THEORY SEMINAR · UNIVERSITÀ DI MILANO-BICOCCA · 19 OCTOBER 2026

Loop-string-hadron formulation for non-Abelian lattice gauge theory and quantum simulation

Jesse Stryker

European Centre for Theoretical Studies in Nuclear Physics and Related Areas (ECT*), Fondazione Bruno Kessler, Trento



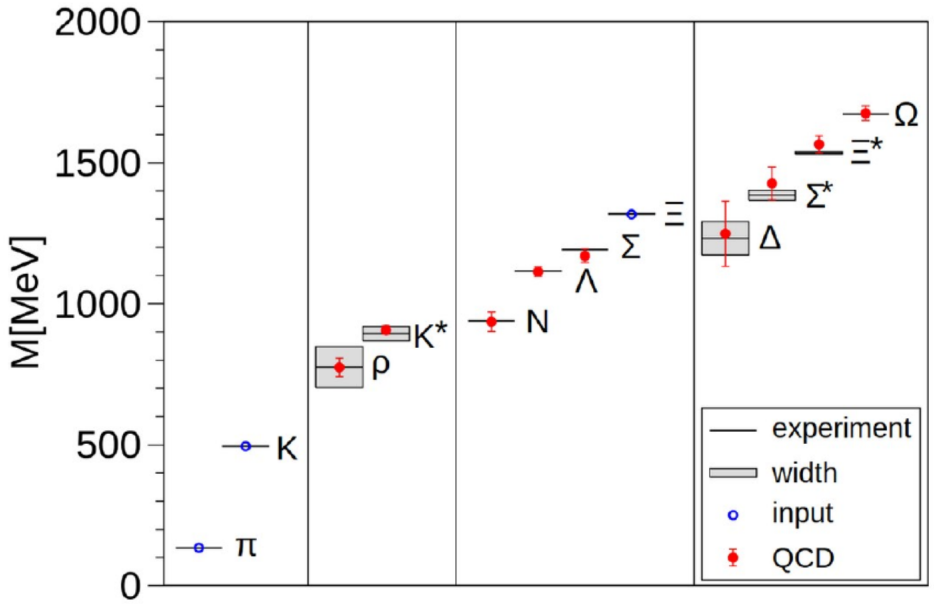
MOTIVATION

Euclidean lattice QCD: Where it shines

Path integral plus Monte Carlo importance sampling is ideal for thermal equilibrium, static properties and the low-lying spectrum.

A few successes

- The light hadron spectrum
  BMW collaboration, Science 322, 1224 (2008)
- The strong coupling αs from first principles
  ALPHA: Bruno, Dalla Brida et al., PRL 119, 102001 (2017)
- Hadronic vacuum polarization for the muon g−2
  e.g. Cè et al., PRD 106, 114502 (2022)
- Thermal QCD up to the electroweak scale
  Giusti, Dalla Brida, Pepe et al., arXiv:2410.12626



Light hadron spectrum vs. experiment (BMW)

MOTIVATION

Euclidean lattice QCD: Where challenges arise

For several decades, three classes of problems have stayed largely out of reach. When the path-integral weight is not real and positive, importance sampling fails.

Real-time dynamics

Hadronization, thermalization after heavy-ion collisions, transport coefficients

Finite density

Dense nuclear matter, neutron-star interiors, the QCD phase diagram at $\mu \neq 0$

Topological terms

θ-vacua and the strong CP problem

One shared obstruction: the sign problem.

MOTIVATION

The Hamiltonian alternative

Hamiltonian lattice gauge theory (Kogut & Susskind, 1975) has no sign problem. Why did we not use it on classical computers long ago?

$$\dim \mathcal{H} = d^V$$

d = local Hilbert-space dimension, V = number of lattice sites. A field theory has a Hilbert space at every point of space.

Classical computers can store wave functions only for small systems.

Two ways to beat the exponential

Tensor networks

Exploit limited entanglement; very effective in 1+1D

Quantum computers

Store the state natively in qubits; time evolution becomes a circuit

Over the last decade, both have driven a renaissance of the Hamiltonian formalism.

What a Hamiltonian calculation must get right

Scope

- real-time and nonequilibrium dynamics (my focus)
- arbitrary chemical potential
- topological terms

Scalability

Cost must not grow exponentially with the volume

Procedures

- state preparation
- time evolution
- measuring observables

Practical choices

- spatial continuum limit
- truncating the Hilbert space
- **choice of formulation and basis**
- method for time evolution

Kogut–Susskind: gauge fields live on links

- Temporal gauge, continuous time
- Each link carries its own Hilbert space; for SU(2), states |j, m, m′⟩
- Left and right electric fields ELa, ERa generate gauge rotations at the two ends of a link
- The link operator U changes the representation: j → j ± 1/2

$$H = H_E + H_B$$

$$H_E = \frac{g^2}{2} \sum_{\text{links}} E^a E^a$$

$$H_B = -\frac{1}{2g^2} \sum_{\text{plaq}} \text{tr}\left(U_P + U_P^\dagger\right)$$

plus mass and hopping terms HM + HI with staggered fermions

Gauge symmetry is a redundancy, so many formulations are possible. **For QCD, a formulation must handle SU(3) and 3+1 dimensions.**

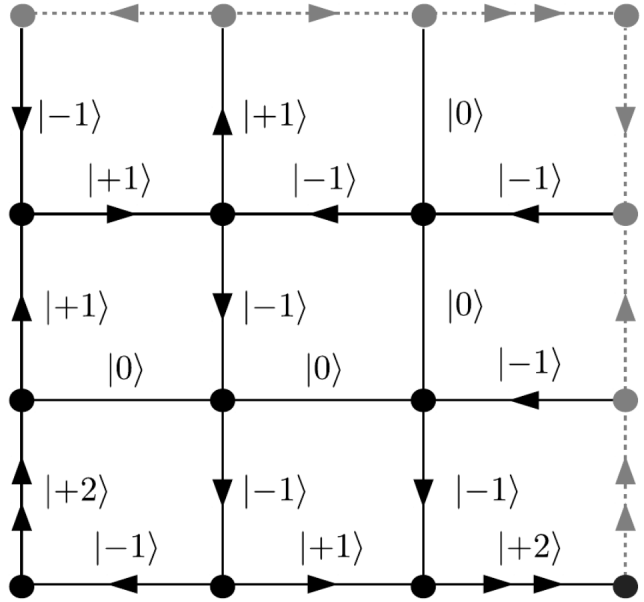
e.g. quantum link models: Chandrasekharan & Wiese, NPB 492, 455 (1997); realized via D-theory: Beard, Pepe, Riederer & Wiese, PRL 94, 010603 (2005)

Gauss’s law: physical states are a tiny subspace

U(1) $\nabla \cdot \mathbf{E} - \rho = 0\,, \quad \rho = \psi^\dagger \psi$

SU(N) $\mathbf{D} \cdot \mathbf{E}^a - \rho^a = 0\,, \quad \rho^a = \psi^\dagger T^a \psi$

- Imposed at every site, one constraint per generator
- Non-Abelian generators do not commute, so the constraints cannot be diagonalized together
- Gauss’s law, gauge invariance and charge conservation are one statement
- Gauge-invariant formulations are expected to pay off (Davoudi, Raychowdhury & Shaw, PRD 2021)



Compact U(1) electric basis: flux on links, balanced at each site

What makes a good simulation candidate

1. A complete, discrete basis for the lattice Hilbert space
2. Constraints known, including how to implement them
3. Matrix elements of the Hamiltonian’s field operators known
4. A clear path to quantum circuits for the Hamiltonian

The loop-string-hadron (LSH) formulation was built to meet all four, for SU(2) and SU(3).

