quobly

quobly PE ToolboX FOXCONN®

Quantum Phase Estimation Toolbox - deep dive

Thibaud Louvet[1], Calvin Ku[2], Yu-Cheng Chen[2], Carlos Ramos Marimon[1], Olivier Gauthé[1], Tristan Meunier[1], Min-Hsiu Hsieh[2] and Benoit Vermersch[1]

[1] *Quobly, Grenoble, France*

[2] *Hon Hai (Foxconn) Research Institute, Taipei, Taiwan*

15 december 2025

quobly PE Toolbox FOXCONN®

quobly

The people

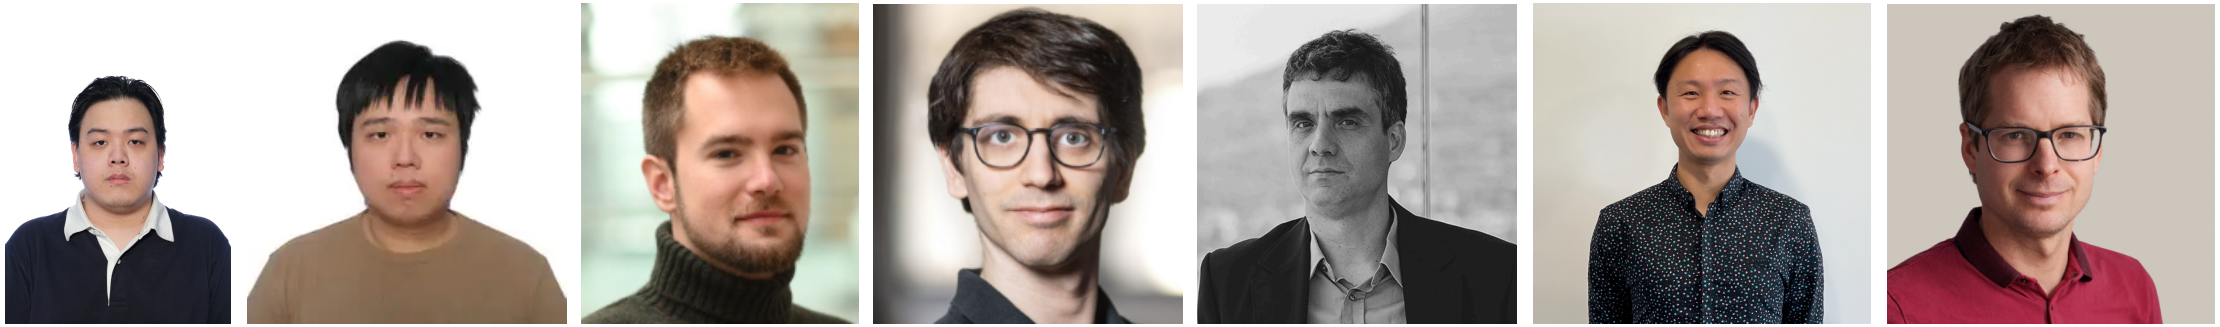

Calvin Ku (Foxconn)

Yu-Cheng Chen (Foxconn)

Carlos Ramos Marimon (Quobly)

Olivier Gauthé (Quobly)

Tristan Meunier (Quobly)

Min-Hsiu Hsieh (Foxconn)

Benoit Vermersch (Quobly)

The project: QPE Toolbox

quobly

Classical preprocessing

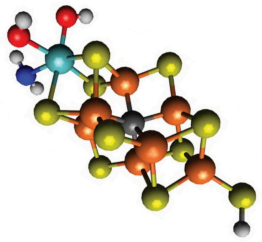

DMRG

MPS to Circuit conversion

Hamiltonian to unitary/gates

Simulator

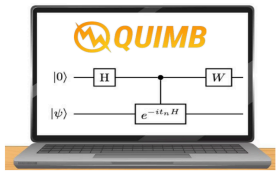

Circuit to Tensor Network

QPU

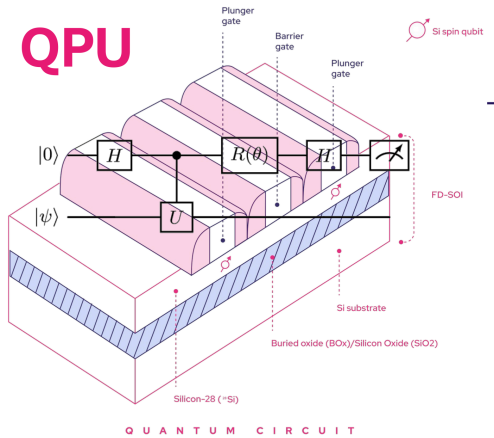

Classical post-processing

0101001010
0101011100
1010111010
1101111011
0101001110
1011010101

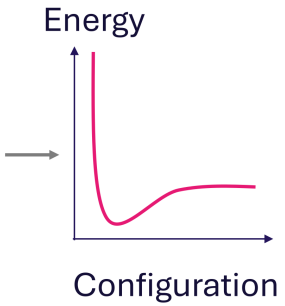

quobly

This talk

1. quantum algorithms
2. chemistry
3. tensor networks
4. our toolbox

Why Quantum Phase Estimation?

There is a variety of quantum algorithms. We chose to focus on one

- The "zoo" of quantum algorithms breaks down into many variations of fewer algorithmic primitives
- We believe the first useful application of a QPU will require **FTQC**
- **QPE is:**
  - a key subroutine in **Shor's algorithm**
  - FTQC: requires **deep** circuits (**hard** to simulate classically)
  - expected to solve problems in **chemistry**, **condensed matter**

Quantum algorithms, quantum advantage

Nielsen and Chuang, *Quantum computation and quantum information* (2010)

> **1.4 Quantum algorithms**
>
> [...]
>
> "Broadly speaking, there are three classes of quantum algorithms which provide an advantage over known classical algorithms"

1. Quantum Fourier Transform (QFT) based algorithms (Phase Estimation, Shor)
2. Quantum search algorithms (Grover)
3. Quantum simulation

1. Quantum Fourier Transform (QFT) based algorithms

- QFT is a subroutine of
  - QPE. QPE is a subroutine of
    - Shor's algorithm: find the prime factors of an integer
- Shor's algorithm runs in polynomial time. **Superpolynomial speedup** compared to the best-known classical algorithm

quobly

1. QFT based algorithms

Origin of Shor's quantum speedup?

Nielsen and Chuang point towards QFT and QPE.

> **5.2 Phase estimation**
>
> The Fourier transform is the key to a general procedure known as phase estimation, which in turn is the key for many quantum algorithms.
> (Nielsen and Chuang, p.221)

1. QFT based algorithms

Quantum Phase Estimation

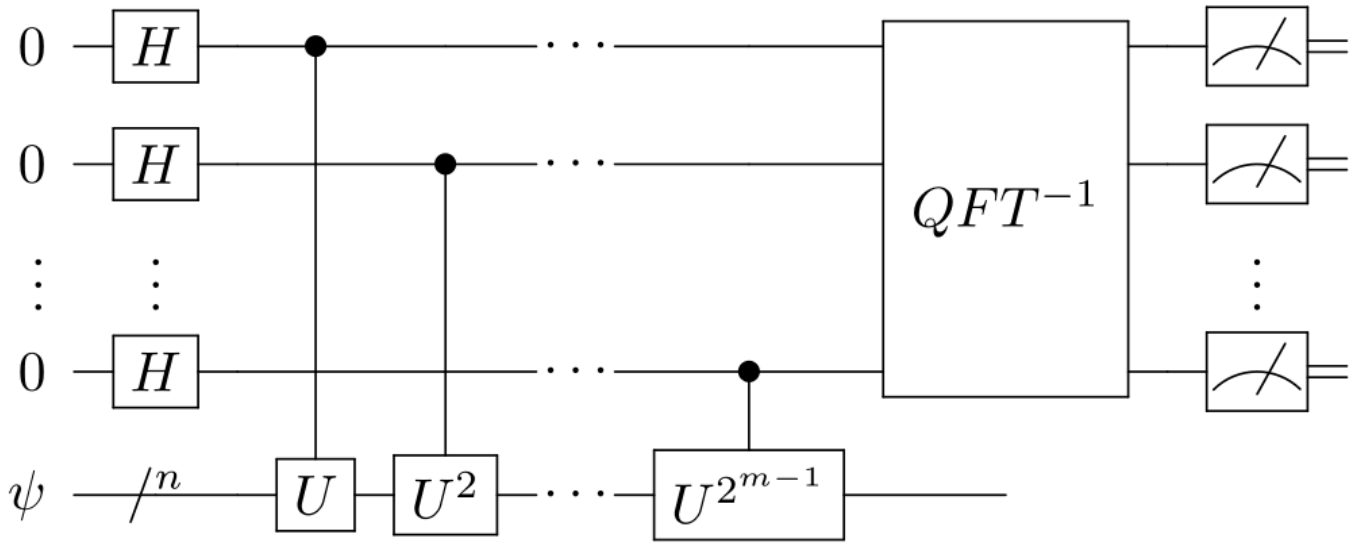

- $\hat{U}|u\rangle = e^{i\theta}|u\rangle$
- Find θ with precision $\sim 1/2^m$

1. QFT based algorithms

Quantum Fourier Transform

$$|j\rangle \rightarrow \frac{1}{\sqrt{2^n}} \sum_{k=0}^{2^n-1} e^{2\pi i j k / 2^n} |k\rangle$$

- Classical Fast Fourier Transform (FFT): cost $\mathcal{O}(n2^n)$

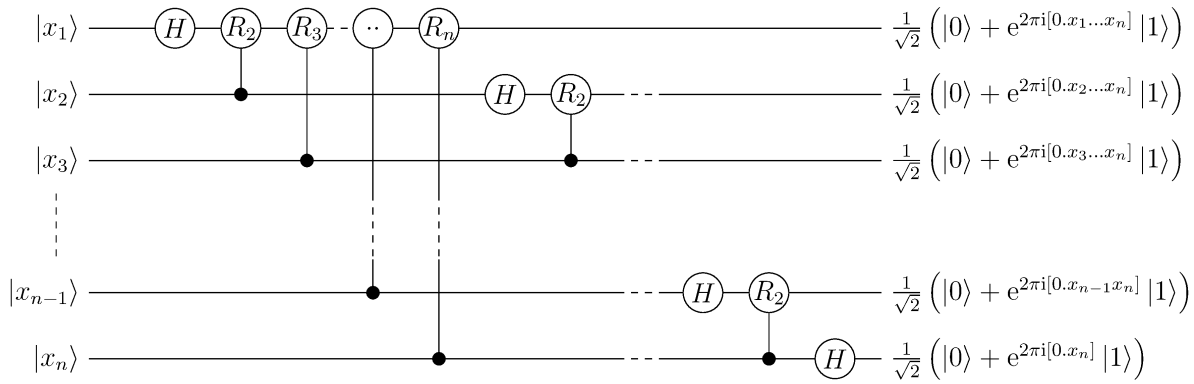

- Quantum Fourier Transform: $\mathcal{O}(n^2)$ (`qpe_toolbox/estimation/qft.py`)

1. QFT based algorithms

But... the QFT has no quantum speedup

- The entanglement in QFT only comes from the reversal of the bits ordering
- The QFT with no bits reversal can be **simulated classically with $\mathcal{O}(n)$ cost**

PRX QUANTUM **4,** 040318 (2023)

Quantum Fourier Transform Has Small Entanglement

Jielun Chen (陈捷伦),[1,2,*] E.M. Stoudenmire,[3] and Steven R. White[1]

[1] *Department of Physics and Astronomy, University of California, Irvine, California 92697-4575, USA*
[2] *Department of Physics, California Institute of Technology, Pasadena, California 91125, USA*
[3] *Center for Computational Quantum Physics, Flatiron Institute, 162 Fifth Avenue, New York, New York 10010, USA*

(Received 2 January 2023; accepted 25 September 2023; published 27 October 2023)

SciPost Phys. 18, 104 (202

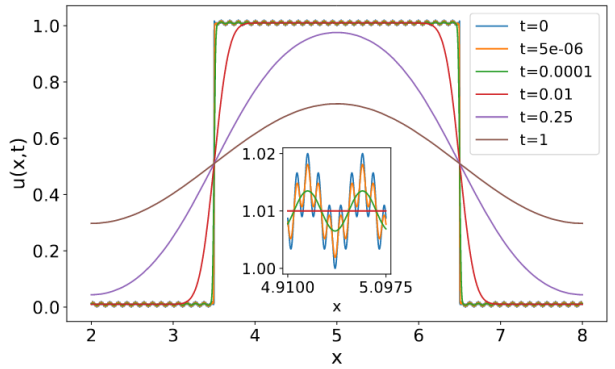

Figure 11: Solution of the heat equation (80) using quantics TCI. The plot shows $u(x,t)$ versus x for different times. We used a 1D grid with $M = 2^{\mathcal{R}}$ points and $\mathcal{R} = 30$, at a computational cost of $\mathcal{O}(\mathcal{R})$. The inset shows a zoom close to $x = 5$.

