University of Oslo

FYS3415/FYS4415 — Quantum computing and quantum information

Problem set 6

This problem set will be discussed on **Mar. 4**.

Problem 1: Measurement in the Circuit Model

a) Measurement commutes with controls. Consider a qubit q_1 that is the control of a controlled-U gate, with target q_2. All measurements are in the computational basis.

i) The following two procedures act on an arbitrary two-qubit state $|\Psi\rangle$:



Show that the two procedures produce the same measurement probabilities and the same post-measurement state for each outcome. Deduce that *measurement commutes with controls*: it does not matter whether the control qubit is measured before or after the controlled operation.

ii) Now suppose the circuit also contains an unconditional single-qubit gate V on q_1 *between* the controlled-U and the measurement. Compare the two procedures:



Does the equivalence from (i) still hold? Work out both procedures and compare the measurement probabilities and post-measurement states.

Hint[1].

Solution:

i) Write $|\Psi\rangle = |0\rangle\,|\phi_0\rangle + |1\rangle\,|\phi_1\rangle$.

Procedure A: The controlled-U gives $|0\rangle\,|\phi_0\rangle + |1\rangle U\,|\phi_1\rangle$. Measuring q_1:

- $m = 0$: probability $\|\,|\phi_0\rangle\|^2$, state $|0\rangle\,|\phi_0\rangle / \|\,|\phi_0\rangle\|$.
- $m = 1$: probability $\|U\,|\phi_1\rangle\|^2 = \|\,|\phi_1\rangle\|^2$, state $|1\rangle U\,|\phi_1\rangle / \|\,|\phi_1\rangle\|$.

Procedure B: Measuring q_1 first:

- $m = 0$: probability $\|\,|\phi_0\rangle\|^2$, state $|0\rangle\,|\phi_0\rangle / \|\,|\phi_0\rangle\|$. No correction applied.
- $m = 1$: probability $\|\,|\phi_1\rangle\|^2$, state $|1\rangle\,|\phi_1\rangle / \|\,|\phi_1\rangle\|$. Apply U to q_2: $|1\rangle U\,|\phi_1\rangle / \|\,|\phi_1\rangle\|$.

The probabilities and post-measurement states are identical in both procedures.

[1] Write $|\Psi\rangle = |0\rangle\,|\phi_0\rangle + |1\rangle\,|\phi_1\rangle$ and work out each procedure separately.

ii) Write $|\Psi\rangle = |0\rangle\,|\phi_0\rangle + |1\rangle\,|\phi_1\rangle$ as before, and let $V = \begin{pmatrix} v_{00} & v_{01} \\ v_{10} & v_{11} \end{pmatrix}$ in the computational basis, so that $V\,|0\rangle = v_{00}\,|0\rangle + v_{10}\,|1\rangle$ and $V\,|1\rangle = v_{01}\,|0\rangle + v_{11}\,|1\rangle$.

Procedure A′: After the controlled-U: $|0\rangle\,|\phi_0\rangle + |1\rangle U\,|\phi_1\rangle$. After V on q_1:

$$V\,|0\rangle\,|\phi_0\rangle + V\,|1\rangle\,U\,|\phi_1\rangle = |0\rangle\big(v_{00}\,|\phi_0\rangle + v_{01}U\,|\phi_1\rangle\big) + |1\rangle\big(v_{10}\,|\phi_0\rangle + v_{11}U\,|\phi_1\rangle\big).$$

Measuring q_1:

- $m = 0$: probability $\|v_{00}\,|\phi_0\rangle + v_{01}U\,|\phi_1\rangle\|^2$, target state $\propto v_{00}\,|\phi_0\rangle + v_{01}U\,|\phi_1\rangle$.
- $m = 1$: probability $\|v_{10}\,|\phi_0\rangle + v_{11}U\,|\phi_1\rangle\|^2$, target state $\propto v_{10}\,|\phi_0\rangle + v_{11}U\,|\phi_1\rangle$.

