $\lceil \log n \rceil$ many qubits, so that the Hamming weight can be output in binary.

2. **A different kind of oracle**

For any Boolean function $f\colon \{0,1\}^n \to \{0,1\}$, recall that we defined in class an oracle O_f, which allows us to apply f unitarily in a quantum circuit. Specifically, we define O_f to be the matrix that has the following behavior on the $(n+1)$-qubit classical basis states:

$$O_f |x\rangle |b\rangle = |x\rangle |f(x) \oplus b\rangle$$

for all $x \in \{0,1\}^n$ and $b \in \{0,1\}$, and where "$\oplus$" is used for the XOR function (alternatively, addition mod 2).

Sometimes however, it will be useful to have a different kind of oracle available to use in a quantum algorithm. We define the *phase oracle* for the Boolean function $f\colon \{0,1\}^n \to \{0,1\}$ to be the matrix that has the following behavior on the n-qubit classical basis states:

$$U_f |x\rangle = (-1)^{f(x)} |x\rangle$$

for all $x \in \{0,1\}^n$.

It turns out that these two oracles are equivalent in the sense that any quantum circuit using one could be simulated by one using the other. Before we show that, let's start with some basics:

(a) Show that O_f is unitary. *Hint: What is the behavior of O_f on the classical basis states? Is the mapping between classical basis states one-to-one?*

(b) Show that U_f is unitary.

To simulate U_f using the O_f oracle, we can use the following circuit:

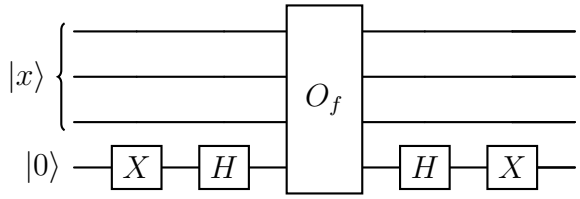

(c) Show that the circuit above implements the U_f unitary. In other words, for all inputs $|x\rangle$, the state of the top n qubits in the circuit above is $(-1)^{f(x)} |x\rangle$. The bottom qubit (called an *ancilla*) starts and ends in the state $|0\rangle$.

To simulate the O_f gate with the U_f gate, let us give ourselves the slightly stronger controlled-U_f gate. Recall that for any n-qubit unitary U, we define the controlled-U gate (denoted C-U) to be the $(n+1)$-qubit gate for which

$$\text{C-}U(|0\rangle \otimes |\psi\rangle) = |0\rangle \otimes |\psi\rangle$$
$$\text{C-}U(|1\rangle \otimes |\psi\rangle) = |1\rangle \otimes U |\psi\rangle$$

for all n-qubit states $|\psi\rangle$.

(d) Show how to implement the O_f gate using a single C-U_f gate. You may use any basic gates we've encountered in the class.