1. QFT based algorithms

Quantum Phase Estimation

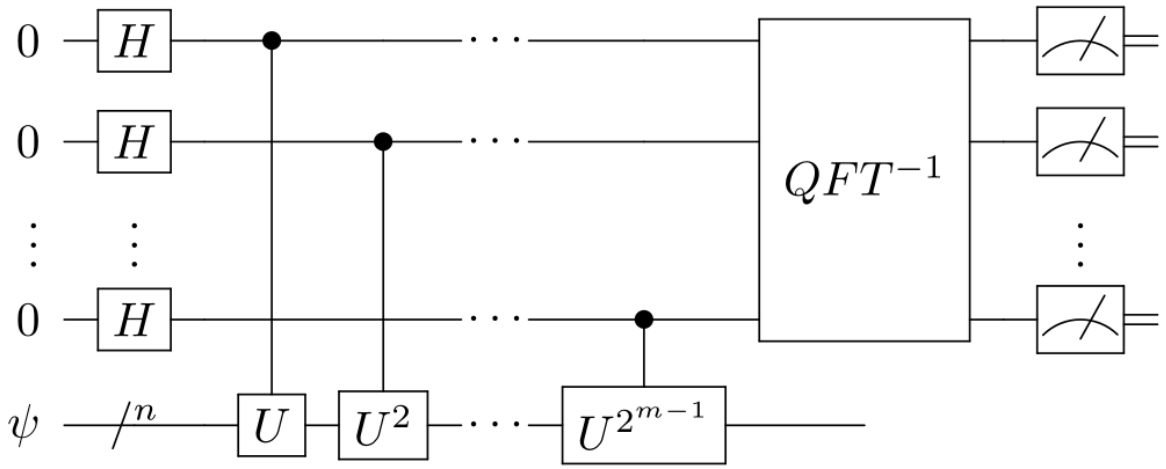

Origin of the speed-up?

- state encoding
- phase encoding

2. Quantum search algorithms

Canonical algorithm: Grover algorithm

- Find the input to a black-box function that produces a particular output in time $\mathcal{O}(\sqrt{N})$
- Quadratic speed-up: best classical algorithm runs in $\mathcal{O}(N)$

2. Quantum search algorithms

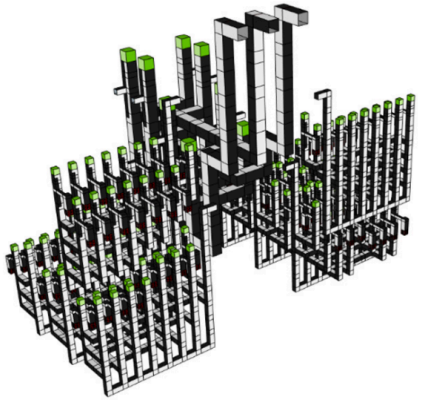

"Quantum NAND"
more than 10 qubit seconds

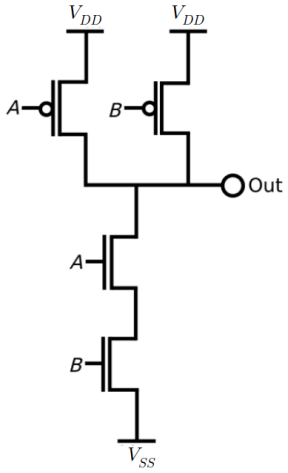

"Classical NAND"
less than 10^{-9}
transistor seconds

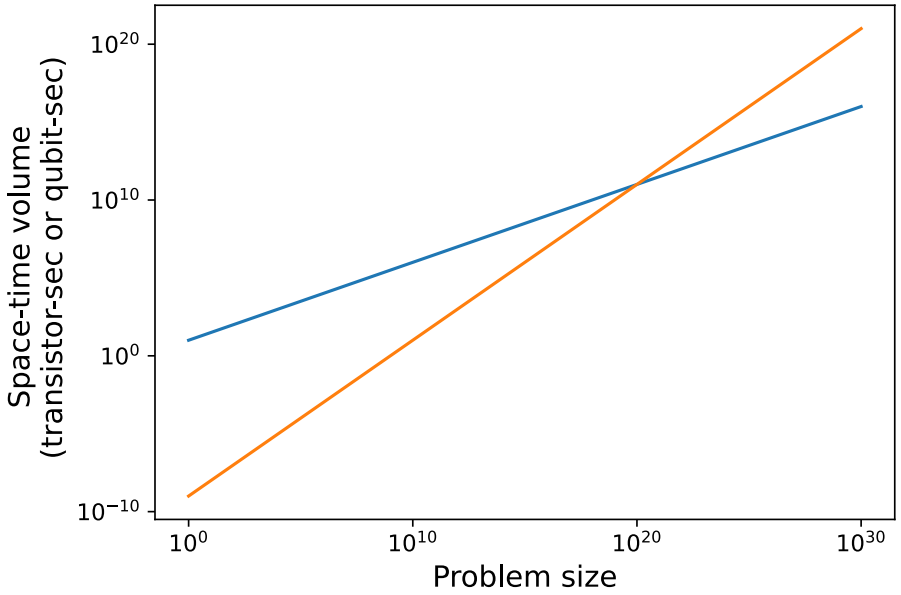

Babbush *et al.*, PRX Quantum 2 (2021): quadratic speed-ups will not enable quantum advantage on early fault tolerant quantum computers

3. Quantum simulation

Nielsen and Chuang, p.39

> "Classical computers have difficulty simulating general quantum systems for much the same reasons they have difficulty simulating quantum computers - the number of complex numbers needed to describe a quantum system generally grows exponentially with the size of the system, rather than linearly, as occurs in classical systems."

Quantum many-body n particles wavefunction $|\psi\rangle$

- Classical computer: $N = 2^n$ complex amplitudes
- Quantum computer: n qubits

quobly

Sum-up: why Quantum Phase Estimation?

Two arguments why QPE might be a good candidate for a useful quantum advantage:

- Shor's algorithm has a proven superpolynomial speedup. QPE is key in Shor
- QPE can be used to measure the energy of a quantum many-body state: "quantum simulation"-type advantage

The quest for a quantum advantage

quobly

A "large Hilbert space" is not enough

Can we solve any problem of size 2^n with just n qubits?
In practice, many factors can limit the quantum advantage of an algorithm.

Data access

- Storing 2^n amplitudes in n qubits (QRAM) has a cost scaling like $\mathcal{O}(2^n)$

Data extraction

- Only $\mathcal{O}(n)$ classical bits of information can be retrieved from n qubits through measurement (Holevo's bound)

The quest for a quantum advantage

Random circuits: quantum "supremacy"... on useless problems

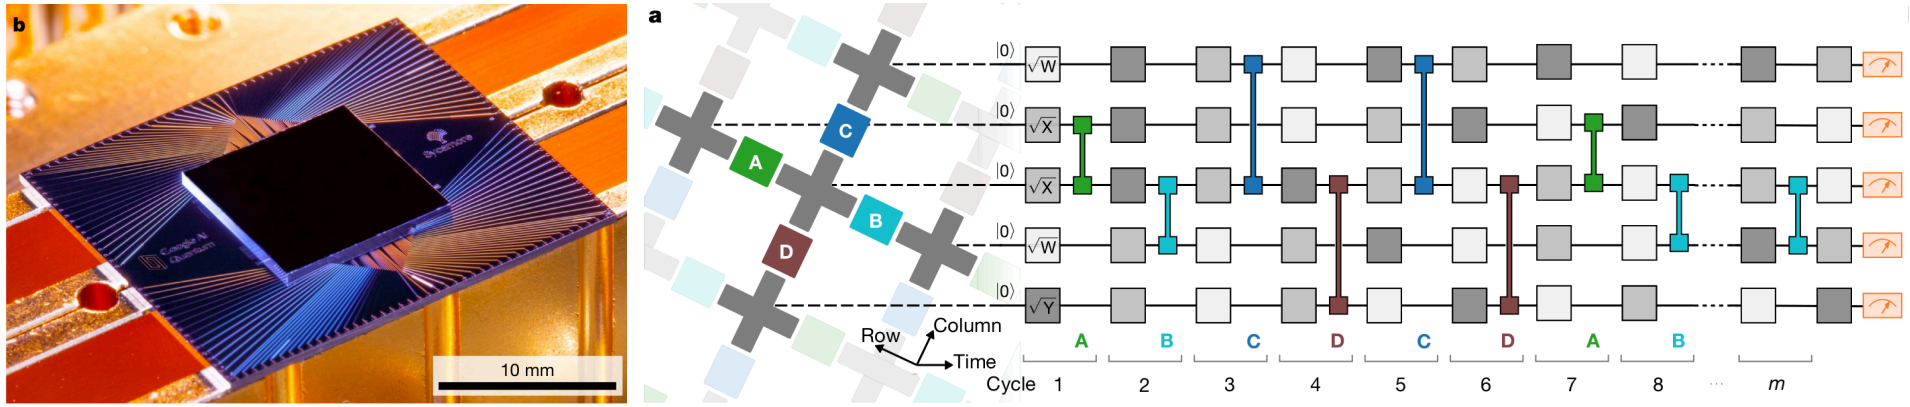

Google's Sycamore "supremacy" experiment, 2019

The quest for a quantum advantage

quobly

Race between NISQ and classical simulations (non exhaustive)

Experiment	qbits	gates	Type	Simulated?	Useful?
Google 2019	53	430	Random Circuit Sampling	Tensor Networks	No
USTC 2023	50	/	Gaussian Boson sampling	Tensor Networks	No
IBM 2023	127	2880	Ising Model	TN, Pauli Path	Yes
Google 2025	65	1000	OTOC measurement	No	No

quobly

The quest for a quantum advantage

- beyond quadratic speedup
- useful task
- quantum "by nature"

$\rightarrow$ focus on chemistry or condensed matter applications

Schrödinger equation:

$$i\hbar\frac{d\psi}{dt} = H\psi$$

quobly

Chemistry

- **Applications:** drugs, fertilizers, batteries
- **Goals:** understand reaction mechanisms $\rightarrow$ find new products/catalysts
- **Arrhenius law:**
  - reaction rate $k \propto e^{-E_a/k_BT}$ (E_a: activation energy)
- **Chemical accuracy:** 1kcal/mol = 1.6mHa = 50meV
- Molecules are **quantum many-body systems**

Molecules

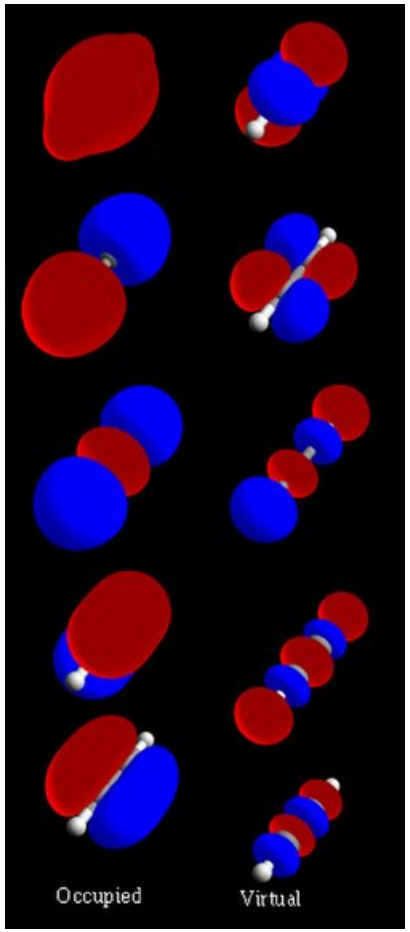

- N electrons, M molecular orbitals
- **Molecular electronic Hamiltonian:**
  $\hat{H} = \hat{T}_{el} + \hat{V}_{nuc-el} + \hat{V}_{el-el} + \hat{V}_{nuc-nuc}$
- Exact diagonalization (FCI): $\dim \hat{H} = 2^M$
  Exponentially hard!
- **Approximate methods:**
  challenged by electronic correlations
  - **Weakly-correlated molecules:**
    electrons can be treated independently
  - **Strongly-correlated molecules:**
    "natural" targets for quantum computers

Quantum chemistry on a quantum computer?

Method	Cost	Correlations
FCI	$\exp(N)$	weak and strong
DFT	N^3	weak
CCSD(T)	N^7	weak to medium
DMRG	$\mathrm{poly}(\chi) = \exp(\mathcal{S})$ (entanglement)	strong

- Weakly-correlated molecules: electrons can be treated independently
- Strongly-correlated molecules: "natural" targets for quantum computers

Chemistry on a QPU: targets?

