Quantum simulation of atomic nuclei using nearly-optimal methods

Speaker: Christopher Kane[1,2]
Collaborators: Zhi-Yuan Wei[2,3], Andrew Childs[2,4], Alexey Gorshkov[2,3], Zohreh Davoudi[1,2]

Appearing on arXiv soon

INT Workshop: Nuclear Hamiltonians for Advancing Nuclear Physics and Beyond

Date: May 14, 2026

[1]Maryland Center for Fundamental Physics and Department of Physics
[2]Joint Center for Quantum Information and Computer Science
[3]Joint Quantum Institute, NIST/UMD
[4]Department of Computer Science and Institute for Advanced Computer Studies

Discussion talking points

What are concrete problems where quantum computers can contribute to nuclear physics?

Near(est) term → IBM roadmap 2030:

- 200 logical qubits and 10^8 gates

Far term → IBM roadmap 2033+:

- 2000 logical qubits and 10^{10} gates

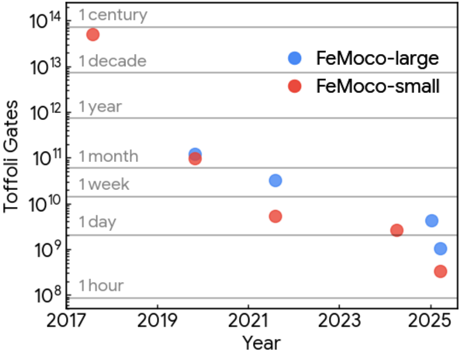

Taken from [R. Babbush et al., 2511.09124]

Where many-body methods stand (my understanding)

Goal: structure and dynamics of nuclei with *ab initio* methods

Static properties (binding energies, radii, ground-state structure)

- Classical many-body methods have long history of success
- Struggle with deformed nuclei/collective modes

Dynamical response (e.g. to weak / EM currents)

- Active and rapidly developing frontier

Can be difficult to quantify “many-body errors”

Concrete frontier: input for DUNE

- Long-baseline ν oscillation experiment, ^{40}Ar target
- Requires input of ν-^{40}Ar cross sections → nuclear response functions across a wide kinematic range

Where quantum computers fit

For static, ground-state properties of quantum chemistry Hamiltonians:

- No expected exponential quantum advantage over classical methods, maybe polynomial?
  [S. Lee *et al.*, Nat. Commun. (2023), arXiv:2208.02199]

For dynamics and precision:

1. Real-time evolution

- Provably efficient simulation of e^{-iHt} even at long times, possible exponential advantage

2. Rigorous error control (the underappreciated case)

- Specify target error ϵ in advance → algorithm cost obeys

$$\text{Cost} \leq f(t, \epsilon, \|H\|, \dots)$$

$\Rightarrow$ in principle, no unquantified many-body systematics

Our work

End-to-end cost analysis performed for three nuclear lattice EFTs:

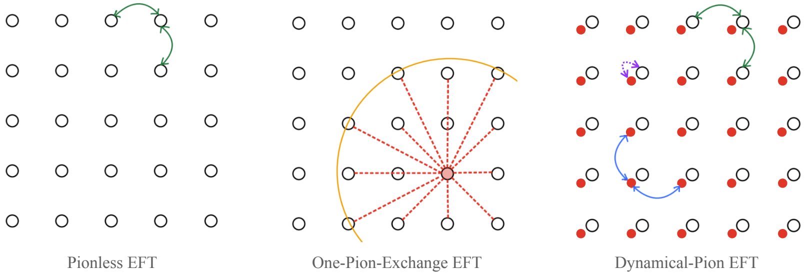

Given: lattice size L, particle number η, error budget ϵ, developed efficient circuits for

- Time-evolution
- Initial state preparation via eigenstate filtering
- Energy estimation via quantum phase estimation

