

Quantum Computing (CST Part II)

Lecture 11: Application 2 of QFT / QPE: Quantum Chemistry

Feynman's 1982 conjecture, that quantum computers can be programmed to simulate any local quantum system, is shown to be correct.

Seth Lloyd

Quantum simulation

Suppose we want to know the state, $|\psi_t\rangle$ of a quantum system at a time, t in the future, given its Hamiltonian, \mathbf{H} and its current state, $|\psi_0\rangle$, then we must solve the Schrödinger equation:

$$|\psi_t\rangle = e^{-i\mathbf{H}t} |\psi_0\rangle$$

Note that this is a *time-invariant* Hamiltonian, and rather than explicitly stating the normalisation by Planck's constant, we have instead implicitly incorporated it into the Hamiltonian itself.

For the purposes of this lecture, it suffices to take the Taylor series as the definition of an exponentiated matrix:

$$e^A = \sum_{n=0}^{\infty} \frac{A^n}{n!}$$

where A^0 is the identity.

Quantum simulation (continued)

One helpful property of many real quantum systems of interest is that they can be decomposed:

$$H = \sum_{k=1}^K H_k$$

where K is sufficiently small, and the physical nature of the system is such that each H_k can be exponentiated (this is because the physical system is dominated by “local” few-body interactions). So we have that:

$$|\psi_t\rangle = e^{-iHt} |\psi_0\rangle = e^{-i\sum_k H_k t} |\psi_0\rangle$$

However, notice we are dealing with *matrix exponentiation*, and in general:

$$e^{-i\sum_k H_k t} \neq \prod_k e^{-iH_k t}$$

so it appears that we cannot directly use the fact that each H_k can individually be exponentiated efficiently.

Trotterisation

However, H_k can still be useful for simulating the quantum system, because of the **Trotter formula**, which is at the heart of quantum simulation:

$$\lim_{n \rightarrow \infty} \left(e^{-iH_1 t/n} e^{-iH_2 t/n} \right)^n = e^{-i(H_1 + H_2)t}$$

To see this, consider that, by definition:

$$\begin{aligned} e^{-iH_1 t/n} &= I - \frac{1}{n} iH_1 t + \mathcal{O}\left(\frac{1}{n^2}\right) \\ \implies e^{-iH_1 t/n} e^{-iH_2 t/n} &= I - \frac{1}{n} i(H_1 + H_2)t + \mathcal{O}\left(\frac{1}{n^2}\right) \\ \implies \left(e^{-iH_1 t/n} e^{-iH_2 t/n} \right)^n &= I + \sum_{l=1}^n \binom{n}{l} \frac{1}{n^l} (-i(H_1 + H_2)t)^l + \mathcal{O}\left(\frac{1}{n}\right) \end{aligned}$$

Proof of the Trotter formula (continued)

Noticing that:

$$\binom{n}{l} \frac{1}{n^l} = \frac{n!}{l!(n-l)!n^l} = \frac{n(n-1)(n-2)\cdots(n-l+1)}{n^l} \frac{1}{l!} = \frac{1}{l!} \left(1 + \mathcal{O}\left(\frac{1}{n}\right)\right)$$

We get that:

$$\begin{aligned} \lim_{n \rightarrow \infty} \left(e^{-iH_1 t/n} e^{-iH_2 t/n} \right)^n &= I + \lim_{n \rightarrow \infty} \sum_{l=1}^n \frac{(-i(H_1 + H_2)t)^l}{l!} \left(1 + \mathcal{O}\left(\frac{1}{n}\right)\right) \\ &= \lim_{n \rightarrow \infty} \sum_{l=0}^n \frac{(-i(H_1 + H_2)t)^l}{l!} \\ &= e^{-i(H_1 + H_2)t} \end{aligned}$$

Following analysis similar to that above, with a small finite time Δt , we get that:

$$e^{-i(H_1 + H_2)\Delta t} = e^{-iH_1 \Delta t} e^{-iH_2 \Delta t} + \mathcal{O}(\Delta t^2)$$

Therefore, if we divide up the duration of the evolution into sufficiently short intervals, we can accurately approximate the overall evolution by evolving each H_k in turn.