- **Benchmark:** energy of strongly correlated molecules within chemical accuracy $\eta_{\text{chem}} = 1.6\text{mHa}$
- Fe-S compounds, e.g. **FeMoCo** (113 electrons, 76x2 orbitals) *active site of nitrogenase*
- Proteins: **cytochromes P450** ($\approx$ 60 electrons, 60x2 orbitals) *drug-metabolizing enzymes*

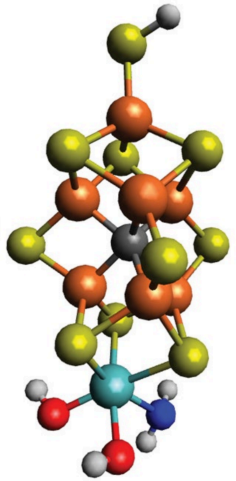

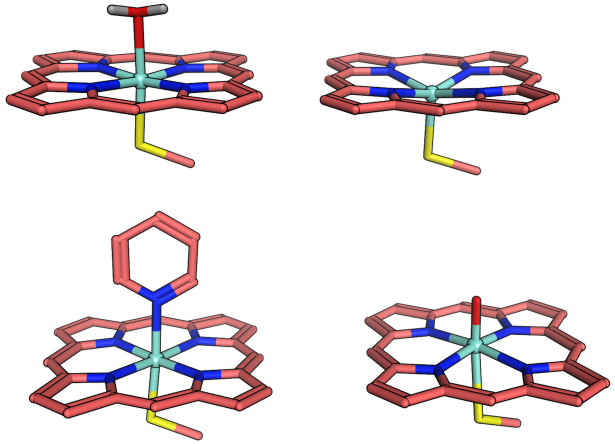

quobly

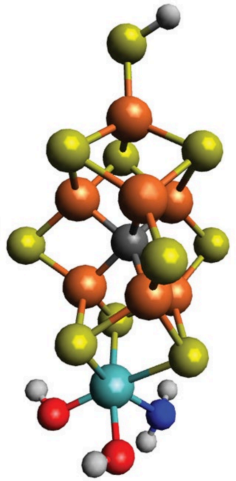

Chemistry on a QPU: targets?

- **FeMoCo:** the "canonical" benchmark
- $N = 113$ electrons, $M = 76 \times 2 = 152$ spin-orbitals
- $\dim \hat{H} = 2^{152} \rightarrow 10^{43}$ bytes
  ("El Capitan" supercomputer: 5 petabytes $= 5 \times 10^{15}$ bytes)
- strongly-correlated

quobly

NISQ?

- FeMoCo: $\dim(\hat{H}) = 2^{152}$
- IBM Heron quantum processor (2023): 156 qubits
- Problem solved?

NISQ algorithms for chemistry

1. Variational Quantum Eigensolver

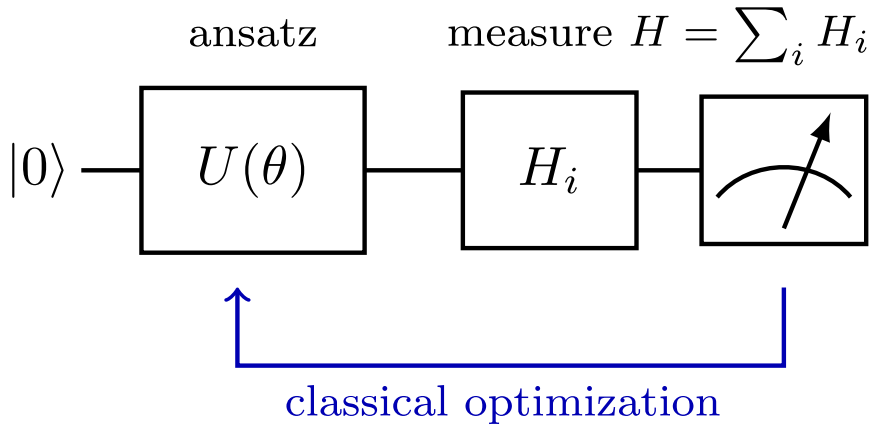

- many measurements, statistical errors
- no guarantee that the optimization converges ("barren plateaus")
- Gonthier *et al.* Phys. Rev. Res. 4 (2022): H_2O (10e) - 2.3 days for a *single* energy
- **TL**, T. Ayral, X. Waintal arXiv:2306.02620: C_6H_6 requires gate error $< 10^{-12}$

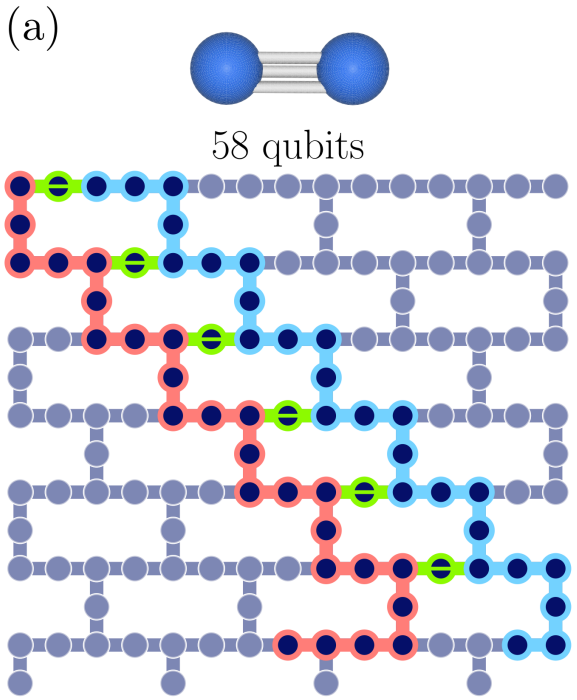

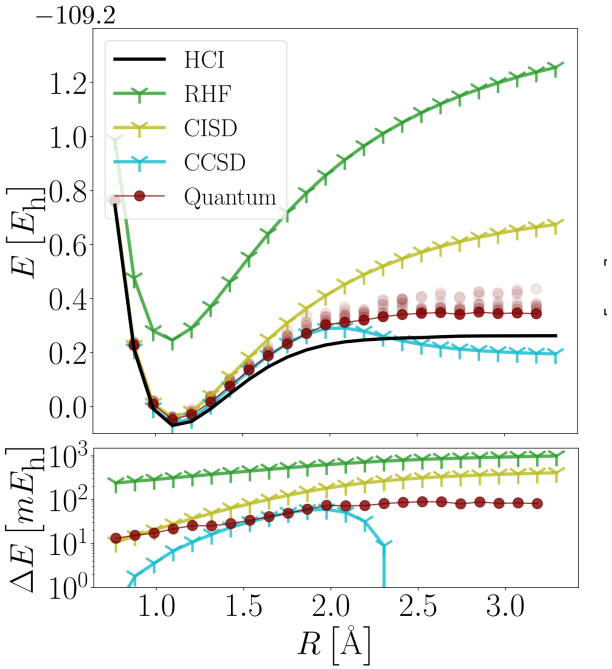

NISQ algorithms for chemistry

2. Sample-Based Quantum Diagonalisation

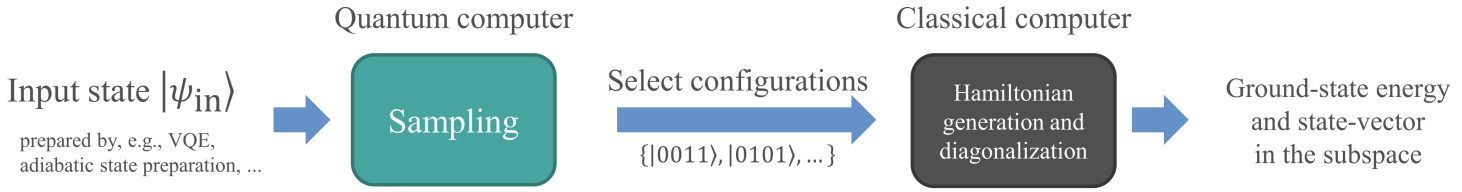

- Idea: Kanno *et al.*, arXiv:2302.11320
- Implementation: IBM group, arXiv:2405.05068
- N_2 (14e, 14o): 10mHa error (chemical accuracy: 1.6mHa)

SQD numerics @ quobly: A. Keita 2025 internship

NISQ algorithms: ansatz and noise

- NISQ algorithms are limited in depth

$$F \approx e^{-\varepsilon N_g}, \qquad \varepsilon = 10^{-3} - 10^{-4}$$

- Noise-induced error:

$$\Delta E = (E_{\text{noise}} - E_{\text{ideal}})(1 - F) \qquad \text{with} \qquad E_{\text{noise}} - E_{\text{ideal}} \gtrsim 1\text{Ha}$$

- Good ansätze are deep
  - Example: chemistry-inspired UCC ansatz
  - Depth: $N_g > N^6$ (triple excitation terms to account for correlations)
  - For N_2 ($N = 14$) chemical accuracy requires $\varepsilon < 10^{-8}$
- Conclusion: **doing chemistry requires fault-tolerant QPUS**

  ***TL**, T. Ayral, X. Waintal arXiv:2306.02620*

Quantum Phase Estimation

- U is a unitary operator related to H, e.g. time evolution $U(t) = \exp(-iHt)$

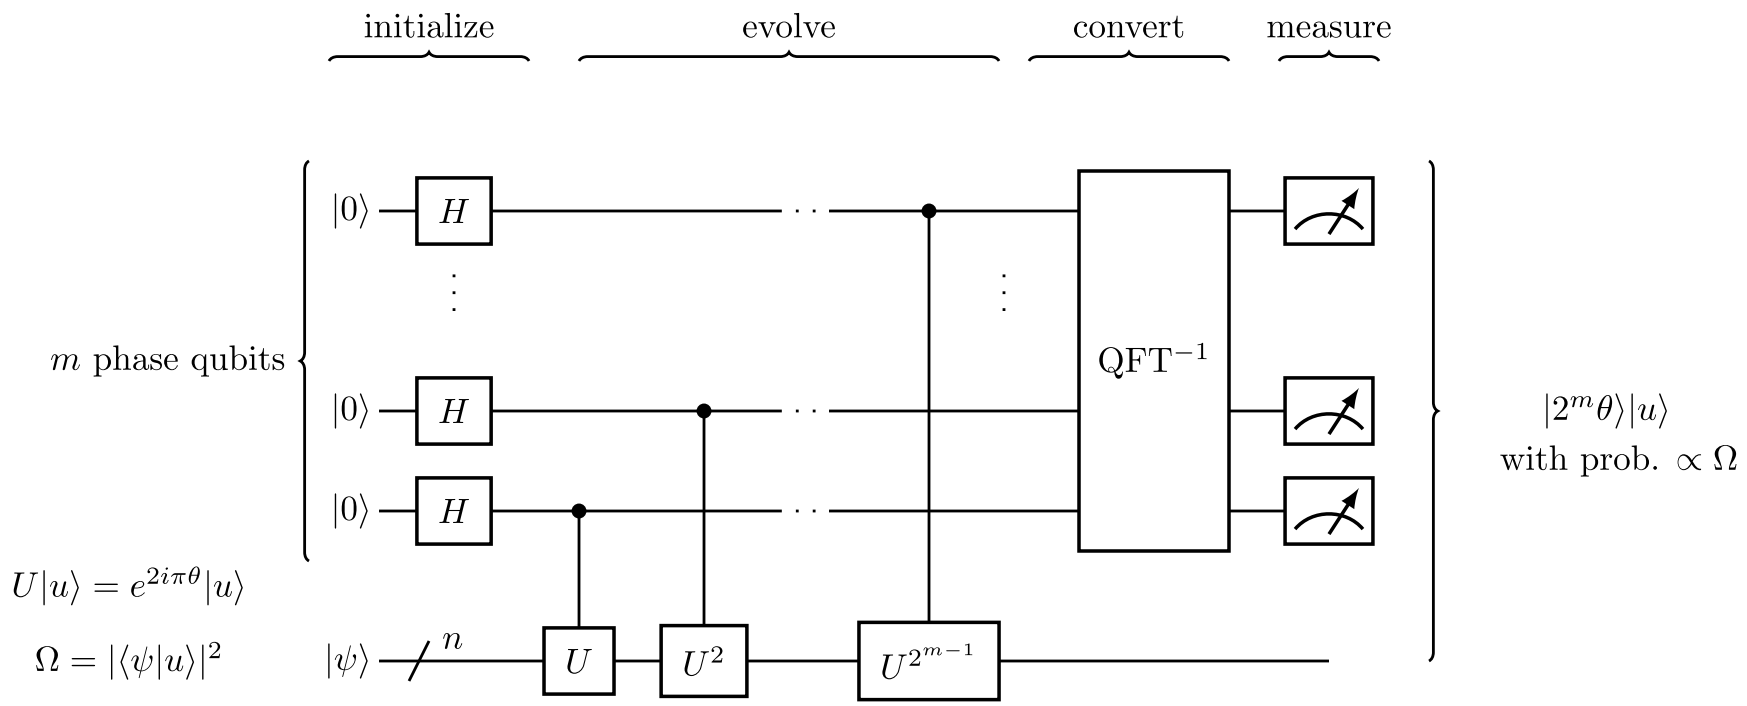

- QPE finds the eigenvalue of U with m bits precision
- Initial state projected on eigenvector with probability $\propto \Omega = |\langle\psi|u\rangle|^2$

qpe_toolbox

quobly

Classical preprocessing

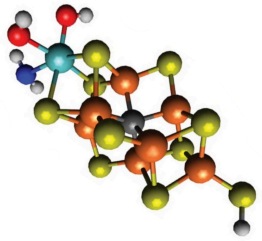

DMRG

MPS to Circuit conversion

Hamiltonian to unitary/gates

Simulator

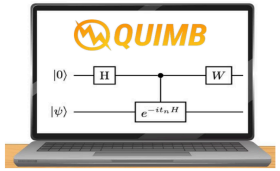

Circuit to Tensor Network

QPU

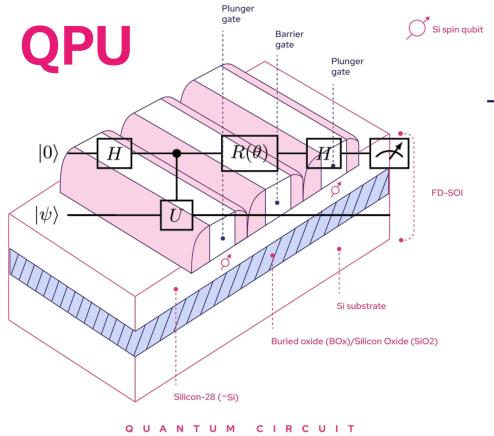

Classical post-processing

0101001010
0101011100
1010111010
1101111011
0101001110
1011010101

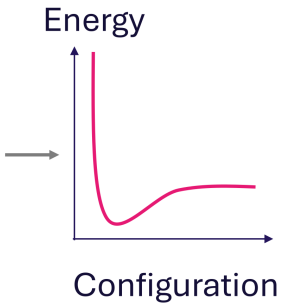

Why tensor networks?