Quantum computing

Calculations on a quantum computer are represented by **quantum circuits**

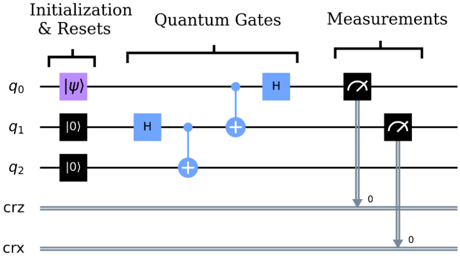

- Information stored in quantum bits (qubits) (two state quantum system)
- State of qubits changed by applying quantum gates (unitary operations)
- Expectation values of observables calculated by measuring qubits

Quantum computing for nuclear physics

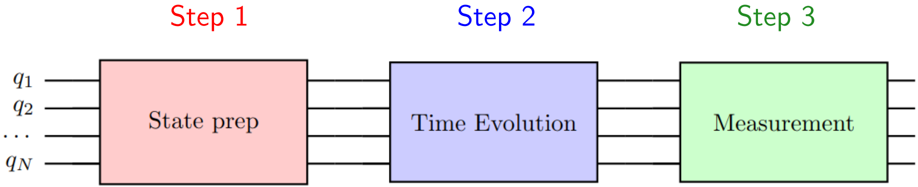

Step 1
State prep

Step 2
Time Evolution

Step 3
Measurement

- Typical complexity of Step 1 $\gtrsim$ complexity of Step 2
- Step 2 not required if only interested in certain eigenstate or thermal state properties

NISQ or Fault-tolerant quantum computers

Class of algorithms we consider depends on devices used

Noisy-intermediate scale era

- Limited number of noisy qubits ~ 100
- Imperfect gates: $\sim 10^3$ entangling gates
- Difficult to control error from noise
- Available now

Fault-tolerant quantum computers

- Error correction
- Access to many logical qubits
- Arbitrarily long coherence time
- Available 2030+

Quantum simulation of lattice EFTs in 3 spatial dimensions
with realistic lattice sizes and controlled errors requires fault-tolerant computers

It therefore makes sense to think about algorithms with (nearly) optimal scaling in the parameters

Broad classes of quantum algorithms

Mimicking "physical" evolution:

- E.g. Trotter formulas, small steps in time
- Simple circuits, no ancillary qubits
- Asymptotic scaling not optimal

Quantum calculator:

- General class of matrix functions of Hamiltonian $f(H)$ (including $f(H) = e^{iH}$)
- More complicated circuits, requires ancillas
- Near-optimal scaling, potentially more efficient

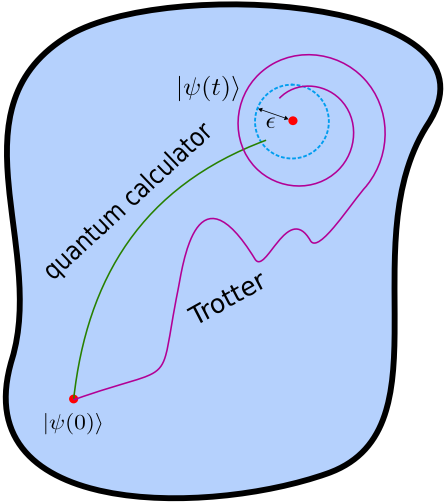

What is near-optimal scaling for time-evolution?

Under general assumptions, rigorous bounds on best scaling with t and ϵ exist

	Product formulas (Trotter)	Quantum Signal Processing
Approximation	$e^{-it(H_1+H_2)} \approx \left(e^{-i\frac{t}{N}H_1} e^{-i\frac{t}{N}H_2}\right)^N$	$e^{-itH} \approx \sum_{k=0}^{d} c_k(t)H^k$
“Cost”	$\mathcal{O}\big(\underbrace{\frac{t^2}{\epsilon}\times}_{\text{sub-optimal}} \lVert[H_1, H_2]\rVert\big)^*$	$\mathcal{O}\big(\lVert H\rVert \underbrace{t + \log(\frac{1}{\epsilon})}_{\text{nearly-optimal}}\big)$

Not obvious a priori which is better

- We are happy with $\epsilon \lesssim 1 - 10\%$ error so not clear which is better *a priori*
- Potentially more important scaling with number of sites L^3 or particle number η
- Need to do the cost analysis to decide for certain