From Schwinger bosons to loops, strings and hadrons: SU(2)

Schwinger bosons (prepotentials)

- A doublet of harmonic oscillators a(L), a(R) at each end of each link
- E and U are bilinears of oscillators
- Occupation-number basis; Clebsch–Gordan coefficients become oscillator factors
- Extra degrees of freedom bring one Abelian constraint per link: NL(x) = NR(x+1)

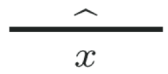
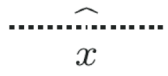
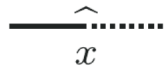
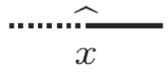


Loop-string-hadron (LSH)

- Contract site-local doublets into SU(2) singlets
- 17 gauge-invariant operators per site: 4 loop, 4 incoming string, 4 outgoing string, 2 hadron, 3 number
- These operators cannot violate color-charge conservation
- Basis per site: |nl, ni, no⟩

I. Raychowdhury & J. R. Stryker, PRD 101, 114502 (2020); PRR 2, 033039 (2020)

Every local excitation has a picture

$\mathcal{L}^{++}(x) \equiv$ 	Create unit of gauge flux.
$\mathcal{L}^{--}(x) \equiv$ 	Destroy unit of gauge flux.
$\mathcal{L}^{+-}(x) \equiv$ 	Change matter-sourced flux direction. $(d > 1)$
$\mathcal{L}^{-+}(x) \equiv$ 	Change matter-sourced flux direction. $(d > 1)$
$\mathcal{H}^{++}(x) \equiv$ 	Create a hadron.
$\mathcal{H}^{--}(x) \equiv$ 	Destroy a hadron.

Loop and hadron operators; string operators, which end flux on quarks, are more involved

Loop segments carry units of gauge flux through a site

String ends terminate a flux line on a quark

On-site hadrons are color singlets formed locally at a site

SU(2) is pseudoreal: flux is unoriented.
U(1) is complex: flux is oriented.

What LSH buys

Non-Abelian Gauss law solved locally

Elementary fields are SU(2)-neutral and local

Only Abelian constraints remain

They commute, so every basis state is definitely allowed or definitely forbidden

A simple Hilbert space

Integer occupation numbers instead of $|j, m, m'\rangle$; the structure looks like U(1)

Circuit-friendly operators

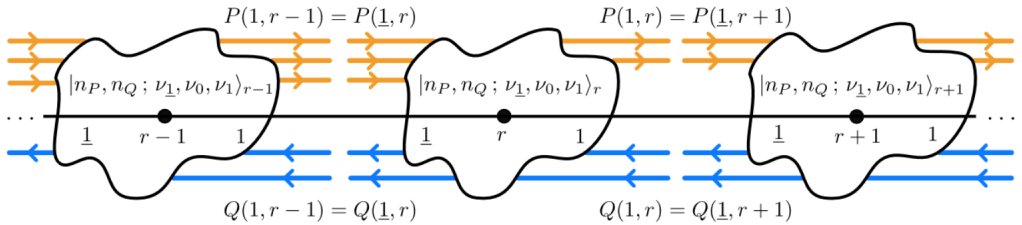
Diagonal matrices times one-sparse ladder operators; first SU(2) physicality circuits built this way

The cost: in d > 1 the magnetic Hamiltonian has many terms and may need more qubits. Developed for 1+1, 2+1 and 3+1 dimensions, with or without staggered fermions.

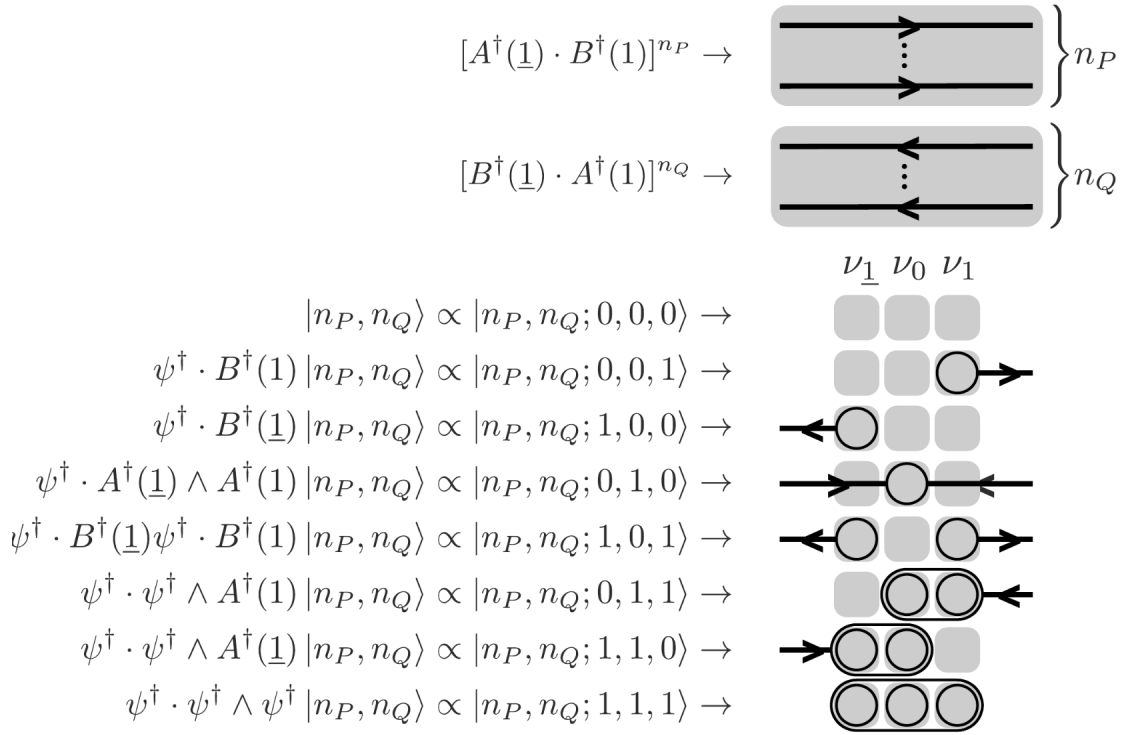
SU(3) in 1+1D: circuit-ready

- Every SU(3) irrep (P, Q) built from one triplet and one antitriplet oscillator: irreducible Schwinger bosons
- A direct generalization of the SU(2) LSH construction follows
- Analytic control on par with SU(2)
- No trivalent vertices needed in 1+1D

S. V. Kadam, I. Raychowdhury & J. R. Stryker, PRD (2023), arXiv:2212.04490; irreducible Schwinger bosons: Anishetty, Mathur & Raychowdhury (2009)