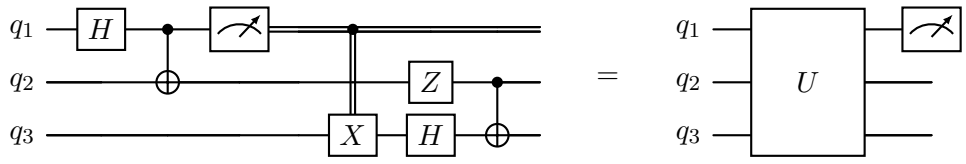
Procedure B′: Apply V first: $V\,|0\rangle\,|\phi_0\rangle + V\,|1\rangle\,|\phi_1\rangle = |0\rangle(v_{00}\,|\phi_0\rangle + v_{01}\,|\phi_1\rangle) + |1\rangle(v_{10}\,|\phi_0\rangle + v_{11}\,|\phi_1\rangle)$. Measuring q_1:

- $m = 0$: probability $\|v_{00}\,|\phi_0\rangle + v_{01}\,|\phi_1\rangle\|^2$, target state $\propto v_{00}\,|\phi_0\rangle + v_{01}\,|\phi_1\rangle$. No U applied.
- $m = 1$: probability $\|v_{10}\,|\phi_0\rangle + v_{11}\,|\phi_1\rangle\|^2$, target state $\propto v_{10}\,|\phi_0\rangle + v_{11}\,|\phi_1\rangle$. Apply U: $\propto v_{10}U\,|\phi_0\rangle + v_{11}U\,|\phi_1\rangle$.

The measurement probabilities on q_1 are *not* in general the same (compare the $m = 0$ terms: A′ has $v_{01}U\,|\phi_1\rangle$ while B′ has $v_{01}\,|\phi_1\rangle$). Neither are the post-measurement target states. The equivalence from (i) breaks down because V mixes the $|0\rangle$ and $|1\rangle$ components of q_1, so the controlled-U and the measurement no longer decouple cleanly. The lesson is that "measurement commutes with controls" applies specifically to the control–target structure; inserting a gate on the control qubit between the controlled gate and the measurement changes the problem.

b) The principle of deferred measurement states that intermediate measurements can always be moved to the end of a quantum circuit.

Consider the following equivalence, where an intermediate measurement and classical control are replaced by a single three-qubit unitary block U followed by a deferred measurement:

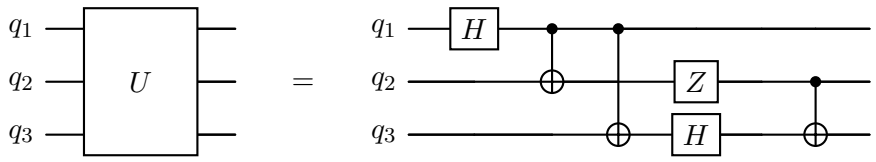


i) Find U. Draw the explicit quantum circuit inside the block U.

ii) Explain why this equivalence holds. What specific rules or previous results guarantee this?

Solution:

i) To construct the unitary U, we replace the classical control with a quantum CNOT gate and slide the measurement to the very end. The internal circuit for U is:



(Note: The measurement is placed outside of U, so it is not drawn inside the block).

ii) The exact equivalence relies on two key principles:

- **Measurement commutes with controls:** As proved in part (a)(i), measuring a control qubit and applying a classically controlled gate is perfectly equivalent to applying a quantum controlled gate and measuring the control qubit afterwards. This justifies swapping the classical wire for a quantum CNOT.
- **Operations on different wires commute:** The Z, H, and final CNOT gates act entirely on q_2 and q_3. Because operators acting on separate subsystems commute, the measurement operator on q_1 does not interact mathematically with them. This justifies safely sliding the measurement past those entangling gates so it can be factored outside the unitary block U.

c) The Principle of Implicit Measurement states that any unterminated quantum wires at the end of a circuit can be assumed to be measured (and their results discarded) without affecting the state of the remaining qubits.

In problem 4 last week, you calculated Bob's reduced density matrix in two different scenarios:

- Scenario 1 (Unterminated wires): Alice does nothing to her qubits. You calculated Bob's state by taking the partial trace over Alice's system before she measures: $\rho^B_{\text{pre}} = \text{Tr}_A(|\Psi_2\rangle\langle\Psi_2|)$.
- Scenario 2 (Measure and discard): Alice measures her qubits, obtaining outcome m with probability $p(m)$, but Bob never receives the classical result. You calculated Bob's averaged state before he received Alice's outcomes as: $\rho^B_{\text{after}} = \sum_m p(m)\rho^B_m$.

