CS 599 P1: Introduction to Quantum Computation — Boston University, Fall 2025
Instructor: Alexander Poremba — **Scribe:** Declan Young

LECTURE #16: VARIATIONAL QUANTUM ALGORITHMS & QUANTUM CHEMISTRY

1 Introduction

The goal of this lecture is to understand how *variational quantum algorithms* work, and how they allow one to solve problems in *quantum chemistry*. We begin by recalling what a Hamiltonian is and why ground-state problems are important. Variational quantum algorithms, and in particular the *Variational Quantum Eigensolver* (VQE), provide heuristic but powerful techniques to approximate low-energy states of Hamiltonians arising in physics, materials science, and chemistry.

- **Hamiltonian $H \in \mathbb{C}^{2^n \times 2^n}$:** A Hermitian operator that represents the total energy of an n-qubit quantum system. Because H is Hermitian, the spectral theorem guarantees that:

$$H = \sum_{i=0}^{2^n-1} E_i \, |\psi_i\rangle\langle\psi_i|,$$

where $\{|\psi_i\rangle\}$ is an orthonormal basis of eigenstates and the $E_i \in \mathbb{R}$ are the corresponding energy eigenvalues. A quantum state always has an associated *average energy*:

$$\langle\psi|H|\psi\rangle,$$

and the *ground state energy* is

$$E_{\min} = \min_i E_i.$$

Computing (or approximating) $E_{\min}$ is a central task in quantum chemistry.

Example: Ising Model Hamiltonian.

The one-dimensional Ising Hamiltonian with nearest-neighbor interactions and an external longitudinal field is

$$H = \sum_{i=1}^{n-1} Z_i Z_{i+1} + \mu \sum_{i=1}^{n} Z_i.$$

Here Z_i is the Pauli-Z operator on qubit i, and μ controls the strength of the magnetic field. The first term encodes pairwise spin coupling, and the second accounts for local fields on each spin.

2 Variational Methods

How can we find or approximate the ground state of a k-local Hamiltonian? A Hamiltonian H acting on n qubits is represented by a $2^n \times 2^n$ Hermitian matrix. Although the eigenvalues of H determine all physically relevant energy levels, computing them directly by classical diagonalization is generally impossible for moderately large n: the matrix size grows exponentially with the number of qubits. For example, even for $n = 50$ qubits, H would be a $2^{50} \times 2^{50}$ matrix, far beyond classical storage capabilities.

This exponential scaling is a fundamental obstacle in classical simulations of quantum many-body physics and quantum chemistry. As a result, we seek alternative methods for identifying low-energy states without explicitly diagonalizing H.

Variational Principle. Fortunately, quantum mechanics provides a remarkably useful tool for approximating ground states: the *variational principle*. It states that for any normalized quantum state $|\psi\rangle$,

$$\langle\psi|H|\psi\rangle \geq E_{\min},$$

where

$$E_{\min} = \min_{\||\psi\|=1} \langle\psi|H|\psi\rangle$$

is the true ground-state energy of the system. In other words, the expected energy of any state is always an upper bound on the ground-state energy. Therefore, if we can search over a family of quantum states whose energies we can compute, minimizing their energy expectation values will drive us closer to the true ground state.

This principle forms the theoretical foundation of classical variational methods (such as the Ritz method in quantum chemistry) and their quantum counterparts, such as the Variational Quantum Eigensolver (VQE).

Ground-State Energy Problem. Given a Hamiltonian H, two computational challenges arise:

I. **State preparation:** Find or approximate a state $|\psi\rangle$ such that

$$\langle\psi|H|\psi\rangle \approx E_{\min}.$$

II. **Energy estimation:** Once such a candidate state is prepared, estimate the energy

$$\langle\psi|H|\psi\rangle$$

efficiently on a quantum device.

Both challenges are nontrivial. The search problem (I) involves navigating an exponentially large Hilbert space. The estimation problem (II) requires extracting physical information through measurements, which are inherently probabilistic.

3 Searching for a Low Energy State (I)

The problem of finding ground states of general Hamiltonians is one of the most difficult in computational physics: it is known to be QMA-complete. In worst-case settings, even quantum computers cannot efficiently solve it.