Abelian Gauss law: P and Q flux each continue across every link

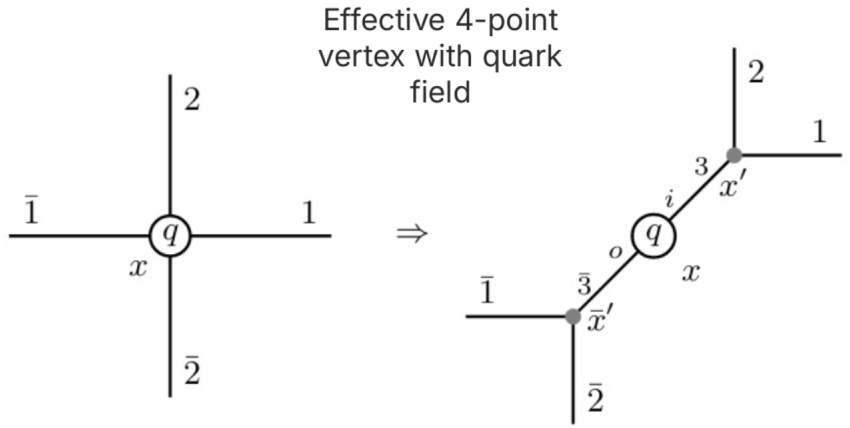


SU(3) beyond one dimension: a seventh quantum number

- Multidimensional lattices reduce to trivalent vertices by point splitting
- SU(2): each coupling occurs at most once. SU(3): Littlewood–Richardson multiplicities can exceed one
- So an SU(3) vertex needs a seventh label beyond its six irrep labels
- LSH states span each sector but are not orthogonal where the multiplicity exceeds one
- **Status:** matrix elements of all site-local operators derived; an analytic orthogonal basis is open

Part I: PRD 111, 074516 (2025), arXiv:2407.19181 · Part II (with A. Naskar, S. V. Kadam & I. Raychowdhury): accepted, in production

Closely related: gauge-invariant SU(3) states via Gram–Schmidt (also not closed-form): Balaji, Draper et al., PRD 112, 054511 (2025); Hidalgo & Draper, arXiv:2607.26445



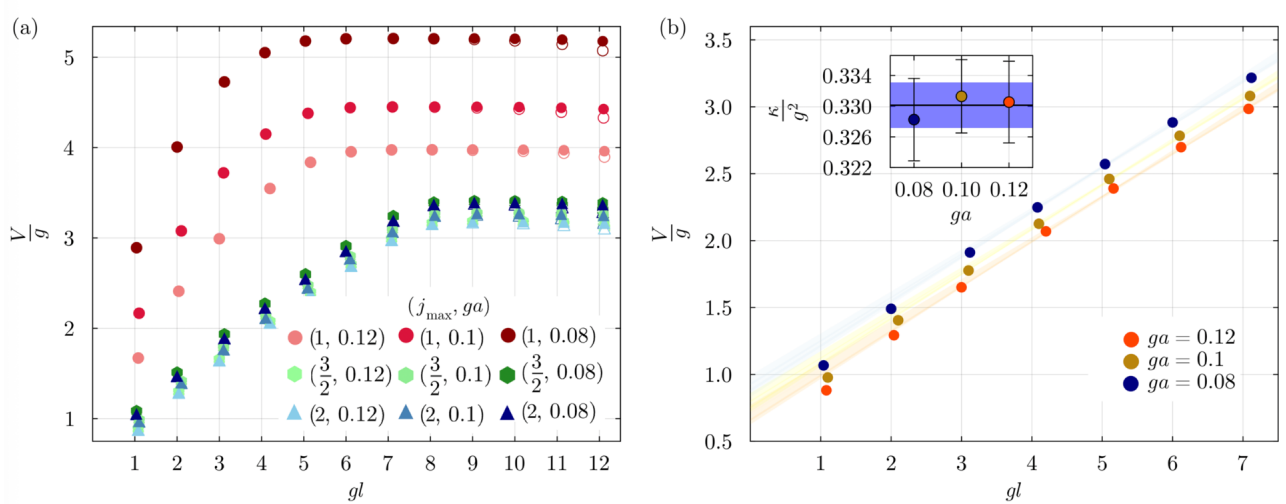
$$(1,1)\otimes(1,1) = (0,0)\oplus 2\,(1,1)\oplus(3,0)\oplus(0,3)\oplus(2,2)$$

$8 \otimes 8$ contains the octet twice

String breaking with tensor networks

An LSH-based tensor-network toolkit for (1+1)D SU(2) with dynamical fermions:

- No non-Abelian constraints to impose, so larger cutoffs and closer to the continuum and infinite volume
- Static strings: continuum string tension κ/g^2 = 0.330(3) at m/g = 0.5
- Quench dynamics: the LSH picture makes string expansion, endpoint splitting and particle production visible



(a) Static potential vs. separation, several cutoffs and lattice spacings. (b) Linear region and continuum string tension (inset).

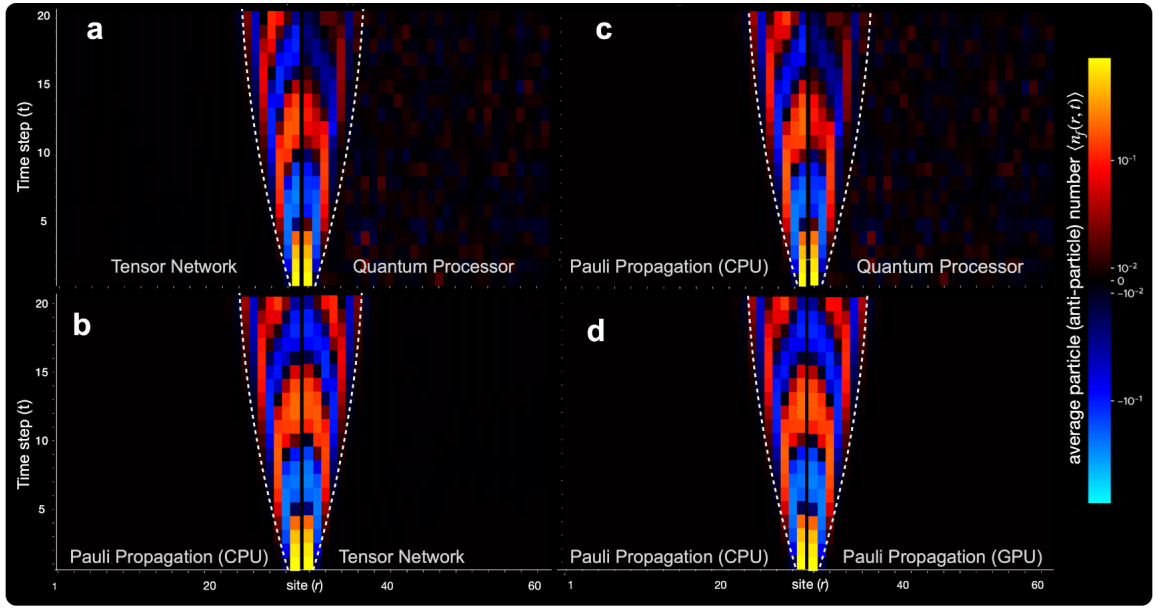
N. Gupta, E. Mathew, S. V. Kadam, J. R. Stryker, A. Bapat, N. Mueller, Z. Davoudi & I. Raychowdhury, arXiv:2603.24698

The formulation is in hand.