You proved that $\rho^B_{\text{pre}} = \rho^B_{\text{after}} = \mathbb{1}/2$. Furthermore, in the general proof of no-signaling, you proved that $\text{Tr}_A(\rho^{AB}) = \sum_m p(m)\rho^B_m$ holds for *any* bipartite state and any measurement. Explain in 2–3 sentences how this specific mathematical identity, $\text{Tr}_A(\rho^{AB}) = \sum_m p(m)\rho^B_m$, perfectly justifies the principle of implicit measurement.

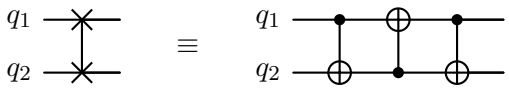
Solution: The left side of the equation, $\text{Tr}_A(\rho^{AB})$, represents the physical act of doing nothing and simply ignoring subsystem A (unterminated wires). The right side, $\sum_m p(m)\rho^B_m$, represents explicitly measuring subsystem A, getting classical outcomes, and discarding that classical information. Because the general no-signaling proof guarantees these two matrices are always equal, there is no mathematical or physical difference for the rest of the circuit. Therefore, adding a measurement that doesn't control any future gates is functionally identical to leaving the wire unterminated.

Problem 2: Circuit decomposition and gate construction

In this problem you will practice a key circuit-design technique — decomposing a gate into simpler pieces and upgrading each piece to its controlled version — and connect the result to the universality of quantum computation.

a) Gate decompositions and controlled gates.

i) Show that the SWAP gate can be decomposed into three alternating CNOTs:



Verify this by computing the action on $|a\rangle\,|b\rangle$ for arbitrary computational basis states $a, b \in \{0, 1\}$.

ii) Apply the "decompose then control" strategy to build a controlled-Z (CZ) gate: First, verify the identity $Z = HXH$. Next, upgrade all three gates in this decomposition to their controlled versions (yielding controlled-H, CNOT, and controlled-H). Finally, argue why the controls on the outer H gates can be safely removed, giving a circuit for CZ using only one CNOT and two unconditional H gates.

Solution:

i) We trace the circuit on each computational basis state (recall that CNOT flips the target if and only if the control is $|1\rangle$):

$$\begin{aligned}
|00\rangle &\xrightarrow{\text{CNOT}_{1\to2}} |00\rangle \xrightarrow{\text{CNOT}_{2\to1}} |00\rangle \xrightarrow{\text{CNOT}_{1\to2}} |00\rangle \\
|01\rangle &\xrightarrow{\text{CNOT}_{1\to2}} |01\rangle \xrightarrow{\text{CNOT}_{2\to1}} |11\rangle \xrightarrow{\text{CNOT}_{1\to2}} |10\rangle \\
|10\rangle &\xrightarrow{\text{CNOT}_{1\to2}} |11\rangle \xrightarrow{\text{CNOT}_{2\to1}} |01\rangle \xrightarrow{\text{CNOT}_{1\to2}} |01\rangle \\
|11\rangle &\xrightarrow{\text{CNOT}_{1\to2}} |10\rangle \xrightarrow{\text{CNOT}_{2\to1}} |10\rangle \xrightarrow{\text{CNOT}_{1\to2}} |11\rangle
\end{aligned}$$

In every case $|a\rangle\,|b\rangle \mapsto |b\rangle\,|a\rangle$, confirming that the three-CNOT circuit acts as SWAP on the computational basis.

To see that this is enough in general: for an arbitrary input

$$|\psi\rangle = \alpha\,|00\rangle + \beta\,|01\rangle + \gamma\,|10\rangle + \delta\,|11\rangle,$$

the table gives

$$|\psi\rangle \;\longmapsto\; \alpha\,|00\rangle + \beta\,|10\rangle + \gamma\,|01\rangle + \delta\,|11\rangle = \text{SWAP}\;|\psi\rangle.$$

ii) First, checking the matrix multiplication for HXH:

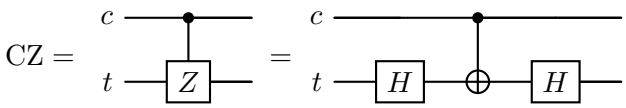
$$HXH = \frac{1}{\sqrt{2}}\begin{pmatrix}1 & 1\\ 1 & -1\end{pmatrix}\begin{pmatrix}0 & 1\\ 1 & 0\end{pmatrix}\frac{1}{\sqrt{2}}\begin{pmatrix}1 & 1\\ 1 & -1\end{pmatrix} = \begin{pmatrix}1 & 0\\ 0 & -1\end{pmatrix} = Z.$$

Following the strategy, we want to build a controlled-Z. We replace Z with HXH, and upgrade every gate to a controlled gate. This gives a circuit with a controlled-H, a CNOT (controlled-X), and a controlled-H.

Crucially, the outer controlled-H gates can be replaced by unconditional H gates. Why? Consider the state of the control qubit c:

- If $c = |1\rangle$: The CNOT fires. The target qubit undergoes $HXH = Z$. This is correct.
- If $c = |0\rangle$: The CNOT does *not* fire (acts as Identity). The target qubit undergoes $H \cdot I \cdot H = H^2 = I$. The target is unchanged, which is also correct!

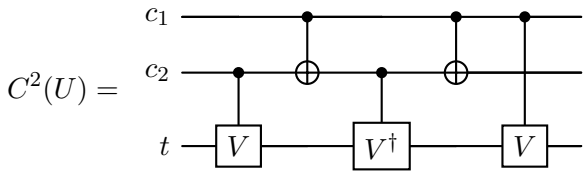
Thus, we can safely drop the controls on the outer H gates because $H^2 = I$:



This technique — upgrading a decomposition and then removing controls from outer self-inverse pairs — is highly useful.

b) Fredkin gate construction Recall that the Fredkin (controlled-SWAP) gate swaps two target qubits if and only if a control qubit is $|1\rangle$.

The following circuit implements a doubly-controlled-U gate ($C^2(U)$) using a gate V with $V^2 = U$:

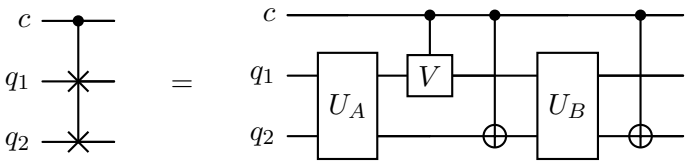


i) Give a quantum circuit which uses three Toffoli gates to construct the Fredkin gate.

ii) Show that the first and last Toffoli gates can be downgraded to plain CNOT gates.

iii) Now replace the middle Toffoli gate with the $C^2(U)$ circuit above to obtain a Fredkin gate construction using only 7 two-qubit gates. Draw the circuit diagram.

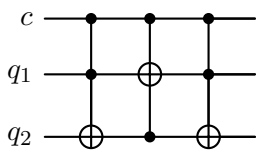
iv) Using explicit commutation arguments, show how to reduce the sequence to only five two-qubit gates:



(The symbol left of the equals sign denotes a controlled-SWAP, i.e. Fredkin gate.) What operations do the U_A and U_B gates contain?

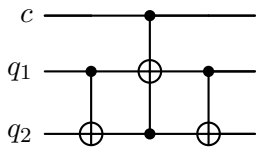
Solution:

i) From part (a), SWAP = CNOT($q_1 \to q_2$)·CNOT($q_2 \to q_1$)·CNOT($q_1 \to q_2$). To make it controlled on c, we upgrade every gate. A controlled-CNOT is exactly a Toffoli gate:

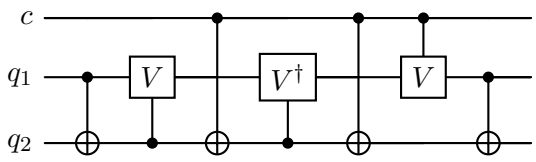


ii) We apply the trick from part a ii). When $c = 0$, the middle gate (Toffoli) does nothing (I). If

we remove the control c from the outer two gates, they revert to standard CNOT($q_1 \to q_2$) gates. When $c = 0$, the circuit acts as CNOT $\cdot I \cdot$ CNOT. Since CNOT is its own inverse, this evaluates to I, which is the correct behavior for $c = 0$. When $c = 1$, the middle gate is CNOT($q_2 \to q_1$), and the full circuit is SWAP, which is also correct. Thus, the outer controls are unnecessary:



iii) The middle gate is a Toffoli (controlled on c, q_2, targeting q_1). This is $C^2(X)$. We use the decomposition block with $U = X$ and $V = \sqrt{X}$ $\left(= \frac{1-i}{2}(I + iX)\right)$. We map the block's wires as follows: $c_1 \to c$, $c_2 \to q_2$, and $t \to q_1$. Inserting this 5-gate sequence between our two outer CNOTs gives a 7-gate circuit:



iv) Look at the first two gates in this sequence: the outer CNOT_{12} (shorthand for CNOT($q_1 \to q_2$)) and the block's first gate CV_{21} (controlled-V from $q_2 \to q_1$). Both of these gates act *exclusively* on the pair of qubits q_1 and q_2. Any sequence of operations on the same two qubits can be mathematically multiplied together to form a single new two-qubit unitary gate. By merging them, the 7-gate sequence reduces to 6 two-qubit gates.

To reduce this to 5 gates, we must be able to merge gates on the right side of the circuit as well. Let's write out the full 7-gate sequence from left to right using our new notation (CV_{ij} for a controlled-V gate with control i and target j):

$$\text{CNOT}_{12} \cdot CV_{21} \cdot \text{CNOT}_{c2} \cdot CV_{21}^\dagger \cdot \text{CNOT}_{c2} \cdot CV_{c1} \cdot \text{CNOT}_{12}$$

We can rearrange this sequence using two fundamental commutation rules:

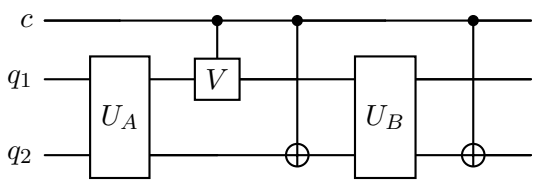
1. CV_{c1} commutes with both CNOT_{c2} (they share the same control but act on different targets) and $CV_{21}^\dagger$ (they share the same target, and their unitaries V and $V^\dagger$ commute). Thus, we can safely slide CV_{c1} to the left, placing it right after CV_{21}.
2. Look at the right end of the circuit. We are now left with CNOT_{c2} followed by CNOT_{12}. Because they share the exact same target (q_2) and their operations on that target (X and X) commute, these two CNOTs commute! We can slide CNOT_{12} to the left, placing it right after $CV_{21}^\dagger$.

After these two slides, our sequence has rearranged to:

$$\underbrace{[\text{CNOT}_{12} \cdot CV_{21}]}_{U_A} \cdot CV_{c1} \cdot \text{CNOT}_{c2} \cdot \underbrace{\left[CV_{21}^\dagger \cdot \text{CNOT}_{12}\right]}_{U_B} \cdot \text{CNOT}_{c2}$$

Notice that the brackets for U_A act exclusively on the pair (q_1, q_2), meaning they combine into a single two-qubit gate. The exact same is true for U_B!

The sequence can now be expressed as the following circuit:



Counting the remaining gates gives us exactly **5 two-qubit gates**.

Problem 3: Irrational Rotations and Approximate Universality

Allowing arbitrary single-qubit gates (a continuous set) together with CNOT gives an exact universal set of gates for quantum computation. We can also consider an approximate universal gate set, when it is possible to implement gates arbitrarily close to any desired unitary using some gate set. A useful way to show approximate universality for single-qubit gates is to find a rotation about some axis by an angle θ which is an *irrational* multiple of 2π. Then the angles $(k\theta) \bmod 2\pi$ obtained by repeating the rotation become dense in $[0, 2\pi)$, giving arbitrarily fine rotation angles about that axis. Together with the ability to rotate about a second, non-parallel axis, this leads to a dense set of single-qubit gates.