What makes it hard to simulate a QPU?

- "Quantum computers allow to manipulate superpositions of $|0\rangle$ and $|1\rangle$"
- n qubits $\rightarrow$ 2^n-dimensional Hilbert space
- **BUT...** do we need the whole Hilbert space?

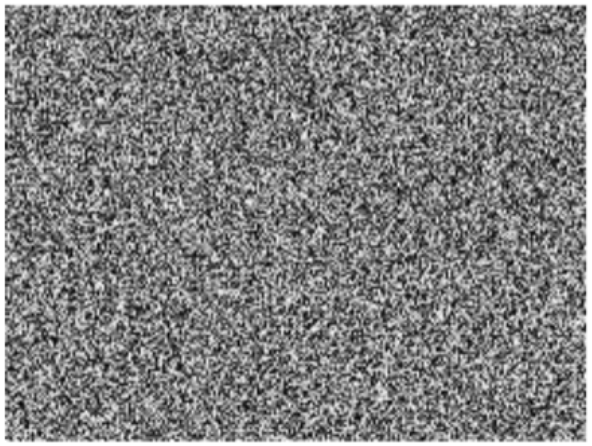

Hard (random)

Easy (structured)

Many-body techniques rely heavily on **structure** (e.g. symmetries)

Tensor Networks

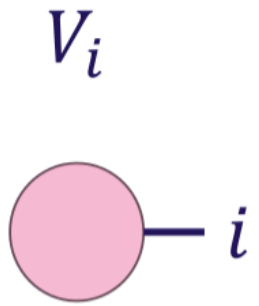

Vector

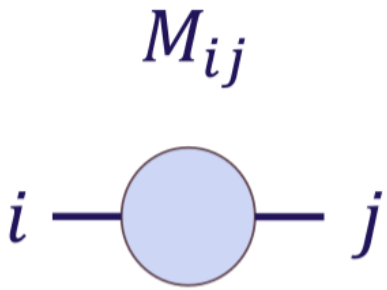

Matrix

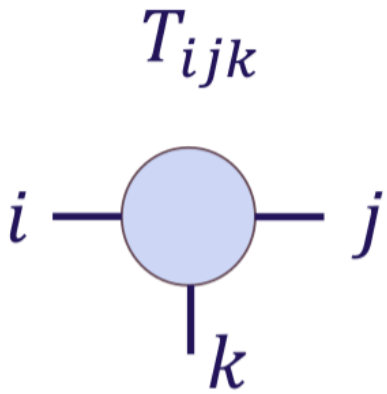

3-leg tensor

$$\sum_i M_{ij} V_i = W_j$$

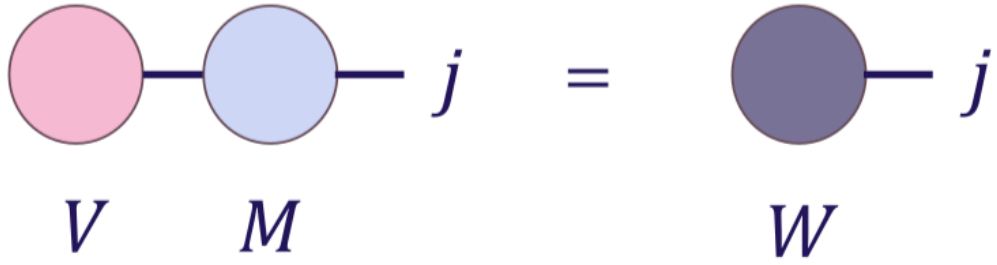

Matrix-Vector multiplication (« contraction »)

Quantum circuits as tensor networks

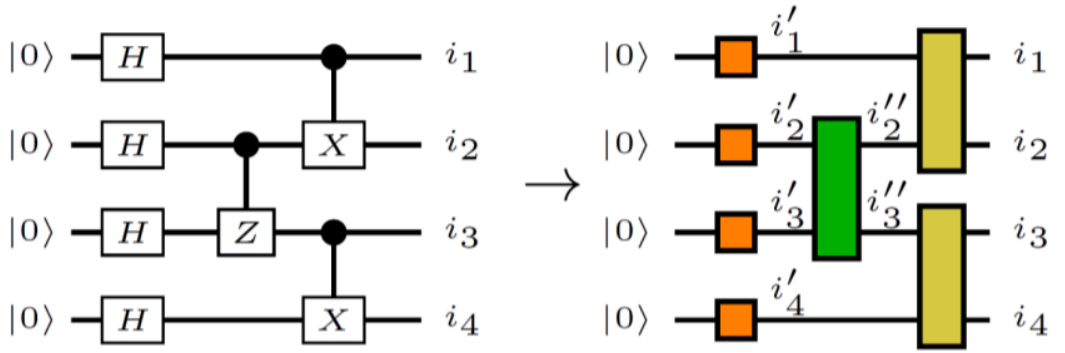

QPE circuit with 4 physical qubits, 2 phase qubits

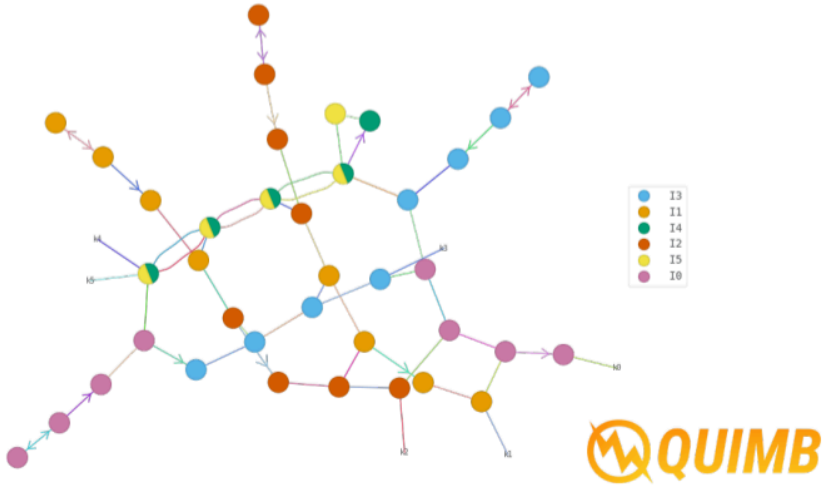

Compression of a quantum state: MPS and SVD

- Matrix Product States

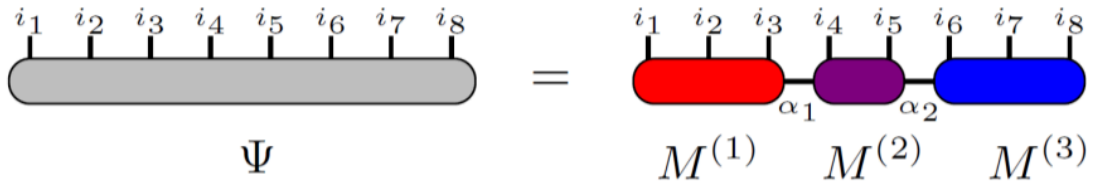

- SVD

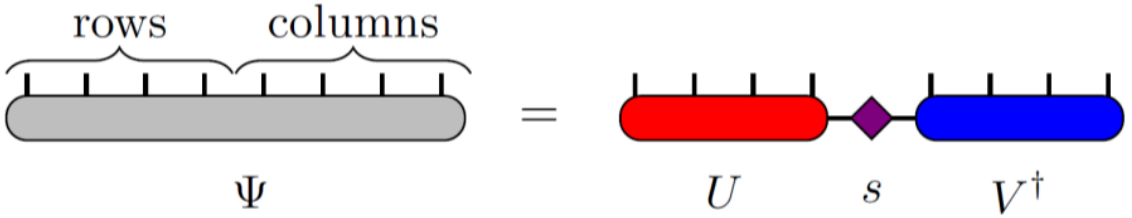

$$\Psi = \sum_{\alpha} U_\alpha s_\alpha V_\alpha^\dagger$$

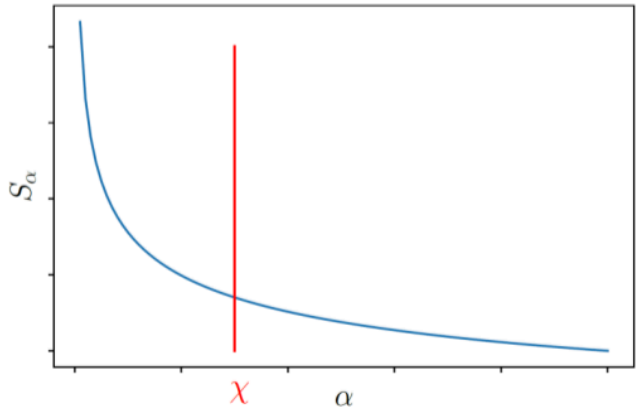

- Compression: bond dimension χ (entanglement)
- Fidelity $F = \left|\langle\Psi|\Psi_{compressed}\rangle\right|^2 = \sum_{\alpha<\chi} s_\alpha^2$

… see examples by C. Ramos Marimon on MPS and hypercontraction

Tensor networks - QUIMB

Quantum
the open journal for quantum science

PAPERS

Hyper-optimized tensor network contraction

Johnnie Gray[1,2] and Stefanos Kourtis[1,3,4]

"in collaboration with NVidia, cotengra & quimb have now been used for a state-of-the-art simulation of the Sycamore chip with cutensor on the Selene supercomputer, producing a sample from a circuit of depth 20 in less than 10 minutes."

QUIMB

NVIDIA

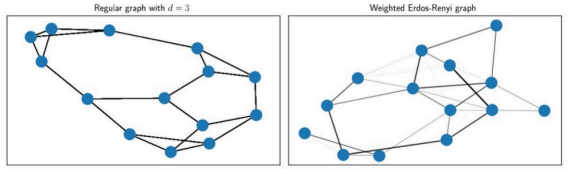

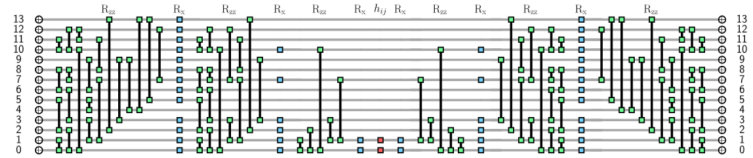

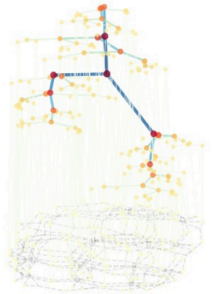

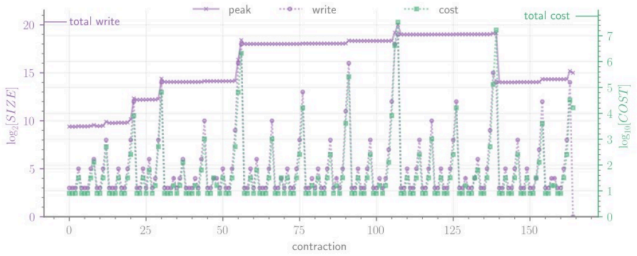

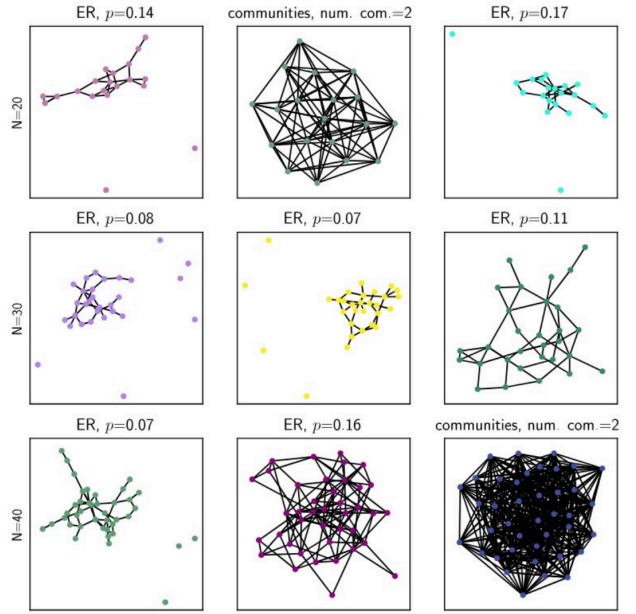

qpe_toolbox, C. Ramos Marimon

Introducing our toolbox

quobly

Objective: Simulate the full QPE pipeline. Get ready for QPE in the long term

Requirements	Features
Include the best classical algorithms	**Component 1:** Classical solution via DMRG
A classical-to-quantum converter	**Component 2:** Tensor-network state preparation
Advanced QPE schemes to reduce quantum resources	**Component 3:** qubitization, single-ancilla
Circuit simulations to quantify the boundary of a quantum advantage, prepare experiments	**Component 4:** Large-scale tensor network simulations and optimization

