

Estimating Energies with the Rodeo Algorithm

A report on K. Choi, D. Lee, J. Bonitati, Z. Qian and J. Watkins,

Rodeo Algorithm for Quantum Computing, Phys. Rev. Lett. **127**, 040505 (2021) [arXiv:2009.04092]

A research note by Sovereign Systems Labs

1. Elevator pitch

The rodeo algorithm is a way of asking a quantum computer a single question — “do you have an eigenstate at energy E ?” — and getting an answer whose reliability improves exponentially with the number of times you ask. You attach one ancilla qubit to the register holding a trial state, put the ancilla in a superposition, and use it to control evolution under the Hamiltonian for a randomly chosen time t . Each eigencomponent of the trial state picks up a phase $e^{-iE_j t}$, and because the evolution is controlled, that phase is kicked back onto the ancilla; a compensating phase gate $P(Et)$ then cancels it *exactly* when the scan parameter E matches the eigenvalue E_j . Measuring the ancilla therefore returns the “success” outcome with certainty on resonance, and with probability $\cos^2[(E_j - E)t/2]$ off resonance. Repeat N times with fresh random times and keep only the runs in which *every* ancilla measurement succeeded: on average each off-resonant eigencomponent has been suppressed by a further factor of $\frac{1}{4}$ per cycle, so the surviving state is dominated by the eigenvectors within a narrow window around E . Sweeping E and plotting the survival probability traces out the spectrum; parking E on a peak and post-selecting *prepares* that eigenstate. Choi et al. show this costs $O[|\log \delta|/(pe)]$ work, which is exponentially cheaper in the target fidelity than quantum phase estimation (QPE) or adiabatic state preparation — and, unlike those methods, it needs no quantum Fourier transform, no slow adiabatic ramp, and (with mid-circuit measurement) a single ancilla qubit. That combination of shallow circuits and exponential convergence is what makes it a credible algorithm for noisy intermediate-scale (NISQ) hardware.

2. Problem and motivation

The question. Given a many-body Hamiltonian H_{obj} , find its energy eigenvalues E_j and prepare the corresponding eigenvectors $|E_j\rangle$ — not just the ground state, but states anywhere in the spectrum. Almost everything one wants downstream (reaction barriers, excitation spectra, response functions, nuclear level schemes, phase diagrams) is a functional of these eigenpairs.

Why it is hard classically. The Hilbert space of n spin- $\frac{1}{2}$ sites has dimension 2^n , so exact diagonalisation is memory-bound at roughly $n \lesssim 20$ –24, and the standard workarounds each buy tractability by giving something up. Krylov methods (Lanczos) avoid storing the matrix but still store the state vector, and converge quickly only to *extremal* eigenvalues — interior levels need shift-and-invert, i.e. repeated large linear solves. Quantum Monte Carlo scales well but suffers the fermionic/frustrated sign problem. DMRG and tensor networks excel in one dimension but assume limited entanglement. None is a general-purpose solver for excited states of an arbitrary Hamiltonian.

Why the obvious quantum answers are also hard. A fault-tolerant machine would use QPE, which is asymptotically excellent but needs an ancilla register, an inverse quantum Fourier transform, and coherent evolution U^{2^k} out to times $\sim 1/\epsilon$ — depths far beyond current devices. Adiabatic state preparation drops the QFT but needs an evolution time set by the inverse minimum gap. Variational algorithms (VQE) go to the other extreme: shallow circuits, but paid for with a hard classical optimisation (barren plateaus, local minima), an $O(1/\epsilon^2)$ measurement cost, no variational certificate for excited states, and no natural way to say “give me the state near $E = -10.5$ ”.

The gap the paper fills. Choi et al. ask for a middle route: the *exponential* convergence of a projective method with the *shallow, repeatable* circuit structure of a NISQ algorithm, valid for interior eigenvalues and returning the whole spectrum. Their answer is a stochastic spectral filter built from interference on a single ancilla.

