Week 8: Phase Estimation Algorithm

COMS 4281 (Fall 2025)

Brief linear algebra review

Eigenvalues

If $M \in \mathbb{C}^{N \times N}$ is a matrix, $|\psi\rangle \in \mathbb{C}^N$ is a vector, and $\lambda \in \mathbb{C}$ satisfying

$$M|\psi\rangle = \lambda|\psi\rangle$$

then we say that $|\psi\rangle$ is an **eigenvector** of M with **eigenvalue** λ.

Eigenvalues of unitary matrices

Fact: The eigenvalues of a unitary matrix U are all of the form $e^{2\pi i\theta}$ for some $\theta \in [0, 2\pi)$.

Eigenvalues of unitary matrices

Fact: The eigenvalues of a unitary matrix U are all of the form $e^{2\pi i\theta}$ for some $\theta \in [0, 2\pi)$.

Proof: Suppose that $U|\psi\rangle = \lambda|\psi\rangle$ for some eigenvector $|\psi\rangle$ and some eigenvalue λ.

Eigenvalues of unitary matrices

Fact: The eigenvalues of a unitary matrix U are all of the form $e^{2\pi i\theta}$ for some $\theta \in [0, 2\pi)$.

Proof: Suppose that $U|\psi\rangle = \lambda|\psi\rangle$ for some eigenvector $|\psi\rangle$ and some eigenvalue λ.

Taking inner products of $\lambda|\psi\rangle$ with itself, on one hand we get

$$(\lambda^* \langle\psi|)(\lambda|\psi\rangle) = |\lambda|^2 \langle\psi|\psi\rangle = |\lambda|^2 .$$

Eigenvalues of unitary matrices

Fact: The eigenvalues of a unitary matrix U are all of the form $e^{2\pi i\theta}$ for some $\theta \in [0, 2\pi)$.

Proof: Suppose that $U|\psi\rangle = \lambda|\psi\rangle$ for some eigenvector $|\psi\rangle$ and some eigenvalue λ.

Taking inner products of $\lambda|\psi\rangle$ with itself, on one hand we get

$$(\lambda^* \langle\psi|)(\lambda|\psi\rangle) = |\lambda|^2 \langle\psi|\psi\rangle = |\lambda|^2 .$$

On the other hand,

$$(\lambda^* \langle\psi|)(\lambda|\psi\rangle) = (\langle\psi| U^\dagger)(U|\psi\rangle) = \langle\psi| U^\dagger U |\psi\rangle = \langle\psi|\psi\rangle = 1$$

because $U^\dagger U = I$ (one of definitions of being unitary).

Eigenvalues of unitary matrices

Fact: The eigenvalues of a unitary matrix U are all of the form $e^{2\pi i\theta}$ for some $0 \leq \theta < 1$.

Proof continued: Therefore

$$|\lambda|^2 = 1$$

and the only such λ's possible are of the form $e^{2\pi i\theta}$.

Some examples

Example: What are the eigenvalues and eigenvectors of

$$Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

Some examples

Example: What are the eigenvalues and eigenvectors of

$$Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

We see that

$$Z |0\rangle = |0\rangle \qquad Z = |1\rangle = - |1\rangle \ .$$

Therefore standard basis are the eigenvectors and ± 1 are corresponding eigenvalues.

Some examples

Example: What are the eigenvalues and eigenvectors of

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} .$$

Some examples

Example: What are the eigenvalues and eigenvectors of

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} .$$

We can compute this by hand, or we can also remember that

$$X|+\rangle = |+\rangle \qquad X|-\rangle = -|-\rangle$$

so the Hadamard basis are the eigenvectors and ± 1 are the corresponding eigenvalues.

Some examples

Example: What are the eigenvalues and eigenvectors of

$$CNOT = \begin{pmatrix} 1 & 0 & & \\ 0 & 1 & & \\ & & 0 & 1 \\ & & 1 & 0 \end{pmatrix}.$$

Some examples

Example: What are the eigenvalues and eigenvectors of

$$CNOT = \begin{pmatrix} 1 & 0 & & \\ 0 & 1 & & \\ & & 0 & 1 \\ & & 1 & 0 \end{pmatrix}.$$

1. $|0,0\rangle$ with eigenvalue 1
2. $|0,1\rangle$ with eigenvalue 1
3. $|1,+\rangle$ with eigenvalue 1
4. $|1,-\rangle$ with eigenvalue -1

Phase Estimation Algorithm

Phase Estimation

Phase Estimation Algorithm (PEA) is one of the most important subroutines in quantum computing.