Wrap-up

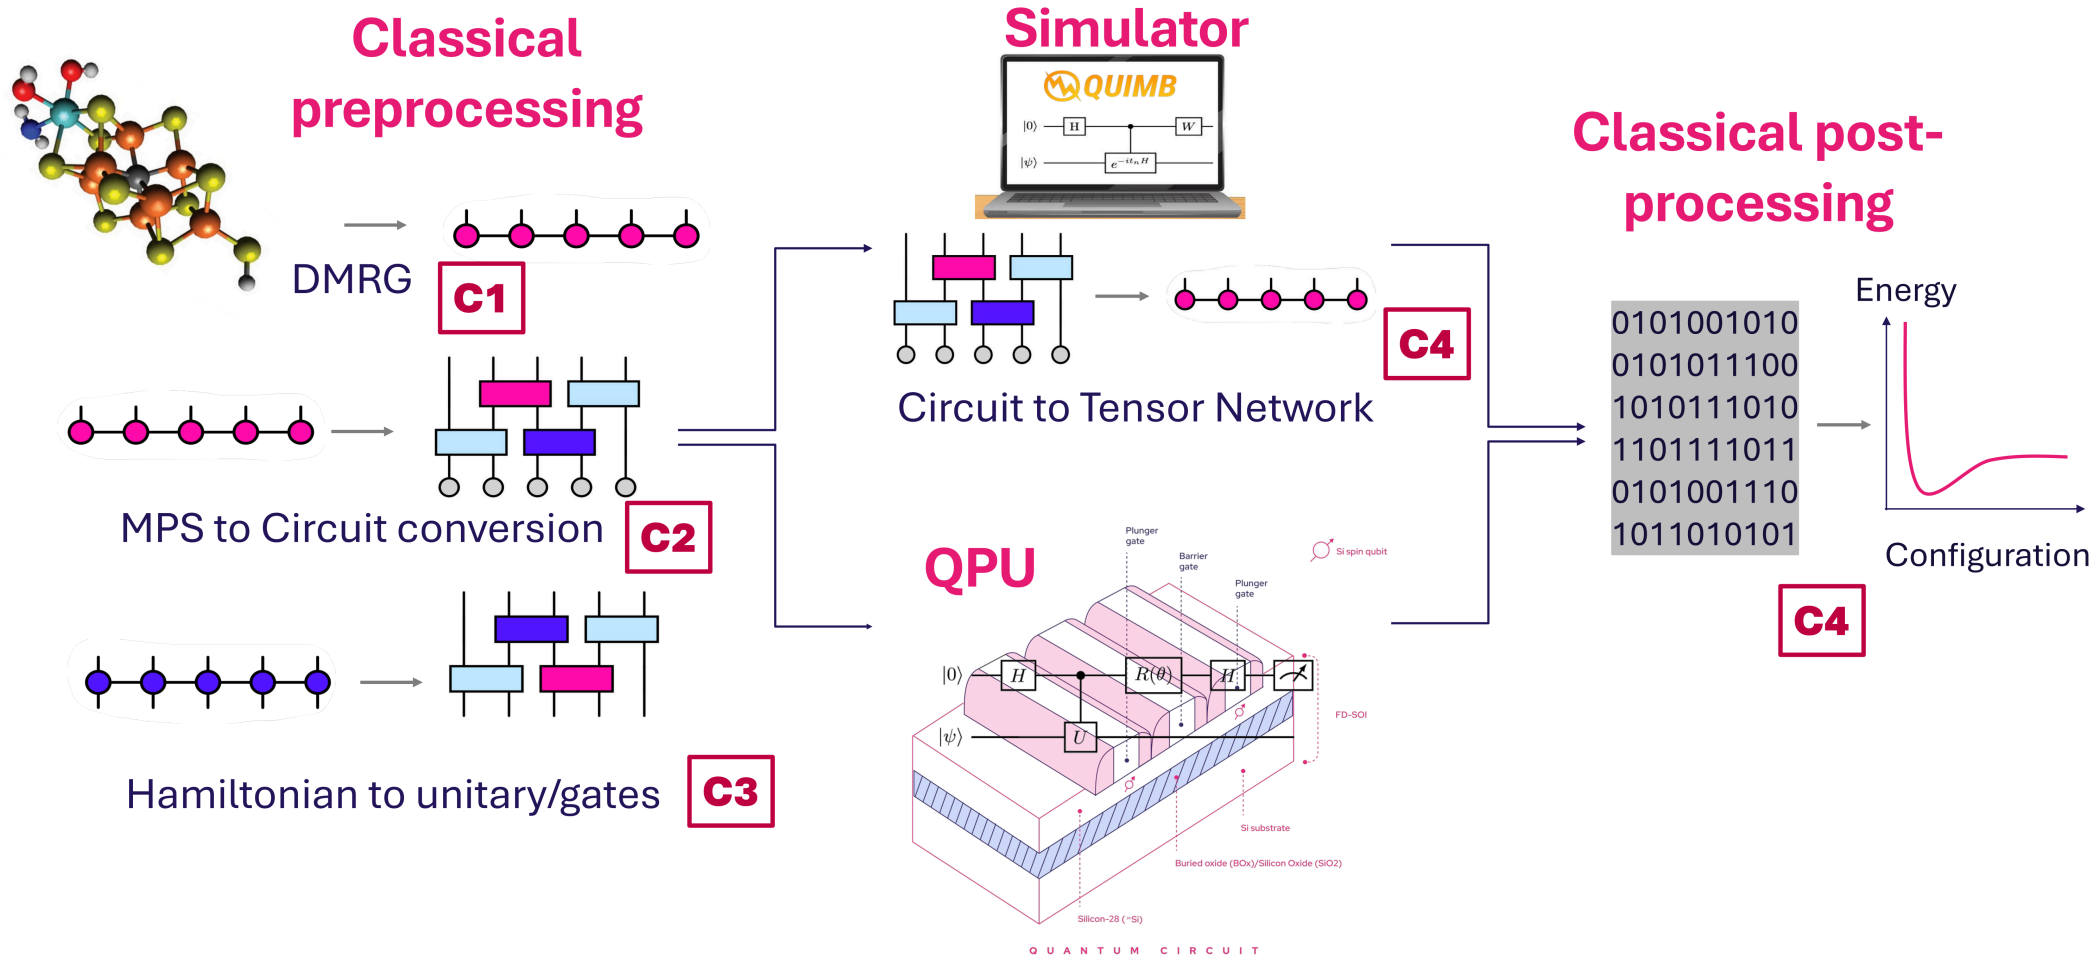
Classical preprocessing
DMRG
C1
MPS to Circuit conversion
C2
Hamiltonian to unitary/gates
C3
Simulator
QUIMB
Circuit to Tensor Network
C4
QPU
Plunger gate
Barrier gate
Plunger gate
Si spin qubit
FD-SOI
Si substrate
Buried oxide (BOx)/Silicon Oxide (SiO2)
Silicon-28 (²⁸Si)
QUANTUM CIRCUIT
Classical post-processing
0101001010
0101011100
1010111010
1101111011
0101001110
1011010101
Energy
Configuration
C4

The "competition" - separate lanes

Components	Competition
C1: Classical solution via DMRG	Block2, Renormalizer... many libraries
C2: Tensor-network state preparation	Quimb, Google, Terra Quantum AG, TU Munich
C3: qubitization and single ancilla scheme	Google (openfermion, Qualtran), Lincoln lab (PyLIQTR), Zapata (Quantum MAMBO), Phasecraft, QC Ware...
C4: Tensor-network simulation and optimization	Flatiron (iTensor), Terra Quantum AG, Multiverse...

qpe_toolbox:contents

```
src/qpe_toolbox/
    |___circuit/
    |    |  controls.py
    |    |  gate_count.py
    |    |  initialization.py
    |    |  parametrized_circuits.py
    |    |  read_write.py
    |
    |___estimation/
    |    |  hadamard_test.py
    |    |  lcu_walk_operator.py
    |    |  phase_estimation.py
    |    |  qft.py
    |    |  robust_phase_estimation.py
    |
    |___hamiltonian/
    |    |  hamiltonian.py
    |    |  pyscf_converter.py
    |
    |___tensor/
    |    |  mpomps_tools.py
```

qpe_toolbox: quick how-to

quobly

1. Chose a system: spin model or molecule. `Hamiltonian` class
2. Prepare guess state: examples `dmrg`, `tn_circuit_optimization`
3. Encode $\hat{H}$ into a unitary:
   3.1. exact time evolution or Trotterization as methods in the `Hamiltonian` class
   3.2. qubitization functions: `lcu_walk_operator`
4. Initialize circuit: `make_circ` from `circuit/initialization`
5. Run QPE: in `estimation/`
   5.1. textbook QPE: `phase_estimation`
   5.2. Robust Phase Estimation (single-ancilla): `robust_phase_estimation`

Phase estimation with qpe_toolbox

`examples/qpe_1_exact_evolution.py`

```python
from qpe_toolbox.circuit import make_circ
from qpe_toolbox.estimation import phase_estimation as pe
from qpe_toolbox.hamiltonian import do_dmrg, heisenberg_hamiltonian

# Define system
n = 4
H = heisenberg_hamiltonian(n)
E0, psi0 = do_dmrg(H) # exact ground state

# guess for energy, search interval
E_target = E0 + 0.2
size_interval = 2.

# initial circuit
n_phase_bits = 5
init_circ = make_circ(n_phase_bits, psi0)

# run qpe (n_steps=0 for exact time evolution)
traces, energy = pe.qpe_energy(H, init_circ, n_steps=0, E_target=E_target, size_interval=size_interval)
```