*Higher order PF: $t^2 \rightarrow t^{1+1/p}$ and $\epsilon^{-1} \rightarrow \epsilon^{-1/p}$ (and more nested commutators)

Review of resource analysis for lattice nuclear EFTs

Prior work: product formulas (Trotter)

2019: Pionless EFT [1]

- time evolution
- 2nd-order product formula

2023: Pionless, One-pion exchange, Dyn. pion [2]

- Time evolution + energy estimation
- Improved product-formula error bounds

This work: Pionless, OPE, Dyn. pion

- Quantum Signal Processing
- Time evolution, energy estimation, state prep
- applied improved boson truncation bound [3]

$$\eta = 40 \text{ nucleons, } 8\times 8\times 8 \text{ box,}$$
$$T = a_L L\sqrt{\frac{M_N}{2E_{\text{kin}}}},\ \epsilon = 0.1$$

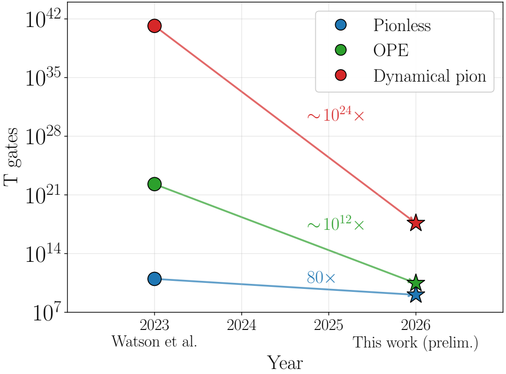

[1] A. Roggero et. al, PRD 2020, , arXiv:1911.06368, [2] J. Watson et. al, arXiv:2312.05344, [3] J. Wang, **CFK**, S. Jabeen, arXiv:2604.24896

Asymptotic gate complexity of time evolution

L^3 : total lattice sites, ℓ: long-range truncation distance, η : particle number, n_b: qubits per bosonic mode, fixed boson truncation

	Trotter [1]	QSP [this work]
Pionless	$\widetilde{\mathcal{O}}\left(\frac{t^{3/2}}{\sqrt{\epsilon}}\,\eta^{1/2}L^3\right)$	$\widetilde{\mathcal{O}}\left(L^6\,t + L^3\log\epsilon^{-1}\right)$
One-pion exchange	$\widetilde{\mathcal{O}}\left(\frac{t^{3/2}}{\sqrt{\epsilon}}\,\eta^{1/2}L^3\ell^3\right)$	$\widetilde{\mathcal{O}}\left(L^6\,t + L^3\log\epsilon^{-1}\right)$
Dynamical pions	$\widetilde{\mathcal{O}}\left(\frac{t^{3/2}}{\sqrt{\epsilon}}\,\eta^{1/2}L^3\,n_b^2\right)$	$\widetilde{\mathcal{O}}\left(n_b\,L^6\,t + n_b\,L^3\log\epsilon^{-1}\right)$

- **Particle number vs Volume:** Trotter cost grows as $\eta^{1/2}L^3$ vs. QSP L^6

[1] J. Watson et. al, arXiv:2312.05344

Gate and qubit cost for time evolution (preliminary)

Setup: $\eta = 40$ particles for time $T_{\text{cross}} = a_L L \sqrt{\frac{M_N}{2E_{\text{kin}}}}$ with error $\epsilon = 0.1$ and $E_{\text{kin}} = 10$ MeV