However, the Hamiltonians arising in real physical systems—such as molecules, materials, condensed matter systems, and lattice models—are far from worst-case. They often exhibit:

- **Locality**: H decomposes into terms acting on only a few qubits (or orbitals).
- **Symmetries**: conservation of particle number, spin, parity, etc.
- **Limited entanglement**: many physically relevant states obey area laws or have efficiently parameterizable structure.

Because of this physical structure, it is reasonable to search for approximate ground states not in the entire Hilbert space, but in some carefully chosen, structured submanifold of states.

Ansatz. A common strategy is to restrict attention to a family of quantum states obtained by applying a parameterized quantum circuit to a fixed reference state. Such a family is called a *variational ansatz*. Formally, we define

$$|\psi(\boldsymbol{\theta})\rangle,$$

where $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_m) \in \mathbb{R}^m$ is a vector of tunable parameters.

The ansatz serves two essential purposes:

1. It greatly reduces the search space from dimension 2^n to a manageable m-dimensional parameter space.
2. It can be designed to incorporate physical intuition: for instance, respecting the symmetries of H or mimicking known approximate wavefunctions (e.g., coupled-cluster states in chemistry).

In practice, different choices of ansatz lead to different algorithmic properties. Hardware-efficient ansätze use shallow circuits well-suited for near-term devices, while chemically inspired ansätze (e.g., unitary coupled cluster) capture the structure of molecular wavefunctions at the cost of deeper circuits. Choosing an appropriate ansatz is therefore a crucial step in any variational quantum algorithm.

Example: Parameterized Quantum State ($m = 2$)

Consider the two-qubit ansatz with two parameters θ_1 and θ_2:

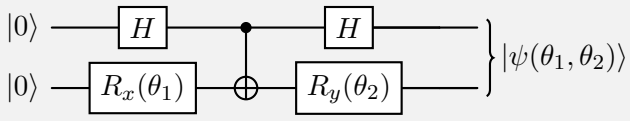

This circuit prepares a tunable entangled state $|\psi(\theta_1, \theta_2)\rangle$ that can be optimized to approximate the ground state.

Relation to Quantum Chemistry. In molecular systems, classical methods such as Coupled Cluster Theory generate structured ansatz families. The same ideas are adapted in quantum chemistry for VQE: one defines an expressive, symmetry-preserving ansatz and optimizes

$$\min_{\boldsymbol{\theta}} \langle \psi(\boldsymbol{\theta}) | H | \psi(\boldsymbol{\theta}) \rangle.$$

4 Calculating the Energy (II)

Suppose we have a candidate state $|\psi(\boldsymbol{\theta})\rangle$. To evaluate its energy, we use

$$H = \sum_{i=1}^{M} H_i, \qquad \text{where each } H_i \text{ is } k\text{-local.}$$

While Quantum Phase Estimation (QPE) is a general method for measuring energies, it is expensive for near-term quantum devices. Instead, VQE relies on Hamiltonians decomposed into Pauli operators.

Pauli Hamiltonians. A Pauli Hamiltonian is of the form $H = \sum_{i=1}^{M} c_i P_i$, where $P_i \in \{I, X, Y, Z\}^{\otimes n}$ is an n-qubit Pauli operator, $c_i \in \mathbb{R}$, M grows polynomially with n for typical molecular Hamiltonians. Then,

$$\langle\psi(\boldsymbol{\theta})|H|\psi(\boldsymbol{\theta})\rangle = \sum_{i=1}^{M} c_i \langle\psi(\boldsymbol{\theta})|P_i|\psi(\boldsymbol{\theta})\rangle.$$

To measure each expectation value, we express P_i in terms of measurement in the computational basis.

Pauli	**Eigendecomposition**	**Relevant Basis Change**
I	$\|0\rangle\langle 0\| + \|1\rangle\langle 1\|$	Identity
X	$\|+\rangle\langle +\| - \|-\rangle\langle -\|$	H
Y	$\|y_+\rangle\langle y_+\| - \|y_-\rangle\langle y_-\|$	$S^\dagger H$
Z	$\|0\rangle\langle 0\| - \|1\rangle\langle 1\|$	Identity

Table 1: Pauli matrices and basis transformations for measurement.