In this problem, we will show that the discrete gate set $\{H, T\}$ achieves this. Recall the general rotation operator $U_{\hat{n}}(\alpha) = \cos(\alpha/2)I - i\sin(\alpha/2)(\hat{n} \cdot \vec{\sigma})$. More precisely,

$$e^{-i\pi/8}T = U_z(\pi/4) = \cos\left(\frac{\pi}{8}\right) I - i \sin\left(\frac{\pi}{8}\right) Z.$$

a) Changing the axis of rotation. Show that wrapping the T-gate in Hadamard gates creates an identical rotation, but around the x-axis. That is, show that up to global phase:

$$HTH \sim U_x(\pi/4)$$

Solution: Ignoring the overall global phase, we may replace T by $U_z(\pi/4) = \cos\left(\frac{\pi}{8}\right) I - i \sin\left(\frac{\pi}{8}\right) Z$. Using $H^2 = I$ and $HZH = X$, we get:

$$HTH \sim H\left(\cos\left(\frac{\pi}{8}\right) I - i \sin\left(\frac{\pi}{8}\right) Z\right) H = \cos\left(\frac{\pi}{8}\right) I - i \sin\left(\frac{\pi}{8}\right) X.$$

This matches the formula for $U_x(\pi/4)$, a rotation of $\pi/4$ around the x-axis.

b) The combined rotation. Now consider the unitary operation U created by applying T followed by HTH:

$$U = (HTH)\,T \sim U_x(\pi/4)\,U_z(\pi/4).$$

Expand this matrix product into the form

$$cI - i\left(v_x X + v_y Y + v_z Z\right).$$

Hint[2].

[2] Use the Pauli matrix multiplication rules, specifically $XZ = -iY$.

Solution: We multiply the two rotation operators:

$$U \sim \left(\cos\frac{\pi}{8}I - i\sin\frac{\pi}{8}X\right)\left(\cos\frac{\pi}{8}I - i\sin\frac{\pi}{8}Z\right) \tag{1}$$

$$= \cos^2\frac{\pi}{8}I - i\sin\frac{\pi}{8}\cos\frac{\pi}{8}Z - i\sin\frac{\pi}{8}\cos\frac{\pi}{8}X - \sin^2\frac{\pi}{8}(XZ). \tag{2}$$

Using the Pauli identity $XZ = -iY$, the last term becomes $+i\sin^2\left(\frac{\pi}{8}\right)Y$. Grouping the terms gives:

$$U \sim \cos^2\left(\frac{\pi}{8}\right)I - i\left[\sin\left(\frac{\pi}{8}\right)\cos\left(\frac{\pi}{8}\right)X - \sin^2\left(\frac{\pi}{8}\right)Y + \sin\left(\frac{\pi}{8}\right)\cos\left(\frac{\pi}{8}\right)Z\right].$$

Thus

$$c = \cos^2\left(\frac{\pi}{8}\right), \qquad v_x = \sin\left(\frac{\pi}{8}\right)\cos\left(\frac{\pi}{8}\right), \qquad v_y = -\sin^2\left(\frac{\pi}{8}\right), \qquad v_z = \sin\left(\frac{\pi}{8}\right)\cos\left(\frac{\pi}{8}\right).$$

c) The irrational angle. The last expression is an element of SU(2), so it must be equivalent to some single rotation

$$U_{\hat{n}}(\theta) = \cos(\theta/2)I - i\sin(\theta/2)(\hat{n}\cdot\vec{\sigma})$$

around a new axis $\hat{n}$.

i) By comparing your result from part (b) with the general formula for $U_{\hat{n}}(\theta)$, find an expression for $\cos(\theta/2)$.

ii) Find an expression for the value of $\cos(\theta/2)$. Use a computer algebra system to determine whether θ is a rational or irrational multiple of 2π. Use the result to explain why this implies that $\{H, T\}$ is an approximately universal single-qubit gate set.

Solution:

i) By comparing the coefficient of the identity matrix I in our result from part (b) with the definition of $U_{\hat{n}}(\theta)$, we read off:

$$\cos\left(\frac{\theta}{2}\right) = \cos^2\left(\frac{\pi}{8}\right).$$

ii) Using $\cos^2(x) = \frac{1+\cos(2x)}{2}$,

$$\cos\left(\frac{\theta}{2}\right) = \frac{1+\cos(\pi/4)}{2} = \frac{1+1/\sqrt{2}}{2} = \frac{2+\sqrt{2}}{4} \approx 0.8536.$$

Numerically this corresponds to $\theta \approx 1.0961$ radians, and one can look up that θ is an *irrational* multiple of 2π.