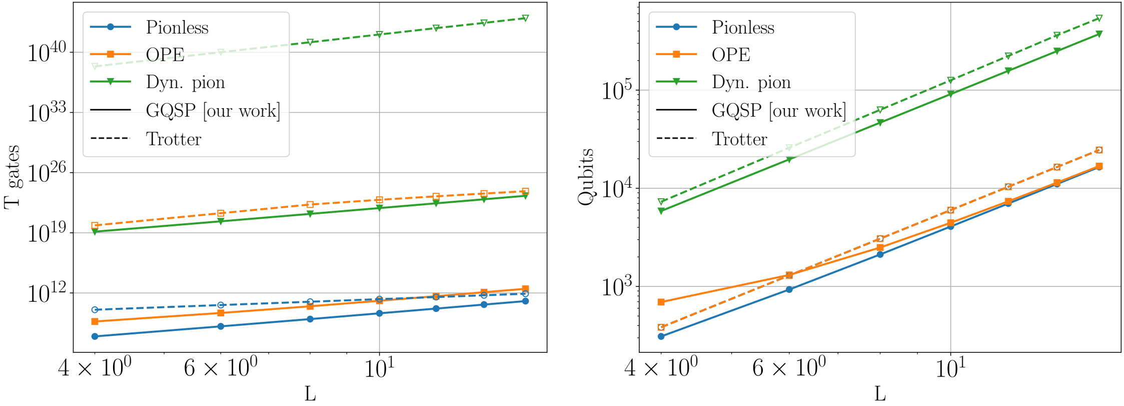

- Trotter results from [1] J. Watson et. al, arXiv:2312.05344
- QSP results from this work

Alternate Hamiltonian representation: second vs first quantization

Second quantization: tracks fermion modes

- Hilbert space spanned by lattice occupation quantum numbers: $|n_0, n_1, \ldots, n_{N-1}\rangle$
- Pionless, OPE, Dynamical pion EFTs Hamiltonian block-diagonal in particle number sectors
  $\rightarrow$ we only use small subspace of full Hilbert space

First quantization: track particle positions

- Hilbert space spanned by position of each particle: $|x_0, x_1, \ldots, x_{\eta-1}\rangle$
- Anti-symmetrize wavefunction $|\psi\rangle$
- Requires less qubits and potentially less gates

Resource analysis done for Pionless lattice EFT in 1st quantization [L. Spagnoli, C. Lissoni, A. Roggero, arXiv:2507.22814].

Pionless lattice EFT: asymptotic complexity for time evolution

L^3: total lattice sites, η: particle number

	Second quantization		**First quantization**
	Trotter [1]	QSP [this work]	QSP [2]
T gates	$\widetilde{\mathcal{O}}\left(\frac{t^{3/2}}{\sqrt{\epsilon}}\,\eta^{1/2}\,L^3\right)$	$\widetilde{\mathcal{O}}\left(L^6\,t + L^3\log\frac{1}{\epsilon}\right)$	$\widetilde{\mathcal{O}}((\eta^2 t + \eta\log\epsilon^{-1})\log L)$
Qubits	$\mathcal{O}\left(L^3\right)$	$\mathcal{O}\left(L^3 + \log(L/\epsilon)\right)$	$\mathcal{O}(\eta\log L + \log(t/\epsilon))$

- Second quant. Trotter: polynomial in η and L
- Second quant. QSP: independent of η and polynomial in L
- First quant. QSP: polynomial in η and logarithmic in L

[1] J. Watson et. al, arXiv:2312.05344, [2] L. Spagnoli, C. Lissoni, A. Roggero, arXiv:2507.22814

Pionless EFT: first vs second quantization

Time evolution in $8 \times 8 \times 8$ box

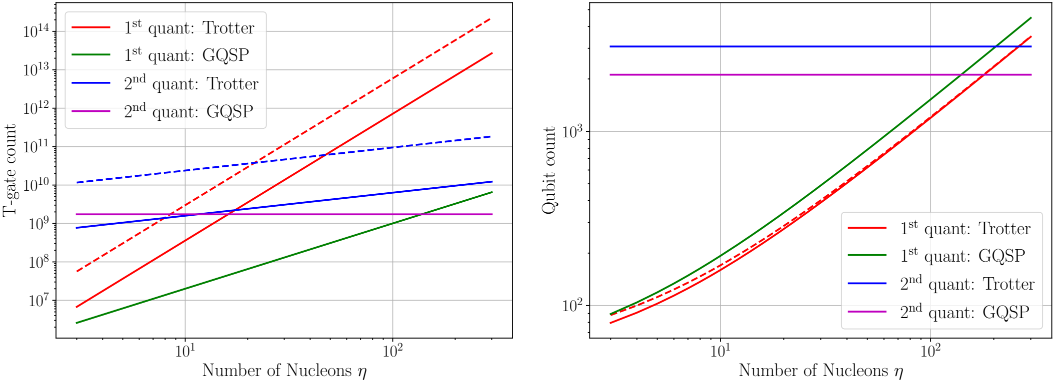

- 2nd quant. Trotter from [1]
- 1st quant. from [2]
- 2nd quant GQSP from this work

