EURO^2

Computational Chemistry in Qiskit: A Practical Overview of Current Approaches, 12 March 2026

Computational Chemistry in Qiskit: A Practical Overview of Current Approaches

Michal Belina

IT4Innovations
VSB - Technical University of Ostrava

12. 3. 2026

VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

Aims of the training

1. Present some of the available approaches
2. Show the implementation in Qiskit
3. Provide the jupyter notebook for testing

Timetable

13:00 – 13:50	Overview of Quantum Computing for Chemistry
13:50 – 14:00	Short break
14:00 – 14:50	Variational Quantum Eigensolver
14:50 – 15:00	Short break
15:00 – 15:50	Adaptive and Improved Variational Methods
15:50 – 16:00	Short break
16:00 – 17:00	Quantum Phase Estimation and Treatment of Excited-State Methods

Michal Belina

University of Chemistry and Technology Prague (2014-2025):

- Physical chemistry
- Computational chemistry

VŠB-Technical University Ostrava (2024 - present):

- Quantum Computing Lab

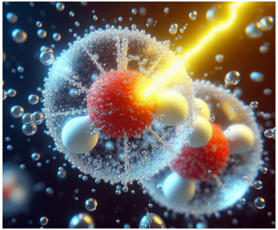

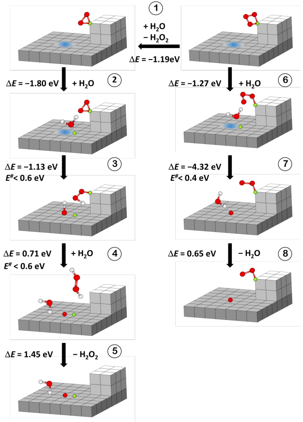

IT4Innovations

- Institute of VSB – Technical University of Ostrava
- Supercomputing and research center
- The most powerful supercomputing systems in the Czech Republic.
- Five research labs
- High-Performance Computing, Data Analysis, Quantum Computing, and Artificial Intelligence

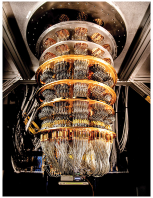

VLQ Quantum computer

- Superconducting qubits in a star-shaped topology
- 24 physical qubits connected to a central resonator
- minimising swap operations
- Integration with EuroHPC infrastructure

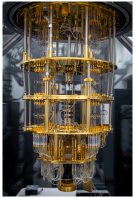

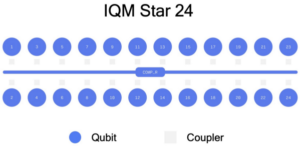

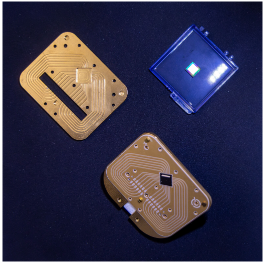

Promised VLQ specification

This is not the result of real-hardware benchmark.

Metric	Value
Qubits	24
Qubit connectivity	one-to-all, star-shape
T1 relaxation time	typically $\sim$40 μs, min. for all qubits: 15 μs
T2 dephasing time	typically $\sim$20 μs, min. for all qubits: 15 μs
1-qubit gate time	< 50 ns
2-qubit gate time	< 120 ns
1-qubit gate fidelity	> 0.997
2-qubit gate fidelity	> 0.99
Readout fidelity	> 0.95

Open call for QPU resources

!!!ONLY FOR RESEARCHERS WITH CZECH AFFILIATION!!!

