QuIC Seminar 4

The Phase Estimation Algorithm

Contents

The phase estimation algorithm is used to determine the eigenvalues of a unitary operator. Conventionally, it is presented after the quantum Fourier Transform (QFT), a subroutine used in phase estimation. Our presentation will use one result from the QFT, which we will derive in a future seminar. The QFT is interesting in that it can perform a basis transformation analogous to a Fourier transform exponentially faster than the best known classical algorithm. We will discuss this result more in the second semester when we introduce (quantum) computational complexity.

For now we want to provide an introduction to the utility of quantum algorithms, which is why we focus on the phase estimation algorithm first. Computing eigenvalues has an enormous number of applications in an enormous number of fields. As such, the phase estimation algorithm is used in a large number of other quantum algorithms from solving linear systems of equations in linear algebra (search for the HHL algorithm) to minimizing the number of features required in machine learning applications (search for quantum principal component analysis).

4.1 Why "Phase Estimation?"

Let's ponder the name of "phase estimation" as we introduce some background notation. The meaning will shortly become clear.

First, we're concerned with an eigenvalue problem, namely an equation of the form

$$A\boldsymbol{x} = \lambda\boldsymbol{x} \tag{4.1}$$

where $A \in \mathbb{C}^{2^m \times 2^m}$, $\boldsymbol{x} \in \mathbb{C}^{2^m}$, and $\lambda \in \mathbb{C}$. Note that we write the dimension as 2^m for convenience, since m qubits imply a state space of size 2^m as we have previously seen. (We're using m instead of the standard n because we'll introduce another n qubits later in the algorithm.)

In the quantum case, we're only going to be concerned with unitary operators, which we normally write as U. Since these operators satisfy $U^\dagger U = I$, any eigenvalue has magnitude one.

Exercise 32: Let λ be an eigenvalue of a unitary matrix U. Prove that $|\lambda| = 1$.

Solution 5: Since $U|x\rangle = \lambda|x\rangle$, we have $\langle x|\lambda^*\lambda|x\rangle = \langle x|U^\dagger U|x\rangle$, hence $|\lambda|^2\langle x|x\rangle = \langle x|x\rangle$, from which the result follows immediately. (Recall eigenvectors are nonzero by definition.)

Since $|\lambda| = 1$, we can write it without loss of generality as

$$\lambda = e^{2\pi i\phi} \tag{4.2}$$

where $0 \leq \phi \leq 1$ is called the *phase*. Hence the term "phase" in "phase estimation." The term "estimation" comes about not from the fact that quantum computation is probabilistic, but rather in the degree of precision that we are going to compute, or estimate, the phase to. The most general case of phase estimation is when the phase cannot be written exactly using n bits of precision. In this lecture, we'll only cover the case where ϕ can be written exactly using n bits.

Before we going about estimating the phase, let's introduce some useful notation that is common in quantum algorithms. The phase ϕ is going to be between zero and one, so we can write it as a decimal in binary notation as follows:

$$\phi = 0.\phi_1\phi_2\cdots\phi_n \tag{4.3}$$

where each ϕ_i is either zero or one. What does this mean exactly? The same thing we mean when we always write decimals, except here we're using binary notation.

Definition 4.1 (Binary decimal notation.). The expression $\phi = 0.\phi_1\phi_2\cdots\phi_n$ is equivalent to

$$\phi = 0.\phi_1\phi_2\cdots\phi_n \iff \phi = \sum_{k=1}^{n} \phi_k 2^{-k}. \tag{4.4}$$

Example 1: Some numbers as binary decimals.
The number 0.5 in decimal is 0.1 in binary, since $0.1 \equiv (1)\cdot 2^{-1} = 1/2 = 0.5$. Note that 0.1 is the same as 0.100000....

The number 0.75 in decimal is 0.11 in binary, since $0.11 \equiv (1)\cdot 2^{-1} + 1\cdot 2^{-2} = 1/2 + 1/4 = 3/4 = 0.75$.

Exercise 33: What is the value of the infinitely repeating binary decimal 0.1111111...?

(If it needed to be proved, the above exercise proves that $0 \leq 0.\phi_1\phi_2\cdots \leq 1$.)

4.2 Action of Controlling an Operator on its Eigenstate

The key to understanding the phase estimation algorithm is what happens in the following circuit.

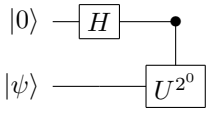

Figure 4.1: Action of a controlled-U^{2^k} operation on the plus state for $k = 0$.