[1] J. Watson et. al, arXiv:2312.05344, [2] L. Spagnoli, C. Lissoni, A. Roggero, arXiv:2507.22814

Summary

End-to-end cost analysis for three nuclear lattice EFTs using QSP in second quantization: time evolution, state preparation, energy estimation.

Cost reduction vs. improved Trotter:

- Pionless EFT: $\sim \mathbf{80}\times$
- One-pion exchange: $\sim \mathbf{10^{12}}\times$
- Dynamical pion: $\sim \mathbf{10^{24}}\times$

Open directions:

- First quantization for OPE / dynamical-pion EFTs
- Alternative algorithms for further reductions?
- Can classical computation accelerate quantum simulation?

Discussion points

What are concrete problems where quantum computers can contribute to nuclear physics?

Near(est) term → IBM roadmap 2030:

- 200 logical qubits and 10^8 gates

Far term → IBM roadmap 2033+:

- 2000 logical qubits and 10^{10} gates

$$\eta = 40 \text{ nucleons}, 8 \times 8 \times 8 \text{ box},$$
$$T = a_L L \sqrt{\frac{M_N}{2E_{\text{kin}}}}, \ \epsilon = 0.1$$

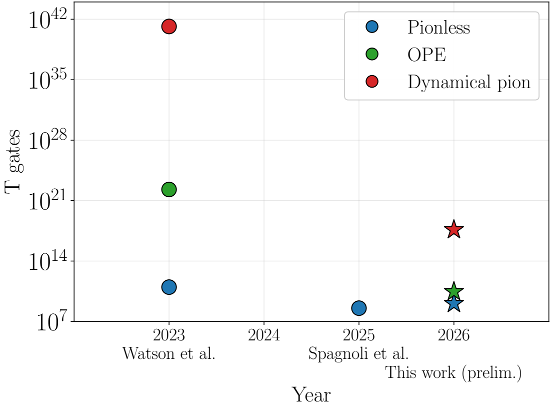

Backup slides

Quantum Signal Processing: high level review

Hamiltonian input model: given $|\psi\rangle \to H|\psi\rangle$, one can implement $f(H)|\psi\rangle$

This is achieved via a **block encoding** of H/λ where $\lambda \geq \|H\|$

$$W_H = \begin{pmatrix} \frac{H}{\lambda} & * \\ * & * \end{pmatrix} \to \begin{pmatrix} \frac{H}{\lambda} & * \\ * & * \end{pmatrix} \begin{pmatrix} |\psi\rangle \\ 0 \end{pmatrix} = \begin{pmatrix} \frac{H}{\lambda}|\psi\rangle \\ * \end{pmatrix}$$

Quantum Signal Processing → d calls to W_H can implement $f(H)|\psi\rangle$ where $\deg(f) \sim d$

[G. Low, I. Chuang, Quantum, arXiv:1606.02685], [G. Low, I. Chuang, Quantum, arXiv:1610.06546], [D. Motlagh, N. Wiebe, PRX Quantum, arXiv:2308.01501]

$$\begin{pmatrix} |\psi\rangle \\ 0 \end{pmatrix} \to W_H \to W_H \to \cdots \to W_H \to \begin{pmatrix} f\left(\frac{H}{\lambda}\right)|\psi\rangle \\ * \end{pmatrix}$$

- Time evolution: $f(\frac{H}{\lambda}) = e^{-iHt}$ has optimal complexity w.r.t. queries to W_H
- State preparation: $f(\frac{H}{\lambda}) = |\psi_0\rangle\langle\psi_0|$ has nearly-optimal complexity w.r.t. queries to W_H
- Energy estimation via phase estimation [R. Babbush et. al, PRX Quantum,, arXiv:1805.03662]