Quantum simulation algorithm

For simulating $|\tilde{\psi}_t\rangle \approx |\psi_t\rangle = e^{-iHt} |\psi_0\rangle$:

1. Initialise $|\tilde{\psi}_0\rangle = |\psi_0\rangle$; $j = 0$
2. $|\tilde{\psi}_{j+1}\rangle \leftarrow U_{\Delta t} |\tilde{\psi}_j\rangle$
3. $j \leftarrow j + 1$; if $j\Delta t < t$ goto step 2
4. Output $|\tilde{\psi}_t\rangle = |\tilde{\psi}_j\rangle$

Where $H = \sum_{k=1}^K H_k$, and:

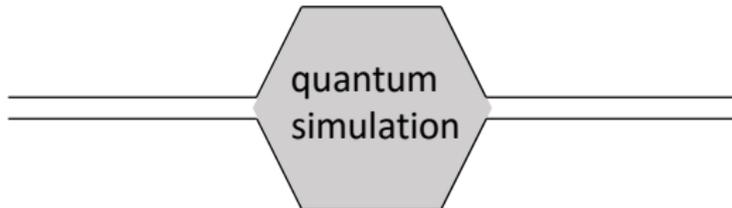
$$U_{\Delta t} = e^{-iH_1\Delta t} e^{-iH_2\Delta t} \dots e^{-iH_K\Delta t}$$

for Δt chosen to be suitably small to achieve the overall desired accuracy.

The previous assertion that each “ H_k can be exponentiated efficiently” can be taken to mean that each unitary $U_k = e^{-iH_k\Delta t}$ can be implemented with only a polynomial (in the number of qubits) number of gates.

Quantum simulation: a solution in search of a problem?

As Feynman asserted, simulation of quantum systems is classically intractable in a fundamental way: because of entanglement it may take an exponential amount of classical memory to even represent the state. So the significance of the discovery of efficient quantum simulation on quantum computers should not be understated. However, from an algorithmic point of view, this really only corresponds to the middle of the double-necked bottle:



In order to find useful application, we need to use quantum simulation to solve a problem with a compact input and output.

Quantum chemistry

Quantum chemistry provides one such application of quantum simulation, the general set-up being:

- A system Hamiltonian, H_s is encoded as a qubit Hamiltonian H_q .
- An important property in computational chemistry is the *ground state energy*, E_0 :

$$E_0 = \min_{|\psi\rangle} \langle \psi | H_q | \psi \rangle$$

which denotes the *measurement of an observable*.

- As E_0 is simply a number, it is clearly a “compact” output. We will see how quantum simulation and quantum phase estimation together enable us to find ground states.

Projective measurements and observables

In quantum mechanics projective measurements are usually used in the context of measuring observables. An observable is a Hermitian operator, say H , which has spectral decomposition:

$$H = \sum_i \lambda_i |e_i\rangle \langle e_i|,$$

where λ_i is the eigenvalue corresponding to eigenvector $|e_i\rangle$ (we let λ_i be ordered from smallest to largest) and we let $P_i = |e_i\rangle \langle e_i|$ denote the projector onto the i^{th} eigenspace. As projectors are always such that $P_i^\dagger P_i = P_i$, we get that the probability of measuring the i^{th} eigenvector as (when measuring some arbitrary state $|\psi\rangle$):

$$p(i) = \langle \psi | P_i | \psi \rangle.$$

Projective measurements and observables (cont.)

If we get the i^{th} eigenvector from the measurement, then we **interpret the measurement outcome as obtaining a numerical value equal to the i^{th} eigenvalue**. This leads to the notion of measuring the expectation of an observable on some state $|\psi\rangle$:

$$\begin{aligned}\mathbb{E}_{|\psi\rangle}(\mathbf{H}) &= \sum_i \lambda_i p(i) \\ &= \sum_i \lambda_i \langle \psi | P_i | \psi \rangle \\ &= \langle \psi | \left(\sum_i \lambda_i P_i \right) | \psi \rangle \\ &= \langle \psi | \left(\sum_i \lambda_i |e_i\rangle \langle e_i| \right) | \psi \rangle \\ &= \langle \psi | \mathbf{H} | \psi \rangle\end{aligned}$$

