OQCL

OPEN QUANTUM COMPUTING LAB

Teacher Notes

Quantum in 7 Sessions: a 20-minute plan and a 75-minute plan for each session, with board examples, misconceptions, discussion prompts, and exit questions.

SESSIONS

1. Qubits and amplitudes
2. Measurement, gates, and interference
3. Entanglement
4. Real hardware: compilation and noise
5. Algorithms: search and phase estimation
6. Near-term methods: Hamiltonians and variational circuits
7. Communication, benchmarking, and your own project

Ali M. Saghiri · William Paterson University
Online version: amsaghiri.github.io/open-quantum-computing-lab/instructor.html
November 2026

Teacher notes — Session 1: Qubits and amplitudes

For instructors | Flipped format

The session in one sentence. A qubit is described by two complex amplitudes, and probabilities come from their squared magnitudes.

The student page for Session 1 is written for self-study: students can read, work the example, run both experiments, and check their own answers without you. These notes are for the class time you do have.

Before class (students, on their own)

- Read Chapter 1.
- Do the worked example and the first experiment on the student page.
- Bring one written prediction and one question.

If you have 20 minutes

MINUTES	WHAT TO DO
0–3	Ask the exit question below cold, and take a show of hands. Do not give the answer yet.
3–10	Board: the worked example from the student page, with students supplying each step.
10–16	Live demo: run the second experiment. Collect predictions first, then run it.
16–20	Return to the exit question, settle it, and assign the project.

If you have a full 75-minute class

MINUTES	ACTIVITY
0–10	Motivating question and show of hands.
10–30	The three points below, with the mathematics on the board.
30–45	Worked example at the board, students supplying steps.

45–55	Misconceptions: put each “common mistake” to the class as a true/false statement.
55–65	Live demo of both experiments, predictions first.
65–75	Project briefing and exit question.

Three points to make

1. A bit has a value. A qubit has two amplitudes, and we only ever see samples.
2. Born rule on the board with one numeric example.
3. Same probabilities does not mean same state: show the two states from the Try-it code.

Board example

Normalize $3|0\rangle + 4i|1\rangle$ and find the probabilities.

1. Squared length: $|3|^2 + |4i|^2 = 9 + 16 = 25$, so the length is 5.
2. Divide by the length: $|\psi\rangle = (3/5)|0\rangle + (4i/5)|1\rangle$.
3. Born rule: $P(0) = (3/5)^2 = 0.36$ and $P(1) = |4i/5|^2 = 0.64$.
4. Check: $0.36 + 0.64 = 1$.

The factor i did not affect either probability. It is a relative phase, and it will matter as soon as another gate is applied.

Misconceptions to raise

- Adding amplitudes and expecting 1. It is the **squared magnitudes** that add to 1.
- Squaring a complex amplitude instead of taking its magnitude first: $|i|^2 = 1$, but $i^2 = -1$.
- Deciding that phase is irrelevant because it does not show up in one measurement.

Discussion prompts

- If we can only observe probabilities, in what sense is the phase “real”?
- What would go wrong if we described a qubit by P(0) and P(1) alone?

Exit question

Ask: Two states both give 50/50 when measured. Are they necessarily the same state?

Looking for: No. They can differ in relative phase, and can even be orthogonal, as $(|0\rangle+i|1\rangle)/\sqrt{2}$ and $(|0\rangle-i|1\rangle)/\sqrt{2}$ are.

Assign

One project from this session, following the predict → baseline → one extension → analyze routine. Project solutions are deliberately not published on this site.

- Q01 Quantum Coin
- Q02 Amplitude and Phase Explorer

Teacher notes — Session 2: Measurement, gates, and interference

For instructors | Flipped format

The session in one sentence. Gates change amplitudes reversibly, measurement samples them, and amplitudes with opposite signs cancel.

The student page for Session 2 is written for self-study: students can read, work the example, run both experiments, and check their own answers without you. These notes are for the class time you do have.

Before class (students, on their own)

- Read Chapter 2.
- Do the worked example and the first experiment on the student page.
- Bring one written prediction and one question.

If you have 20 minutes

MINUTES	WHAT TO DO
0–3	Ask the exit question below cold, and take a show of hands. Do not give the answer yet.
3–10	Board: the worked example from the student page, with students supplying each step.
10–16	Live demo: run the second experiment. Collect predictions first, then run it.
16–20	Return to the exit question, settle it, and assign the project.

