nature physics

Supplementary information

https://doi.org/10.1038/s41567-025-02797-w

Simulating two-dimensional lattice gauge theories on a qudit quantum computer

In the format provided by the authors and unedited

CONTENTS

I. Quantum Computing for Particle Physics 2

II. Qudit Advantage 2

III. VQE Circuit for Pure-Gauge QED with Periodic Boundary Conditions 3

IV. Experimental Details 4
- IVa. Optimal Encoding and Cross-Talk Suppression 4
- IVb. Gate Count Optimization 5
- IVc. Measurement and Readout of Qudits 5

V. Noise Model 5

VI. Results for Alternative Lattice QED Model 6

References 8

Supplementary Note I: Quantum Computing for Particle Physics

Predictions in particle physics that involve the strong force can in general not be performed with perturbative methods and require therefore computer simulations. So far, lattice gauge theory (LGT) is the only general solution available in this endeavor. The standard approach of performing LGT simulations using the action formalism and Markov Chain Monte Carlo (MCMC) calculations has proven to be extremely successful but breaks down in situations that involve, for example, finite fermionic chemical potentials (i.e. finite matter densities) and real-time dynamics [20]. This limitation motivates an intensive ongoing search for new computational methods. Since quantum computing does not face the sign-problems that currently impede MCMC-based approaches, LGT quantum simulations offer an exciting opportunity to explore new avenues for gauge theory calculations in high energy physics [28–30]. In particular, quantum computers could allow making inroads into the two big problem areas outlined below.

Equilibrium physics – A concrete example of a problem that is out of reach for current MCMC-based methods is the phase diagram of quantum chromodynmics (QCD), which describes quarks and gluons at different fermionic chemical potentials and temperatures [86]. Forming the basis for understanding which states of matter are possible in Nature, the QCD phase diagram informs important open questions in nuclear physics, particle physics and cosmology, and could be explored by future LGT quantum simulations at finite densities.

Dynamics – Another specific example of a problem area that is currently inaccessible to MCMC calculations is the real-time evolution in collider physics [25]. Since there is no fundamental roadblock to study dynamics with quantum computers, LGT quantum simulations could open the door to a better understanding of particle collisions, i.e. the modelling and simulation of how the different degrees of freedom evolve in time during scattering processes observed at collider experiments at CERN and other facilities.

Traditional methods cannot even begin addressing the problem areas outlined above, as MCMC calculations do not converge and the method is unsuitable from the outset. Using quantum computers as a new tool for investigating the above examples, will involve a number of challenges including the following two key steps:

1. Simulating LGTs in 3D, which requires representing gauge degrees of freedom with three and more levels.

2. Simulating LGTs with non-Abelian symmetry groups, i.e. including color-degrees of freedom.

Supplementary Note II: Qudit Advantage

We now compare our native qudit implementation to a standard *qubit* implementation of the circuits presented in the main text. Here we follow the efficient paradigm outlined in Ref. [41], where high-dimensional gauge fields are encoded into qubits using a one-hot encoding: we map each d-level gauge field onto d qubits. All qubits are in the up state, except for one; which qubit is in the down state indicates which of the d levels is occupied. This has the advantage of greatly simplifying the required entangling gates and thus reducing the circuit depth, at the cost of an increase in qubit number. One may also consider a binary encoding, where we map a qudit onto $\lceil \log_2 d \rceil$ qubits. This optimizes the required register size at the cost of more complicated entangling gates. For the qudit platform, we make the assumption that an embedded CNOT gate (i.e. a CNOT gate acting on a qubit subspace of the qudit Hilbert space) can be performed with roughly equal fidelity as the same gate in a pure qubit system. This amounts to the assumption that the gate fidelity is independent of the presence or absence of unaffected spectator levels. We note that, while this assumption is reasonable in the trapped-ion context [9] it might not hold in all cases. Hence, the estimated fidelities in Tab. SI and Tab. SII should be understood in this context.

We now consider the scaling of circuit complexity with increasing gauge-field truncation to capture the physics more accurately. Starting with the full gauge theory of one plaquette in 2D, as shown in the main text Fig. 1, we propose a physically motivated generalization of the VQE circuit, which requires a linear increase in gate count with gauge field dimension. In particular, we show in Fig. S1 the generalization of the circuit when using a ququint truncation for the gauge field. More generally the qudit implementation of the circuit requires three CNOT for each of the qubit A gates, and two CNOT for each of the C-ROT for a total of $18 + 4(d - 1)$ CNOT gates. In the case of the qubit implementation studied in Ref. [54] $18 + 36(d - 1)$, since the implementation of a control-rotation, in this case, requires six CNOT. This shows that the scaling of the qudit platform is about an order of magnitude better than a comparable qubit platform, already in this simplest instance of the problem. Note that we use the CNOT-based gate count for cross-platform comparability, while our experiment is using a different decomposition of the gates, specific to trapped ions. The results are summarized in Tab. SI.