Here we let U be a unitary operator and $|\psi\rangle$ an eigenstate[1] with eigenvalue $\lambda = e^{2\pi i 0.\phi_1}$. We first perform a Hadamard gate on the first qubit to get the state

$$|+\rangle \otimes |\psi\rangle = |0\rangle|\psi\rangle + |1\rangle|\psi\rangle \tag{4.5}$$

[1]Why are we starting with an eigenstate? Doesn't this defeat the whole purpose? After all, we could just do $U|\psi\rangle$ and see what it's effect was. Answer: here, we're just using the eigenstate for pedagogical purposes. As we'll see, the general phase estimation algorithm does not require this.

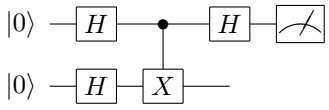

Figure 4.2: Example of using the phase estimation algorithm to compute the eigenvalues of Pauli-X.

where we have intentionally omitted the normalization factor of $1/\sqrt{2}$ for clarity[2]. We then perform a controlled U operation, which we have written as U^{2^0} for reasons that will become clear shortly. The action of this gate is to produce the new state

$$|0\rangle|\psi\rangle + |1\rangle U|\psi\rangle = |0\rangle|\psi\rangle + e^{2\pi i 0.\phi_1}|1\rangle|\psi\rangle \tag{4.6}$$

$$= (|0\rangle + e^{2\pi i 0.\phi_1}|1\rangle) \otimes |\psi\rangle. \tag{4.7}$$

Note what happened: The second qubit register containing $|\psi\rangle$ hasn't changed. We shouldn't expect it to, since $|\psi\rangle$ is an eigenstate of U. Thus, no matter how many times we apply U to this register, nothing happens to $|\psi\rangle$.

That's rather odd though—what's the point of applying U then? The effect was that it *wrote* some information about the eigenvalue into the relative phase of the first qubit. Namely, the entire effect was to map

$$|0\rangle + |1\rangle \mapsto |0\rangle + e^{2\pi i 0.\phi_1}|1\rangle \tag{4.8}$$

How can we read out this information from the quantum state? Consider the effect of applying another Hadamard transformation on the first qubit, which will produce

$$H(|0\rangle + e^{2\pi i 0.\phi_1}|1\rangle) = (1 + e^{2\pi i 0.\phi_1})|0\rangle + (1 - e^{2\pi i 0.\phi_1})|1\rangle. \tag{4.9}$$

where we have again ignored the normalization factor of $1/2$.

Exercise 34: Verify (4.9).

Now, ϕ_1 can only be zero or one. In the case that $\phi_1 = 0$, $e^{2\pi i 0.\phi_1} = 1$, hence the state is exactly $|0\rangle$.

Exercise 35: Prove that the RHS of (4.9) becomes exactly $|1\rangle$ up to global phase. (Recall again that (4.9) omits the normalization factor of $1/2$.)

Thus, we measure with certainty (i.e., not probabilistically) a state that tells us exactly what the phase, and hence the eigenvalue, is.

Example 2: Phase estimation on Pauli-X.

First, prove that the eigenvalues of X are -1 and 1 with eigenvectors $|-\rangle$ and $|+\rangle$ respectively.

Now, let's see how we can do this using the phase estimation algorithm. The first thing we need to do is prepare an eigenstate, which we can do by performing the Hadamard gate on the $|0\rangle$ state in the second register to get the plus state. Then, we control on the first qubit in uniform superposition to implement a controlled-X gate, otherwise known as CNOT. Finally, we implement H again on the top qubit, then measure. The complete circuit for this state is shown in Fig. 4.2.

The complete state after applying both Hadamard gates to both qubits is given by the uniform superposition

$$|+\rangle \otimes |+\rangle = |00\rangle + |01\rangle + |10\rangle + |11\rangle. \tag{4.10}$$

Note that this state remains unchanged under the CNOT. After performing the second Hadamard on the top qubit, the top qubit is in the state $|0\rangle$. Hence, measurement always returns the outcome $\phi_1 = 0$, and so we conclude that the corresponding eigenvalue is $\lambda = e^{2\pi i \cdot 0} = 1$, as we must.

Exercise 36: Repeat the above example but for the $|-\rangle$ eigenstate and verify that measuring the top qubit always yields $\phi_1 = 1$, thus $\lambda = -1$. (*Hint:* How can the circuit in Fig. 4.2 be modified to produce the $|-\rangle$ state?)

4.3 Controlling Higher Powers

The key idea of the phase estimation algorithm is to keep applying the same controlled-U operations, but at successively higher powers of two. The final algorithm introduces n qubits in the top register and implements $C(U^{2^k})$ between the kth qubit and the bottom register. Let's build up to this with one more example where $n = 2$.

The complete circuit for the $n = 2$ case is shown below. (Note that, although we're only writing one wire for the bottom register, this really consists of m qubits. That is, $|\psi\rangle$ is an m-qubit state and U is an m-qubit unitary.)