Therefore, powers U^k never repeat exactly and generate a dense set of rotation angles about the axis $\hat{n}$. Since the gate set $\{H, T\}$ also gives rotations about different (non-parallel) axes, the subgroup it generates is dense in SU(2) (up to global phase), meaning any single-qubit rotation can be approximated arbitrarily well. This is precisely the sense in which $\{H, T\}$ is *approximately* universal for single-qubit gates (it does not generate all rotations *exactly*, but it can get arbitrarily close).

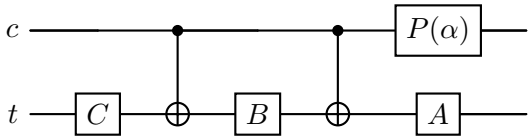
Bonus insight: If we had used the Phase gate $S = U_z(\pi/2)$ instead of T, we would get $\cos(\theta/2) = \cos^2(\pi/4) = 1/2$. This gives $\theta/2 = \pi/3 \implies \theta = 2\pi/3$, which is a rational fraction

of a circle, so repeated applications would cycle through finitely many points. This illustrates why $\{H, S\}$ is not universal.

d) From Single-Qubit to Multi-Qubit Universality. To achieve universal quantum computation, we must extend our single-qubit control to multiple qubits using an entangling gate. A crucial step is showing that we can construct a *controlled-U* gate for *any* single-qubit unitary U using only CNOTs and single-qubit gates.

Any single-qubit unitary can be factored as $U = e^{i\alpha}AXBXC$, where A, B, and C are single-qubit unitaries satisfying $ABC = I$, X is the Pauli-X gate, and $e^{i\alpha}$ is a global phase.

Consider the following circuit, where $P(\alpha) = \left(\begin{smallmatrix} 1 & 0 \\ 0 & e^{i\alpha} \end{smallmatrix}\right)$ is a phase gate:



i) Verify that this circuit correctly implements a controlled-U gate by tracing the evolution of the target qubit when the control is $|0\rangle$ and when it is $|1\rangle$.

ii) A fundamental result in quantum circuit theory establishes that any n-qubit unitary operation can be decomposed into a sequence of single-qubit gates and two-qubit controlled-U gates. In 2–3 sentences, explain how combining this exact decomposition, the circuit above, and the properties of the $\{H, T\}$ gate set from the previous steps shows that $\{H, T, \text{CNOT}\}$ is universal for quantum computation in the *approximate* sense.

Solution:

i) We follow the operations applied to the target qubit t based on the control qubit c, reading the circuit from left to right:

- **Control** $= |0\rangle$**:** The CNOTs do not fire (acting as the identity I), and $P(\alpha)$ acts as the identity on the $|0\rangle$ control branch. The overall operation on the target is $A\,B\,C$, and since $ABC = I$, the target is unchanged. This is exactly what a controlled gate should do when $c = 0$.
- **Control** $= |1\rangle$**:** The CNOTs both fire (applying X), and $P(\alpha)$ contributes a phase $e^{i\alpha}$ on the $|1\rangle$ control branch. The phase $e^{i\alpha}$ is not global but it can be viewed as being part of the control qubit or the target qubit, meaning that we can write the final state as

$$\begin{aligned}\big(P(\alpha)\,|1\rangle\big) \otimes \big(AXBXC\,|\psi\rangle\big) &= \big(e^{i\alpha}\,|1\rangle\big) \otimes \big(AXBXC\,|\psi\rangle\big) = |1\rangle \otimes (e^{i\alpha}AXBXC\,|\psi\rangle) \\ &= |1\rangle \otimes \big(U\,|\psi\rangle\big).\end{aligned}$$

  This is exactly what a controlled-U should do when $c = 1$.

ii) The circuit above shows how any controlled-U can be constructed using only CNOTs and single-qubit gates. Since $\{H, T\}$ can approximate any single-qubit gate arbitrarily well, we can approximate the single-qubit gates appearing in such decompositions; with CNOT available, this implies we can approximate any n-qubit unitary arbitrarily well, i.e. $\{H, T, \text{CNOT}\}$ is universal in the approximate sense.