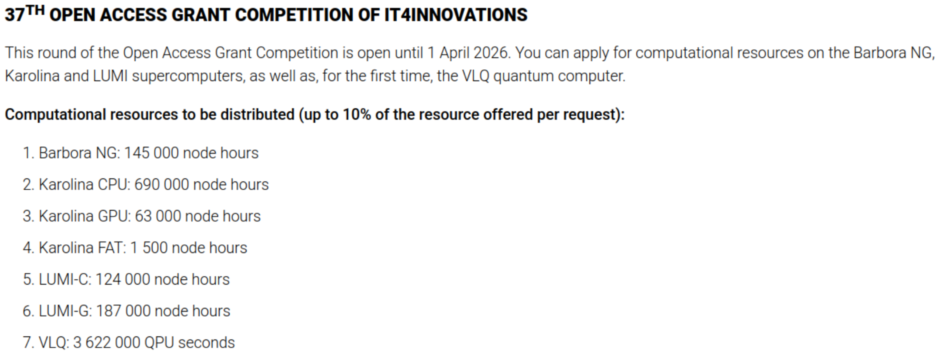

37TH OPEN ACCESS GRANT COMPETITION OF IT4INNOVATIONS

This round of the Open Access Grant Competition is open until 1 April 2026. You can apply for computational resources on the Barbora NG, Karolina and LUMI supercomputers, as well as, for the first time, the VLQ quantum computer.

Computational resources to be distributed (up to 10% of the resource offered per request):

1. Barbora NG: 145 000 node hours
2. Karolina CPU: 690 000 node hours
3. Karolina GPU: 63 000 node hours
4. Karolina FAT: 1 500 node hours
5. LUMI-C: 124 000 node hours
6. LUMI-G: 187 000 node hours
7. VLQ: 3 622 000 QPU seconds

Who are the participants?

Please, vote in the ZOOM survey. Voluntarily introduce yourself.

Preferably in chat:
Something like name, institute, field of main expertise and knowledge of quantum mechanics

Questions and comments are welcome at any time and in any way!

Computational chemistry bottlenecks: scaling of key electronic-structure methods

Method	Typical cost	What is captured well	Main limitation
Hartree–Fock (HF)	$\mathcal{O}(N^4)$	Mean-field ground state	No dynamical correlation
DFT	$\sim \mathcal{O}(N^3\text{–}N^4)$	Large systems at moderate cost	Functional-dependent accuracy
CCSD(T)	$\mathcal{O}(N^7)$	"Gold standard" for many molecules	Too costly for large systems; weak for multireference cases
CASSCF	exponential in active space	Correlation in selected orbitals	Active space remains small in conventional implementations
Full CI (FCI)	combinatorial / exponential	Exact solution in a chosen basis	Feasible only for very small problems

”Logical” qubits: public demonstrations

Company	**Platform**	**Where**	**What to say on the slide**
Quantinuum	Trapped ions	ArXiv (25.02.2026)	48 error-corrected logical qubits; 64 error-detected logical qubits reported.
Microsoft + Atom	Neutral atoms	ArXiv (09.06.2025)	24 entangled logical qubits; computation on 28 logical qubits.
Google Quantum AI	Superconducting	NATURE (09.12.2024)	Willow showed below-threshold error correction; scalable logical-memory milestone.
QuEra	Neutral atoms	NATURE (10.11.2025)	Public fault-tolerance claims; logical-qubit demonstrations with partners.

Status: We already have logical qubits that can outperform physical qubits, but not yet large-scale universal logical qubits that behave as nearly ideal computational objects.

”Logical” qubits: roadmaps and future targets

Company	Platform	Target year	Public target / roadmap
IBM	Superconducting	2029	Starling: 200 logical qubits.
IonQ	Trapped ions	2030	40,000–80,000 logical qubits in roadmap language.
Pasqal	Neutral atoms	2029–2030	100 logical qubits by 2029; 200 by 2030.
Alice & Bob	Cat qubits	2030	100 high-fidelity logical qubits.
PsiQuantum	Photonic	long-term	Utility-scale FTQC; emphasis on million-qubit-scale architecture.
Microsoft (Majorana 1)	Topological	long-term	Scalable topological path; no public large logical-qubit count yet.