3. Method

3.1 One rodeo cycle

Let $|\psi_I\rangle = \sum_j c_j |E_j\rangle$ be an initial (trial) state of the n -qubit *object* register. Cycle n uses one ancilla prepared in $|1\rangle$ and consists of: a Hadamard, a controlled evolution $e^{-iH_{\text{obj}}t_n}$, a phase gate $P(Et_n)$ (which multiplies $|1\rangle$ by e^{iEt_n}), a second Hadamard, and a measurement. Because H_{obj} commutes with every gate in the cycle, the action factorises over eigenvectors, and the cycle implements the (non-unitary, measurement-induced) map

$$|\psi_I\rangle \mapsto \frac{1}{2} \left[\mathbb{I} + e^{-i(H_{\text{obj}}-E)t_n} \right] |\psi_I\rangle \quad (\text{conditioned on the "success" outcome}). \quad (1)$$

This is the whole idea: the ancilla measurement projects onto a *cosine filter* centred on E . The success probability for a single eigencomponent is

$$P_1(E_j) = \cos^2 \left[(E_j - E) \frac{t_n}{2} \right], \quad \text{so after } N \text{ cycles} \quad P_N(E_j) = \prod_{n=1}^N \cos^2 \left[(E_j - E) \frac{t_n}{2} \right]. \quad (2)$$

3.2 Why randomness matters

If the t_n were all equal, Eq. (2) would revive whenever $(E_j - E)t/2$ is a multiple of π : the filter would be periodic and pass a whole comb of false energies. Drawing t_n from a Gaussian of root-mean-square width t_{RMS} destroys those revivals. On resonance every factor is exactly 1 whatever t_n is, so the target survives untouched; off resonance the phases are effectively uniform and the geometric mean of $\cos^2 \theta$ is $\frac{1}{4}$, giving

$$P_N(E_j) \sim 4^{-N} \quad (E_j \neq E), \quad (3)$$

i.e. **exponential suppression in the number of cycles**. Two practical rules follow from the paper: the resonance peak has width $O[1/(\sqrt{N} t_{\text{RMS}})]$, so the energy resolution is set by $t_{\text{RMS}} \propto 1/\epsilon$; and the asymptotic 4^{-N} rate is reached only once $\theta_{\text{RMS}} \equiv (E_j - E)t_{\text{RMS}}/2 \gtrsim 1$, i.e. once t_{RMS} exceeds roughly twice the inverse level spacing. (When N is smaller than the number of populated eigenstates the effective rate is the arithmetic mean 2^{-N} instead — the two regimes are the estimators F_A and F_G of the paper.)

3.3 Two modes of operation

- **Spectroscopy.** Scan E across an interval, and for each E estimate the all-success probability from repeated shots. The resulting curve is the *initial-state spectral function* $S(E) = \sum_j |\langle E_j | \psi_I \rangle|^2 \delta_{E, E_j}$, broadened by the filter: peak positions give the eigenvalues, peak heights give the overlaps. Resolution is refined by re-scanning a narrower interval with a proportionally larger t_{RMS} ; $O(\log \epsilon)$ such zoom levels give total cost $O[(\log \epsilon)^2 / (p\epsilon)]$, close to the $O(1/\epsilon)$ Heisenberg limit.
- **Eigenstate preparation.** Fix $E = E_j$, hold t_{RMS} fixed, and post-select on all-success. After N cycles the residual orthogonal component has magnitude $\Delta \sim 2^{-N/2}$ or $4^{-N/2}$, so $N = O(\log \Delta)$ cycles suffice; including the $1/p$ post-selection overhead the cost is $O(\log \Delta / p)$.

3.4 Resources

Only one ancilla is needed if the device supports mid-circuit measurement and reset; otherwise N ancillas are measured at the end. The dominant cost is the controlled evolution; Table 1 collects the scalings quoted in the paper.