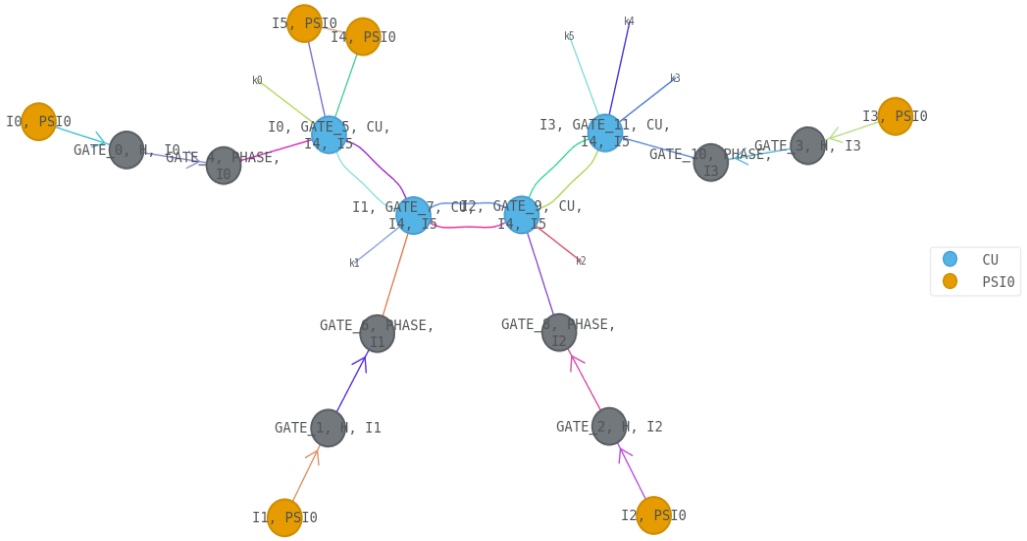

1D Heisenberg with 4 spins

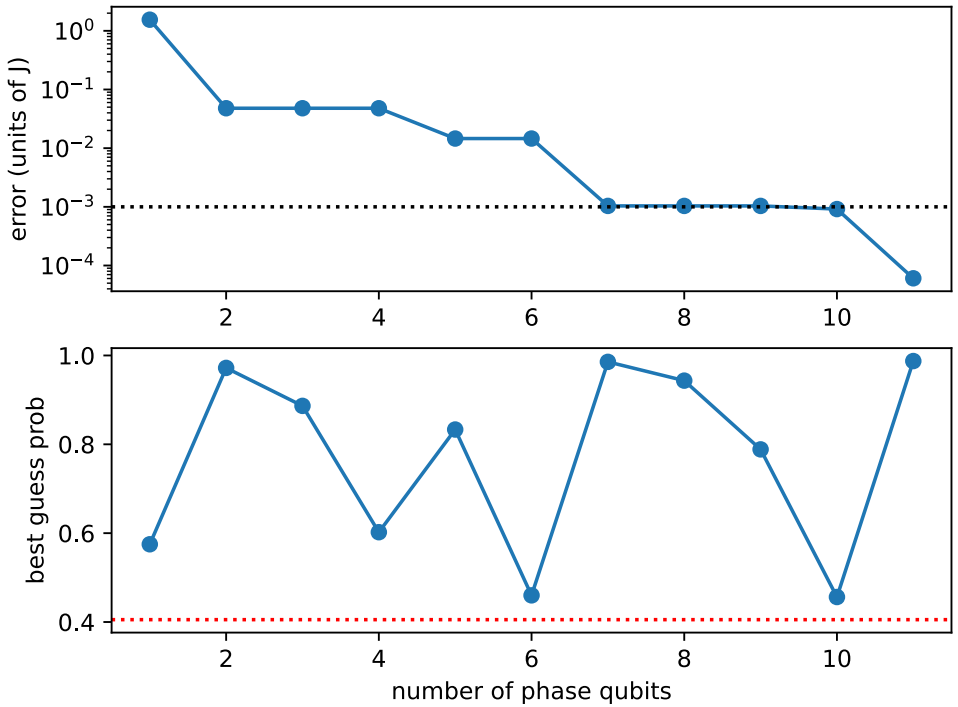

qpe_toolbox

```
bitstr |k>  theta   proba
-----------------------------------------
10011  |19> 0.59375 0.875252366065979
10100  |20> 0.625 0.0548088438808918
10010  |18> 0.5625 0.024422215297818184
10101  |21> 0.65625 0.010917366482317448
10001  |17> 0.53125 0.007346082013100386

energy = -0.7375
E0 = -0.75
size_interval / 2**n_phase_bits = 0.0625
```

Getting ready for early-FTQC: challenges

- First generation FTQC: 100 logical qubits, 10^5 logical gates
- Reduce logical qubit count and Toffoli gate count

Naïve implementation (C3)

$$\hat{V}_{el-el} = \frac{1}{2} \sum_{p,q,r,s=1}^{M} V_{pqrs} \hat{a}_p^\dagger \hat{a}_q \hat{a}_r^\dagger \hat{a}_s$$

- $\mathcal{O}(M^4)$ terms
- Trotterization: $e^{-iH\delta t} \approx e^{-iH_0\delta t} e^{-iH_1\delta t} \ldots \rightarrow \mathcal{O}(M^4)$ gates
- Total time: $t \propto 2^m (E_{\max} - E_{\min})$. Chemical accuracy: $m > 10$.

  $\rightarrow$ **total cost:** $\mathcal{O}(2^m (E_{\max} - E_{\min}) M^4)$
- $E_{\max} - E_{\min}$ increases with M
- Additional costs: statistical error, state preparation...

Trotter error

- First-order Trotter formula: $e^{-i(H_1+H_2)t} = (e^{-iH_1\delta t}e^{-iH_2\delta t})^r + \mathcal{O}(t^2/r)$
- Second-order: error $\mathcal{O}(t^3/r^2)$

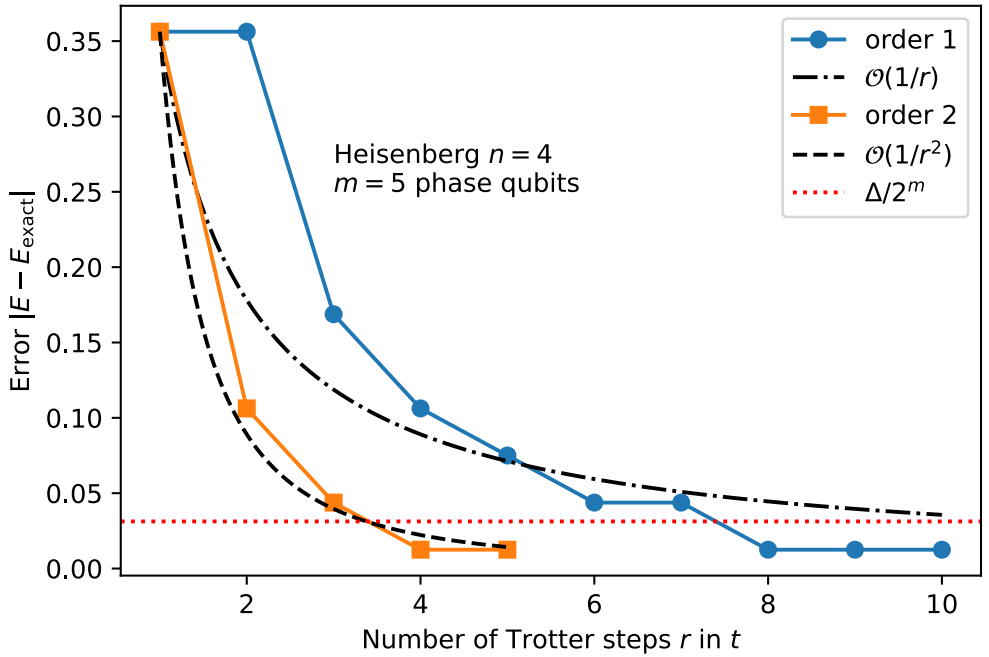

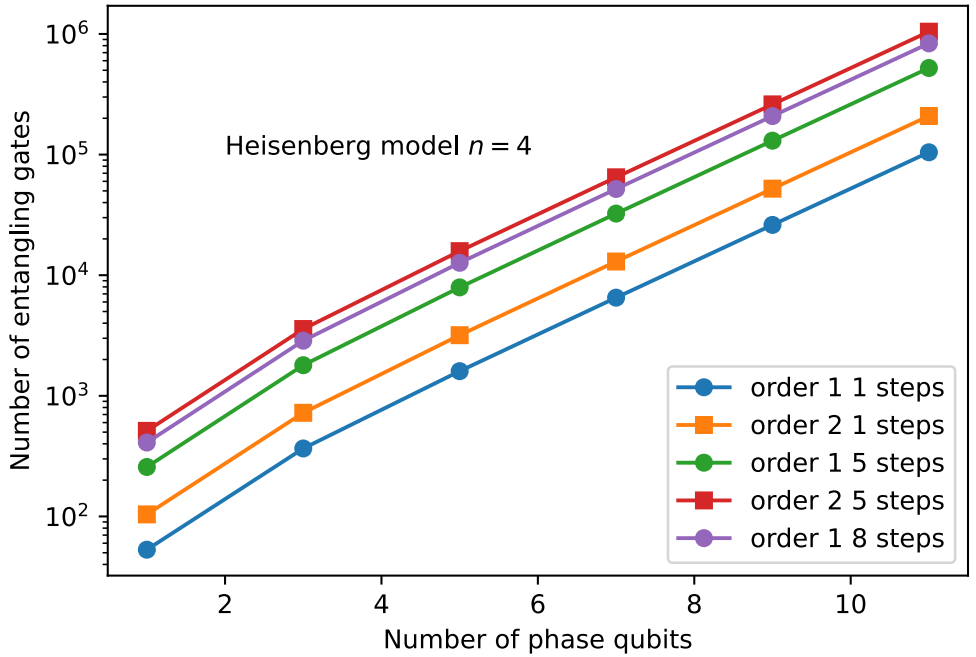

Progress towards early-FTQC

Year	Reference	Primary algorithmic innovation	Space complexity	Toffoli or T complexity
2005	Aspuru-Guzik *et al.* [4]	First algorithm (no compilation or bounds)	$\mathcal{O}(N)$	$\mathcal{O}[\mathrm{poly}(N/\epsilon)]$
2010	Whitfield *et al.* [11]	First compilation (no Trotter bounds)	$\mathcal{O}(N)$	$\mathcal{O}[\mathrm{poly}(N/\epsilon)]$
2012	Seeley *et al.* [44]	Use of Bravyi-Kitaev transformation	$\mathcal{O}(N)$	$\mathcal{O}[\mathrm{poly}(N/\epsilon)]$
2013	Wecker *et al.* [45]	First chemistry-specific Trotter bounds	$\mathcal{O}(N)$	$\widetilde{\mathcal{O}}(N^{10}/\epsilon^{3/2})$
2013	Toloui *et al.* [46]	Use of first quantization	$\mathcal{O}(\eta \log N)$	$\widetilde{\mathcal{O}}(\eta^2 N^8/\epsilon^{3/2})$
2014	Hastings *et al.* [47]	Better compilation and multiresolution Trotter	$\mathcal{O}(N)$	$\widetilde{\mathcal{O}}(N^8/\epsilon^{3/2})$
2014	Poulin *et al.* [48]	Tighter Trotter bounds and ordering	$\mathcal{O}(N)$	$\widetilde{\mathcal{O}}(N^6/\epsilon^{3/2})$
2014	McClean *et al.* [25]	Exploiting Hamiltonian sparsity with Trotter	$\mathcal{O}(N)$	$\widetilde{\mathcal{O}}(N^4 S/\epsilon^{3/2})$
2014	Babbush *et al.* [49]	Tighter system-specific Trotter bounds	$\mathcal{O}(N)$	$\widetilde{\mathcal{O}}(N^2 S/\epsilon^{3/2})$
2015	Babbush *et al.* [50]	Use of Taylor series (database method)	$\mathcal{O}(N)$	$\widetilde{\mathcal{O}}(N^4 \lambda_V/\epsilon)$
2015	Babbush *et al.* [50]	Use of Taylor series (on-the-fly method)	$\mathcal{O}(N)$	$\widetilde{\mathcal{O}}(N^5/\epsilon)$
2015	Babbush *et al.* [51]	Use of Taylor series with first quantization	$\mathcal{O}(\eta \log N)$	$\widetilde{\mathcal{O}}(\eta^2 N^3/\epsilon)$
2016	Reiher *et al.* [23]	First T count and tighter Trotter bounds	$\mathcal{O}(N)$	$\widetilde{\mathcal{O}}(N^2 S/\epsilon^{3/2})$
2018	Motta *et al.* [30]	Use of low-rank factorization with Trotter	$\mathcal{O}(N)$	$\widetilde{\mathcal{O}}(N^4 \Xi/\epsilon^{3/2})$
2018	Campbell [52]	Use of randomized compiling with Trotter	$\mathcal{O}(N)$	$\widetilde{\mathcal{O}}(\lambda_V^2/\epsilon^2)$
2019	Berry *et al.* [9]	Use of qubitization (sparse method)	$\widetilde{\mathcal{O}}(N+\sqrt{S})$	$\widetilde{\mathcal{O}}[(N+\sqrt{S})\lambda_V/\epsilon]$
2019	Berry *et al.* [9]	Use of qubitization (single factorization)	$\widetilde{\mathcal{O}}(N^{3/2})$	$\widetilde{\mathcal{O}}(N^{3/2}\lambda_{\mathrm{SF}}/\epsilon)$
2019	Kivlichan *et al.* [53]	Better randomized compiled phase estimation	$\mathcal{O}(N)$	$\widetilde{\mathcal{O}}(\lambda_V^2/\epsilon^2)$
2020	von Burg *et al.* [10]	Use of qubitization (double factorization)	$\widetilde{\mathcal{O}}(N\sqrt{\Xi})$	$\widetilde{\mathcal{O}}(N\lambda_{\mathrm{DF}}\sqrt{\Xi}/\epsilon)$
2020	Present work	Use of tensor hypercontraction	$\widetilde{\mathcal{O}}(N)$	$\widetilde{\mathcal{O}}(N\lambda_\zeta/\epsilon)$

1. Costs scalings - from Lee *et al.* (Google), PRXQuantum 2 2021
   ($S \sim \mathcal{O}(N^4)$ and $\lambda \sim \mathcal{O}(N) - \mathcal{O}(N^3)$)

Progress towards early-FTQC

Year	Innovation	FeMoco-54 [37]			FeMoco-76 [38]		
		Qubits	Toffolis	Reference	Qubits	Toffolis	Reference
2017	First resource estimate by Trotterization [37]	111	5.0×10^{13}	[37]	-	-	-
2019	Qubitization of Single-Factorization [17]	3320	9.5×10^{10}	[7]	3628	1.2×10^{11}	[7]
2020	Qubitization of Double-Factorization (DF) [9]	3600	2.3×10^{10}	[9]	6404	5.3×10^{10}	[7]
2020	Tensor-Hyper-Contraction (THC) [7]	2142	5.3×10^{9}	[7]	2196	3.2×10^{10}	[7]
2024	Symmetry compression of DF [39]	1994	2.6×10^{9}	[39]	-	-	-
2025	Symmetry compression of THC [8]	-	-	-	1512	4.3×10^{9}	[8]
This work	Spectrum amplification & DFTHC	1137	3.41×10^{8}		1459	9.99×10^{8}	
Improvement of this work over [39] and [8][a]		1.8×	7.0×		1.0 ×	4.3 ×	

2. Resources for FeMoCo - from Low *et al.* (Google team), PRX 15 2025

quobly

Progress towards early-FTQC: sum-up

- Between 2017 and 2025:
  - Scalings reduced from $\mathcal{O}(N^6/\epsilon^{3/2})$ to $\mathcal{O}(N^4/\epsilon)$
  - FeMoCo costs reduced from $5 \cdot 10^{13}$ to $3.4 \cdot 10^8$ Toffolis (increasing qubits count from 111 to 1137)
- an order of magnitude reduction every two years for Toffoli count!
- still some way to go for early-FTQC (100 logical qubits, 10^5 Toffolis)

Advanced QPE schemes for early-FTQC (C3)

Two classes of advanced variations on QPE:

- **Qubitization:** clever heuristics for unitary encoding of H
  (quantum linear algebra)
  See: Google team, many others...
- **Single-ancilla schemes:** "space-time" tradeoff
  less ancilla qubits - many circuit repetitions - classical post-measurement signal processing
  See: Lin Lin, Aram H. Harrow...