Phase Estimation

Phase Estimation Algorithm (PEA) is one of the most important subroutines in quantum computing.

Goal of PEA:

- Ability to run controlled versions of U^k for $k = 1, 2, \ldots$.
- An **eigenstate** $|\psi\rangle$ where $U|\psi\rangle = e^{2\pi i\theta}|\psi\rangle$,

estimate θ.

Question: The eigenvalue $e^{2\pi i\theta}$ looks like a global phase... how can you possibly estimate it?

Question: The eigenvalue $e^{2\pi i\theta}$ looks like a global phase... how can you possibly estimate it?

Answer: It becomes a **relative** phase once you run the controlled-U gate in superposition:

$$\begin{aligned} cU|+\rangle|\psi\rangle &= \frac{1}{\sqrt{2}}(|0\rangle|\psi\rangle + |1\rangle U|\psi\rangle) \\ &= \frac{1}{\sqrt{2}}(|0\rangle|\psi\rangle + e^{2\pi i\theta}|1\rangle|\psi\rangle) \\ &= \frac{1}{\sqrt{2}}(|0\rangle + e^{2\pi i\theta}|1\rangle)|\psi\rangle \end{aligned}$$

Warmup towards Phase Estimation

Let U be a unitary with an eigenvector $|\psi\rangle$ whose corresponding eigenvalue is either $+1$ or -1. How to tell which is the case, given one copy of $|\psi\rangle$ and the ability to apply controlled versions of U?

Warmup towards Phase Estimation

Let U be a unitary with an eigenvector $|\psi\rangle$ whose corresponding eigenvalue is either $+1$ or -1. How to tell which is the case, given one copy of $|\psi\rangle$ and the ability to apply controlled versions of U?

A "baby" form of phase estimation:

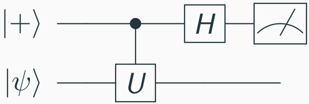

When $|\psi\rangle$ is a $+1$-eigenvector of U, the output is $|0\rangle$. When it is a -1-eigenvector, the output is $|1\rangle$.

Warmup to Phase Estimation

What if the phase were $\exp(2\pi i\theta)$ for some $0 \leq \theta < 1$?

Warmup to Phase Estimation

What if the phase were $\exp(2\pi i\theta)$ for some $0 \leq \theta < 1$? We can analyze the same circuit:

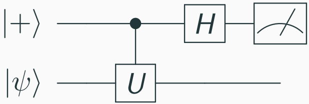

Warmup to Phase Estimation

What if the phase were $\exp(2\pi i\theta)$ for some $0 \leq \theta < 1$? We can analyze the same circuit:

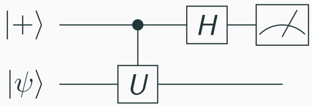

The state of top qubit before measurement is:

$$\left(\frac{1 + e^{2\pi i\theta}}{2}\right)|0\rangle + \left(\frac{1 - e^{2\pi i\theta}}{2}\right)|1\rangle \ .$$

Warmup to Phase Estimation

What if the phase were $\exp(2\pi i\theta)$ for some $0 \le \theta < 1$? We can analyze the same circuit:

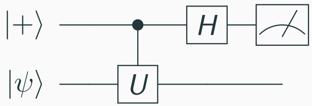

The state of top qubit before measurement is:

$$\left(\frac{1+e^{2\pi i\theta}}{2}\right)|0\rangle + \left(\frac{1-e^{2\pi i\theta}}{2}\right)|1\rangle \ .$$

Measuring this qubit yields

$$\Pr[|0\rangle] = \left|\frac{1+e^{2\pi i\theta}}{2}\right|^2 = \cdots \text{ high school trig } \cdots = \cos^2(\pi\theta) \ .$$

The state $|\psi\rangle$ is undisturbed after running the circuit. So we can repeat it multiple times with fresh ancilla qubits to get an estimate of θ.

By repeating the phase estimation circuit $O(1/\epsilon)$ times, can obtain an estimate of $\cos^2(\pi\theta) \pm \epsilon$. Does this uniquely identify θ?

The state $|\psi\rangle$ is undisturbed after running the circuit. So we can repeat it multiple times with fresh ancilla qubits to get an estimate of θ.

By repeating the phase estimation circuit $O(1/\epsilon)$ times, can obtain an estimate of $\cos^2(\pi\theta) \pm \epsilon$. Does this uniquely identify θ?

No: There is ambiguity between θ and $1 - \theta$:

$$\cos^2(\pi\theta) = \cos^2(\pi(1-\theta)) \ .$$

In other words, this estimation procedure cannot distinguish between whether θ is smaller or bigger than $\frac{1}{2}$.