The results are even more striking in the case of a pure gauge theory as discussed in Fig. 2 of the main text. Here, the qudit platform shows no scaling with gauge field truncation, as the added complexity is only in the local qudit operations. In the qubit platform, on the other hand, the gate count scales as $48 + 12d$, and the register size scales as $3d$, see Tab. SII. Considering an encoding optimized for register size, this growth can be reduced to $3\lceil \log_2(d) \rceil$, at the cost of a gate count that

	Qudit encoding			Qubit encoding		
Dimension d	3	5	7	3	5	7
Register size	5	5	5	7	9	11
CNOT count	26	34	42	90	162	234
CNOT fidelity	99%					
Approx. circ. fid.	77%	71%	66%	40%	20%	10%
CNOT fidelity	99.5%					
Approx. circ. fid.	88%	84%	81%	64%	44%	31%

TABLE SI. **Circuit complexity of a single plaquette of the full gauge theory in 2D**. The complexity is measured in a hardware-agnostic way by the gate count when all entangling gates (qubit and qudit) are decomposed into two-level CNOT gates under the assumption that these gates work equally well in qudit as in qubit systems. More efficient decompositions might exist by using custom hardware-efficient gates.

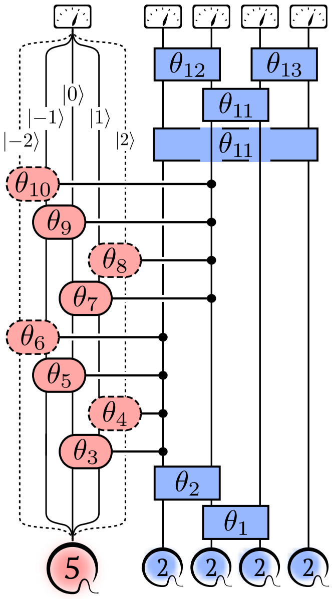

FIG. S1. Variational ansatz for simulating a single plaquette with open boundary conditions when the gauge field is represented by a ququint. The gates used in the circuit are defined in Methods MIV I.

scales in general quadratically with the dimension.

Finally, one might suspect that the price to pay for the reduced circuit complexity in the qudit approach is a less favorable runtime. While it is generally true that the qudit circuits take longer due to long readout times, we show in Tab. SIII that even the tradeoff between qudit readout ($(d-2)4$ ms) and qubit gate complexity (0.2 ms per gate) is slightly in favor of the qudit system for the problems studied here.

	Qudit encoding			Qubit encoding		
Dimension d	3	5	7	3	5	7
Register size	3	3	3	9	15	21
CNOT count	8	8	8	84	108	132
CNOT fidelity	99%					
Approx. circ. fid.	92%	92%	92%	43%	34%	27%
CNOT fidelity	99.5%					
Approx. circ. fid.	96%	96%	96%	66%	58%	52%

TABLE SII. **Circuit complexity of a pure gauge plaquette with periodic boundary conditions**. The complexity is measured in a hardware-agnostic way by the gate count when all entangling gates (qubit and qudit) are decomposed into two-level CNOT gates under the assumption that these gates work equally well in qudit as in qubit systems. More efficient decompositions might exist by using custom hardware-efficient gates.

	Qudit encoding			Qubit encoding		
Dimension d	3	5	7	3	5	7
Register size	3	3	3	9	15	21
Rel. runtime open	1	1.52	2.04	1.48	2.26	3.04
Rel. runtime periodic	1	1.54	2.08	1.76	2.08	2.41

TABLE SIII. **Runtimes of one execution of the qubit and qudit implementations**. These times are calculated for the device used in the main text and may differ in other architectures.

Supplementary Note III: VQE Circuit for Pure-Gauge QED with Periodic Boundary Conditions

We explain here the essential points of the design of the circuit used to study the pure-gauge QED model producing the results shown in Fig. 2d. A more detailed discussion is found in Ref. [54].