Ground state energy

Recall from Lecture 2 that the eigenvectors of a Hermitian operator form an orthonormal basis. Therefore we can write out an arbitrary state $|\psi\rangle$ as a superposition over this basis.

$$|\psi\rangle = \sum_i a_i |e_i\rangle$$

Thus we can see that:

$$\langle\psi|\mathbf{H}|\psi\rangle = \left(\sum_i a_i^* \langle e_i|\right) \left(\sum_i \lambda_i |e_i\rangle \langle e_i|\right) \left(\sum_i a_i |e_i\rangle\right) = \sum_i |a_i|^2 \lambda_i$$

Noting that $\sum_i |a_i|^2 = 1$, we can see that this is minimised when $a_0 = 1$ (recalling that we have ordered λ_i from smallest to largest).

So the ground state is obtained when $|\psi\rangle = |e_0\rangle$, and the ground state energy itself is equal to λ_0 – the smallest eigenvalue of \mathbf{H} , that is

$$E_0 = \lambda_0$$

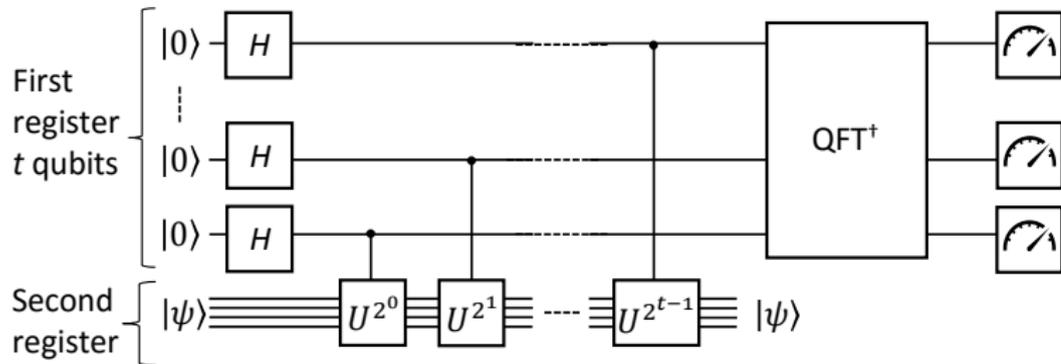
Using quantum phase estimation to find the ground state

Finding the ground state energy is a computationally hard problem in general. Quantum phase estimation can be used to find the ground state energy of a Hamiltonian H . First, we define the unitary:

$$U = e^{2\pi i H}$$

Which we use in the QPE algorithm, for now we treat the contents of the second register as an arbitrary state expressed as a superposition over the basis formed by the eigenstates of H :

$$|\psi\rangle = \sum_i a_i |e_i\rangle$$



Using QPE to find the ground state (cont.)

It is important to note that the eigenvectors of H are also the eigenvectors of $U = e^{2\pi i H}$:

$$\begin{aligned}U |e_i\rangle &= e^{2\pi i H} |e_i\rangle \\&= \sum_{k=0}^{\infty} \frac{1}{k!} (2\pi i)^k H^k |e_i\rangle \\&= \sum_{k=0}^{\infty} \frac{1}{k!} (2\pi i)^k \lambda_i^k |e_i\rangle \\&= e^{2\pi i \lambda_i} |e_i\rangle\end{aligned}$$

i.e., the eigenvalue of U corresponding to $|e_i\rangle$ is $e^{2\pi i \lambda_i}$. So the state *before* the inverse QFT is:

$$\frac{1}{\sqrt{2^t}} \sum_i a_i \sum_{j=0}^{2^t-1} e^{2\pi i \lambda_i j} |j\rangle |e_i\rangle$$