How to uniquely identify θ?

How to uniquely identify θ?

Suppose, in addition to having a good estimate of $\cos^2(\pi\theta)$, we also knew (a good estimate of)

$$\cos^2(\pi\theta + \frac{\pi}{4}) \ .$$

This is enough to recover θ! (Proof by picture on board).

Thus, after estimating $\cos^2(\pi\theta)$ using the first circuit, we can run a different circuit to get an estimate of $\cos^2(\pi\theta + \frac{\pi}{4})$:

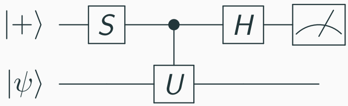

where $S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}$. You can show that the probability of getting $|0\rangle$ after measurement is

$$\Pr[|0\rangle] = \cos^2(\pi\theta + \frac{\pi}{4})$$

as desired.

In general, if get charged \$1 each time we query controlled U, we can obtain $\pm\epsilon$ approximations of θ by spending $O(1/\epsilon)$ dollars. This is fine for many applications, but for Shor's factoring algorithm, we need something much, much cheaper.

In general, if get charged \$1 each time we query controlled U, we can obtain $\pm\epsilon$ approximations of θ by spending $O(1/\epsilon)$ dollars. This is fine for many applications, but for Shor's factoring algorithm, we need something much, much cheaper.

If we have the ability to query controlled U^k for arbitrarily large k for \$1, then we can get $\pm\epsilon$ approximations of θ using $O(\log 1/\epsilon)$ dollars. Exponentially cheaper!

In general, if get charged \$1 each time we query controlled U, we can obtain $\pm\epsilon$ approximations of θ by spending $O(1/\epsilon)$ dollars. This is fine for many applications, but for Shor's factoring algorithm, we need something much, much cheaper.

If we have the ability to query controlled U^k for arbitrarily large k for \$1, then we can get $\pm\epsilon$ approximations of θ using $O(\log 1/\epsilon)$ dollars. Exponentially cheaper!

Main idea: estimate θ bit-by-bit.

Phase Estimation Algorithm

Assume for simplicity that θ can be represented using exactly t bits. In other words the binary representation of θ looks like

$$\theta = 0.\theta_1\theta_2\cdots\theta_t$$

where $\theta_1, \theta_2, \ldots \in \{0, 1\}$. This is equivalent to

$$\theta = \frac{\theta_1}{2} + \frac{\theta_2}{2^2} + \cdots + \frac{\theta_t}{2^t}.$$

Phase Estimation Algorithm

First we will estimate $\theta_t \in \{0, 1\}$. Let $k = 2^{t-1}$. Since $U|\psi\rangle = e^{2\pi i\theta}|\psi\rangle$, we have

$$U^k|\psi\rangle = e^{2\pi ik\theta}|\psi\rangle \ .$$

Phase Estimation Algorithm

First we will estimate $\theta_t \in \{0,1\}$. Let $k = 2^{t-1}$. Since $U|\psi\rangle = e^{2\pi i\theta}|\psi\rangle$, we have

$$U^k|\psi\rangle = e^{2\pi i k\theta}|\psi\rangle \ .$$

But notice that

$$k\theta = \frac{k\theta_1}{2} + \frac{k\theta_2}{2^2} + \cdots + \frac{k\theta_t}{2^t} = \underbrace{2^{t-2}\theta_1 + \cdots + \theta_{t-1}}_{\text{integer}} + \frac{\theta_t}{2} \ .$$

Phase Estimation Algorithm

First we will estimate $\theta_t \in \{0, 1\}$. Let $k = 2^{t-1}$. Since $U|\psi\rangle = e^{2\pi i\theta}|\psi\rangle$, we have

$$U^k|\psi\rangle = e^{2\pi i k\theta}|\psi\rangle \ .$$

But notice that

$$k\theta = \frac{k\theta_1}{2} + \frac{k\theta_2}{2^2} + \cdots + \frac{k\theta_t}{2^t} = \underbrace{2^{t-2}\theta_1 + \cdots + \theta_{t-1}}_{\text{integer}} + \frac{\theta_t}{2} \ .$$

Therefore

$$e^{2\pi i k\theta} = e^{2\pi i \frac{\theta_t}{2}} \in \{+1, -1\} \ .$$

If we run this circuit

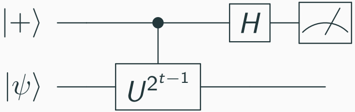

the final qubit will be $|\theta_t\rangle$. We have learned one bit about θ!