As for the case of 2D-QED with matter discussed in the main text, the form of the circuit is inspired by the Hamiltonian that the VQE solves, providing us with a scalable strategy that works also for higher truncations of the gauge field, and also with two different variational circuits for the electric and magnetic representation respectively. Let us consider the electric Hamiltonian given in Eqs. (MIII6), its ground state is found via the variational ansatz given by the circuit in Fig. 2c. After creating superpositions of all states for rotator $\hat{R}_2$, there is a layer of C-ROT gates that reflects the coupling between rotators given by $\hat{H}_E^{(e)}$, where we have used the symmetry between rotator $\hat{R}_1$ and $\hat{R}_3$ to reduce the number of independent variational parameters. For larger values of g^{-2}, the system tends to occupy states where the three rotators occupy the same electric field eigenstate, therefore we spread the population of rotators 1 and 3 controlling suitably on the corresponding state of rotator 2.

The circuit for the magnetic Hamiltonian given in Eqs. (MIII8) and depicted in Fig. S2 is designed with similar arguments.

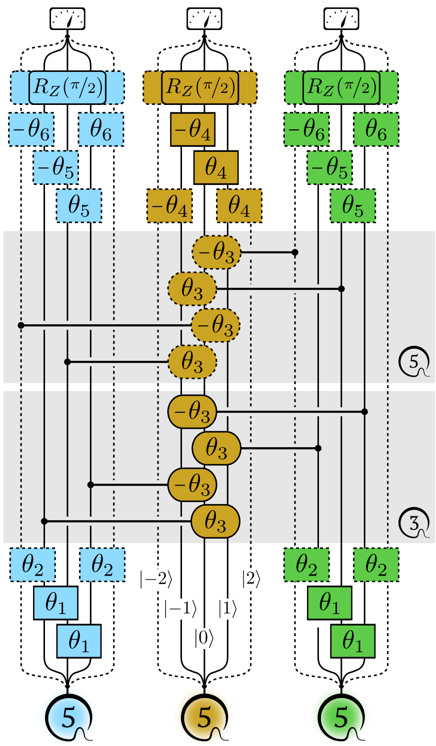

FIG. S2. Variational circuit for the magnetic representation for the qutrit (solid lines) and the ququint (all, except box marked with qutrit) truncation. The explicit form of the gates employed is given in Methods MIV I. The last gate in the circuit is defined as $\hat{R}_Z(\theta) = \exp\left(-i\theta\hat{Z}\right)$, where $\hat{Z} = \mathrm{diag}(-1, 0, 1)$ for the qutrit, and $\hat{Z} = \mathrm{diag}(-2, -1, 0, 1, 2)$ for the ququint truncation.

Supplementary Note IV: Experimental Details

IVa. Optimal Encoding and Cross-Talk Suppression

Employing embedded C-ROTs requires careful treatment of the frequency spectrum and the structure of the quantum circuit to optimize the encoding of qudits in trapped ions. From the pure gauge variational ququint circuit shown in Fig. 2c, it is evident that only few states participate in the C-ROTs, namely the states $|\pm 2\rangle$ of control ququint and $|0\rangle$, $|\pm 1\rangle$ of the two target qudits. Optimal encoding yields the strongest coupling of the qudit states with the auxiliary ground state $|g\rangle$ on the blue sideband while simultaneously minimizing accumulating phase errors as a consequence of imperfect AC Stark shift compensation. The latter is achieved by restricting transitions, which lead to large Stark shifts, to static angles of $\pm\pi$, making the shift constant and easier to compensate. Figure S3 shows the chosen encodings for control and target qudits for the ququint circuit shown in Fig. 2c for the electric representation. The states with the largest coupling to $|g\rangle$ are highlighted in blue. Note, that in the magnetic representation, a different encoding is chosen.

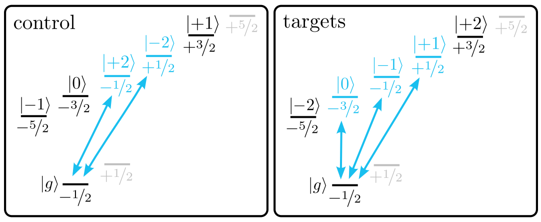

FIG. S3. Optimal encodings of control and target qudits in the $D_{5/2}$ manifold of $^{40}\mathrm{Ca}^+$ applied for the pure gauge ququint circuit in Fig. 2c – the encoding is chosen such that relevant states for the C-ROTs have the strongest coupling to the blue sideband. We seek to minimize phase errors by allowing only π rotations on all transitions, which induce large Stark shifts on the spectator levels.