Break

Software stack: Qiskit libraries

Library	Version	Comment
`qiskit`	1.4.5	Core SDK: circuits, operators, transpilation, primitives.
`qiskit-aer`	0.16.4	Local simulation: statevector, shots, noisy backends.
`qiskit-algorithms`	0.4.0	High-level algorithms such as VQE and eigensolvers.
`qiskit-nature`	0.7.2	Chemistry layer: electronic structure problems, mappings, ansätze.
`qiskit-ibm-runtime`	0.34.0	Interface to IBM Runtime primitives and hardware workflows.

Main idea:

- `qiskit` = core framework
- `nature` = chemistry problem construction

Supporting libraries and version choice

Library	Version	Comment
`PySCF`	2.8.0	Classical chemistry backend for molecular integrals and reference data.
`matplotlib`	latest	Plotting convergence, energies, and comparisons.
`pandas`	latest	Tables and structured result analysis.
`ipywidgets`	latest	Optional interactive notebook controls.

Why pin versions?

- Improves compatibility in Google Colab.
- Reduces API mismatch between Qiskit subpackages.
- Makes the notebook reproducible for training and demos.

Target environment: Google Colab, Runtime Version 2025.07

Variational Quantum Eigensolver in one slide

Aim: smallest eigen value of a Hamiltonian

$$E_0 = \min_{|\psi\rangle} \langle\psi| H |\psi\rangle ,$$

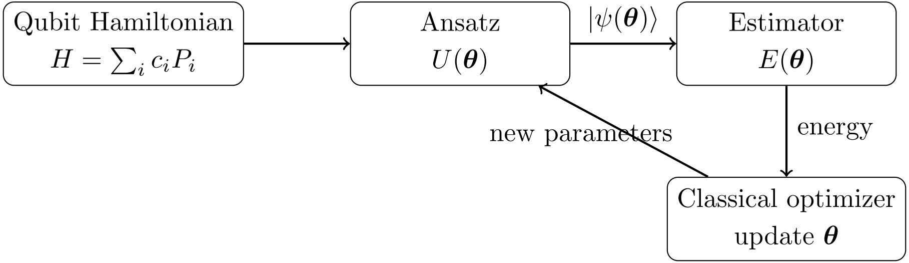

Qubit Hamiltonian $H = \sum_i c_i P_i$

Ansatz $U(\boldsymbol{\theta})$

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

Estimator $E(\boldsymbol{\theta})$

energy

Classical optimizer update $\boldsymbol{\theta}$

new parameters

- Quantum part: prepare $|\psi(\boldsymbol{\theta})\rangle$ and estimate $\langle H\rangle$.
- Classical part: update parameters $\boldsymbol{\theta}$ to lower the energy.
- Result: a hybrid quantum–classical optimization loop.

From electronic structure to second quantization

Electronic Hamiltonian

$$\hat{H}_e = \sum_i \left(-\frac{1}{2}\nabla_i^2 - \sum_A \frac{Z_A}{|\mathbf{r}_i - \mathbf{R}_A|} \right) + \sum_{i<j} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|}$$

Acts on the many-electron wavefunction

$$\Psi(\mathbf{r}_1, \mathbf{r}_2, \ldots, \mathbf{r}_N)$$

Main difficulty: antisymmetry and exponential growth of the many-body space.

Second quantization

Choose a spin-orbital basis $\{\phi_p\}$ and rewrite the problem in occupation-number language:

$$\hat{H} = \sum_{pq} h_{pq}\, a_p^\dagger a_q + \frac{1}{2} \sum_{pqrs} h_{pqrs}\, a_p^\dagger a_q^\dagger a_r a_s$$

- h_{pq}: one-electron integrals
- h_{pqrs}: two-electron integrals
- $a^\dagger, a$: enforce fermionic statistics

This is the standard starting point for Jordan–Wigner / Bravyi–Kitaev mapping.