Example I: Measuring the expectation of a single-qubit Pauli

Suppose we wish to estimate the expectation value

$$\langle\psi|Y|\psi\rangle.$$

If we apply the unitary $S^\dagger H$ *before measuring*, then a standard Z-basis measurement will behave *as if* we measured Y on the original state. Hence, we can proceed as follows:

1. Given $|\psi\rangle$, apply the basis-change circuit $S^\dagger H$ to the qubit.
2. Measure the qubit in the computational (Z) basis. This yields an outcome $b \in \{0, 1\}$.
3. Convert the bit into an eigenvalue of Z, i.e. $0 \mapsto +1$ and $1 \mapsto -1$.
4. Repeat this process N times and take the average:

$$\frac{1}{N}\sum_{s=1}^{N}(-1)^{b^{(s)}} \approx \langle\psi|Y|\psi\rangle.$$

Example II: Measuring the expectation of a two-qubit Pauli product $(X \otimes Y)$

Now consider the expectation value of a *two-qubit* observable:

$$\langle\psi|X \otimes Y|\psi\rangle.$$

Because we can naturally only measure each qubit in the Z basis, we will once again use a basis change to convert X and Y into Z. The required changes of basis are:

$$X = HZH, \qquad Y = S^\dagger H\, Z\, (S^\dagger H)^\dagger.$$

This tells us exactly which gates to apply before measurement. Therefore, it sufices to do the following:

1. Given the two-qubit state $|\psi\rangle$, apply H to the *first* qubit (turns X into Z).
2. Apply $S^\dagger$ then H to the *second* qubit (turns Y into Z).
3. Measure both qubits in the Z basis. This yields a pair of classical bits:
$$(b_1, b_2) \in \{0, 1\}^2.$$
4. Convert each bit into a Z eigenvalue, i.e. $b_i = 0 \mapsto +1, b_i = 1 \mapsto -1$.
5. Multiply the two eigenvalues:
$$(-1)^{b_1} \cdot (-1)^{b_2}.$$
This product gives the eigenvalue of $Z \otimes Z$ for this measurement shot.
6. Repeat N times and average:
$$\frac{1}{N}\sum_{s=1}^{N}(-1)^{b_1^{(s)}}(-1)^{b_2^{(s)}} \approx \langle\psi|X \otimes Y|\psi\rangle.$$

5 Variational Quantum Eigensolver (VQE)

The Variational Quantum Eigensolver (VQE) is one of the most prominent quantum algorithms for near-term quantum devices. Unlike algorithms such as Quantum Phase Estimation, which require deep coherent quantum circuits, VQE is specifically designed to operate within the depth and noise limitations of current hardware. It accomplishes this by combining shallow parameterized quantum circuits with classical optimization techniques.

The Basic Idea

VQE leverages the *variational principle*, which states that for any normalized trial state $|\psi\rangle$,

$$\langle\psi|H|\psi\rangle \geq E_{\min},$$

where $E_{\min}$ is the ground-state energy. Thus, if we can produce a family of states

$$|\psi(\boldsymbol{\theta})\rangle,$$

parameterized by a vector $\boldsymbol{\theta} \in \mathbb{R}^m$, we can attempt to approximate the ground state by minimizing the corresponding expected energy.

Objective Function. VQE formulates ground-state search as the optimization problem

$$\min_{\boldsymbol{\theta}} \langle \psi(\boldsymbol{\theta}) | H | \psi(\boldsymbol{\theta}) \rangle.$$

This expectation value serves as a *loss function* for the classical optimizer.

The Hybrid Quantum–Classical Loop

Unlike purely quantum algorithms, VQE is *hybrid*. The quantum computer is responsible only for preparing quantum states and estimating energy expectation values, while all optimization is performed on a classical machine. The two systems interact repeatedly in a feedback loop.

Algorithm: Variational Quantum Eigensolver (VQE)

1. **Classical optimizer proposes parameters.** The classical computer chooses a vector $\boldsymbol{\theta}^{(t)}$ of parameters at iteration t. This choice may be random (initially) or based on optimization rules such as Nelder–Mead, COBYLA, CMA-ES, or gradient-based methods.