Can it work on real quantum hardware?

What would full QCD cost?

LSH on hardware: two complementary demonstrations



A confined meson spreading on a 60-site lattice: IBM processor vs. tensor networks and Pauli propagation

First: scale

Ilčić, Majumdar, Mathew, Ali, Earnest-Noble & Raychowdhury, arXiv:2602.18080

- 60 staggered sites on 120 qubits of a 156-qubit IBM processor
- Weak coupling, up to 25 Trotter steps, over 17,000 two-qubit gates
- Confined meson light cone, benchmarked classically
- Link occupations fixed by the Abelian Gauss law: **fermionic qubits only**

Next: generality (this talk)

Smaller, but the gauge field kept as its own multiqubit register, beyond j = 1/2: the ingredient d > 1 will need

SU(2) in 1+1D, with the hard parts kept in

Beyond the Schwinger model

Same minimal setting, one flavor and one spatial dimension, but a non-Abelian gauge group: one step closer to SU(3)

Keep the gauge fields

In 1+1D they can be integrated out, and usually are. In d ≥ 2 they cannot, so we keep the local bosonic degrees of freedom

Go beyond j = 1/2

The smallest truncation cannot reach the continuum limit. Higher truncations need multiqubit registers

T. Kaji, J. R. Stryker, Y. Iiyama, L. Nagano, E. Rinaldi, C. W. Bauer & K. Terashi

Data: Quantinuum System Model H2 (H2-2), May–June 2026



Trapped-ion quantum simulation of (1+1)D SU(2) lattice gauge theory via loop-string-hadron formulation

Toshiaki Kaji[1,*] Jesse R. Stryker[2,3] Yutaro Iiyama[1] Lento Nagano[4,1] Enrico Rinaldi[5] Christian W. Bauer[2,6] and Koji Terashi[1]

[1] International Center for Elementary Particle Physics (ICEPP), The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan.
[2] Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA.
[3] European Centre for Theoretical Studies in Nuclear Physics and Related Areas (ECT*), Fondazione Bruno Kessler, Strada delle Tabarelle, 286, 38123 Trento (TN), Italy.
[4] Graduate School of Science and Technology, Keio University, Yokohama, Kanagawa 223-8522, Japan
[5] Quantinuum, Partnership House, Carlisle Place, London SW1P 1BX, United Kingdom.
[6] Leinweber Institute for Theoretical Physics, Berkeley, CA 94720, USA.

(Dated: October 16, 2026)

We present a quantum simulation of the SU(2) lattice gauge theory in 1+1 dimensions with one flavor of fermions, performed using the trapped-ion Quantinuum System Model H2. This work is based on the loop-string-hadron (LSH) formulation of the gauge theory, and is distinguished by using qubits to represent both bosonic-gauge and fermionic-matter degrees of freedom. Furthermore, we access the $j = 1$ electric-flux sector beyond the $j_{\max} = 1/2$ truncation commonly used in SU(2) quantum simulations. To that end, we digitize an electric-field degree of freedom using a multiqubit register for representing the bosonic quantum number. Our simulation therefore synthesizes several key ingredients relevant to quantum simulations in higher spatial dimensions, in particular explicit local gauge-field degrees of freedom represented here by a multiqubit bosonic register, while relying solely on circuit-construction optimizations that are expected to scale polynomially with lattice size. This work demonstrates that modern quantum hardware is increasingly able to accommodate the features of complicated but physically important lattice gauge theories while successfully reproducing the expected Trotterized dynamics.

I. INTRODUCTION

One of the key applications for quantum computers has been identified as the simulation of lattice gauge theories [1–11]. Lattice gauge theory is the only available numerical tool to make predictions for interacting field theories, in particular those describing the interactions of elementary particles, without the need to rely on perturbation theory. Most lattice gauge theory calculations are based on evaluating the path integral using Monte Carlo integration. Since Monte Carlo integration requires a positive definite integrand, it is generally applied to theories defined in Euclidean spacetime in order to avoid the sign problem [12, 13]. The requirement for a positive-definite integrand limits the kinds of problems that can be studied with Monte Carlo integration, and in particular does not allow to directly study the dynamics of the field theory. Recently, approaches based on the Hamiltonian formulation [14] have attracted considerable attention, driven in large part by the ability to efficiently perform the required calculations using quantum computers. This is expanding the scope of theoretical investigations possible [15–18] and in particular allowing physicists to study the dynamics of fundamental processes of nature [19–22].

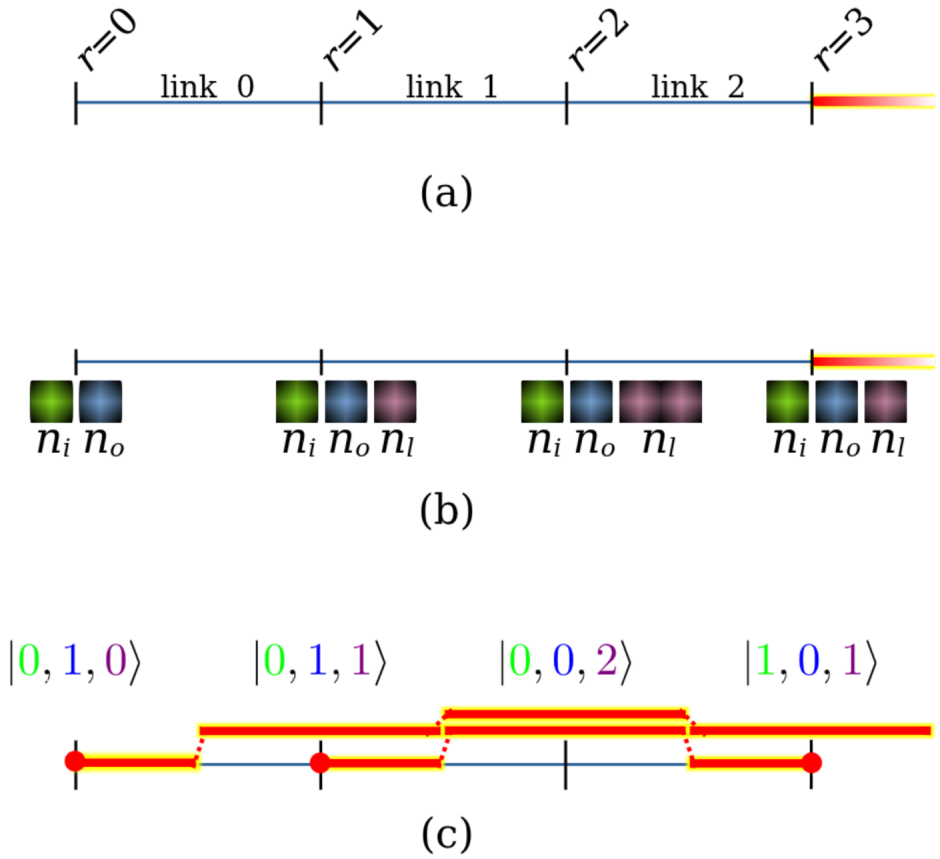
Our current best understanding of nature at small scales, the Standard Model, has multiple layers of complexity. The strong interaction sector, corresponding to the SU(3)-color gauge symmetry of the Standard Model, is especially complex because of its nonperturbative coupling at scales larger than approximately one femtometer. A proper representation of the physics of hadrons should account for the SU(3) gauge fields, multiple flavors of massive chiral fermions, and 3+1 dimensional spacetime. When calculations are based on lattice field theory, one must also take limits such as zero lattice spacing (i.e., the continuum limit), infinite volume, and infinite Hilbert space dimension [23–26]. In practice, simulations cannot start with all this complexity built in. Instead, toy models of increasing complexity are studied, such that more and more features of QCD or the Standard Model are increasingly understood in terms of their practical implementation. The simplest toy model exhibiting several key features of more complicated gauge theories—QCD included—is the theory with a single flavor of fermion, the simpler gauge group of U(1), and only one spatial dimension. This is the Schwinger model, describing quantum electrodynamics in 1+1 dimensions [15, 27]. In this work, we take one step in the direction of QCD by generalizing the gauge group of the Schwinger model from U(1) to SU(2).