4. Application to the project Hamiltonian

The project applies the algorithm to the spin- $\frac{1}{2}$ Heisenberg model in a uniform field — the same model used in the paper, but with $n = 5$ sites instead of 10:


$$H_{\text{obj}} = J \sum_{\langle j,k \rangle} \vec{\sigma}_j \cdot \vec{\sigma}_k + h \sum_j \sigma_j^z, \quad J = 1.0, \quad h = 3.0, \quad E \in [-13, -9] \text{ a.u.} \quad (4)$$

Exact diagonalisation of the 32×32 matrix (my own check, to have a reference for the quantum result) gives the levels in Table 2. The number of levels inside the window depends on the boundary


Table 1: Computational cost (from Choi et al.). $p = |\langle E_j | \psi_I \rangle|^2$ is the initial-state overlap, ϵ the target resolution, Δ the residual error of the prepared eigenvector. The “exponential” advantage is in Δ , i.e. in *state-preparation fidelity*; for eigenvalue estimation the rodeo algorithm is merely comparable to QPE.

	Rodeo	Phase est.	Adiabatic	Direct $\langle H \rangle$
Eigenvalue to accuracy ϵ	$O[(\log \epsilon)^2 / (p\epsilon)]$	$O[1/(p\epsilon)] + \text{QFT}$	—	$O(1/\epsilon^2)$ shots
Eigenstate to residual Δ	$O(\log \Delta / p)$	$O[1/(p\Delta)]$	$O(1/\Delta) \times f(p)$	—
Ancillas	1 (mid-circuit) or N	$O(\log 1/\epsilon)$	0	0
Needs QFT / long ramp	no / no	yes / no	no / yes	no / no
Gives excited states	yes, any E_j	yes	hard	via constraints

(a) One rodeo cycle (repeat N times with fresh t_n)



(b) Three ways to build the shaded box



(c) Workflow




Figure 1: Structure of the rodeo algorithm. (a) A single cycle: the ancilla interferometer converts the phase $e^{-iE_j t_n}$ kicked back by the controlled evolution into a measurement probability $\cos^2[(E_j - E)t_n/2]$, which equals 1 exactly on resonance. (b) The controlled evolution is the only expensive ingredient, and is the object of the three implementations compared in this project. (c) The outer loop: scan, post-select, read off peaks, then zoom in. (*Original figure.*)

condition, which the brief does not fix: a closed ring has five bonds and a doubly degenerate ground state at $-9 - \sqrt{5}$, whereas an open chain has four bonds and four non-degenerate levels in the window. Either way the search interval brackets the low-energy part of the spectrum. The smallest gap inside it is 0.47 for the ring and 0.29 for the open chain, so the rule $t_{\text{RMS}} \gtrsim 2/\Delta E$ suggests $t_{\text{RMS}} \approx 4-7$ and roughly $N \approx 8-10$ cycles.

Table 2: Exact eigenvalues of Eq. (4) for $n = 5$ inside the search window $E \in [-13, -9]$, computed by direct diagonalisation as a benchmark for the quantum calculation.

Boundary condition	Eigenvalues in $[-13, -9]$	Closed forms
Ring (5 bonds)	$-11.2361 (\times 2), -10.4721 (\times 2), -10.0000$	$-9 - \sqrt{5}, -6 - 2\sqrt{5}, -10$
Open chain (4 bonds)	$-12.2361, -11.0000, -10.7115, -10.2361$	$-10 - \sqrt{5}, -11, -, -8 - \sqrt{5}$

The initial state is the real design choice. Since $[H_{\text{obj}}, \sum_j \sigma_j^z] = 0$ and every gate in a cycle commutes with H_{obj} , the rodeo algorithm *cannot* move weight between magnetisation sectors: it only sharpens whatever is already in $|\psi_I\rangle$. This has a concrete consequence here. The Néel state $|01010\rangle$ — the natural analogue of the $|0101010101\rangle$ used in the paper — has *zero* overlap with every level in $[-13, -9]$, because $h = 3$ pushes the low-lying states into negative-magnetisation sectors. (Reproducing the paper’s Table S1 for $n = 10$ confirms the same effect there: the Néel state reaches only -18.1 ,

the lowest state of the $S^z = 0$ sector, while the true ground state of that ring lies at -23.9 .) A product state such as $|11011\rangle$ gives $p \approx 0.40$ on the ground doublet, and a three-term superposition drawn from the $S^z = -5, -3, -1$ sectors illuminates all three window levels at once — this is what is used in Fig. 2.