The pure-gauge qutrit circuit can be constructed in such a way that spatially neighboring ions store the qutrit in different sets of Zeeman levels, with each set separated by frequencies of a few MHz. Assigning integers to each ions in the string from left to right, ions with even indices encode the qutrit in Zeeman states with positive magnetic quantum numbers $|-1'\rangle = D_{+1/2}$, $|0'\rangle = D_{+3/2}$ and $|+1'\rangle = D_{+5/2}$, coupled via the second ground state $|g'\rangle = S_{+1/2}$ as shown in Fig. S4. With this spectroscopic decoupling scheme, next-neighbor cross-talk is effectively eliminated [12] and higher axial trap frequencies can be used. However, spectroscopic decoupling is not possible for ququints, as no bi-partition of the $D_{5/2}$ manifold exists in $^{40}\mathrm{Ca}^+$ for qudits of dimension $d > 3$; here one would require different ion species or isotopes to apply this scheme. A second approach to reduce next-neighbor cross-talk is to increase the spatial separation between ions encoding the qudits. This can be achieved in two ways, either lowering the axial confinement by reducing the end-cap voltage of the ion trap or by introducing additional *buffer* ions, which are transferred to a *hidden* state which does not couple with the laser pulses driving the actual quantum circuit. The former strategy is less suited for our specific circuits due to significant *crowding*

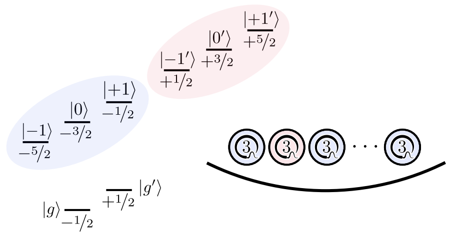

FIG. S4. Next-neighbor cross-talk suppression in qutrits by proper choice of encoding – spatially neighboring ions encode the qutrit in different sets of Zeeman states with each set separated by a few MHz. This encoding pattern is well suited to suppress cross-talk [12], allowing higher axial trap frequencies, reducing the inter-ion distance.

of the mode spectrum, making it difficult to faithfully implement the C-ROT operations. Thus, for the $d = 5$ pure gauge circuit as well as the gauge-matter circuit we chose the second method and introduced a pair of extra ions - in the pure gauge circuits, these buffers are located between each qudit, while the open plaquette circuit requires a buffer between each of the matter qubits, while we order the ion string in such a way that the gauge field qudit is physically located in the center of the string as shown in Fig. S5. This re-ordering ensures optimal performance of the Mølmer-Sørensen gates on the matter qubits at the very end of the circuit. Additionally to the

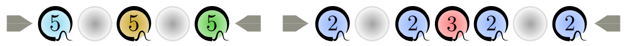

FIG. S5. Artificially increasing the inter-qudit distance by introducing *buffer* ions. Buffer ions do not contribute to the circuit and do not couple with the light fields used to control the qudits. This strategy is employed for the $d = 5$ pure gauge system, as well as for the circuits describing dynamical matter. For the latter case, we also chose to re-order the qudits inside the ion string to enhance the performance of the Mølmer-Sørensen entangling gates on the qubits.

two approaches presented in this section, we also employ composite rotations [12] when transferring the buffer ions to their hidden state.

IVb. Gate Count Optimization

The structure of the quantum circuits enables us to implement shortcuts, reducing the number of entangling gates; these shortcuts are best explained by the circuit in Fig. 1d. First, if the input state of one of the qudits is known, the number of interactions with the phonon mode $B(\theta, \phi)$ can be reduced to $n_{\mathrm{B}(\theta,\phi)} = 2$ compared to the general case of a C-ROT being composed of $n_{B(\theta,\phi)} = 5$ BSB pulses. This strategy is applied to the matter *qubits* in the gauge-matter-circuit for the angles θ_1, θ_2 and θ_7, where the state of the *'left'* qubit entering each operation is exactly known. Second, if two C-ROTs are executed successively on the same control state as it would be for the pairs (θ_3, θ_4) and (θ_5, θ_6) in Fig. 1d, two successive $B(\theta, \phi)$ cancel out as it can be imagined by extending the sequence in Extended Data figure M2b by a second C-ROT. Third, we implement the SWAP operations with angles θ_7, θ_8 and θ_9 by sequences of three Mølmer-Sørensen (MS) gates [84] and local single-qubit rotations, which, in this case, is more favorable than the blue-sideband approach in terms of gate count.