Another important issue is the encoding of gauge degrees of freedom. It is well known that gauge theories

arXiv:submit/8242668 [hep-lat] 16 Oct 2026

In the arXiv queue: appears tomorrow morning (submit/8242668, 16 Oct 2026)

A minimal system that needs a bosonic register



- 4 staggered sites, open boundaries; one unit of external flux (j = 1/2) at the right edge
- Each site: fermionic ni, no $\in$ {0, 1} and bosonic nl $\in$ {0, 1, 2, …}; each box is one qubit
- Abelian Gauss law then allows nl(2) $\in$ {0, 1, 2}: a genuine two-qubit register, reaching the j = 1 flux sector
- Initial state (c): maximal electric energy, primed to disperse
- Observable: electric energy density HE(r), diagonal, read straight off the bitstrings

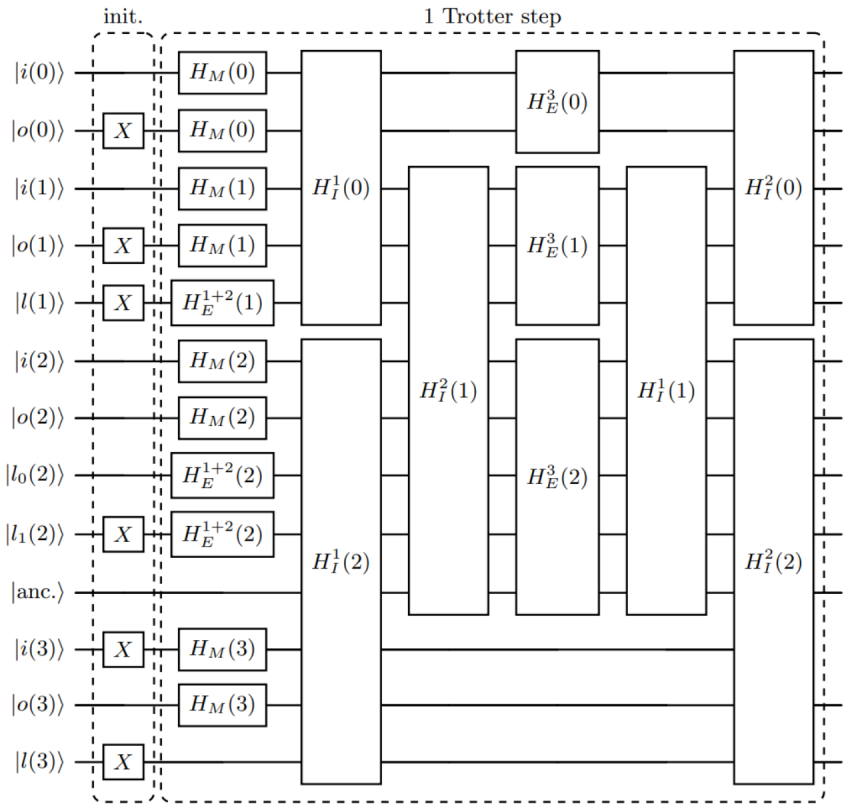
One Trotter step on trapped ions

Quantinuum H2

- 56 fully connected 171Yb+ qubits; ions shuttled around a race-track trap
- Native two-qubit gate: parameterized ZZPhase
- Typical infidelity: 1e-3 two-qubit, 3e-5 one-qubit

Circuit

- First-order Trotter; terms ordered so ladder operators cancel between neighboring hopping terms
- One step pre-optimized, compiled with pytket as a unit, then repeated
- **226 two-qubit gates per step**, up to six steps (t = 7.5)
- 100 to 540 shots per circuit, more for deeper circuits



Two error-mitigation strategies

Depolarizing rescaling

Model the noise on each observable as depolarizing, with strength p calibrated on test circuits with known answers; then invert

$$\langle O\rangle_{\text{noisy}} = (1-p)\,\langle O\rangle_{\text{ideal}} + p\,o_{\text{tr}}\,, \quad o_{\text{tr}} := \text{Tr}[O]/2^n$$

$$\langle O\rangle_{\text{mit}} = \frac{\langle O\rangle_{\text{raw}} - o_{\text{tr}}}{1-p} + o_{\text{tr}}$$

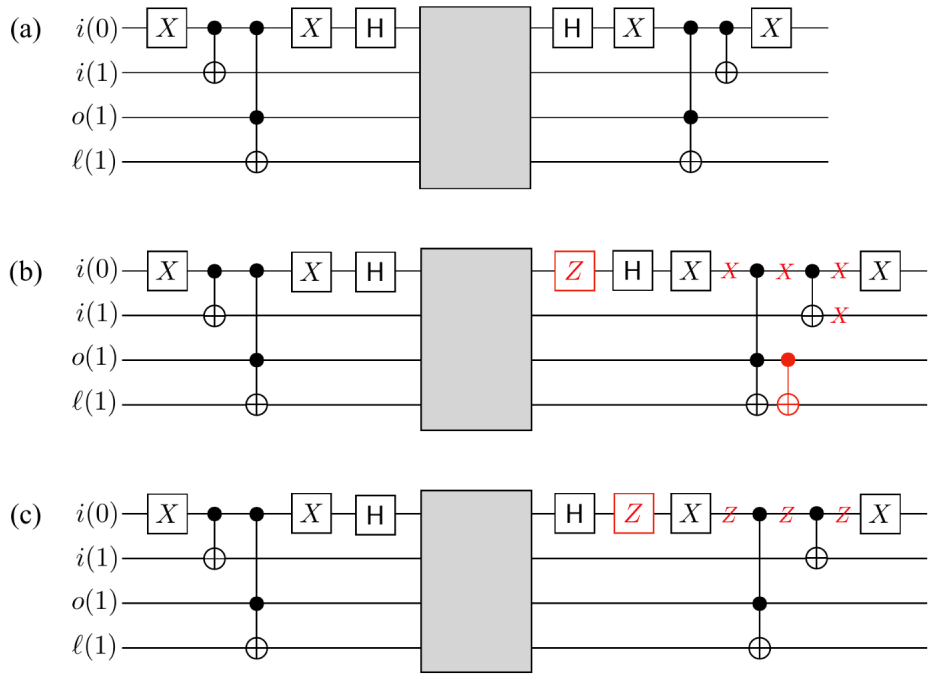
Gauss's-law post-selection

The Abelian constraints depend only on measured bitstrings: discard any shot that breaks one. No noise model needed