Pauli Strings in Qubit Hamiltonians

From Fermions to Qubits

Electronic Hamiltonian in second quantization:

$$H = \sum_{pq} h_{pq} a_p^\dagger a_q + \frac{1}{2} \sum_{pqrs} h_{pqrs} a_p^\dagger a_q^\dagger a_r a_s$$

After fermion-to-qubit mapping (Jordan–Wigner, Bravyi–Kitaev):

$$H = \sum_i c_i P_i$$

where

- c_i — real coefficients
- P_i — **Pauli strings** (I, X, Y, Z tensor products)

Example:

$$H = -1.05I + 0.39Z_0 - 0.39Z_1 - 0.01Z_0Z_1 + 0.18X_0X_1 + 0.18Y_0Y_1$$

Physical Meaning of Pauli Strings

Each Pauli string corresponds to a measurable observable.

Pauli term	Origin in fermionic Hamiltonian	Physical meaning
Z_i	$a_i^\dagger a_i$	orbital occupation
$Z_i Z_j$	$n_i n_j$	electron–electron interaction
$X_i X_j + Y_i Y_j$	$a_i^\dagger a_j + a_j^\dagger a_i$	electron hopping
long Z strings	JW parity operators	fermionic antisymmetry

Example:

$$0.18 X_0 X_1 + 0.18 Y_0 Y_1$$

represents electron excitation between orbitals 0 and 1.

Commutation of Pauli Strings

Two Pauli strings either **commute** or **anticommute**.
Rule:

- If the number of qubits where Pauli operators anticommute is **even** → strings commute
- If it is **odd** → strings anticommute

Example:

$$P_1 = X_0 X_1 I I$$

$$P_2 = Z_0 Z_1 I I$$

Qubit comparison:

qubit	P_1	P_2	relation
0	X	Z	anticommute
1	X	Z	anticommute

Two anticommutes ⇒ even ⇒

$$P_1 P_2 = P_2 P_1$$

The operators commute.

Expectation Value Measurement

Energy evaluation in VQE:

$$E = \langle \psi | H | \psi \rangle = \sum_i c_i \langle P_i \rangle$$

Each Pauli string expectation value is estimated from measurements.
Example measurement results:

bitstring	count	$Z_0 Z_1$ eigenvalue
00	420	+1
01	300	-1
10	200	-1
11	80	+1

Expectation value:

$$\langle Z_0 Z_1 \rangle = \frac{420 - 300 - 200 + 80}{1000} = 0$$

The key idea: the variational principle

Variational statement

For any normalized trial state $|\psi(\boldsymbol{\theta})\rangle$,

$$E(\boldsymbol{\theta}) = \langle\psi(\boldsymbol{\theta})|\, H \,|\psi(\boldsymbol{\theta})\rangle \geq E_0.$$

- If our ansatz family can represent a state close to the true ground state, minimizing the expectation value can approximate E_0.
- Therefore, VQE does not directly diagonalize H.
- Instead, it searches over a parametrized family of quantum states prepared by a circuit $U(\boldsymbol{\theta})$.
- The whole method stands or falls with three practical questions:
    1. Is the ansatz expressive enough?
    2. Can we estimate the energy accurately enough?
    3. Can the optimizer find good parameters efficiently?

Main components of a VQE calculation

Quantum ingredients

- **Operator:** the qubit Hamiltonian.
- **Reference state:** often Hartree–Fock.
- **Ansatz:** parametrized circuit creating trial states.
- **Estimator:** expectation-value evaluation.

Classical ingredients

- **Optimizer:** COBYLA, SLSQP, SPSA, . . .
- **Initialization:** initial parameter vector.
- **Stopping rule:** max iterations, tolerance, convergence test.
- **Diagnostics:** callback, energy trace, parameter history.

Important practical message

VQE is not a single algorithmic choice, but a **framework** whose performance depends strongly on all of the design decisions above.