Advanced QPE schemes for early-FTQC

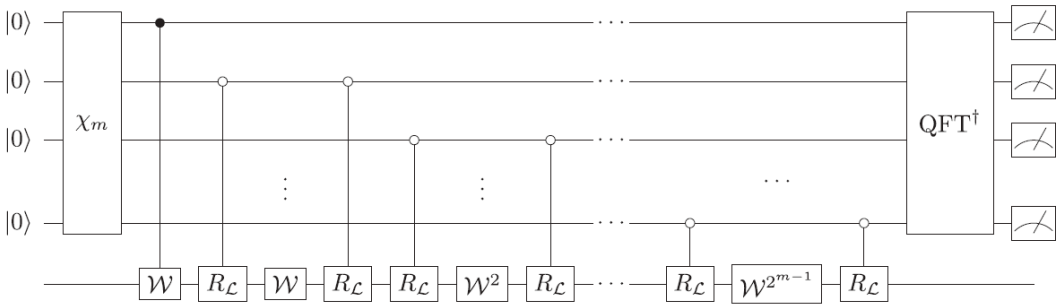

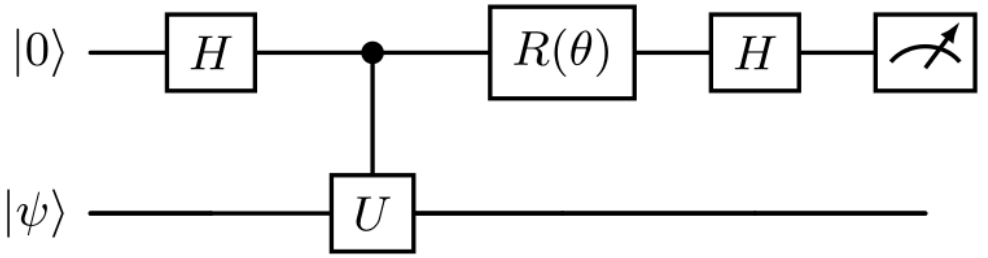

- **In the toolbox** we introduce (**C3**):
  - i. **Linear Combination of Unitaries** and **Walk operator** for Qubitization
  - ii. The single-ancilla **Robust Phase Estimation** algorithm

Advanced QPE schemes I - Qubitization

quobly

see `lcu_walk_operator` example

(inspired from Lin Lin's lecture notes and Babbush, PRX **8**, 041015 (2018))

- Linear Combination of Unitaries: write $\hat{H} = \sum_{\ell=0}^{L-1} w_\ell H_\ell$, with $w_\ell > 0$ and $H_\ell^\dagger H_\ell = \mathbb{1}$.
- **1-norm** of coefficients: $\lambda = \sum_\ell |w_\ell|$
- **Walk operator:** $\mathcal{W}[\hat{H}] = e^{-i \arccos(\hat{H}/\lambda)}$
- **Exact** unitary encoding of the spectrum (no Trotter error)
- Apply QPE on $\mathcal{W}[\hat{H}]$ to get $\arccos(E/\lambda)$
- Total cost ($C_\mathcal{W} \equiv$ qubitization cost):

$$\mathcal{O}(C_\mathcal{W} \frac{\lambda}{\Omega \varepsilon})$$

Qubitization circuits: the oracles

Babbush (Google), PRX **8**, 041015 (2018)

$$H = \sum_{\ell=0}^{L-1} w_\ell H_\ell, \qquad \lambda = \sum_\ell w_\ell$$

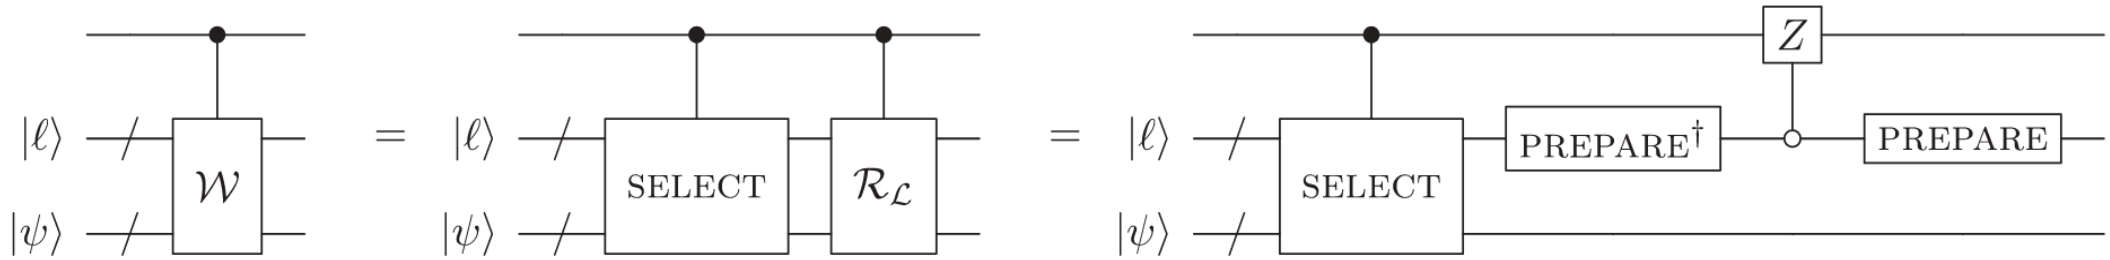

$$\text{PREPARE} \equiv \sum_\ell \sqrt{\frac{w_\ell}{\lambda}} |\ell\rangle\langle 0| \;\rightarrow\; \text{PREPARE}|0\rangle = \sum_\ell \sqrt{\frac{w_\ell}{\lambda}} |\ell\rangle \equiv |\mathcal{L}\rangle$$

$$\text{SELECT} \equiv \sum_\ell |\ell\rangle\langle\ell| \otimes H_\ell$$

Qubitization circuits in the qpe_toolbox (C3)

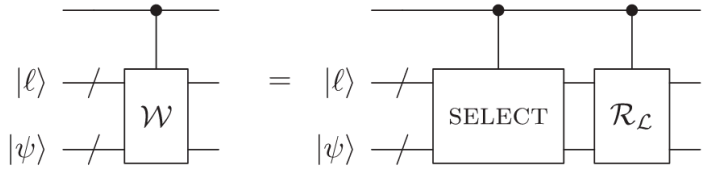

- SELECT and PREPARE circuits are **hard**. Involve **multi-controlled qubits**, non-trivial subroutines for classical data loading (**"QROM"**)
- Current implementation (2025):
  - SELECT with multi-controlled gates
  - $\mathcal{R}_L$ as a Matrix Product Operator (no explicit gate decomposition)
- Future prospects (**Work in progress 2026**):
  get a gate decomposition for PREPARE via explicit circuit or MPO-to-circuit encoding

Qubitization - cost reduction

quobly

- Reduce the Hamiltonian 1-norm λ
- Dominant contribution: two-body term $\hat{V}_{el-el} = \frac{1}{2} \sum_{p,q,r,s=1}^{M} V_{pqrs} \hat{a}_p^\dagger \hat{a}_q \hat{a}_r^\dagger \hat{a}_s$
- Naïve LCU: $\lambda = \mathcal{O}(M^4)$
- Different **contraction schemes** for the V_{pqrs} tensor
  - Berry *et al.*, Quantum **3**, 208 (2019) - Single-Factorization
  - Von Burg *et al.*, PRR **3**, 033055 (2021) - Double-Factorization
  - Lee *et al.*, PRX Quantum **2**, 030305 (2021) Tensor Hypercontraction
- SF and DF Hamiltonians implemented by C. Ku and K. Chen (Foxconn)
- **Work in Progress (2026)**: get SF and DF circuits for $\mathcal{W}$

Advanced QPE schemes II - Robust Phase Estimation

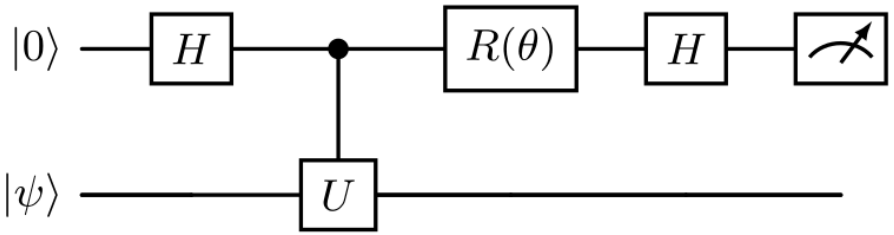

- Inspired from J. Günther *et al.* (Harrow's group) arXiv:2503.05647
- **Hadamard test:** measure $\langle\psi|U|\psi\rangle$, $U = \exp(-iHt)$, for $t = 2^k$, $k = 1, \ldots, R$
- Find ϕ_k, an estimate of $2^k E \bmod 2\pi$ within measurement and Trotter error
- For each k, select the estimate closest to the previous one: improve precision on E by one bit
- Measure E up to $O(1/2^R)$ error with R circuits

Robust Phase Estimation with the qpe_toolbox (C3)

