

Exercises for everyone (try these first)

1. Show that unitaries cannot “delete” information: there is no 1-qubit unitary U that maps $|\phi\rangle \mapsto |0\rangle$ for every 1-qubit state $|\phi\rangle$.
2. Show that surrounding the target qubit of a CNOT gate with Hadamard gates turns it into a controlled- Z gate, i.e., $(I \otimes H)\text{CNOT}(I \otimes H) = \text{CZ}$.
3. Suppose we have one copy of the state $\frac{1}{\sqrt{2}}(|0\rangle|\phi\rangle + |1\rangle|\psi\rangle)$, where $|\phi\rangle$ and $|\psi\rangle$ are unknown normalized quantum states with the same number of qubits. Suppose we apply a Hadamard gate to the first qubit and then measure that first qubit in the computational basis. Give the probability of measurement outcome 0, as a function of the states $|\phi\rangle$ and $|\psi\rangle$.

Exercises to choose from

4. This question is about parallelizing a variant of quantum phase estimation. Suppose $U = \begin{pmatrix} 1 & 0 \\ 0 & e^{2\pi i\phi} \end{pmatrix}$ with unknown phase $\phi = \{0, 1/4, 1/2, 3/4\}$, so 2 bits suffice to represent ϕ .
 - (a) Show how to prepare a qubit $\frac{1}{\sqrt{2}}(|0\rangle + e^{2\pi i\phi}|1\rangle)$ using one application of U and some elementary gate(s).
 - (b) Show how to prepare a qubit $\frac{1}{\sqrt{2}}(|0\rangle + e^{2\pi i2\phi}|1\rangle)$ using two parallel applications of U and some elementary gate(s). You’re allowed to use an auxiliary qubit.
Hint: Set up an EPR-pair.
 - (c) Give a circuit that acts on 3 qubits, that uses $U^{\otimes 3}$ and no other applications of U , and that returns ϕ with probability 1.
5. Suppose we have a 2-bit input $x = x_0x_1$ and a phase-query that maps

$$O_x : |b\rangle \mapsto (-1)^{x_b}|b\rangle \text{ for } b \in \{0, 1\}.$$

- (a) Suppose we run the 1-qubit circuit HO_xH on initial state $|0\rangle$ and then measure. What is the probability distribution on the output bit, as a function of x ?
- (b) Now suppose the query leaves some workspace in a second qubit, which is initially $|0\rangle$:

$$O'_x : |b, 0\rangle \mapsto (-1)^{x_b}|b, b\rangle \text{ for } b \in \{0, 1\}.$$

Suppose we just ignore the workspace and run the algorithm of (a) on the first qubit with O'_x instead of O_x (and $H \otimes I$ instead of H , and initial state $|00\rangle$). What is now the probability distribution on the output bit (i.e., measuring the 1st of the 2 qubits)?

Comment: This exercise illustrates why it’s important to “clean up” (i.e., set back to $|0\rangle$) workspace qubits of some subroutine before running it on a superposition of inputs: the unintended entanglement between the address and workspace registers can thwart the intended interference effects.

6. Show how we can delay intermediate measurement until the end of the computation, using one auxiliary qubit for each 1-qubit measurement that needs to be delayed.