IVc. Measurement and Readout of Qudits

Measurements on a qudit with dimension d are implemented by a series of single-qudit operations followed by a sequence of $d-1$ detection events and local π rotations. Our optimal encoding of the qudits requires an additional reordering of the quantum states to match a W-pattern as shown in Fig. S6. In this arrangement any state $|i\rangle$ directly couples with $|i \pm 1\rangle$, which enables a straightforward QR-decomposition-based compilation of the measurement operator into a sequence of single qudit rotations [9]. After the measurement operation has

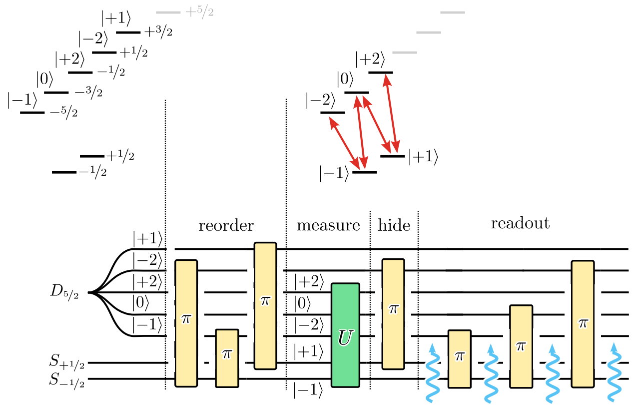

FIG. S6. Measurement and readout of ququints – the ququint is initially encoded in the $D_{5/2}$ manifold of $^{40}Ca^+$ similar to Fig. S3 (left). It is then reordered by a sequence of π-pulses, such that the states are arranged in the shape of the letter W (right). The coupling between states $|i \pm 1\rangle \leftrightarrow |i\rangle$ is required to faithfully implement the measurement operation U in green. After hiding $|+1\rangle$ in the $D_{5/2}$ manifold, the readout sequence is executed as a series of detection and re-cooling & optical pumping pulses indicated by blue arrows. Additional π-pulses rotate the desired state to the ground state manifold, such that photons can be scattered on the short-lived $S_{1/2} \leftrightarrow P_{1/2}$ transition (see Extended Data figure M2a).

been applied, the readout is implemented by a sequence of $d-2$ local π-flips and $d-1$ detection events. Each event is composed of a 1 ms window, during which the fluorescence of each ion is recorded on a sensitive EMCCD camera, followed by 2 ms of Doppler cooling and 0.5 ms of polarization gradient cooling; this re-cooling pattern is not required after the last detection. We mention, that the duration of the readout in our experiment scales as $\approx (d-2) \cdot 3.5$ ms $+ 1$ ms for any qudit dimension d – this is on the order of the duration of the ground state cooling sequence at the beginning of each experiment run.

Supplementary Note V: Noise Model

While systems consisting of qubits only can often be accurately described by generic noise models such as simple depolarising noise, we found that for our qudit systems this is too reductive. Therefore, we construct a more realistic noise model based on the physics of how the gates of our circuits are realized. In particular, our heuristic model includes imperfect precision in the variational parameters and dephasing of the prepared quantum state. We find good agreement between the experimental results and the simulated data.

Amplitude Fluctuations. We expect that the precision in the gate angles we aim to apply throughout the experiment suffers from fluctuations of the addressing laser amplitude. A gate $\hat{U}(\theta)$ is hence replaced by $\hat{U}(\theta + \delta\theta)$, where $\delta\theta$ is a stochastic variable extracted from a Gaussian $\mathcal{N}(0, \sigma_{\mathrm{a}}^2)$ with mean $\mu = 0$ and variance σ_{a}^2. Given the higher fidelity with which single qudit gates can be realized compared to entangling gates, we use in our model a variance σ_{a}^2 for parametrized entangling gate, and a variance $\sigma_{\mathrm{a}}^2/10$ for single qudit gates.

Phase Errors. In trapped ion systems, while single qudit operations can be carried out with a fidelity close to unity, two-qudit entangling gates and controlled excitation of the ions' motion are the dominant sources of errors. C-ROT operations in qudits introduce phase noise due to AC Stark shifts of different magnitude (see Methods MI): (1) exciting the motion on a chosen control state $|i\rangle$ introduces a (small) phase error ϕ_i on that state due to imperfect compensation and (2) spectator states that are not involved in the interaction ($|j\rangle$, with $j \neq i$) accumulate a phase Φ_j due to imperfect compensation of AC stark shifts. Each C-ROT is a composition of single qudit operations as shown in Extended Data figure M2b—for any blue sideband operation exciting the motion via the state $|i\rangle$, we identify the mapping $c_i |i\rangle + \sum_{j\neq i} c_j |j\rangle \mapsto e^{i\phi_i} c_i |i\rangle + \sum_{j\neq i} e^{i\Phi_j} c_j |j\rangle$. As can be deduced by counting how many times each level in Extended Data figure M2b is involved in a sideband pulse, the control qudit is mapped to the state in Eq. (S1), while the target acquires phases according to Eq. (S2).