If you have a full 75-minute class

MINUTES	ACTIVITY
0–10	Motivating question and show of hands.
10–30	The three points below, with the mathematics on the board.
30–45	Worked example at the board, students supplying steps.

45–55	Misconceptions: put each “common mistake” to the class as a true/false statement.
55–65	Live demo of both experiments, predictions first.
65–75	Project briefing and exit question.

Three points to make

1. Measurement is sampling: show 10 shots versus 1,000 shots.
2. HH = I on the board, tracking the signs.
3. Insert Z: no probability changes, yet the answer flips. That is interference.

Board example

Follow the amplitudes through H, Z, H starting from $|0\rangle$.

1. $H|0\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$.
2. Z flips the sign of $|1\rangle$: $(|0\rangle - |1\rangle)/\sqrt{2}$. The probabilities are still 50/50.
3. Apply H to each part: $H|0\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ and $H|1\rangle = (|0\rangle - |1\rangle)/\sqrt{2}$.
4. Combine: $\frac{1}{2}[(|0\rangle + |1\rangle) - (|0\rangle - |1\rangle)] = \frac{1}{2}[2|1\rangle] = |1\rangle$.

The two $|0\rangle$ contributions cancel and the two $|1\rangle$ contributions add. Without the Z, the signs are the other way round and the result is $|0\rangle$.

Misconceptions to raise

- Reading H as “makes the qubit random”. A random bit stays random after a second H; a superposition does not.
- Thinking more shots change the state. Shots only improve your estimate of the same probabilities.
- Assuming Z “does nothing” because it leaves the probabilities of $|0\rangle$ and $|1\rangle$ unchanged.

Discussion prompts

- How is a superposition different from “we just do not know the value yet”?
- Which experiment could convince a sceptic that the qubit after one H is not simply a random bit?

Exit question

Ask: Why does H followed by H always give 0, when one H gives a random result?

Looking for: The two paths to $|1\rangle$ carry opposite signs and cancel, while the two paths to $|0\rangle$ add.

Assign

One project from this session, following the predict → baseline → one extension → analyze routine. Project solutions are deliberately not published on this site.

- Q03 X, H, and Interference
- Q08 Shot-Count Convergence

Teacher notes — Session 3: Entanglement

For instructors | Flipped format

The session in one sentence. Entangled qubits have joint results that cannot be explained by giving each qubit its own state.

The student page for Session 3 is written for self-study: students can read, work the example, run both experiments, and check their own answers without you. These notes are for the class time you do have.

Before class (students, on their own)

- Read Chapter 3.
- Do the worked example and the first experiment on the student page.
- Bring one written prediction and one question.

If you have 20 minutes

MINUTES	WHAT TO DO
0–3	Ask the exit question below cold, and take a show of hands. Do not give the answer yet.
3–10	Board: the worked example from the student page, with students supplying each step.
10–16	Live demo: run the second experiment. Collect predictions first, then run it.
16–20	Return to the exit question, settle it, and assign the project.

If you have a full 75-minute class

MINUTES	ACTIVITY
0–10	Motivating question and show of hands.
10–30	The three points below, with the mathematics on the board.
30–45	Worked example at the board, students supplying steps.
45–55	Misconceptions: put each “common mistake” to the class as a true/false statement.

55–65	Live demo of both experiments, predictions first.
65–75	Project briefing and exit question.

Three points to make

1. Four basis states for two qubits; 2^n in general.
2. Build the Bell state on the board in three lines.
3. Correlation alone is not entanglement: show the X-basis measurement.

Board example

Build the Bell state step by step, writing states as |first qubit, second qubit⟩.

1. Start: $|00\rangle$.
2. H on the first qubit: $(|00\rangle + |10\rangle)/\sqrt{2}$.
3. CNOT with the first qubit as control flips the second qubit only in the $|10\rangle$ term: $(|00\rangle + |11\rangle)/\sqrt{2}$.

