Quantum Algorithms 2/2/2023

Lecture 4: Quantum Circuits

Instructor: Dieter van Melkebeek

We continue our discussion of intermediate measurements and ways to eliminate them. Incorporating them leads us to the notion of mixed state and the model of a quantum circuit that we will use in the rest of the course. The folklore way of eliminating them involves cloning of basis states and leads to the notion of purification. We also show that cloning of arbitrary states is not possible.

1 Recap

1.1 Quantum Computation

Recall that in a quantum system, a state is described by a superposition

$$|\psi\rangle = \sum_{s\in\{0,1\}^m} \alpha_s |s\rangle$$

with amplitudes $\alpha_s \in \mathbb{R}$ (or $\mathbb{C}$) and $\| |\psi\rangle \|_2 \doteq \sqrt{\sum_s |\alpha_s|^2} = 1$. Computations on these systems go through an initialization phase, with a basis state (e.g., $|x0^{m-n}\rangle$), proceed through a sequence of quantum gates, and then terminate, where $|\psi\rangle$ is measured. This measurement collapses the superposition into a basis state $|s\rangle$, from which output y is extracted. The probability of obtaining any given basis state $|s\rangle$ is given by

$$\Pr[\text{obtain } s] = |\alpha_s|^2.$$

1.2 Simulating Probabilistic Computations and Partial Measurements

Recall that probabilistic computations can be simulated by applying the Hadamard gate H to $|0\rangle$ and observing the component. This provides a basic coin-flip mechanism. However, this requires a measurement that is partial (only part of the qubits) and intermediate. We don't want to measure the entire system all the time because doing so would collapse the system and eliminate interference, losing the power of quantum computations. Partial measurements can be viewed as sequences of single-qubit measurements.

Another way to think about partial measurements is in geometric terms. For a single-qubit measurement, we can project the state $|\psi\rangle$ to be measured onto two subspaces, with one spanned by all the basis vectors where the qubit to be measured is in the zero state, and the other being an orthogonal space where the qubit is in the one state. A partial measurement is then a projection of the state onto one of the two subspaces. The probability of ending up in a particular subspace is the length squared of the projection onto that subspace. The fact that these probabilities add up to one corresponds geometrically to the Pythagorean theorem (as the two subspaces are orthogonal). After the projection, we need to renormalize in order to make the 2-norm one again. If P_b foro $b \in \{0,1\}$ denotes the projection onto the subspace where the measured component is b, with probability $\|P_b |\psi\rangle\|^2$ the new state becomes $P_b |\psi\rangle / \|P_b |\psi\rangle\}$. This interpretation extends in the

natural way to the effect of measuring multiple components simultaneously (which is equivalent to measuring them individually in any order).

2 Exercise #2 – Deferring Measurements

Measurement can be deferred to the end of a computation by using ancilla qubits. Consider the following two circuits starting from same state $|\psi\rangle \doteq \sum_{s\in\{0,1\}^m} \alpha_s |s\rangle$. The circuit on the left does a (full) measurement in between two unitaries U_1 and U_2. The circuit on the right does not have an intermediate measurement but instead applies for each of the m qubits a CNOT controlled by the qubit and acting on a fresh ancilla. Both circuits perform a measurement of the qubits involved in $|\psi\rangle$ at the end.

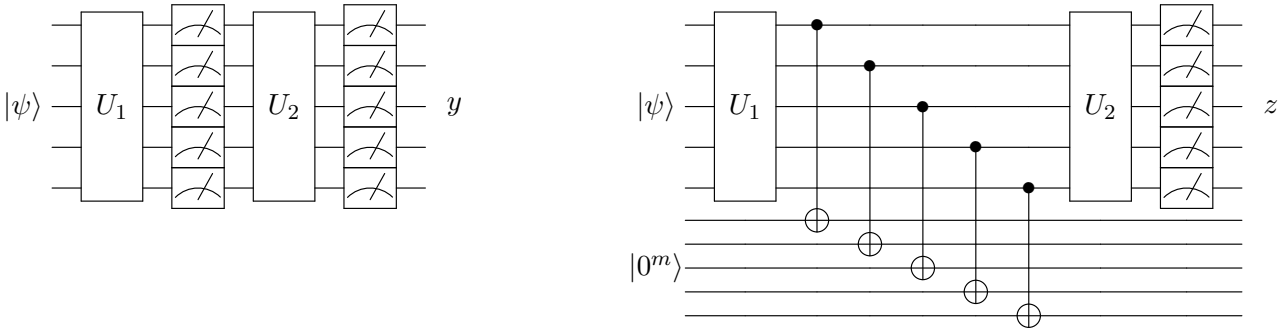

Recall the questions:

(a) Show that the distributions of y and z is the same.

(b) What if intermediate measurements on the left are removed, or equivalently, the CNOTs on the right are removed?

For the circuit on the left, starting from the state $|\psi\rangle$, the state immediately after U_1 is $U_1|\psi\rangle$, which can be written as $U_1|\psi\rangle = \sum_s \beta_s |s\rangle$. This intermediate state is then measured to be in state $|s\rangle$ with a probability of $|\beta_s|^2$. Passing $|s\rangle$ through U_2 yields a state $U_2|s\rangle = \sum_t \gamma_{st}|t\rangle$, which when measured yields the the output $y = t$ with probability $|\gamma_{st}|^2$. By the law of total probability, the overall probability of outputting $y = t$ is given by summing the contributions of all intermediate states s:

$$\sum_s |\beta_s|^2 |\gamma_{st}|^2. \tag{1}$$

For the circuit on the right, the system now starts in a state $|\psi\rangle|0^m\rangle$, and after U_1, is in a state $U_1|\psi\rangle|0^m\rangle$, which, given $U_1|\psi\rangle = \sum_s \beta_s|s\rangle$, equals $\sum_s \beta_s|s\rangle|0^m\rangle$. After the CNOT gates but before U_2, the state becomes $\sum_s \beta_s|s\rangle|s\rangle$. After applying U_2 to the top m qubits, we get that the final state is $\sum_s \beta_s(U_2|s\rangle)|s\rangle = \sum_s \beta_s(\sum_t \gamma_{st}|t\rangle)|s\rangle$, which can be rewritten as $\sum_{s,t} \beta_s\gamma_{st}|t\rangle|s\rangle$. The final measurement then yields the output $z = t$ with probability $\| \sum_s \beta_s\gamma_{st}|s\rangle \|^2$, which equals

$$\sum_s |\beta_s\gamma_s t|^2. \tag{2}$$

Expressions (1) and (2) are equivalent, so the two probability distributions are the same. Thus, replacing the intermediate measurement by the CNOTs with fresh ancillas effectively defers the intermediate measurement to the end of the computation.

Finally, consider what happens when we drop the intermediate measurement on the left are removed, or equivalently, the CNOTs on the right are removed. Note that both are indeed equivalent as the bottom m qubits in the circuit on the right are not involved in the computation. Thus, we have the following circuit:

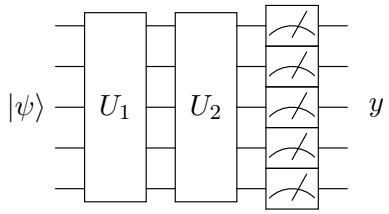

If a measurement is not taken in the middle of U_1 and U_2, then the state right after U_2 is simply $U_2U_1\,|\psi\rangle = \sum_s \beta_s(\sum_t \gamma_{st}\,|t\rangle) = \sum_t(\sum_s \beta_s\gamma_{st})\,|t\rangle$. Then, the probability of outputting $y = t$ is $|\sum_s \beta_s\gamma_{st}|^2$, which is not equivalent to the previous probabilities. In fact, Exercise #1 gives an example in which the probability distributions are different, namely $m = 1$ and $U_1 = U_2 = H$.