$$\hat{N}^L(r) = \hat{N}^R(r) \;\text{ on every link } r$$

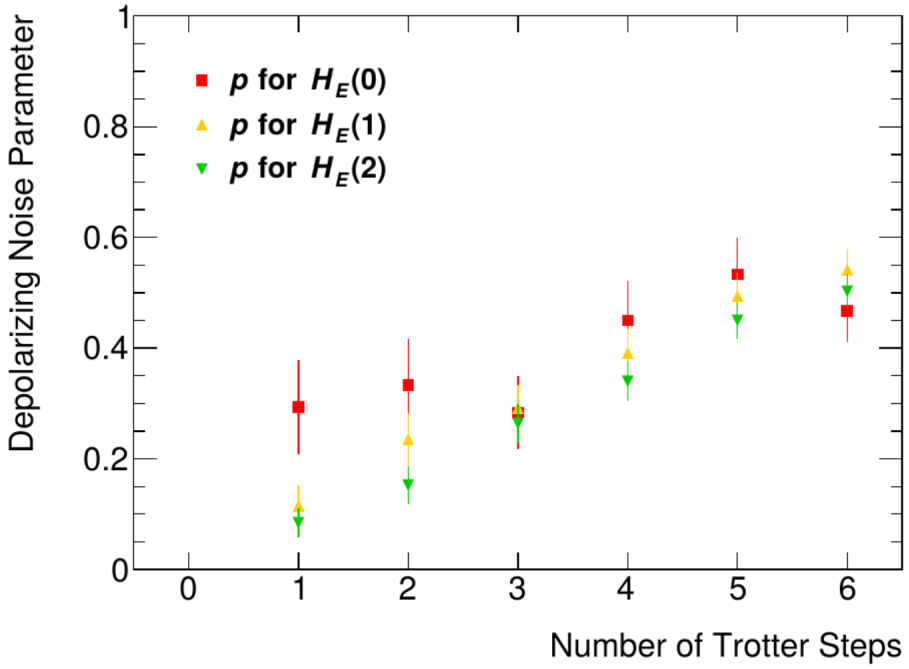
$$\hat{N}^L(r) = \hat{n}_l(r) + \hat{n}_o(r)[1-\hat{n}_i(r)]\,, \quad \hat{N}^R(r) = \hat{n}_l(r{+}1) + \hat{n}_i(r{+}1)[1-\hat{n}_o(r{+}1)]$$

Readout mitigation, dynamical decoupling and zero-noise extrapolation were tested and set aside

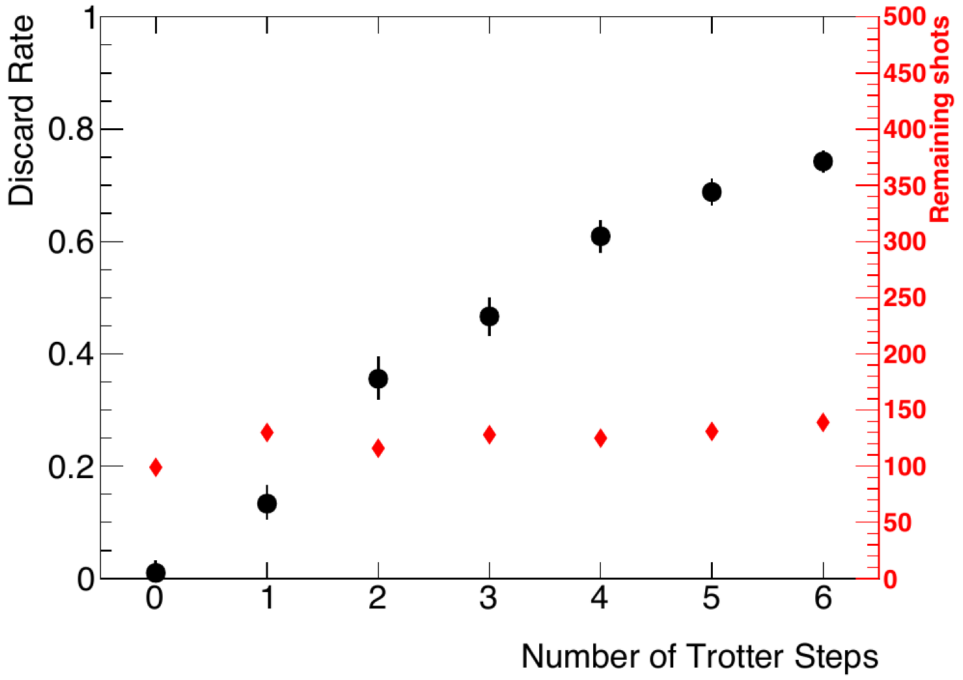


Post-selection also catches phase errors: a Z error between the H gates becomes an X error that trips the checks (b); one outside them may not (c)

Calibration, and the price of post-selection



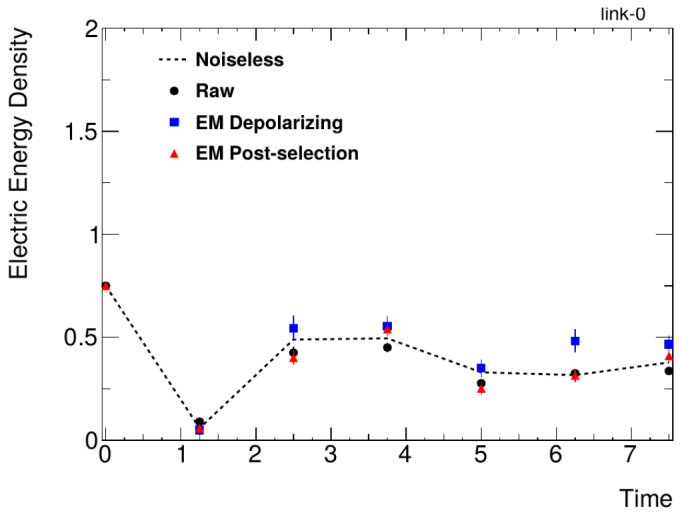
Effective depolarizing parameter grows roughly linearly with depth



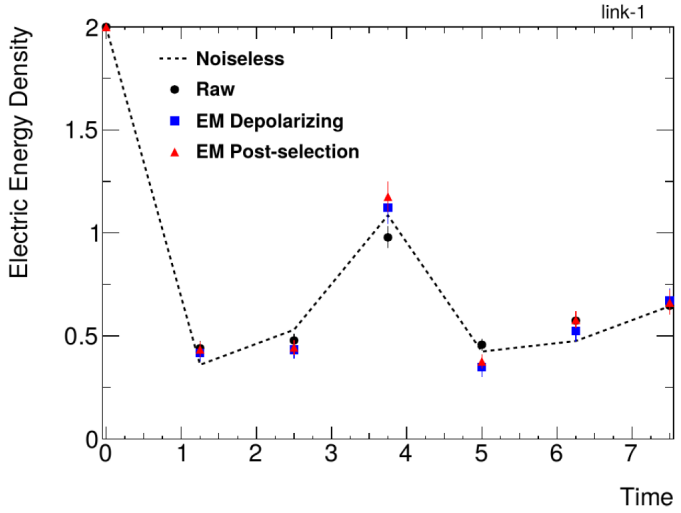
Discard rate reaches 0.74 at six steps; shots budgeted so that about 100–150 survive