2. **Quantum computer prepares a state.** A parameterized quantum circuit (the *ansatz*) prepares the state
$$|\psi(\boldsymbol{\theta}^{(t)})\rangle.$$

3. **Quantum computer evaluates the energy.** The Hamiltonian is decomposed into measurable terms:
$$H = \sum_{i=1}^{M} c_i P_i,$$
where each P_i is a tensor product of Pauli operators. The quantum device measures the expectation of each term
$$\langle \psi(\boldsymbol{\theta}^{(t)}) | P_i | \psi(\boldsymbol{\theta}^{(t)}) \rangle$$
using basis changes and repeated measurements. The energy estimate is assembled as
$$E(\boldsymbol{\theta}^{(t)}) = \sum_{i=1}^{M} c_i \, \langle \psi(\boldsymbol{\theta}^{(t)}) | P_i | \psi(\boldsymbol{\theta}^{(t)}) \rangle.$$

4. **Classical optimizer updates parameters.** The energy value $E(\boldsymbol{\theta}^{(t)})$ is returned to the classical machine, which uses it to determine a new parameter vector $\boldsymbol{\theta}^{(t+1)}$.

5. **Repeat until convergence.** The loop continues until the energy converges to within a chosen tolerance, or until a maximum iteration count is reached.

Why Hybridization Helps

A key insight is that quantum computers are well-suited for representing and manipulating high-dimensional quantum states, while classical computers excel at numerical optimization. VQE allows each side to do what it does best:

- **Quantum side:** Efficiently prepares states in a 2^n-dimensional Hilbert space and measures energy terms.
- **Classical side:** Performs nonlinear optimization in an m-dimensional parameter space (where m is typically polynomial in n).

This division of labor avoids the need for deep circuits or full quantum coherence over long periods.

Characteristics of the VQE Workflow

Shallow Circuits. Ansätze are chosen to be implementable on current devices (few layers of gates, limited entanglement).

Stochastic Objective Evaluation. Each energy estimate is based on finite sampling, so

$$E(\boldsymbol{\theta}) \approx \widehat{E}(\boldsymbol{\theta})$$

with statistical noise of order $1/\sqrt{N}$ where N is the number of measurement shots.

Local Minima and Optimization Challenges. Because the landscape may contain many local minima or "barren plateaus," the performance depends strongly on the choice of ansatz and optimization method.

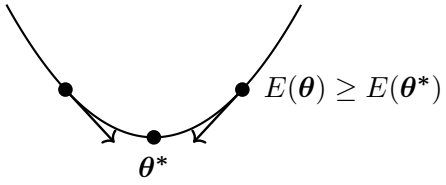

Figure 1: Finding the minimum value θ^* of the energy loss function $E(\theta)$.

6 VQE for the Hydrogen Molecule H_2

To understand how VQE is applied in practice, it is helpful to study the simplest nontrivial molecule: the hydrogen molecule H_2. This example illustrates the entire workflow used in quantum chemistry simulations on a quantum computer, without overwhelming chemical detail. The goal is to use VQE to approximate the ground-state energy of H_2 for a fixed internuclear distance R.

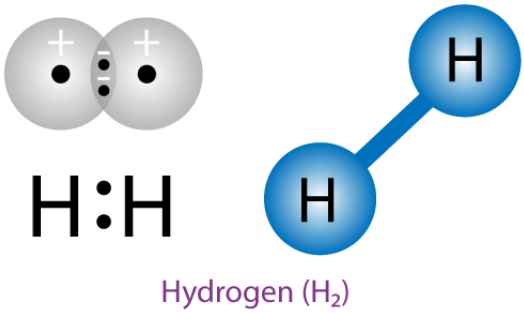

Figure 2: The Hydrogen molecule. Source: https://byjus.com/chemistry/hydrogen-gas/.

6.1 From the molecular problem to a qubit Hamiltonian

Every molecule is governed by an electronic Hamiltonian derived from Coulomb interactions. In full generality this Hamiltonian is complicated, but for quantum computing we do not work with the full continuum model. Instead, we follow three standard steps:

1. **Choose a finite orbital basis.** For pedagogical purposes, we use the minimal STO–3G basis for H_2, which contains only two spatial orbitals.
2. **Map orbitals to qubits.** Each spin–orbital (spatial orbital + spin label) can be either occupied or unoccupied, which is naturally represented by a qubit. After symmetry reductions and removing irrelevant subspaces, the H_2 problem reduces to *two qubits*.
3. **Rewrite the Hamiltonian in Pauli form.** Using a fermion-to-qubit mapping such as Jordan–Wigner or Bravyi–Kitaev, the molecular Hamiltonian becomes

$$H_{H_2} = c_0 I + c_1 Z_0 + c_2 Z_1 + c_3 Z_0 Z_1 + c_4 X_0 X_1 + c_5 Y_0 Y_1,$$

where the coefficients c_i depend on the internuclear distance R.