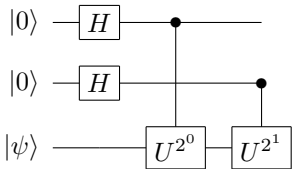

Figure 4.3: Phase estimation for $n = 2$ qubits in the top register.

We've already computed what the first control does to the top register. The wavefunction at this point in the circuit is

$$(|0\rangle + e^{2\pi i 0.\phi_1\phi_2}|1\rangle \otimes (|0\rangle + |1\rangle) \otimes |\psi\rangle. \tag{4.11}$$

Note that the relative phase in the first qubit now has two digits because we assumed that $\phi = 0.\phi_1\phi_2$ consists of two digits.

The same exact operation now happens on the second qubit, except for one key difference: the power of U. In particular, U is now squared. What effect does this have?

Exercise 37: Verify that

$$U^2|\psi\rangle = e^{2\pi i(2\phi)}|\psi\rangle \tag{4.12}$$

by applying the unitary twice to the eigenstate.

Note an interesting thing from (4.12), namely that

$$2\phi = 2 \cdot 0.\phi_1\phi_2 = 2(\phi_1 2^{-1} + \phi_2 2^{-2}) = \phi_1 + 2\phi_2 2^{-1} = \phi_1.\phi_2. \tag{4.13}$$

The effect is that the decimal moves one place to the right. (This should not be too surprising based on common arithmetic in base ten.) What's interesting and convenient, however, is what happens in the exponent:

$$e^{2\pi i(2\phi)} = e^{2\pi i(\phi_1 + 0.\phi_2)} = e^{2\pi i\phi_1} e^{2\pi i 0.\phi_2} = e^{2\pi i 0.\phi_2}. \tag{4.14}$$

The last inequality follows because ϕ_1 is an integer and so the exponential is unity.

Exercise 38: Prove the general case that

$$e^{2\pi i(2^k\phi)} = e^{2\pi i 0.\phi_k\phi_{k+1}\cdots}. \tag{4.15}$$

Thus, we have proven that the state after applying the $C(U^2)$ to the wavefunction (4.11) is

[2] This is a common, though probably not great, convention that many authors follow.

$$(|0\rangle + e^{2\pi i 0.\phi_1\phi_2}|1\rangle \otimes (|0\rangle + e^{2\pi i 0.\phi_2}|1\rangle) \otimes |\psi\rangle. \tag{4.16}$$

How do we read out this information? This time it's not so easy as a simple Hadamard transform. The missing ingredient is what we said we were going to assume was true in the introduction: the quantum Fourier transform. In particular, the key fact that we'll be using is the following:

Definition 4.2 (Quantum Fourier transform action.). The quantum Fourier transform is a unitary change of basis with the following effect:

$$\mathrm{QFT}\,(|\phi_1\rangle|\phi_2\rangle\cdots|\phi_n\rangle) = 2^{-n/2}(|0\rangle + e^{2\pi i 0.\phi_1}|1\rangle) \otimes (|0\rangle + e^{2\pi i 0.\phi_1\phi_2}|1\rangle) \otimes \cdots \otimes (|0\rangle + e^{2\pi i 0.\phi_1\phi_2\cdots\phi_n}|1\rangle). \tag{4.17}$$

The state on the right looks exactly like what we end up with from phase estimation! (Except with the order of the bits reversed, but this is not problem.) So, what we need to apply then to read out the information is the inverse Fourier transform $\mathrm{QFT}^\dagger$. (The QFT is unitary, as every quantum algorithm must be, therefore it has an inverse).

Thus, if we apply $\mathrm{QFT}^\dagger$ to the state in (4.16), we get a product state containing the information we want about the phase:

$$\mathrm{QFT}^\dagger(|0\rangle + e^{2\pi i 0.\phi_1\phi_2}|1\rangle \otimes (|0\rangle + e^{2\pi i 0.\phi_2}|1\rangle) \otimes |\psi\rangle = |\phi_2\rangle \otimes |\phi_1\rangle \otimes |\psi\rangle. \tag{4.18}$$

4.4 General Phase Estimation Algorithm

We now present the general n-qubit case for the phase estimation algorithm. The circuit follows immediately from our work in the previous two sections:

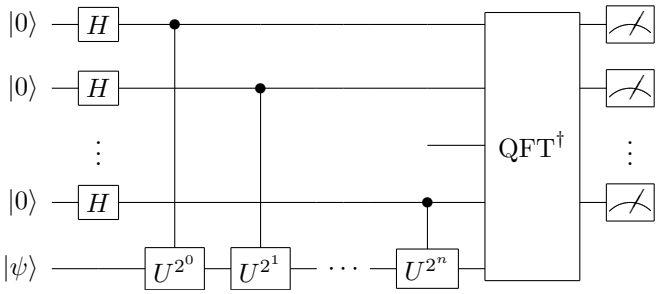

Figure 4.4: Full algorithm for quantum phase estimation. If the phase ϕ can be written exactly in n bits, this circuit computes ϕ exactly. If it requires more bits, this circuit computes a "good" approximation to ϕ.