Ansatz: choosing a trial wave function

- A parametrized circuit prepares

$$|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})\,|\psi_{\mathrm{ref}}\rangle\,.$$

- In chemistry, common choices include:
  - hardware-efficient ansätze,
  - UCC-inspired ansätze (especially UCCSD),
  - problem-tailored or symmetry-preserving constructions.
- Trade-off:
  - more expressive ansatz $\Rightarrow$ potentially better accuracy,
  - but also more parameters, deeper circuits, and harder optimization.
- A chemically motivated reference state often improves convergence.

A good ansatz should balance expressivity, circuit depth, symmetry preservation, and hardware feasibility.

Chemistry-motivated reference states and ansätze

- **Hartree–Fock reference:** a natural starting point because it approximates the mean-field electronic configuration.
- **UCCSD idea:**

  $$|\psi(\boldsymbol{\theta})\rangle \approx e^{T(\boldsymbol{\theta})-T^\dagger(\boldsymbol{\theta})}\,|\mathrm{HF}\rangle\,,$$

  where T contains excitation operators.
- UCCSD is chemically interpretable but may become costly.
- Hardware-efficient circuits are shallower, but may be harder to interpret and optimize.

Energy estimation is not a single measurement

- Because

$$H = \sum_i c_i P_i,$$

  the energy expectation value is

$$E(\boldsymbol{\theta}) = \sum_i c_i \langle P_i \rangle_{\boldsymbol{\theta}}.$$

- The Pauli terms generally do **not** all commute.
- Therefore, we cannot usually measure the full Hamiltonian in one shot.
- Instead we:
  1. group commuting terms,
  2. rotate into appropriate measurement bases,
  3. sample circuits repeatedly,
  4. combine results into an energy estimate.
- This measurement overhead is one of the central bottlenecks of VQE.

Classical optimization loop

- Starting from $\boldsymbol{\theta}^{(0)}$, the optimizer iteratively updates the parameters.
- At each step, the cost function is the estimated energy

$$E(\boldsymbol{\theta}) = \langle\psi(\boldsymbol{\theta})|\, H \,|\psi(\boldsymbol{\theta})\rangle \,.$$

- Common choices:
  - **COBYLA:** derivative-free, simple, popular for small demonstrations,
  - **SLSQP:** gradient-based constrained optimization,
  - **SPSA:** robust option for noisy settings.
- What matters in practice:
  - number of energy evaluations,
  - robustness to shot noise and hardware noise,
  - sensitivity to initial parameters,
  - optimizer–ansatz compatibility.

Minimal algorithmic pseudocode

```
1 Build molecular Hamiltonian in second quantization.
2 Map the fermionic operator to a qubit Hamiltonian.
3 Choose reference state and ansatz $U(\boldsymbol{\theta})$.
4 Initialize parameters $\boldsymbol{\theta}^{(0)}$.
5 Repeat until convergence:
    1 prepare $|\psi(\boldsymbol{\theta})\rangle$,
    2 estimate $E(\boldsymbol{\theta})$,
    3 update $\boldsymbol{\theta}$ with a classical optimizer.
6 Return best energy and corresponding parameters.
```

Output of interest

Not only the approximate ground-state energy, but also the optimized variational state and observables derived from it.

Why VQE became so important

- It is compatible with relatively shallow circuits compared with algorithms designed for fault-tolerant machines.
- It naturally combines quantum state preparation with classical optimization.
- It is flexible enough to accommodate chemistry-inspired wave functions.
- It remains a standard benchmark for new ideas in:
  - ansatz design,
  - measurement reduction,
  - error mitigation,
  - excited-state and adaptive extensions.

Pedagogical value

Even when a later workflow outperforms plain VQE, understanding VQE is essential because many modern hybrid algorithms modify only one part of the same basic loop.

Main limitations of plain VQE

Algorithmic challenges

- difficult ansatz selection,
- local minima / flat landscapes,
- parameter sensitivity,
- scalability of optimization.

Quantum-hardware challenges

- finite sampling error,
- device noise,
- limited coherence,
- measurement overhead.

- These limitations motivate the next topics in the training:
    - adaptive and improved variational methods,
    - phase-estimation-based approaches for the fault-tolerant regime.