Why it is not two separate qubits. A product state looks like $(a|0\rangle + b|1\rangle)(c|0\rangle + d|1\rangle) = ac|00\rangle + ad|01\rangle + bc|10\rangle + bd|11\rangle$. To match the Bell state we need $ad = 0$ and $bc = 0$, but also $ac \neq 0$ and $bd \neq 0$. If $ad = 0$ then $a = 0$ or $d = 0$, which makes ac or bd zero. No choice works, so the state cannot be split.

Misconceptions to raise

- Saying “the qubits are both 0 or both 1, we just do not know which”. That describes a classical shared coin, not a Bell state.
- Believing that measuring one qubit sends a signal to the other. The correlations cannot be used to communicate.
- Reading Qiskit bit strings left to right. Qiskit prints qubit 0 on the **right**.

Discussion prompts

- Why can entanglement not be used to send a message faster than light?
- What does it mean physically that the Bell state cannot be factored?

Exit question

Ask: A Bell pair and a pair of glued classical coins both give only 00 and 11. How do you tell them apart?

Looking for: Measure both in a second basis (H before measurement). The Bell state stays perfectly correlated; the classical pair does not.

Assign

One project from this session, following the predict → baseline → one extension → analyze routine. Project solutions are deliberately not published on this site.

- Q04 Bell Correlations
- Q06 GHZ Scaling
- Q05 Bell State Under Noise

Teacher notes — Session 4: Real hardware: compilation and noise

For instructors | Flipped format

The session in one sentence. The device runs a compiled version of your circuit, and the extra gates it needs are a main source of error.

The student page for Session 4 is written for self-study: students can read, work the example, run both experiments, and check their own answers without you. These notes are for the class time you do have.

Before class (students, on their own)

- Read Chapter 4 and Chapter 7.
- Do the worked example and the first experiment on the student page.
- Bring one written prediction and one question.

If you have 20 minutes

MINUTES	WHAT TO DO
0–3	Ask the exit question below cold, and take a show of hands. Do not give the answer yet.
3–10	Board: the worked example from the student page, with students supplying each step.
10–16	Live demo: run the second experiment. Collect predictions first, then run it.
16–20	Return to the exit question, settle it, and assign the project.

If you have a full 75-minute class

MINUTES	ACTIVITY
0–10	Motivating question and show of hands.
10–30	The three points below, with the mathematics on the board.
30–45	Worked example at the board, students supplying steps.

45–55	Misconceptions: put each “common mistake” to the class as a true/false statement.
55–65	Live demo of both experiments, predictions first.
65–75	Project briefing and exit question.

Three points to make

1. Logical versus physical circuit; native gates and connectivity.
2. Live demo: star-shaped GHZ on a line doubles the two-qubit gates.
3. Same state as a chain: no overhead. Design with the layout in mind.

Board example

Estimate how much routing costs. Suppose every two-qubit gate succeeds with probability 0.99, and ignore every other error.

1. Logical circuit, 3 two-qubit gates: $0.99^3 \approx 0.970$.
2. Compiled circuit, 6 two-qubit gates: $0.99^6 \approx 0.941$.
3. Routing alone moved the estimate from about 97% to about 94%, before counting single-qubit gates, readout, or idle time.

This is a rough model, not a prediction for any particular device. Its use is to compare two versions of the same circuit before you spend hardware time.

Misconceptions to raise

- Reporting the depth of the circuit you drew instead of the circuit that ran.
- Assuming all qubits on a device are equally good and all pairs are connected.
- Blaming every difference from the ideal result on “noise” without separating sampling error, compilation, and device error.

Discussion prompts

- Who should adapt, the algorithm or the compiler? When is each the better choice?
- Which numbers would you insist on seeing before believing a hardware result?

Exit question

Ask: Two circuits prepare the same state. One compiles to 3 two-qubit gates and one to 6. Which do you run, and why?

Looking for: The one with 3. Every additional two-qubit gate adds error, so fewer gates means a result closer to ideal.

Assign

One project from this session, following the predict → baseline → one extension → analyze routine. Project solutions are deliberately not published on this site.

- Q09 Physical Qubit Mapping
- Q10 Compilation Tradeoffs
- Q23 Cost-Aware Hardware Plan
- Q14 Syndrome Decoding

Teacher notes — Session 5: Algorithms: search and phase estimation

For instructors | Flipped format

The session in one sentence. Grover marks the answer with a sign and then uses interference to turn that sign into a high probability.