The general calculation is also an immediate generalization of our previous work. Namely, the action of the circuit, ignoring the $|\psi\rangle$ register that never changes, is

$$|0\rangle^{\otimes n} \longmapsto \frac{1}{2^{n/2}}(|0\rangle + |1\rangle) \otimes (|0\rangle + |1\rangle) \otimes \cdots \otimes (|0\rangle + |1\rangle) \qquad \text{(after the Hadamards)}$$

$$\longmapsto \frac{1}{2^{n/2}}(|0\rangle + e^{2\pi i 0.\phi_1\phi_2\cdots\phi_n}|1\rangle) \otimes (|0\rangle + e^{2\pi i 0.\phi_2\phi_3\cdots\phi_n}|1\rangle) \otimes \cdots \otimes (|0\rangle + e^{2\pi i 0.\phi_n}|1\rangle) \qquad \text{(after the } C(U)\text{'s)}$$

$$\longmapsto |\phi_n\rangle \otimes |\phi_{n-1}\rangle \otimes \cdots \otimes |\phi_1\rangle \qquad \text{(after the } \mathrm{QFT}^\dagger\text{)}.$$

By measuring the final state, we obtain a description of the phase, and hence the eigenvalue, exactly.

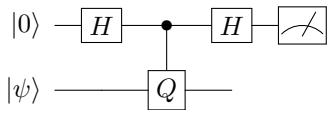

Figure 4.5: Hadamard test algorithm. This circuit can be used to compute $\mathrm{Re}\,[\langle\psi|Q|\psi\rangle]$ as described in Example 4.5.

Exercise 39: Prove that the phase estimation algorithm works even if we don't know an eigenstate $|u\rangle$ of U by inputting an arbitrary state $|\psi\rangle$ and writing it as a linear combination of the eigenstates of U,

$$|\psi\rangle = \sum_u c_u |u\rangle. \tag{4.19}$$

4.5 Closing Remarks

Throughout, we have assumed implicitly that ϕ can be written exactly with n bits. The case where ϕ is irrational (requires infinitely many bits) is a bit more subtle and requires more work. For a complete discussion, we refer the reader to Nielsen and Chuang Section Chapter 5.

Example 3: The Hadamard test.
The Hadamard test is an algorithm for computing the expectation value of an operator that is similar to the phase estimation algorithm. Let $|\psi\rangle$ be a state on m qubits and Q a unitary operator on m qubits, and consider the algorithm described by Fig. 4.5. The action of this circuit is

$$\begin{aligned}
|0\rangle|\psi\rangle &\longmapsto |0\rangle|\psi\rangle + |1\rangle|\psi\rangle && \text{(first Hadamard)}\\
&\longmapsto |0\rangle|\psi\rangle + |1\rangle Q|\psi\rangle && \text{(controlled } Q\text{)}\\
&\longmapsto |0\rangle|\psi\rangle + |1\rangle|\psi\rangle + |0\rangle Q|\psi\rangle - |1\rangle Q|\psi\rangle && \text{(second Hadamard)}\\
&= |0\rangle(I+Q)|\psi\rangle + |1\rangle(I-Q)|\psi\rangle.
\end{aligned}$$

Putting in the omitted factor of $1/2$ (each Hadamard contributes $1/\sqrt{2}$), the probability of measuring zero on the first qubit is

$$p(0) = \frac{1}{4}\langle\psi|(I+Q)^\dagger(I+Q)|\psi\rangle = \frac{1}{4}\langle\psi|I+Q+Q^\dagger+Q^\dagger Q|\psi\rangle. \tag{4.20}$$

Similarly, the probability of measuring one on the first qubit is

$$p(1) = \frac{1}{4}\langle\psi|(I-Q)^\dagger(I-Q)|\psi\rangle = \frac{1}{4}\langle\psi|I-Q-Q^\dagger+Q^\dagger Q|\psi\rangle. \tag{4.21}$$

Subtracting these yields

$$p(0) - p(1) = \frac{1}{2}\langle\psi|Q+Q^\dagger|\psi >= \mathrm{Re}\,[\langle\psi|Q|\psi\rangle], \tag{4.22}$$

since if $z = a + bi$, $z + z^* = (a+bi) + (a-bi) = 2a = 2\mathrm{Re}\,a$.

Exercise 40: Prove that implementing an S gate before the final Hadmard on the first qubit allows one to compute the imaginary part of the expectation value $\mathrm{Im}\,[\langle\psi|Q|\psi\rangle]$. Recall, $S|0\rangle = |0\rangle$ and $S|1\rangle = i|1\rangle$.