$$c_i |i\rangle + \sum_{j\neq i} c_j |j\rangle \mapsto e^{i2\phi_i} c_i |i\rangle + \sum_{j\neq i} e^{i2\Phi_j} c_j |j\rangle , \tag{S1}$$

$$\begin{aligned} & c_i |i\rangle + c_j |j\rangle + \sum_{k\neq i,j} c_k |k\rangle \mapsto \\ & e^{i(2\phi_i+\Phi_i)} c_i |i\rangle + e^{i(2\Phi_j+\phi_j)} c_j |j\rangle + \sum_{k\neq i,j} c_k e^{i3\Phi_k} |k\rangle . \end{aligned} \tag{S2}$$

The second-beam technique [87] applied to cancel the AC Stark shift on the *actively* driven transition $|g\rangle \leftrightarrow |i\rangle$ is robust to intensity noise, as both beams are affected equally by fluctuations. Thus, the phase error ϕ_i is considered to be small compared to Φ_j and can be neglected for simplicity. Furthermore, we make the simplifying assumption that the passively (in software) compensated spectator shifts yield the same phase error for all levels $\Phi_j \mapsto \Phi$, and thus the mappings simplify to Eq. (S3) and Eq. (S4).

$$c_i |i\rangle + \sum_{j\neq i} c_j |j\rangle \mapsto c_i |i\rangle + e^{i2\Phi} \sum_{j\neq i} c_j |j\rangle , \tag{S3}$$

$$\begin{aligned} & c_i |i\rangle + c_j |j\rangle + \sum_{k\neq i,j} c_k |k\rangle \mapsto \\ & e^{i\Phi} c_i |i\rangle + e^{i2\Phi} c_j |j\rangle + e^{i3\Phi} \sum_{k\neq i,j} c_k |k\rangle . \end{aligned} \tag{S4}$$

Consecutive C-ROTs controlled by the same qudit and same state (see for example circuit in Fig. 1d) can be contracted, reducing the number of blue sideband pulses. In this specific case, the pattern on the target *qutrit* is then given by

$$\sum_{j\in\{-1,0,+1\}} c_j |j\rangle \mapsto e^{i3\Phi} c_{-1} |-1\rangle + e^{i2\Phi} c_0 |0\rangle + e^{i3\Phi} c_{+1} |+1\rangle , \tag{S5}$$

while the control qudit acquires $e^{i\Phi}$ on all spectator levels. We model the phase error Φ with a Gaussian distribution $\mathcal{N}(0, \sigma_{\mathrm{p}}^2)$ with mean $\mu = 0$ and variance σ_{p}^2.

Results. In Fig. S7 we compare the experimental VQE results with simulated data for the open plaquette (first row) and for the pure gauge theory (second and third row). In the latter case the blue (orange) markers correspond to the electric (magnetic) representation. For the simulation, we consider variational parameters obtained without noise, and we calculate the expectation values of the relevant observables when the variational circuit is affected by the noise with parameters estimated as $\sigma_{\mathrm{a}}^2 = 0.0.047$, $\sigma_{\mathrm{p}}^2 = 0.073$. We can see good qualitative agreement between the experimental and the simulated data.

Supplementary Note VI: Results for Alternative Lattice QED Model

In this work, we consider compact lattice QED (cQED) with staggered fermions [51]. There are two different models that are both referred to as cQED and that are both currently discussed in the recent literature. The main text of this paper focuses on the first of these two models. As explained in Methods MII, this model is obtained by discretising space in two-dimensional QED [74, 76]. In this section, we focus on the second version of cQED. This model has been studied in the literature on quantum simulations of lattice gauge theories [52, 54, 61, 63], especially in condensed matter physics [64, 88, 89]. Both models respect a local U(1) gauge symmetry, but the kinetic parts of the Hamiltonian have different phases, which leads to different theories in the continuum limit.