Hardware follows the flux dynamics

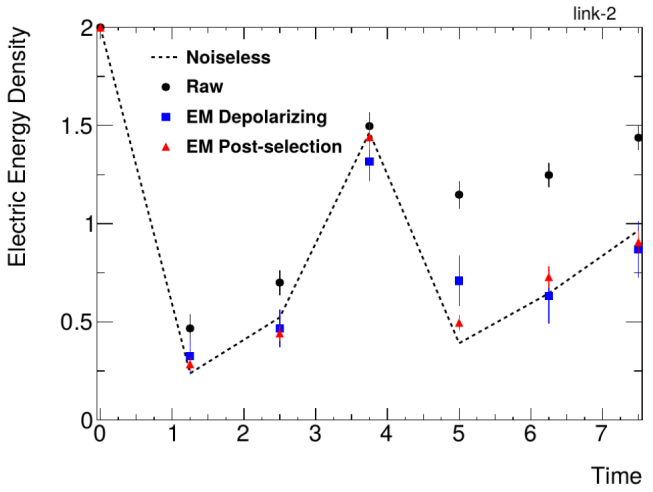
Link 0



Link 1



Link 2: the two-qubit register

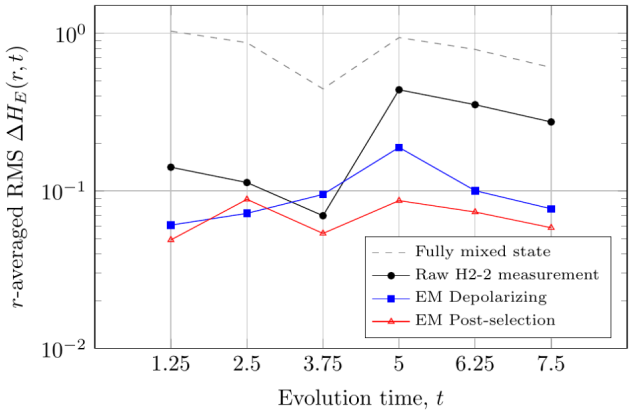


Six first-order Trotter steps, Δt = 1.25. Dotted: noiseless Trotterized baseline. Raw data (black) already tracks it on links 0 and 1; on link 2, the most noise-sensitive observable, mitigation pulls the late-time points back.

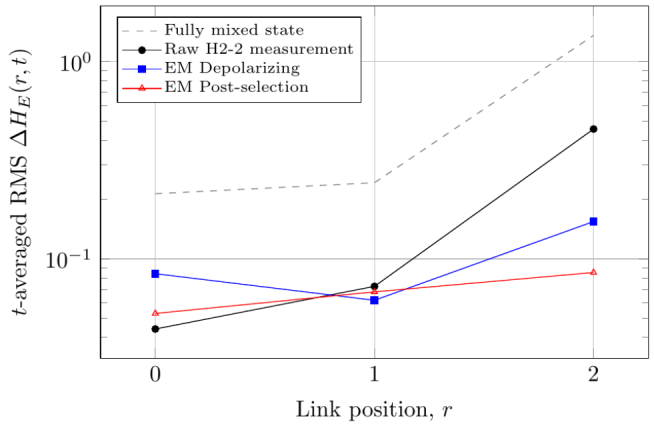
Mitigation cuts the deviation by 60–75%

RMS deviation from the noiseless baseline, all links and times

Fully mixed state	**0.806**
Raw H2-2 data	**0.267**
Depolarizing rescaling	**0.108**
Post-selection	**0.067**



Averaged over links, vs. time



Averaged over times, vs. link

Link 2 gains the most. Where raw data are already close, depolarizing rescaling can overcorrect (t = 3.75, link 0).

What we take away

- Fermions and a multiqubit gauge-field register can be implemented together and meaningfully probed on current hardware
- Raw data already carry real signal: about 67% below the fully mixed deviation
- Gauss’s-law checks catch some phase errors too: the constraints do a lot of work
- Post-selection works best, but its pass rate falls exponentially with lattice size; depolarizing rescaling is the scalable fallback (for now).
- Everything else scales polynomially: locality, circuit optimization, calibration. Next: more sites, higher truncations, then d > 1 and SU(3)

Product formulas and the inherited cost estimates

$$H = \sum_{\gamma=1}^{\Gamma} H_\gamma , \qquad \| V_2(T;s) - e^{-iTH} \| \leq \frac{\alpha_2 T^3}{s^2}$$

Each subevolution exp(−itHγ) needs its own circuit; non-Abelian ones are expensive, so the number of summands Γ matters

Byrnes & Yamamoto (2006)

Pauli decomposition in the electric basis; the number of Pauli terms grows as

$O\left(\Lambda^{8(N_c^2-1)}\right)$ and many terms do not commute

Kan & Nam (2021)

Gate count improves to $O(\Lambda \operatorname{polylog} \Lambda)$

but about 1e15 subevolutions per SU(3) plaquette

"Efficient" in the complexity-theory sense. Prohibitive in practice.

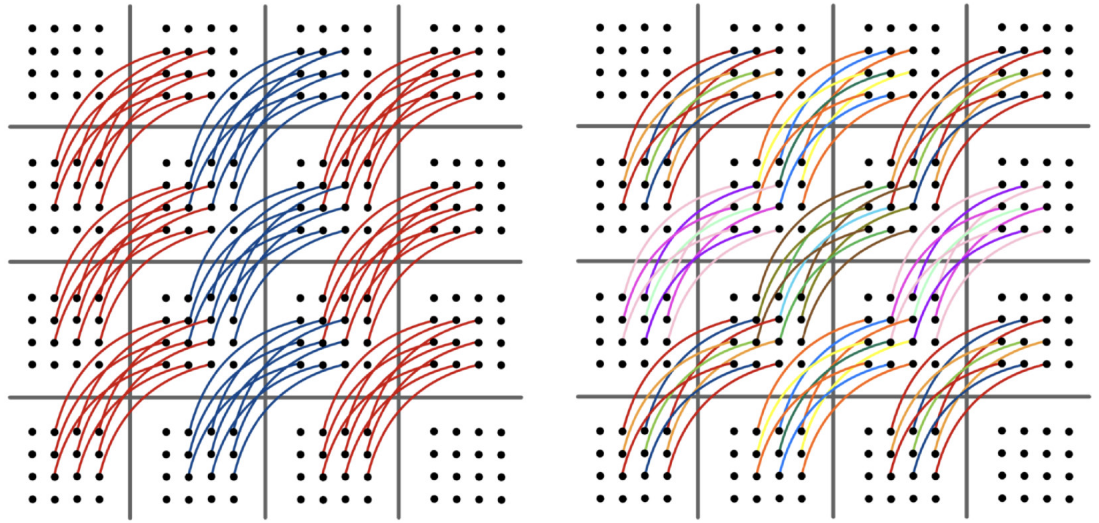
Split one ladder operator, not all of them

- Goal: summands that only rotate within 2×2 blocks, so each subevolution can be diagonalized and circuitized in closed form
- Kan & Nam split every ladder operator into even and odd parts; each split multiplies the count
- Our observation: one even-odd split of a single ladder factor is enough, using SVD-based diagonalization (Davoudi, Shaw & Stryker, Quantum 7, 2023)

Result: a factor $O\left(2^{4(N_c^2-1)}\right)$

fewer subevolutions per step

independent of cutoff, volume and simulation time



Graph picture of the U(1) plaquette Hamiltonian: one splitting needs 2 colors (a); four splittings need 2^4 (b). Each color is one subevolution.