Figure 2: Simulated rodeo scan of the search window for the $n = 5$ ring of Eq. (4), with $t_{\text{RMS}} = 4$, averaged over 40 sets of Gaussian random times and initial state $|\psi_I\rangle \propto |11011\rangle + |01011\rangle + |11111\rangle$. With $N = 3$ cycles the filter is too broad to separate the levels; by $N = 9$ the three peaks sit on the exact eigenvalues and their heights converge to the exact overlap probabilities (open circles). (*Original figure.*)

The three evolution methods and their CX cost. The controlled evolution dominates the two-qubit gate count, and the three routes trade accuracy against depth in different ways (Table 3). Two implementation notes that matter for the notebook:

- *Trotterisation biases the answer, it does not blur it.* With a fixed step δ , the circuit exactly implements $e^{-iH_{\text{eff}}\delta}$ for an effective Hamiltonian differing from H_{obj} by $O(\delta)$ at first order, so the rodeo peaks appear at eigenvalues of H_{eff} . It is therefore better to fix δ and set $r_n = \text{round}(t_n/\delta)$ than to fix r and let δ vary with t_n — otherwise each cycle filters a slightly different Hamiltonian and the peaks broaden rather than merely shift.
- *Controlling everything is wasteful.* Each controlled two-qubit rotation costs about twice its uncontrolled version. Since $\vec{\sigma}_j \cdot \vec{\sigma}_k$ decomposes into three mutually commuting Pauli-pair rotations, the uncontrolled bond gate needs 6 CX (3 with an optimal $SU(4)$ decomposition) but the controlled one needs ≈ 12 .

Table 3: Estimated CX count for *one* controlled evolution (5-site ring; 5 bonds, 5 field terms) using the textbook decompositions.

Method	CX per controlled cycle	Error behaviour
Exact unitary $e^{-iH_{\text{obj}}t_n}$	$\lceil \frac{1}{4}(4^6 - 19) \rceil = 1020$ (generic 6-qubit synthesis)	none (machine precision)
First-order Trotter, r steps	$r(5 \times 12 + 5 \times 2) = 70r$	systematic $O(\delta)$ peak shift
Constrained depth D	D , fixed by construction	variational (size of fit residual)

5. Discussion: strengths, limitations, outlook

5.1 Where it wins

Against classical methods: it stores an n -site state in $n + 1$ qubits rather than 2^n amplitudes, has no sign problem, and — unusually — reaches *interior* eigenvalues as easily as extremal ones, since the filter is centred wherever you point it; classically that is the expensive shift-and-invert regime.

Against QPE: no QFT, no U^{2^k} ladder, one ancilla, and a short cycle repeated N times — exactly the structure that mid-circuit measurement and error mitigation handle well. For eigenstate preparation the gain is genuinely exponential, $O(\log \Delta)$ versus $O(1/\Delta)$. **Against VQE:** no classical optimiser, hence no barren plateaus or local minima, plus direct access to excited states. **Against adiabatic evolution:** no long coherent ramp — and the two combine, since the paper reports a short adiabatic “preconditioner” ($t_{\text{AE}} = 5$) raising p from 0.11 to 0.83, a sevenfold gain that feeds straight into the $1/p$ cost.