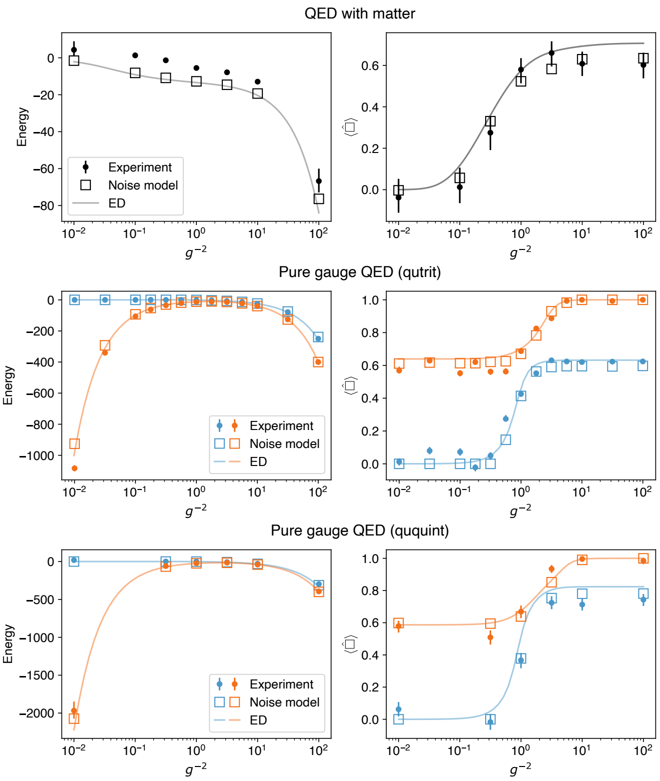

FIG. S7. **Simulated noise model for the experiment.** Energy (first column) and plaquette expectation value (second column) for the two models studied via experimental VQE, namely 2D-QED with matter (first row), and an instance of pure gauge 2D-QED (second and third row, corresponding to qutrit and ququint truncations respectively). For the pure gauge model, blue (orange) lines and markers correspond to the simulations performed with the electric (magnetic) representation. After obtaining variational parameters in the absence of noise, the simulated noisy data (squares) is obtained via the noise model detailed in the text, which includes amplitude fluctuations and phase errors. The simulated noisy data are in good agreement with the measured results from experimental quantum computations (dots) shown with one standard deviation of statistical uncertainty from Monte-Carlo resampling around the measured value, averaged over 150 (300) repetitions for qutrits (ququints).

The Hamiltonian in the second case is given by [61]

$$\hat{H}_E = \frac{1}{2}\sum_{\mathbf{n}}\left(\hat{E}^2_{\mathbf{n},\mathbf{e}_x} + \hat{E}^2_{\mathbf{n},\mathbf{e}_y}\right), \tag{S1a}$$

$$\hat{H}_B = -\frac{1}{2}\sum_{\mathbf{n}}\left(\hat{P}_{\mathbf{n}} + \hat{P}^\dagger_{\mathbf{n}}\right), \tag{S1b}$$

$$\hat{H}_m = \sum_{\mathbf{n}}(-1)^{n_x+n_y}\hat{\phi}^\dagger_{\mathbf{n}}\hat{\phi}_{\mathbf{n}}, \tag{S1c}$$

$$\hat{H}_k = \sum_{\mathbf{n}}\sum_{\mu=x,y}\left(\hat{\phi}_{\mathbf{n}}\hat{U}^\dagger_{\mathbf{n},\mathbf{e}_\mu}\hat{\phi}^\dagger_{\mathbf{n}+\mathbf{e}_\mu} + \text{H.c.}\right). \tag{S1d}$$

Since the two models only differ by sign choice in the kinetic term (compare Eq. (MIII2d)), they coincide for a pure gauge theory, as considered in Sec. IV of the main text. After eliminating the gauge fields as shown in Ref. [54] and applying the Jordan-Wigner transformation given in Eq. (MIII1) the resulting qubit Hamiltonian is given by:

$$\begin{aligned}\hat{H}_E = \frac{1}{4}(8\hat{E} + 2\hat{E}\,(-\hat{\sigma}^z_1 - 2\hat{\sigma}^z_2 + 3\hat{\sigma}^z_3 + 2) \\ - \hat{\sigma}^z_2 + \hat{\sigma}^z_3 + \hat{\sigma}^z_1\hat{\sigma}^z_2 + 3),\end{aligned} \tag{S2a}$$

$$\hat{H}_B = -\frac{1}{2}\left(\hat{U} + \hat{U}^\dagger\right), \tag{S2b}$$

$$\hat{H}_m = \frac{1}{2}\left(\hat{\sigma}^z_1 - \hat{\sigma}^z_2 + \hat{\sigma}^z_3 - \hat{\sigma}^z_4\right), \tag{S2c}$$

$$\hat{H}_k = \hat{\sigma}^+_1\hat{\sigma}^-_2 + \hat{\sigma}^+_2\hat{U}^\dagger\hat{\sigma}^-_3 + \hat{\sigma}^+_4\hat{\sigma}^-_3 - \hat{\sigma}^+_1\hat{\sigma}^-_4 + \text{H.c.}, \tag{S2d}$$