If we run this circuit

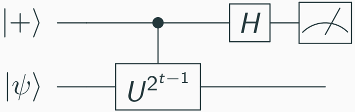

the final qubit will be $|\theta_t\rangle$. We have learned one bit about θ!

Consider the unitary

$$V = e^{-2\pi i \frac{\theta_t}{2^t}} U \ .$$

which has eigenvector

$$V |\psi\rangle = e^{-2\pi i \frac{\theta_t}{2^t}} U |\psi\rangle = e^{2\pi i (\theta - \frac{\theta_t}{2^t})} |\psi\rangle \ .$$

If we run this circuit

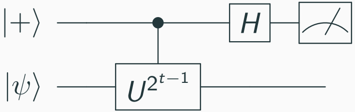

the final qubit will be $|\theta_t\rangle$. We have learned one bit about θ!

Consider the unitary

$$V = e^{-2\pi i \frac{\theta_t}{2^t}} U \ .$$

which has eigenvector

$$V |\psi\rangle = e^{-2\pi i \frac{\theta_t}{2^t}} U |\psi\rangle = e^{2\pi i (\theta - \frac{\theta_t}{2^t})} |\psi\rangle \ .$$

Notice that

$$\theta - \frac{\theta_t}{2^t} = \frac{\theta_1}{2} + \frac{\theta_2}{4} + \cdots + \frac{\theta_{t-1}}{2^{t-1}} \ .$$

We can try to learn θ_{t-1} next by doing phase estimation on

$$V^{2^{t-2}} = e^{-2\pi i \frac{\theta_t}{4}} U^{2^{t-2}}.$$

using the following circuit:

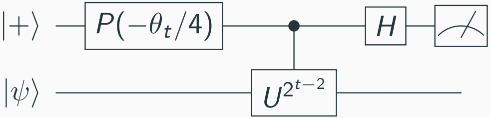

where

$$P(\alpha) = \begin{pmatrix} 1 & 0 \\ 0 & e^{2\pi i \alpha} \end{pmatrix} .$$

We can continue in this manner until we learn all the bits of θ.

The number of iterations is t, which is the number of bits of precision of θ.

The number of iterations is t, which is the number of bits of precision of θ.

Since t bits of precision translates to $\pm 2^{-t}$ error, this means that to get $\pm\epsilon$ error we have $O(\log 1/\epsilon)$ iterations.

Phase Estimation Algorithm Analysis

Question: What if the phase θ cannot be exactly expressed as t bits?

Phase Estimation Algorithm Analysis

Question: What if the phase θ cannot be exactly expressed as t bits?

Answer: If we use $t + k$ ancilla qubits, and measure only the first t ancilla qubits, we will get the best t-bit approximation $\widetilde{\theta}$ of θ with probability $1 - 2^{-k}$.

Phase Estimation Algorithm Analysis

Question: What happens if $|\psi\rangle$ is not an eigenvector of U?

Phase Estimation Algorithm Analysis

Question: What happens if $|\psi\rangle$ is not an eigenvector of U?

Answer: The set $\{|\phi_j\rangle\}$ of eigenvectors of U forms a basis for $\mathbb{C}^{2^n}$ (if U is n-qubit unitary). We can write $|\psi\rangle$ as

$$|\psi\rangle = \sum_j \alpha_j |\phi_j\rangle$$

for some coefficients α_j.

Running a “coherent version” of Phase Estimation on $|\psi\rangle$ with ancilla qubits $|0\cdots 0\rangle$ yields a state that is close to

$$\approx \sum_j \alpha_j |\phi_j\rangle \otimes |\widetilde{\theta}_j\rangle$$

where $\widetilde{\theta}_j$ is an approximation of the eigenphase θ_j, i.e. $U|\phi_j\rangle = e^{2\pi i\theta_j}|\phi_j\rangle$.

Running a "coherent version" of Phase Estimation on $|\psi\rangle$ with ancilla qubits $|0\cdots 0\rangle$ yields a state that is close to

$$\approx \sum_j \alpha_j |\phi_j\rangle \otimes |\widetilde{\theta}_j\rangle$$

where $\widetilde{\theta}_j$ is an approximation of the eigenphase θ_j, i.e. $U|\phi_j\rangle = e^{2\pi i \theta_j}|\phi_j\rangle$.

Measuring the last register yields $\widetilde{\theta}_j$ with probability $|\alpha_j|^2$.

Next time

RSA, Order Finding, Shor's algorithm