A couple remarks:

- When copying the intermediate state with CNOT gates, the resultant state is $\sum_s \beta_s\,|s\rangle\,|s\rangle$, not simply $U_1\,|\psi\rangle\,U_1\,|\psi\rangle = \big(\sum_s \beta_s\,|s\rangle\big)\big(\sum_t \beta_t\,|t\rangle\big)$, as copying might imply. The first state is entangled, whereas the second one is not. The two states behave differently, as can be discovered by first measuring the first m qubits and then the second m qubits. In the first state, if the first measurement yields s, then the second one is guaranteed to give s, as well. In the second state, if the first measurement yields s, the second one yields s with the same probability as the first one, so not 100% unless $|\beta_s| = 1$. Thus, individual basis basis states are being cloned, but not arbitrary superpositions. At the end of the lecture we will prove that cloning arbitrary superpositions is, in fact, not possible (no cloning theorem).

- In case of no measurement or CNOTs in the middle, the various computation paths that lead to basis state $|t\rangle$ can interfere, as is reflected in the summation expression for the final amplitude $\sum_s \beta_s\gamma_{s,t}$. Both the intermediate measurement and the introduction of the CNOTs prevent the interference from happening; the intermediate measurement do this by turning the state into a probability distribution of unit-2 superpositions that cannot interact any further; the CNOTs do it by tagging each basis state $|s\rangle$ with $|s\rangle$, so if they both $|s_1\rangle$ and $|s_2\rangle$ yield contributions to a given basis state $|t\rangle$ for the first register, those contributions cannot interfere as the first one yields $|t\rangle\,|s_1\rangle$ whereas the second one yields $|t\rangle\,|s_2\rangle$.

- The above process of deferring measurements requires fresh ancilla qubits for each deferral. With current technology, adding these additional qubits is infeasible. On the other hand, taking a measurement is slower and less reliable than elementary unitary operations.

3 Intermediate Measurements

So far, we've only been doing measurements at the end to extract some classical information out of the quantum state, in which case the intermediate states of the system can be described by superpositions of basis states. Allowing intermediate measurements necessitates a richer description of intermediate states of the system. Instead of only superpositions of states, we now have to consider a probability distribution over superpositions of 2-norm 1. In probabilistic computations, a probability distribution over probability distributions over deterministic states can be collapsed into a single probability distribution over deterministic states, and thus into a single equivalent superposition. This, however, is not the case in quantum computations.

Some notes on state terminology:

- A *basis state* is $|s\rangle$ for some $s \in \{0,1\}^m$, which represents a deterministic configuration of the system.
- A *pure state* is a superposition $|\psi\rangle = \sum_s \alpha_s |s\rangle$ of basis states $|s\rangle$, each with an amplitude α_s and with $\sum_s |\alpha_s|^2 = 1$.
- A *mixed state* is a probability distribution ρ over superpositions, where $\rho = \{(p_i, |\psi_i\rangle)\}_i$ with $p_i \geq 0$ and $\sum_i p_i = 1$. Each element of ρ is a pair, where p_i is the probability of the system being in the pure state $|\psi_i\rangle$.

As an example of a mixed state, consider the measurement of the first component of

$$\frac{1}{\sqrt{3}}(|01\rangle + |10\rangle - |11\rangle)$$

Last lecture, we showed that:

$$\Pr[\text{measure 0}] = \frac{1}{3}, \text{ new state: } |01\rangle$$
$$\Pr[\text{measure 1}] = \frac{2}{3}, \text{ new state: } \frac{1}{\sqrt{2}}(|10\rangle - |11\rangle)$$

This measurement thus transforms the pure state:

$$|\psi\rangle = \frac{1}{\sqrt{3}}(|01\rangle + |10\rangle - |11\rangle)$$

into the mixed state:

$$\rho = \left\{ \left(\frac{1}{3}, |01\rangle \right), \left(\frac{2}{3}, \frac{1}{\sqrt{2}}(|10\rangle - |11\rangle) \right) \right\}$$

Allowing intermediate measurements enables us to efficiently simulate probabilistic computations. The also naturally arise when combining quantum subroutines; the final measurement of the subroutine becomes an intermediate measurement in the overall computation. Thus, intermediate measurements are convenient in algorithm design. However, they also complicate theoretical analysis. As we have seen, these intermediate measurements can be avoided by entangling the measured components with fresh ancillas using CNOTs. This process gives us a pure state $|\phi\rangle$ on $m + l$ qubits, where m denotes the original number of qubits and l the number of ancillas, or equivalently, the number of intermediate single-qubit measurements. In the circuit on the right of Exercise #2,