quobly

Example `robust_phase_estimation`

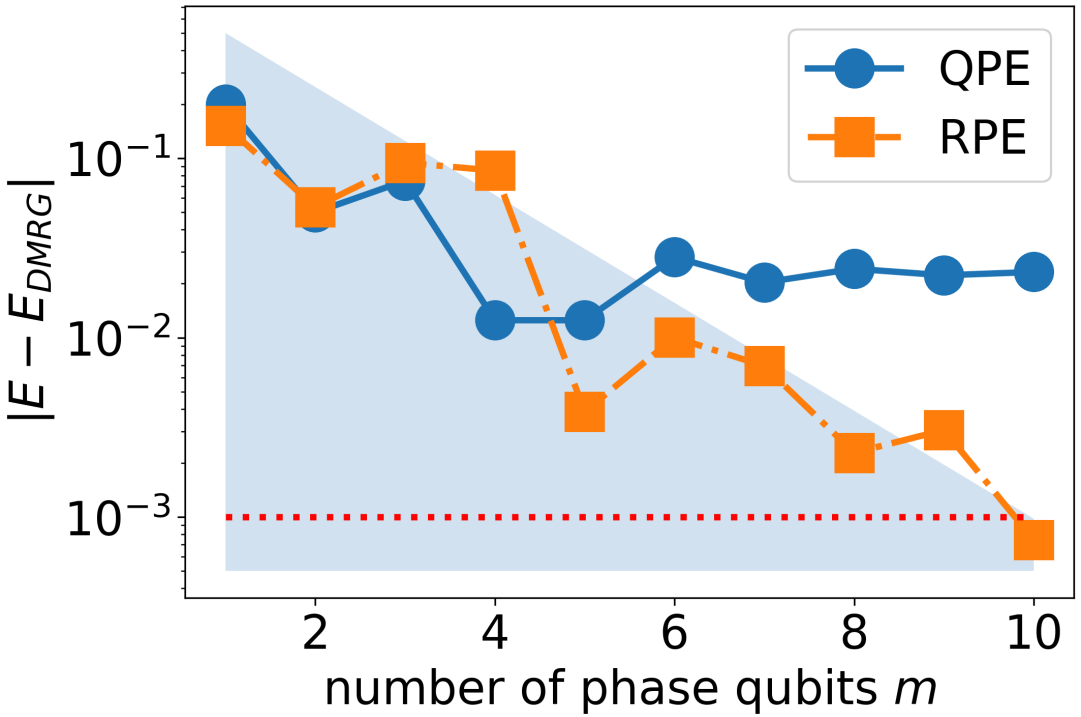

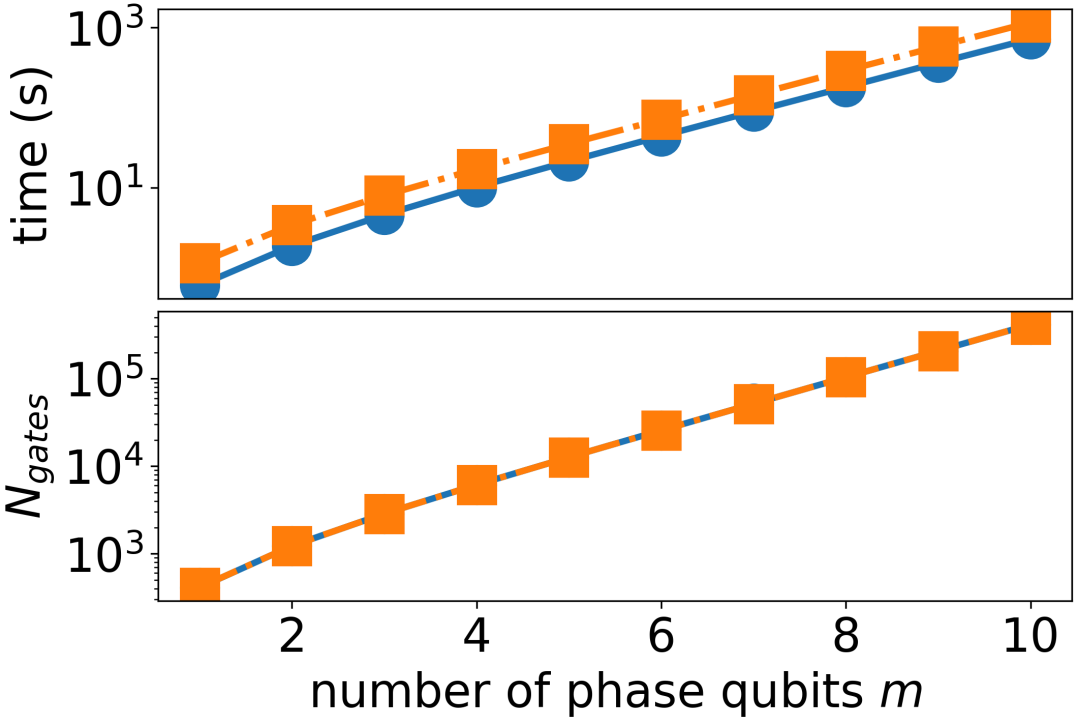

- Heisenberg model with $n = 4$ spins, 2^{d}-order Trotter with $dt = t/4$
- Textbook QPE limited by Trotter error. RPE is not

State preparation - challenge (C1)

- Success probability $\propto \Omega = |\langle\psi|\psi_0\rangle|^2 \rightarrow$ cost $\propto 1/\Omega$
- **Orthogonality catastrophe:** Generically, $\Omega \propto \exp(-N)$

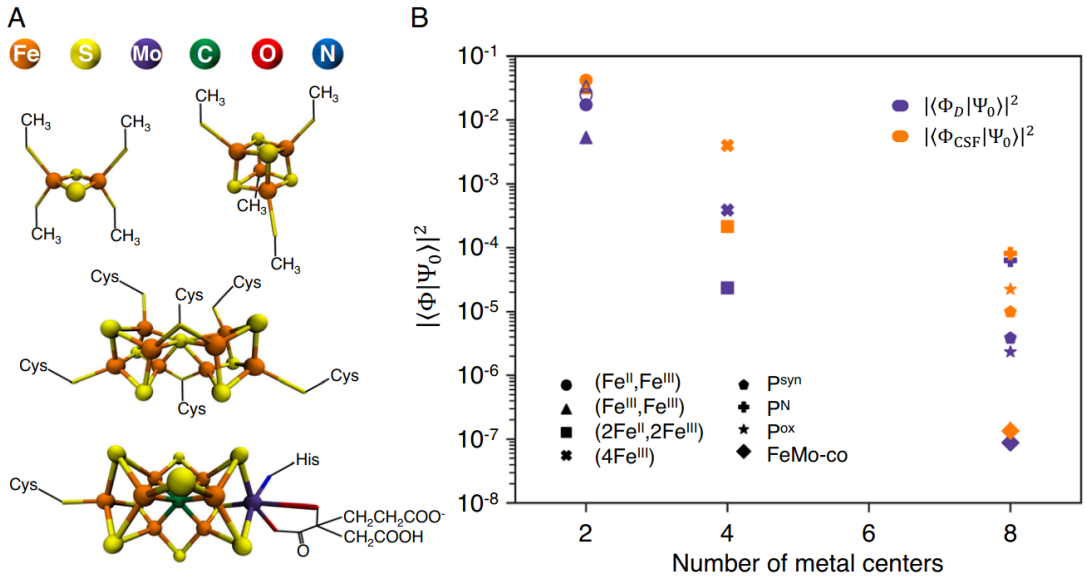

"Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry", Lee *et al.*, Nature Communications **14**, 1952 (2023)

State preparation - MPS (C1-C2)

VOLUME 69, NUMBER 19 PHYSICAL REVIEW LETTERS 9 NOVEMBER 1992

Density Matrix Formulation for Quantum Renormalization Groups

Steven R. White

The density-matrix renormalization group in the age of matrix product states

Ulrich Schollwöck *

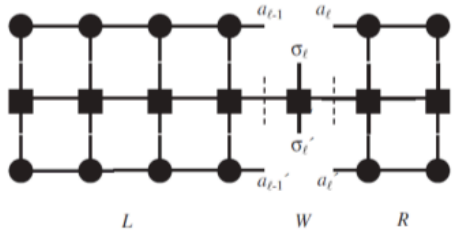

Rapid initial state preparation for the quantum simulation of strongly correlated molecules

Dominic W. Berry,[1,*] Yu Tong,[2,3,4] Tanuj Khattar,[5] Alec White,[5] Tae In Kim,[6] Sergio Boixo,[5] Lin Lin,[7] Seunghoon Lee,[6,8] Garnet Kin-Lic Chan,[8] Ryan Babbush,[5] and Nicholas C. Rubin[5,†]

System	Estimated $\lvert\langle MPS \lvert \psi_0 \rangle\rvert$	Bond Dimension	Spatial Orbitals	MPS Toffolis	qubits
$Fe_2(III)Fe_2(II)$	0.88	1000	36	42 200 000	359
$Fe_4(III)$	0.92	1000	36	42 200 000	359
FeMoco [43] MPS1	0.99	6000	76	1 360 000 000	833
FeMoco [43] MPS2	0.95	4000	76	733 000 000	682
FeMoco [43] MPS3	0.98	4000	76	733 000 000	682

DMRG might be the best candidate for state preparation

- First stage: small-scale DMRG (H_2, O_2, H_2O) in `dmrg` example (**C1**)
- Prospects 2026 (O. Gauthé): state-of-the-art DMRG for chemistry (**C1**), MPS circuit encoding (**C2**)

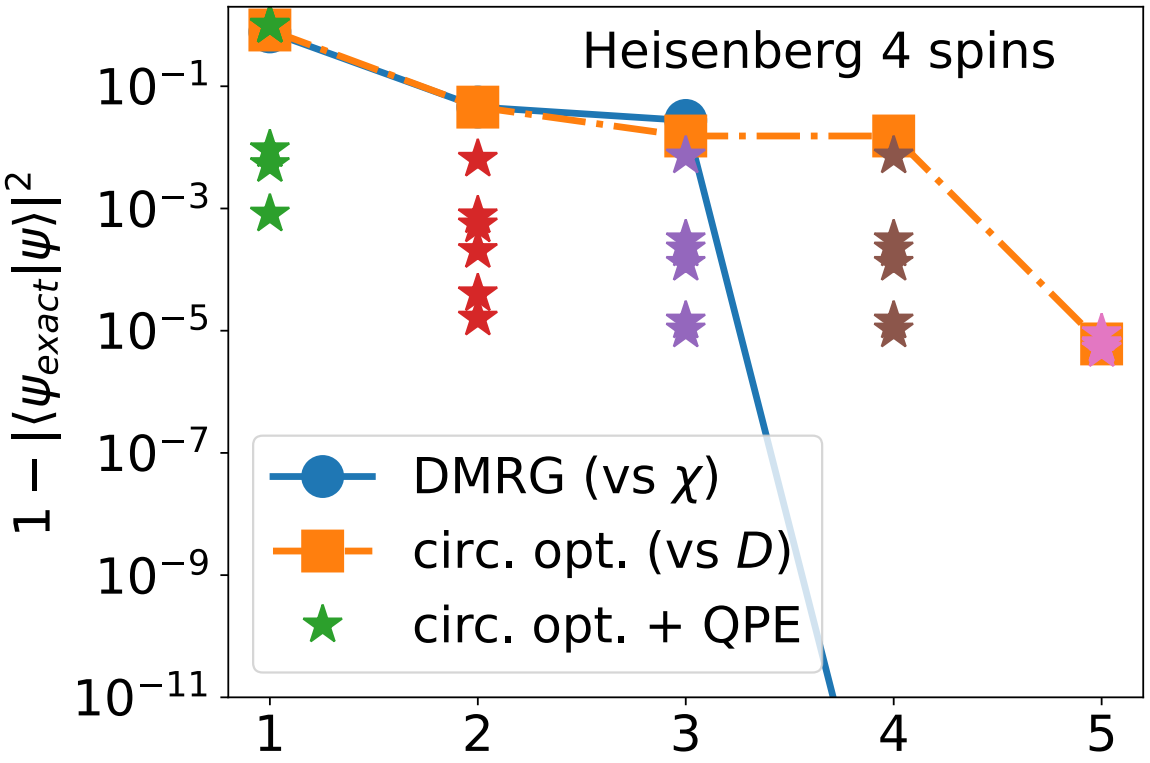

State preparation

- circuit optimization

`tn_circuit_optimization` example (**C1-C2**)

- DMRG: find the MPS wavefunction which minimizes the energy
- We take a shortcut and optimize the parameters of a simple ansatz circuit to minimize the energy
- QPE: project onto g.s. $|\psi_{\text{exact}}\rangle$

quobly

Summary

- **We develop an open-source numerical toolbox**
  - Based on advanced tensor-network techniques
  - Synthetize and explore circuits in a practical sense
- **First stage release - "My pedagogical toolbox" (coming soon)**
  - Initialization: DMRG or circuit optimization
  - Hamiltonian encoding: Trotterization or naïve qubitization
  - Phase estimation: textbook QPE or Robust Phase Estimation

Future release - "My TN optimized toolbox" (2026)

- Initialization: state-of-the-art DMRG and MPS to circuit encoding (O. Gauthé)
- Hamiltonian encoding:
  - Single-Factorization, Double-Factorization... and advanced LCU/qubitization circuits (Foxconn, TL)
  - MPO to circuit encoding
- Circuit simulation (C. Ramos Marimon):
  - Pauli truncation
  - State-of-the-art tensor network techniques e.g. Belief Propagation...

Release plan

Components	First release - "My pedagogical toolbox"	Next release - "My TN optimized toolbox"
C1: DMRG	Small scale molecules	State-of-the art
C2: TN state preparation	"Perfect" (no gates)	TN optimization
C3: advanced QPE	LCU via Pauli strings, RPE	LCU with SF and DF
C4: TN simulation and optimization	Introduce Quimb features	Large-scale simulations

QPE Toolbox

quobly

Classical preprocessing

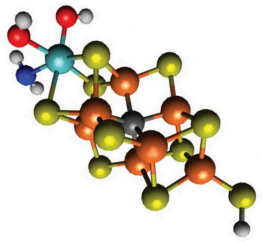

DMRG **C1**

MPS to Circuit conversion **C2**

Hamiltonian to unitary/gates **C3**

Simulator

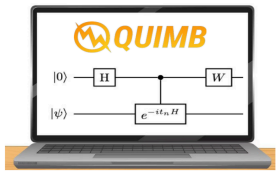

Circuit to Tensor Network **C4**

QPU

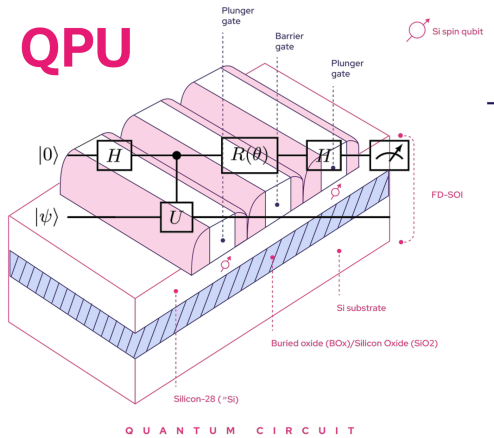

QUANTUM CIRCUIT

Classical post-processing

0101001010
0101011100
1010111010
1101111011
0101001110
1011010101

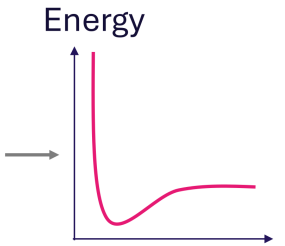

C4