Take-home messages before the notebook session

1. VQE estimates the ground-state energy by minimizing an expectation value over a parametrized quantum state.
2. The algorithm is hybrid: quantum processor for state preparation/measurement, classical computer for optimization.
3. The crucial practical design choices are the Hamiltonian representation, ansatz, estimator strategy, and optimizer.
4. For chemistry, physically motivated references and ansätze are often essential.
5. Plain VQE is a starting point, not the final word.

Break

Motivation for adaptive ansätze

- Chemistry demands accuracy, but near-term hardware limits depth and measurement budget.
- A fixed ansatz can be:
  - over-parameterized (wastes depth/measurements), or
  - under-expressive (misses correlation).
- **ADAPT-VQE idea:** grow the ansatz *only as needed*, guided by energy gradients.

Key promise

Compact, problem-tailored circuits with a built-in convergence criterion.

Ansatz structure (adaptive growth)

ADAPT-VQE builds an ansatz as a product of exponentials from an **operator pool**:

$$\left|\psi(\vec{\theta})\right\rangle = \prod_{k=1}^{m} e^{\theta_k A_k} \, |\phi_0\rangle ,$$

where each A_k is chosen from a pool $\mathcal{P} = \{A_1, \ldots, A_{|\mathcal{P}|}\}$.

- Start from a reference state (often).
- Iteratively append one (or a batch of) operator(s) that most improve the energy.
- Re-optimize all parameters after each growth step.

How does ADAPT choose the next operator?

For each candidate $A_i \in \mathcal{P}$, estimate the **energy gradient** of adding $e^{\theta A_i}$ at $\theta = 0$:

$$g_i \equiv \left| \frac{\partial E}{\partial \theta} \right|_{\theta=0} \quad \propto \quad \left| \langle \psi | \left[H, A_i \right] | \psi \rangle \right| .$$

Then pick the operator with the largest g_i:

$$A_\star = \arg \max_{A_i \in \mathcal{P}} g_i .$$

Stopping criterion

Stop when $\max_i g_i < \varepsilon$ (or $\|\vec{g}\| < \varepsilon$): the pool has no strong “missing directions”.

Operator pools (chemistry vs hardware)

Fermionic (chemistry-inspired) pools

- Excitation operators (singles/doubles, spin-adapted variants)
- Often accurate/compact for molecules
- Mapping to qubits (JW / Parity / BK) yields Pauli sums

Qubit (hardware-friendly) pools

- Pauli-string generators directly in qubit space
- Can dramatically reduce depth
- Known variants: *qubit-ADAPT-VQE*

Practical note

Pool choice strongly affects depth, measurement cost, and convergence speed.

ADAPT-VQE loop (pseudo-code)

Algorithm sketch

```
1 Initialize $|\psi\rangle \leftarrow |\phi_0\rangle$, ansatz list $\mathcal{A} \leftarrow [\,]$.
2 **Repeat:**
    1 For all $A_i \in \mathcal{P}$, estimate gradient $g_i$ on current state $|\psi\rangle$.
    2 Pick $A_\star = \arg\max_i g_i$ (or a batch of top-$b$ operators).
    3 Append: $\mathcal{A} \leftarrow \mathcal{A} \cup [A_\star]$.
    4 Re-optimize all parameters $\vec{\theta}$ to minimize $E(\vec{\theta})$.
    5 Update $|\psi\rangle \leftarrow \prod_{A_k \in \mathcal{A}} e^{\theta_k A_k} |\phi_0\rangle$.
3 Until $\max_i g_i < \varepsilon$ (or iteration budget reached).
```

Measurement overhead: the big catch

- Each iteration needs gradients for *all* pool operators:

$$\text{cost per iter} \sim |\mathcal{P}| \times (\text{shots per gradient}).$$

- Plus energy evaluations during classical optimization.