At this point the chemistry has been translated into a form perfectly suited for quantum computation: a weighted sum of Pauli operators acting on two qubits.

6.2 Building an ansatz for the ground state

Because of symmetry, the true ground state of H_2 lies in the two-dimensional subspace spanned by $|01\rangle$ and $|10\rangle$. Thus, a minimal and physically motivated variational ansatz is

$$|\psi(\theta)\rangle = \cos\theta\ |01\rangle + \sin\theta\ |10\rangle\,.$$

A simple two-qubit circuit that prepares this family of states is:

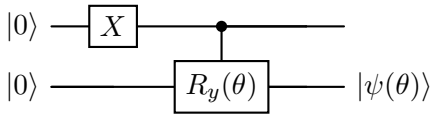

This ansatz contains a single parameter θ, yet is expressive enough to capture the exact ground state of H_2 in the minimal basis.

6.3 Running the VQE loop

The complete VQE procedure for H_2 is:

1. The classical optimizer chooses an initial parameter value $\boldsymbol{\theta^{(0)}}$.
2. The quantum computer prepares $|\psi(\boldsymbol{\theta^{(t)}})\rangle$ and measures all Pauli expectations for H_{H_2}.
3. The measured expectations are combined to estimate
$$E(\boldsymbol{\theta^{(t)}}) = \langle\psi(\boldsymbol{\theta^{(t)}})|H_{\mathrm{H}_2}|\psi(\boldsymbol{\theta^{(t)}})\rangle.$$
4. The classical optimizer updates the parameter:
$$\theta^{(t+1)} = \mathrm{OptimizerStep}\big(\boldsymbol{\theta^{(t)}}, E(\boldsymbol{\theta^{(t)}})\big).$$
5. Repeat until convergence.

This iterative process continues until the energy estimate stabilizes near the minimum. For each atomic distance (i.e., internuclear separation) R, the same procedure yields an approximate ground-state energy $E_{\min}(R)$. Plotting $E_{\min}(R)$ as a function of the atomic distance R produces the *potential energy curve* of the molecule and reveals the equilibrium bond length.

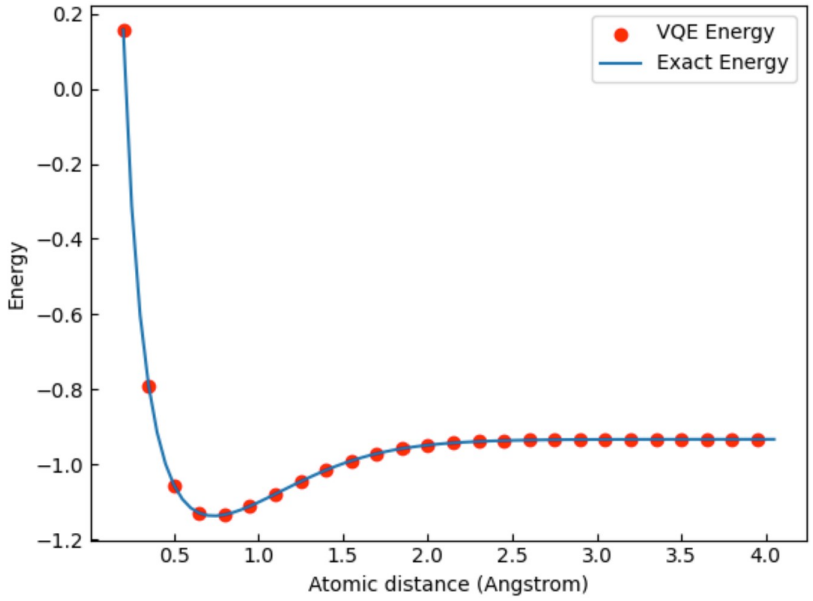

Figure 3: Potential energy surface of the hydrogen molecule. Source: [QX23].

References

[QX23] Maomin Qing and Wei Xie. Use vqe to calculate the ground energy of hydrogen molecules on ibm quantum, 2023. 9