The student page for Session 5 is written for self-study: students can read, work the example, run both experiments, and check their own answers without you. These notes are for the class time you do have.

Before class (students, on their own)

- Read Chapter 5 and Chapter 6.
- Do the worked example and the first experiment on the student page.
- Bring one written prediction and one question.

If you have 20 minutes

MINUTES	WHAT TO DO
0–3	Ask the exit question below cold, and take a show of hands. Do not give the answer yet.
3–10	Board: the worked example from the student page, with students supplying each step.
10–16	Live demo: run the second experiment. Collect predictions first, then run it.
16–20	Return to the exit question, settle it, and assign the project.

If you have a full 75-minute class

MINUTES	ACTIVITY
0–10	Motivating question and show of hands.
10–30	The three points below, with the mathematics on the board.
30–45	Worked example at the board, students supplying steps.

45–55	Misconceptions: put each “common mistake” to the class as a true/false statement.
55–65	Live demo of both experiments, predictions first.
65–75	Project briefing and exit question.

Three points to make

1. Oracle flips a sign: nothing visible yet.
2. Diffusion as “reflect about the average”, worked on the board for four states.
3. One round too many spoils it: run the second experiment live.

Board example

Grover on four states, by hand. Suppose the marked state is $|11\rangle$.

1. After the first H gates, all four amplitudes are ½.
2. Oracle: the marked amplitude becomes −½. The amplitudes are (½, ½, ½, −½).
3. Average: (½ + ½ + ½ − ½)/4 = ¼.
4. Diffusion replaces each amplitude a by 2 × average − a. Unmarked: 2 × ¼ − ½ = 0. Marked: 2 × ¼ − (−½) = 1.

The final amplitudes are (0, 0, 0, 1), so the marked state is measured with certainty after one round.

Misconceptions to raise

- Saying Grover “tries all answers at once”. It rotates amplitude toward the answer step by step.
- Running more rounds to be safe. Too many rounds rotate the state past the answer and the success probability falls.
- Forgetting that the oracle changes only a sign. On its own it changes no probability.

Discussion prompts

- Where exactly does the speed-up come from, if the oracle is called one item at a time?
- Why is a quadratic speed-up for unstructured search less dramatic than it sounds for real databases?

Exit question

Ask: After the oracle, what is the probability of measuring the marked state among four?

Looking for: Still ¼. The oracle changes only the sign of the amplitude.

Assign

One project from this session, following the predict → baseline → one extension → analyze routine. Project solutions are deliberately not published on this site.

- Q11 Grover Search: Four States
- Q13 Phase Estimation
- Q12 Grover Under Noise

Teacher notes — Session 6: Near-term methods: Hamiltonians and variational circuits

For instructors | Flipped format

The session in one sentence. Variational algorithms tune circuit angles with a classical loop to minimize a measured average.

The student page for Session 6 is written for self-study: students can read, work the example, run both experiments, and check their own answers without you. These notes are for the class time you do have.

Before class (students, on their own)

- Read Chapter 8 and Chapter 9 and Chapter 10.
- Do the worked example and the first experiment on the student page.
- Bring one written prediction and one question.

If you have 20 minutes

MINUTES	WHAT TO DO
0–3	Ask the exit question below cold, and take a show of hands. Do not give the answer yet.
3–10	Board: the worked example from the student page, with students supplying each step.
10–16	Live demo: run the second experiment. Collect predictions first, then run it.
16–20	Return to the exit question, settle it, and assign the project.

If you have a full 75-minute class

MINUTES	ACTIVITY
0–10	Motivating question and show of hands.
10–30	The three points below, with the mathematics on the board.
30–45	Worked example at the board, students supplying steps.

45–55	Misconceptions: put each “common mistake” to the class as a true/false statement.
55–65	Live demo of both experiments, predictions first.
65–75	Project briefing and exit question.

Three points to make

1. Expectation value from counts: P(0) − P(1) on the board.
2. The hybrid loop: quantum estimates, classical decides.
3. Always state the classical baseline next to the quantum result.

Board example

From counts to an expectation value. The state $Ry(\theta)|0\rangle$ is $\cos(\theta/2)|0\rangle + \sin(\theta/2)|1\rangle$.