$l = m$ and $|\phi\rangle$ corresponds to the state of the entire system right before the final measurement. The state $|\phi\rangle$ has the property that measuring and eliminating the last l qubits yields the mixed state ρ on m qubits. Any state $|\phi\rangle$ with the latter property (for any choice of l) is called a *purification of* ρ. Purification allows us to work with pure states instead of mixed states.

The folklore way of eliminating intermediate measurements requires one fresh ancilla qubit per measurement, which may result in an additive space overheard of $O(t)$ for a computation that runs in time t. As explained before, this is undesirable. Recently, more efficient approximate simulations were developed. The following table represents the state-of-the-art simulations without intermediate measurements of quantum computations running in time t and space s. We indicate their respective time and space complexity as a function of t and s. The folklore simulation is exact; the other two are to within an error of ϵ; in the table we are suppressing the dependency on the error ϵ.

time	space	
$O(t)$	$O(s+t)$	[folklore]
$\text{poly}(t, 2^s)$	$O(s + \log t)$	[FR21], [GRZ21]
$t\,\text{poly}(s)$	$O(s \log t)$	[GR22]

The first method is time-optimal. It is the one we developed of deferring measurements by entangling with fresh ancillas using CNOT gates. The second one is space-optimal. The last one describes an interesting middle ground.

4 Quantum Circuits

In our definition of quantum circuits, we are going to allow intermediate measurements, as they are useful and give additional power to algorithms. We are also going to allow non-classical inputs and outputs to these circuits. While our algorithms as a whole need to start and end in a classical state, we can use circuits with non-classical inputs and outputs like subroutines—bits and pieces of algorithms where some superposition or mixed state is transformed into another. An example of a subroutine like this would be amplitude amplification, which is a procedure transforming one pure state into another.

Definition 1. *A quantum circuit is a system acting on m components (thought of like a register of m qubits) and wires. Operations can be quantum gates acting on a constant number of components, or single-component measurements. Operations on disjoint components can act in parallel.*

A special type of quantum circuits are *unitary circuits*, which are circuits with no measurements and only quantum gates. These types of circuits make analysis simpler, as they are reversible and their transition matrices can be combined. With a unitary circuit of k gates, the reverse can also be realized in k gates, namely by applying the reverse of each unitary gate in the reverse order. The transition matrix of a reverse gate is the complex conjugate transpose of the original transition matrix.

All states in unitary circuits are pure because there are no measurements. Another property of unitary circuits is that their effect is fully specified by the output for each basis state. Each pure state is just a superposition of basis states, and the result of the circuit on this linear combination is the corresponding linear combination of the effects of the circuit on the basis states.

The above properties of unitary circuits are not true in general for all quantum circuits. These circuits can involve measurement, may not be reversible, may induce mixed states, and may not be fully specified by their action on the basis states.

Exercise #3: Give an example of two quantum circuits on the same number of qubits that behave the same on all basis states but not on all pure states.

5 No Cloning

We have been doing some sort of "copying" with CNOTs. We're copying basis states, and as such, we can transform $|\psi\rangle|0^m\rangle \mapsto |\psi\rangle|\psi\rangle$ for every basis state $|\psi\rangle = |s\rangle$. However, this does not hold in general for all pure states.

Theorem 1. *There does not exist a quantum circuit realizing $|\psi\rangle|0^m\rangle \mapsto |\psi\rangle|\psi\rangle$ for every pure state $|\psi\rangle$ on m qubits.*

In fact, the following proof shows that the set of states $|\psi\rangle$ for which a given quantum circuit can succeed in copying, can only consist of states that are pairwise parallel or orthogonal.

Proof. We prove this by beginning with unitary quantum circuits, and then generalize to general quantum circuits. Suppose there exists a unitary quantum circuit that realizes

$$|\psi\rangle|0^m\rangle \mapsto |\psi\rangle|\psi\rangle \tag{3}$$

for every pure state $|\psi\rangle$ on m qubits.