Subevolutions per term: before and after

Term	Kan & Nam (quoted)	(tightened)	This work
U(1) hopping	2	2	1
U(1) plaquette	16	16	2
SU(2) hopping	64	64	8
SU(2) plaquette	1,048,576	524,288	256
SU(3) hopping	13,872	9,248	75
SU(3) plaquette	1.47e15	2.90e14	1,062,882

Electric (irrep) basis. “Tightened”: a careful recount of strictly necessary summands within the Kan & Nam scheme.

Z. Davoudi & J. R. Stryker, arXiv:2608.21258

~1e14

fewer T gates for QCD with a second-order product formula

regardless of cutoff, system size or simulation time

Honest comparison: algorithms are a live frontier

Near-optimal post-Trotter (Rhodes, Kreshchuk & Pathak)

About 1e3 fewer terms and 1e6 lower cost than our product formula, plus better scaling in time and error

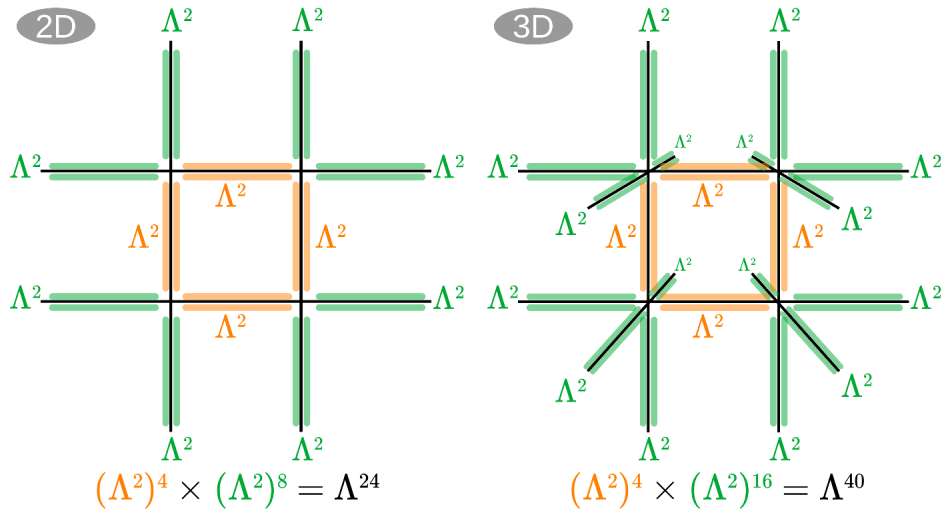
Local-multiplet basis (Ciavarella, Klco & Savage)

Compact qubit encoding, but plaquette cost grows as Λ^24 on square and Λ^40 on cubic lattices

Basis introduced for SU(2) in the first hardware simulation of a truncated non-Abelian gauge theory: Klco, Savage & Stryker, PRD 101, 074512 (2020)

Still missing for an end-to-end QCD cost

Truncation errors, continuum and infinite-volume limits, state preparation, measurement



Nonzero plaquette matrix elements in the local-multiplet basis: orange links on the plaquette, green external links

OUTLOOK

Open problems

An orthogonal SU(3) basis in d > 1

A seventh Casimir with analytically constructible eigenstates

Non-Abelian dynamics beyond 1+1D on hardware

Larger lattices, larger registers, higher truncations

String breaking and particle production

Tensor networks in richer theories, toward hadronization

Cheaper time evolution

Product formulas vs. post-Trotter methods; an end-to-end QCD cost

Hamiltonian meets Euclidean

Cross-checks where both methods work (statics); real-time dynamics, where only the Hamiltonian approach applies

Tools and partners in place

- LSH formulations for SU(2) and SU(3), with public code
- LSH tensor-network toolkit
- Collaboration with Quantinuum
- Hardware access through Berkeley Lab and NERSC

Each problem can be attacked analytically, with classical numerics, or on hardware.

SUMMARY

Three takeaways

1. Hamiltonian lattice gauge theory complements what Euclidean methods offer: especially for real time, finite density, topological terms.

2. Loop-string-hadron formulations make non-Abelian gauge invariance local and simple, for SU(2) and SU(3), and they already work on classical and quantum hardware.

3. The road to QCD is long but concrete: a real-time SU(2) simulation with a multiqubit gauge register on Quantinuum H2, and a ~1e14 cut in product-formula cost for QCD.

Thank you

THANKS

Acknowledgments

I’ve been lucky to work with many talented scientists based in the USA and around the world.

Soon my collaborator list will grow to include some based in Italy!



U.S. DEPARTMENT OF ENERGY | Office of Science

Backup

Degenerate sectors and a seventh Casimir

- Naive basis |ℓ12 ℓ23 ℓ31; ℓ21 ℓ32 ℓ13; t⟩ from loop operators L†IJ = A†(I)·B†(J) and trilinears TA†, TB†
- In the sector (p1,q1,p2,q2,p3,q3) = (1,1,1,1,1,1), the six irrep labels do not fix the state
- Counting LSH states evaluates SU(3) Littlewood–Richardson coefficients
- Gram–Schmidt always works, but is not analytic and gives no insight into the seventh label

One candidate

$$C_T \equiv (T_A T_B)^\dagger T_A T_B$$

Commutes with the (pi, qi); in the (1,1,1,1,1,1) sector:

$$\mathrm{Spec}(C_T) = \left\{0,\ \tfrac{80}{3}\right\}$$

$$|\phi_{1,2}\rangle = |111;000;0\rangle \mp |000;111;0\rangle$$

Goal: a Casimir whose eigenstates can be built analytically, like |j, m⟩ for SU(2)

PRD 111, 074516 (2025), arXiv:2407.19181; Mathematica notebook public on GitHub

Reduction factors by gauge group

Gauge group	Fewer magnetic-term subevolutions	Fewer T gates (2nd order)
U(1)	8	≈ 20
SU(2)	4.096e3	≈ 1e5
SU(3)	1.383e9	≈ 1e14

- Relative to Kan & Nam; independent of system size, simulation time and electric cutoff
- Our product formula and Kan & Nam both scale as O(Λ polylog Λ) in the cutoff; local-multiplet plaquettes as Λ^24 (2D) and Λ^40 (3D)
- Draper (2026) combines the near-optimal and local-multiplet approaches; a comparison is future work