1. $P(0) = \cos^2(\theta/2)$ and $P(1) = \sin^2(\theta/2)$.
2. Z gives +1 for outcome 0 and −1 for outcome 1, so $\langle Z\rangle = P(0) - P(1) = \cos^2(\theta/2) - \sin^2(\theta/2) = \cos\theta$.
3. Suppose 1,000 shots give 250 zeros and 750 ones. Then $\langle Z\rangle \approx 0.25 - 0.75 = -0.5$.
4. $\cos\theta = -0.5$ means $\theta = 2\pi/3$, or 120°.

This is how an energy is estimated on real hardware: from counts, not from the state itself.

Misconceptions to raise

- Treating a variational result as exact. It is an estimate with sampling error, and the optimizer may stop in a local minimum.
- Skipping the classical baseline. At teaching sizes the exact answer is always available, so compare against it.
- Calling a small successful demonstration “quantum advantage”.

Discussion prompts

- What would count as honest evidence that a variational method is useful?
- Why might an optimizer fail on hardware even when the circuit is correct?

Exit question

Ask: 1,000 shots give 250 zeros and 750 ones. What is ⟨Z⟩?

Looking for: 0.25 − 0.75 = −0.5.

Assign

One project from this session, following the predict → baseline → one extension → analyze routine. Project solutions are deliberately not published on this site.

- Q16 Ising Trotter Study
- Q18 VQE Energy Estimation
- Q19 QAOA Max-Cut
- Q20 Variational Quantum Classifier

Teacher notes — Session 7: Communication, benchmarking, and your own project

For instructors | Flipped format

The session in one sentence. BB84 detects eavesdropping because measuring in the wrong basis disturbs the state, and honest results need recorded evidence.

The student page for Session 7 is written for self-study: students can read, work the example, run both experiments, and check their own answers without you. These notes are for the class time you do have.

Before class (students, on their own)

- Read Chapter 11 and Chapter 12.
- Do the worked example and the first experiment on the student page.
- Bring one written prediction and one question.

If you have 20 minutes

MINUTES	WHAT TO DO
0–3	Ask the exit question below cold, and take a show of hands. Do not give the answer yet.
3–10	Board: the worked example from the student page, with students supplying each step.
10–16	Live demo: run the second experiment. Collect predictions first, then run it.
16–20	Return to the exit question, settle it, and assign the project.

If you have a full 75-minute class

MINUTES	ACTIVITY
0–10	Motivating question and show of hands.
10–30	The three points below, with the mathematics on the board.
30–45	Worked example at the board, students supplying steps.

45–55	Misconceptions: put each “common mistake” to the class as a true/false statement.
55–65	Live demo of both experiments, predictions first.
65–75	Project briefing and exit question.

Three points to make

1. Two bases, random choices, keep only matching rounds.
2. Why Eve causes 25% errors: two halves multiplied, on the board.
3. What a benchmark number does and does not tell you.

Board example

How much error does an eavesdropper cause? Look only at rounds where Alice and Bob used the same basis.

1. Eve guesses the basis. She is right half the time, and then causes no error.
2. The other half, she measures in the wrong basis and sends on a state in that wrong basis. Bob, measuring in Alice’s basis, now gets a random result and is wrong half the time.
3. Error rate in the sifted key: ½ × ½ = ¼, or 25%.
4. If Alice and Bob compare k bits, Eve escapes detection on each with probability ¾. For k = 20 that is $(¾)^{20} \approx 0.003$, about 0.3%.

Misconceptions to raise

- Thinking the basis announcement leaks the key. Bases are revealed only after the measurements, and the bit values never are.
- Keeping the bits used for the error check. They were announced publicly and must be discarded.
- Treating a simulated protocol as a security proof, or a single benchmark run as a general statement about a device.

Discussion prompts

- What does BB84 assume about the devices that a real attacker might violate?
- What would you need to see before trusting a published hardware benchmark?

Exit question

Ask: In BB84 with an intercept-and-resend eavesdropper, what error rate appears in the sifted key?

Looking for: About 25%: she picks the wrong basis half the time, and then Bob is wrong half the time.

Assign

One project from this session, following the predict → baseline → one extension → analyze routine. Project solutions are deliberately not published on this site.

- Q21 BB84 and Eavesdropping
- Q22 Reproducible Processor Benchmark
- Q24 Reproduce, Extend, Publish