Common mitigation strategies (high-level)

- **Batched ADAPT**: add multiple operators per iteration to reduce iterations.
- **Pool design**: smaller/structured pools, qubit pools, symmetry pruning.
- **Measurement reduction**: grouping/commuting structure, smarter gradient estimation.

Takeaways

- ADAPT-VQE replaces “choose an ansatz” with “choose a pool + convergence rule”.
- Strength: compact, system-tailored circuits with interpretable growth.
- Tradeoff: gradient screening can be measurement-heavy; pool design matters.
- Variants (qubit pools, batched growth, measurement mitigation) target NISQ practicality.

Good default for chemistry demos

HF reference + chemistry-inspired excitation pool + modest batching + careful shot budgeting.

References (starter set)

- Grimsley *et al.*, “An adaptive variational algorithm for exact molecular simulations on a quantum computer” (ADAPT-VQE).
  `arXiv:1812.11173` / PRX (2019).
- Tang *et al.*, “qubit-ADAPT-VQE” (hardware-efficient adaptive pools).
  `arXiv:1911.10205` / PRX Quantum (2021).
- Qiskit Nature How-to: “Find ground state energy using AdaptVQE”.
  `qiskit-community.github.io/qiskit-nature/howtos/adapt_vqe.html`

Break

Why Quantum Phase Estimation?

- Goal: extract eigenvalues of a Hermitian operator from a quantum computer.
- In quantum chemistry, the operator of interest is the electronic Hamiltonian $\hat{H}$.
- QPE targets the phase of a unitary operator

$$U \left|\psi_k\right\rangle = e^{2\pi i \phi_k} \left|\psi_k\right\rangle ,$$

  where $\left|\psi_k\right\rangle$ is an eigenstate of U.
- If we encode $U = e^{-i\hat{H}t}$, then the phase contains the energy:

$$\hat{H} \left|\psi_k\right\rangle = E_k \left|\psi_k\right\rangle \quad \Rightarrow \quad U \left|\psi_k\right\rangle = e^{-iE_k t} \left|\psi_k\right\rangle .$$

- Therefore, estimating the phase ϕ_k lets us recover the energy E_k.

Key message: QPE is a route from *time evolution* to *energy estimation*.

From molecular Hamiltonian to a phase

- Start from the electronic Hamiltonian in second quantization:

$$\hat{H} = \sum_{pq} h_{pq}\, a_p^\dagger a_q + \frac{1}{2}\sum_{pqrs} h_{pqrs}\, a_p^\dagger a_q^\dagger a_r a_s.$$

- Map fermionic operators to qubits:

$$\hat{H} \;\longrightarrow\; H_q = \sum_j c_j P_j,$$

  where P_j are Pauli strings.

- Build a unitary from the qubit Hamiltonian:

$$U = e^{-iH_q t}.$$

- If $|\psi_k\rangle$ is an eigenstate of H_q with energy E_k, then

$$U\,|\psi_k\rangle = e^{-iE_k t}\,|\psi_k\rangle\,.$$

Interpretation: chemistry becomes an eigenphase-estimation problem.

Core idea of QPE

- Use two registers:
  - **phase register** with m qubits,
  - **system register** prepared close to an eigenstate.
- Apply Hadamards on the phase register.
- Apply controlled powers of the unitary:

$$U,\ U^2,\ U^4,\ \ldots,\ U^{2^{m-1}}.$$

- Apply the inverse quantum Fourier transform ($\mathrm{QFT}^\dagger$).
- Measure the phase register to obtain an m-bit estimate of ϕ.

Standard QPE circuit structure

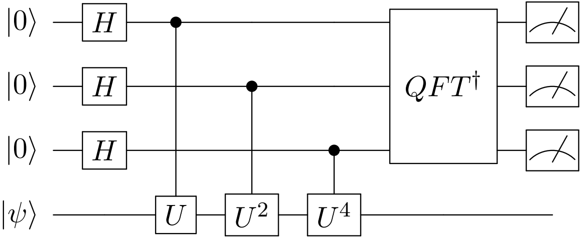

- The bottom register should ideally be an eigenstate $|\psi\rangle$.
- In practice, we often prepare an approximation with nonzero overlap with the target eigenstate.
- Measurement returns the target phase with probability proportional to that overlap.