And thus the state *after* the inverse QFT is:

$$\sum_i a_i |\tilde{\lambda}_i\rangle |e_i\rangle$$

Notes on using QPE to find the ground state

- $\tilde{\lambda}_i$ is a t -bit approximation of the eigenvalue λ_i .
- There is a $|a_0|^2$ probability of collapsing into the desired state $|e_0\rangle$ (and so obtaining an estimate of $\lambda_0 = E_0$). This in turn tells us that **we should prepare the initial state such that its sufficiently dominated by $|e_0\rangle$** . There are various ways to do this, and later in the course we will study one of them: adiabatic state preparation.
- Noting that $U^{2^j} = (e^{2\pi i H})^{2^j} = e^{2\pi i H 2^j}$ we can see that each controlled unitary is a time evolution of U , and thus we can use the quantum simulation algorithm to achieve these.
- Overall we have a quantum algorithm that is asymptotically efficient: that is, it only requires a circuit of depth which is polynomial in the number of qubits. However, in practise the circuit depth is prohibitive for near-term quantum computers.

Near-term quantum chemistry

Performing QPE for quantum chemistry requires a *full-scale fault-tolerant quantum computer*, however even in the absence of such a device the principle that quantum simulation on classical computers is intractable still holds. Therefore much current research concerns **hybrid quantum-classical algorithms** which aim to execute only **shallow-depth quantum circuits**, in which an **unmanageable amount of error** is not expected to occur.

The most famous and promising of these hybrid algorithms is the **variational quantum eigensolver**.

The variational quantum eigensolver

VQE addresses the problem of ground state energy evaluation directly, and relies on the Rayleigh-Ritz variational principle:

$$\langle \psi(\theta) | \mathbf{H} | \psi(\theta) \rangle \geq E_0$$

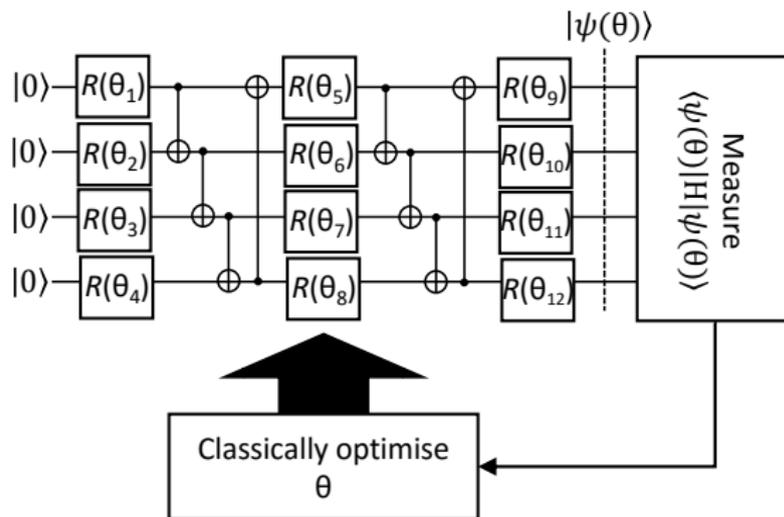
where $|\psi(\theta)\rangle$ is a quantum state parameterised by θ . This implies we can find the ground state energy by finding the value of parameters that minimise $\langle \psi(\theta) | \mathbf{H} | \psi(\theta) \rangle$.

VQE then iterates the following:

1. Run a shallow-depth quantum circuit $U(\theta) : |0\rangle \rightarrow |\psi(\theta)\rangle$ to prepare $|\psi(\theta)\rangle$.
2. Projectively measure to get $E(\theta) = \langle \psi(\theta) | \mathbf{H} | \psi(\theta) \rangle$.
3. Perform classical optimisation to update the parameter values θ .

After sufficiently many iterations VQE converges on $E(\theta) \approx E_0$.

An illustration of VQE



Here we can see that the quantum circuit takes the form of a **parameterised quantum circuit** (PQC), or **variational quantum circuit**. There are many forms of PQC, in this sketch we show one where entangling **CNOT** gates are interspersed with parameterised rotation gates. PQCs are sometimes described as quantum analogues of neural networks, and are also used in **quantum machine learning**.

Summary

In this lecture we have covered:

- Quantum simulation, including Trotterisation.
- Quantum phase estimation for ground state energy estimation in quantum chemistry.
- The variational quantum eigensolver: a near-term quantum chemistry algorithm.