5.2 Where it does not

- **The overlap p is the real bottleneck.** Every cost in Table 1 carries a $1/p$, and for a generic many-body Hamiltonian the overlap of an efficiently preparable trial state with a specific eigenvector shrinks exponentially with system size. The advantage is in the convergence parameter Δ , *not* in n : the algorithm inherits the same state-preparation wall as QPE, so the brief’s phrase “exponential speedup compared to QPE” should be read with that caveat.
- **Post-selection throws away most runs.** The accepted fraction is $\approx pP_N$, so deep filtering and poor overlap multiply into a large shot budget.
- **“Shallow” is relative.** The total evolution time is $\sum_n |t_n| \sim N t_{\text{RMS}}$ (≈ 29 a.u. for $N = 9$, $t_{\text{RMS}} = 4$), which Table 3 turns into hundreds or thousands of CX gates on a 5-qubit problem — which is exactly why the controlled evolution, not the interferometer, is what the project measures.
- **Degeneracy is not resolved.** The filter selects an energy, not a state: for the ring, whose window levels are doubly degenerate, it prepares whichever vector in the degenerate subspace $|\psi_I\rangle$ happens to select; extra symmetry labels are needed to separate them.
- **Approximate evolution shifts the answer.** Trotter and constrained-depth circuits return the spectrum of an *effective* Hamiltonian; that bias does not average away with more shots.
- **It is not variational.** Despite the project heading, this is a projective filter with no parameters to optimise; its relationship to VQE is complementary — use a variational ansatz to make a high- p initial state, then let the rodeo algorithm project, certify and refine it.

5.3 Outlook

Two developments since the paper attack the limitations above directly. Qian et al. ran the algorithm on IBM’s *Casablanca* using mid-circuit measurement on a single ancilla, recovering a one-qubit spectrum to 0.08% relative error and beating direct expectation-value estimation on the same device. Bee-Lindgren et al. replaced the controlled evolution with *controlled reversal gates*: a reversal gate R is a product of single-qubit Paulis anticommuting with (part of) H_{obj} , so $Re^{-iH_{\text{obj}}t}R^\dagger = e^{+iH_{\text{obj}}t}$, and controlling R alone toggles time forwards or backwards rather than on or off. This yields the same interference pattern (with $t \rightarrow 2t$) while controlling only a few single-qubit gates — a fivefold two-qubit gate reduction in their example, with larger gains at scale; here the relevant gates would be $\otimes_i Z_{2i}$ and $\otimes_i X_{2i}$. Beyond that: better initial states, and replacing product formulas with qubitisation once hardware allows.

References

- [1] K. Choi, D. Lee, J. Bonitati, Z. Qian and J. Watkins, *Rodeo Algorithm for Quantum Computing*, Phys. Rev. Lett. **127**, 040505 (2021); arXiv:2009.04092.
- [2] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th Anniv. Ed., CUP (2010) — Ch. 4 (gate decompositions), Ch. 5 (phase estimation).
- [3] M. Cerezo et al., *Variational Quantum Algorithms*, Nature Reviews Physics **3**, 625 (2021) — review of VQE, its cost model and barren plateaus.
- [4] J. Preskill, *Quantum Computing in the NISQ Era and Beyond*, Quantum **2**, 79 (2018).
- [5] S. Lloyd, Science **273**, 1073 (1996); H. F. Trotter, Proc. Amer. Math. Soc. **10**, 545 (1959); M. Suzuki, Commun. Math. Phys. **51**, 183 (1976) — product-formula Hamiltonian simulation.
- [6] V. V. Shende, S. S. Bullock and I. L. Markov, IEEE Trans. CAD **25**, 1000 (2006) — source of the $\lceil \frac{1}{4}(4^n - 3n - 1) \rceil$ CNOT count in Table 3.
- [7] Z. Qian et al., *Demonstration of the Rodeo Algorithm on a Quantum Computer*, Eur. Phys. J. A **60**, 151 (2024); arXiv:2110.07747.
- [8] M. Bee-Lindgren et al., *Rodeo Algorithm with Controlled Reversal Gates*, arXiv:2208.13557; *Controlled gate networks*, Eur. Phys. J. A (2025).
- [9] N. Wiebe and N. S. Babcock, New J. Phys. **14**, 013024 (2012) — adiabatic cost quoted in Table 1.