How the energy is recovered

Assume

$$U = e^{-iHt}, \qquad H\left|\psi_k\right\rangle = E_k\left|\psi_k\right\rangle.$$

Then

$$U\left|\psi_k\right\rangle = e^{-iE_kt}\left|\psi_k\right\rangle.$$

Write the phase in the QPE convention as

$$e^{-iE_kt} = e^{2\pi i\phi_k}.$$

Hence

$$\phi_k = -\frac{E_kt}{2\pi} \pmod{1}.$$

So after estimating ϕ_k, we infer the energy as

$$E_k \approx -\frac{2\pi}{t}\,\phi_k,$$

up to the usual phase-wrap ambiguity.

Important practical point:

- the evolution time t must be chosen so that the relevant energies can be resolved,
- too large t may introduce phase wrapping,
- too small t reduces spectral resolution.

Algorithmic workflow for chemistry

1. Define the molecule, basis, charge, and spin.
2. Build the fermionic electronic Hamiltonian.
3. Map it to a qubit Hamiltonian.
4. Prepare an approximate eigenstate:
    - Hartree–Fock state,
    - classically optimized VQE state,
    - other chemically motivated ansatz/state preparation.
5. Implement controlled time evolution e^{-iHt}.
6. Run QPE and estimate the phase.
7. Convert phase to energy.

Why state preparation matters: if the prepared state is

$$|\psi_{\text{in}}\rangle = \sum_k c_k |\psi_k\rangle ,$$

QPE collapses onto eigenvalue E_k with probability $|c_k|^2$.

Precision and resource scaling

- With m phase qubits, the phase is resolved to about m binary digits.
- Better precision requires:
  - more phase qubits,
  - longer coherent controlled evolutions,
  - more accurate implementation of e^{-iHt}.
- The most expensive part is usually controlled time evolution.
- This is why QPE is often viewed as a fault-tolerant, long-circuit algorithm.

Main cost drivers

- number of phase qubits
- circuit depth of U^{2^k}
- Hamiltonian simulation method
- overlap with target state
- noise / decoherence

Takeaway: QPE can be very accurate, but hardware demands are much stronger than for VQE.

QPE vs. VQE

Feature	**QPE**	**VQE**
Target quantity	Eigenphase / energy	Variational energy
Circuit depth	High	Lower
Classical optimization	No outer loop	Yes
State preparation sensitivity	High	High
Precision scaling	Favorable	Sampling-limited
Best hardware regime	Fault-tolerant	NISQ / early utility

- VQE is typically more practical on near-term devices.
- QPE is the more direct eigenvalue algorithm and the long-term high-precision route.
- In practice, a VQE state can serve as an input state for QPE.

Iterative QPE and practical variants

- **Standard QPE**: uses a full phase register and inverse QFT.
- **Iterative QPE (IQPE)**:
    - uses one ancilla qubit repeatedly,
    - estimates phase bits sequentially,
    - reduces qubit count,
    - still requires controlled time evolution.
- Practical chemistry demonstrations often simplify one or more of:
    - system size,
    - ansatz / initial state,
    - Hamiltonian simulation,
    - noise model.

QPE take-home messages

- QPE estimates eigenvalues by measuring phases of a unitary operator.
- For chemistry, choose
$$U = e^{-iHt},$$
so that measured phase encodes molecular energy.
- Accuracy improves systematically with phase resolution.
- Success probability depends on overlap with the target eigenstate.
- Main bottleneck: deep, controlled Hamiltonian time evolution.
- Conceptually elegant, practically demanding, and central for fault-tolerant quantum chemistry.

Q&A and Closing the training

Thank you for attention!