We studied this model experimentally using the same VQE circuit structure as in Fig. 1d (see Ref. [54]). The gates that entangle the qubits and the qudit are still controlled on the qubits adjacent to the dynamical gauge field, which in this case are qubits ii and iii, as shown in the inset of Fig. S8. The experimentally measured plaquette expectation values in the ground state found via VQE as a function of g^{-2} are displayed in Fig. S8 for $\Omega = 5$ and $m = 0.1$, and show good agreement with the theoretical prediction.

In the large coupling regime ($g^{-2} \ll 1$) the electric field energy term $\hat{H}_E$ dominates the Hamiltonian of Eq. (S2), favoring the ground state $|vvvv, 0\rangle$ with $\langle\hat{\square}\rangle = 0$. In the weak coupling regime, on the other hand, the magnetic field energy term $\hat{H}_B$ dominates the Hamiltonian, favoring a positive vacuum plaquette expectation value ($\langle\hat{\square}\rangle = \frac{1}{\sqrt{2}} \approx 0.707$ for the chosen truncation). In the intermediate regime where $g^{-2} \approx 1$, we encounter a competition between the two field-energy terms and $\hat{H}_k$. The presence of the kinetic term $\hat{H}_k$ makes it energetically favorable to create pairs, which in turn produce a background magnetic flux that opposes the flux present in the ground state preferred by the term $\hat{H}_B$. On its own, the term $\hat{H}_k$ leads to a ground state with a negative plaquette expectation value ($\langle\hat{\square}\rangle \approx -0.642$ in the considered case). The presence of dynamical matter, therefore, results in a reduction of the plaquette expectation value in the ground state, as shown in Fig. S8. Notably, even with our small system size we already clearly observe the dip in the plaquette expectation value that originates from pair creation processes in the ground state, see Ref. [54] for details.

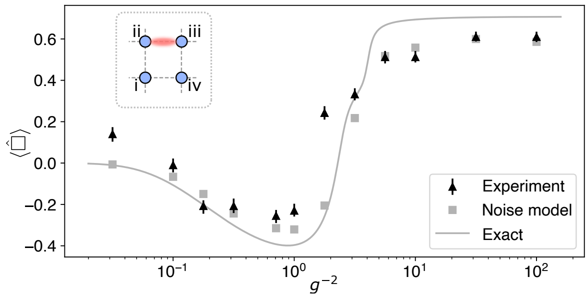

FIG. S8. Experimental VQE results for the plaquette expectation value for the model described in Eqs. (S2). Errorbars indicate one standard deviation of statistical uncertainty from Monte-Carlo resampling around the measured value, averaged over 150 repetitions.

[86] S. Bogdanov, E. Fonseca, R. Kashyap, A. Kurkela, J. M. Lattimer, J. S. Read, B. S. Sathyaprakash, H. T. Cromartie, T. Dietrich, A. Dhani, T. Dolch, T. Gorda, S. Guillot, W. C. G. Ho, R. Huxford, F. K. Lamb, P. Landry, B. W. Meyers, M. C. Miller, J. Nättilä, R. Paatelainen, C. Prescod-Weinstein, S. Säppi, I. H. Stairs, N. Stergioulas, I. Tews, A. Vuorinen, Z. Wadiasingh, and A. L. Watts, Snowmass 2021 Cosmic Frontier White Paper: The Dense Matter Equation of State and QCD Phase Transitions (2022), arXiv:2209.07412 [astro-ph, physics:gr-qc, physics:hep-th, physics:nucl-ex, physics:nucl-th].

[87] H. Häffner, C. F. Roos, and R. Blatt, Quantum computing with trapped ions, Phys. Rep. **469**, 155 (2008).

[88] I. Affleck and J. B. Marston, Large-n limit of the Heisenberg-Hubbard model: Implications for high-${T}_{c}$ superconductors, Physical Review B **37**, 3774 (1988).

[89] W. Rantner and X.-G. Wen, Electron Spectral Function and Algebraic Spin Liquid for the Normal State of Underdoped High ${T}_{c}$ Superconductors, Physical Review Letters **86**, 3871 (2001).