Consider two arbitrary pure states, $|\psi\rangle = |\psi_1\rangle$ and $|\psi\rangle = |\psi_2\rangle$. We take the inner product of the left-hand side of (3) for these two states. We know that the inner product is preserved for unitary operations, so we set this equal to the inner product of the right-hand side, yielding the equation

$$|\langle\psi_1|\psi_2\rangle \cdot \langle 0^m|0^m\rangle| = |\langle\psi_1|\psi_2\rangle \cdot \langle\psi_1|\psi_2\rangle|$$

Setting $a \doteq |\langle\psi_1|\psi_2\rangle|$, we find that $a = a^2$, as $\langle 0^m|0^m\rangle = 1$. This implies that a can only be either 0 or 1 to satisfy the relationship.

In the case that $a = 0$, this implies that $|\psi_1\rangle$ and $|\psi_2\rangle$ are orthogonal, while when $a = 1$, $|\psi_1\rangle$ and $|\psi_2\rangle$ must be parallel. Clearly, two arbitrary pure states $|\psi_1\rangle$ and $|\psi_2\rangle$ are not guaranteed to be orthogonal or parallel to each other, so this is a contradiction.

For the general case, we consider a purification of a general circuit since our circuit can now have arbitrary intermediate measurements, say l single-qubit measurements. A purification with with l extra ancilla qubits allows us to defer these measurements to the end. The purification realizes

$$|\psi\rangle|0^m\rangle|0^l\rangle \mapsto |\psi\rangle|\psi\rangle|g(\psi)\rangle \tag{4}$$

where the new ancilla qubits start in the zero state, and at the end are in some garbage state $|g(\psi)\rangle$. We want it to be the case that after observing the ancillas, regardless of their state, we always get the superposition $|\psi\rangle|\psi\rangle$ in the upper qubits. This means that the state $|g(\psi)\rangle$ cannot be entangled with $|\psi\rangle|\psi\rangle$.

Consider again two arbitrary pure states, $|\psi\rangle = |\psi_1\rangle$ and $|\psi\rangle = |\psi_2\rangle$. We take the inner product of the left-hand side of (4) for these two states. Since purification of the general circuit yields a

unitary circuit, the inner product is preserved, so we again set this equal to the inner product of the right-hand side, this time including the extra ancilla bits.

$$|\langle\psi_1|\psi_2\rangle \cdot \langle 0^m|0^m\rangle \cdot \langle 0^l|0^l\rangle| = |\langle\psi_1|\psi_2\rangle \cdot \langle\psi_1|\psi_2\rangle \cdot \langle g(\psi_1)|g(\psi_2)\rangle|$$

Again taking $a \doteq |\langle\psi_1|\psi_2\rangle|$, we find that $a = a^2 \cdot |\langle g(\psi_1)|g(\psi_2)\rangle|$. Since both a and $|\langle g(\psi_1)|g(\psi_2)\rangle|$ are in $[0, 1]$, we again find that a must be either 0 or 1 (and $|\langle g(\psi_1)|g(\psi_2)\rangle|$ must be 1 if $a = 1$). This implies that $|\psi_1\rangle$ and $|\psi_2\rangle$ must be either orthogonal or parallel, which contradicts that they are arbitrary pure states. Therefore, such a general circuit cannot exist. □

References

[FR21] Bill Fefferman and Zachary Remscrim. Eliminating intermediate measurements in space-bounded quantum computation. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 1343–1356. ACM, 2021.

[GR22] Uma Girish and Ran Raz. Eliminating intermediate measurements using pseudorandom generators. In Mark Braverman, editor, *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, January 31 - February 3, 2022, Berkeley, CA, USA*, volume 215 of *LIPIcs*, pages 76:1–76:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[GRZ21] Uma Girish, Ran Raz, and Wei Zhan. Quantum logspace algorithm for powering matrices with bounded norm. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference)*, volume 198 of *LIPIcs*, pages 73:1–73:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.