

Recovering What Breaks in a Trapped-Ion Quantum Simulation of a Lattice Gauge Theory

Hassan Ugail

h.ugail@bradford.ac.uk

University of Bradford

Research Article

Keywords: lattice gauge theory, trapped ions, open quantum systems, Gauss's law, symmetry verification, post-selection, error mitigation

Posted Date: October 19th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11353424/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

RECOVERING WHAT BREAKS IN A TRAPPED-ION QUANTUM SIMULATION OF A LATTICE GAUGE THEORY

Hassan Ugail
Centre for Visual Computing and Intelligent Systems
University of Bradford
United Kingdom

October 12, 2026

ABSTRACT

A recent trapped-ion experiment realised a lattice gauge theory with the gauge field in the internal states of the ions and bosonic matter in their vibrations. We showed earlier that its noise damage can be predicted. Here, we show that most of it can be undone. The conservation law of the theory is a parity of each vibrational mode, which heating flips, so a heating event leaves a mark that can be read without disturbing the physics. We keep only the runs in which this mark never changed. With equal heating rates and continuous monitoring, the kept runs reproduce the heating-free dynamics exactly, with a closed-form yield. Reading the mark at finite intervals introduces a small bias linear in the interval, which we derive and bound. At the reported rates, a single final check removes most of the heating and preparation damage while accepting eight runs out of ten.

Keywords lattice gauge theory · trapped ions · open quantum systems · Gauss’s law · symmetry verification · post-selection · error mitigation

1 Introduction

A lattice gauge theory ties the motion of matter to a field that lives on the links between lattice sites. The local conservation law that expresses this tie is Gauss’s law, and a quantum simulation of a gauge theory is judged by how well it respects it [1, 2, 3]. Over the past decade, the simplest discrete example, the $\mathbb{Z}_2$ lattice gauge theory coupled to dynamical matter, has become a laboratory object. It has been realised with trapped ions [4], with ultracold atoms in Floquet-engineered double wells [5], and with superconducting qubits in which both matter and gauge field are spins [6]. Most recently, it has been realised on a trapped-ion processor in which the gauge field is encoded in the internal state of each ion and the matter in the quantised vibrations of the crystal [7, 8]. In that experiment, two ions close a two-link loop. The link qubits are prepared in Bell states that fix a $\mathbb{Z}_2$ flux through the loop, and a phonon either tunnels freely or is frozen in place by Aharonov–Bohm interference [9], depending on the flux.

The hybrid qubit–oscillator encoding gives the matter a feature that spin and hard-core encodings lack. A site can hold any number of excitations, and the Gauss generator at a site is the product of the phonon parity of that site and the link qubits that meet it. The local conservation law is therefore a parity, and a parity can be measured without learning anything else about the state. Three noise processes dominate the device. The vibrational modes heat, their motional coherence decays, and the qubits suffer errors. In a companion paper, we showed that the damage each process does to each observable class of the experiment can be predicted before the machine is run, as certified bounds and first-order laws derived from the reported noise rates [10]. The central fact of that paper is that heating violates Gauss’s law by flipping a parity at a rate that grows with the number of quanta already in the mode.

The present paper turns that fact around. A heating event flips a parity, and the parity can be read without disturbing the gauge-invariant dynamics, so a heating event leaves a mark. One can read the mark at intervals and keep only the runs in which it never changed. This is symmetry verification by post-selection, and it has a substantial literature. It was proposed as a low-cost error mitigation strategy for variational algorithms [11, 12] and has since been applied to

gauge theories in several forms. Sparse error detection with a final post-selection on the gauge charges has been studied for Hamiltonian simulation [13]. Dynamical post-selection on mid-circuit Gauss-law checks has been worked out for non-Abelian theories on qudits [14]. Full stabiliser codes built from Gauss's law have been constructed for hard-core $\mathbb{Z}_2$ matter [15]. The older idea of pinning the dynamics to the gauge-invariant subspace by an engineered measurement or dissipation is due to [16], and the reliability of the gauge-invariant dynamics under coherent and incoherent errors has been analysed in general [17, 18, 19]. All of this work treats matter that is a spin or a hard-core boson. For such matter under Pauli noise, the jump operators are unitary. The evolution conditioned on no error is then the ideal evolution by construction, and post-selection removes every detected error exactly. For bosonic matter, the jump operators are not unitary. The evolution conditioned on no error carries a backaction that in general distorts the surviving state, and post-selection on a symmetry is then only approximate, however well the errors are detected.

The first result of this paper is that for bosonic matter under heating, the backaction is nevertheless trivial, so the post-selected estimator is exact. The mechanism itself is not new. A dissipative channel whose total jump rate depends only on a conserved quantity produces a common decay factor on each sector of that quantity and leaves the normalised state inside the sector untouched. In the bosonic-code literature, this is why two-mode codes with a fixed total excitation number are unaffected by the no-jump evolution of uniform loss [20], and the construction has recently been formalised as uniformly decaying subspaces for error-mitigated computation [21]. What this paper adds is the recognition that the gauge-invariant dynamics of the ion platform supplies such a conserved quantity, the total phonon number, and that the Gauss generators supply a syndrome that reads the heating history without disturbing it. In essence, they give an exact post-selected estimator with a closed-form yield on any geometry and at any electric field. The result holds for any heating model in which the sum of the gain and loss rates is the same on every mode.

The rest of the paper works out what happens when the ideal conditions are relaxed, because the experiment will relax them, and this is where the main analytic content lies. The syndrome cannot be read continuously. A finite reading interval admits runs in which two heating events on the same mode cancelled between readings. Such return events are excluded by hypothesis in the uniformly decaying subspace construction [21]. The heating channel produces them at any occupation because an empty mode can gain a quantum and lose it again, whereas pure loss needs at least two quanta on a mode. We show that the resulting bias grows linearly in the interval. We derive its coefficient and the accompanying excess yield from the heating-free dynamics alone, and we give a certificate for the bias of any bounded observable. For a state of definite phonon number the certificate needs only the measured fraction of accepted runs and the closed-form yield. The heating rates of the modes are not equal in practice, and unequal rates tilt the conditional state by an amount with a closed form at zero electric field. The matter is prepared on a thermal background rather than in a Fock state. A syndrome measurement handles this too, because projecting on the nominal syndrome fixes the parity of every mode and so removes most of the thermal background along with the heating. Finally, we state what a destructive Gauss-law measurement at the end of each run can test, with the readout it requires, and what would need a non-demolition measurement of the generators during the run. The methodology follows the certification programme we developed for spectral diagnostics in numerical relativity [22] and for noisy quantum neural networks [23, 24]. Every predictive law is accompanied by a certified bound, and every bound is tested on an exact computation.

2 Model, noise and the Gauss-law syndrome

We use the model and notation of [10], which follow [7, 8]. A graph has sites m carrying oscillators with ladder operators $\hat{a}_m$, and links $\ell = (i, j)$ carrying qubits with Pauli operators $\hat{\sigma}_\ell^x$, $\hat{\sigma}_\ell^y$ and $\hat{\sigma}_\ell^z$. The Hamiltonian of the $\mathbb{Z}_2$ lattice gauge theory on this graph is


$$\hat{H} = \sum_{\ell=(i,j)} \left[J(\hat{a}_j^\dagger \hat{\sigma}_\ell^z \hat{a}_i + \text{h.c.}) + h \hat{\sigma}_\ell^x \right], \quad (1)$$

where J is the gauge-invariant tunnelling and h is the electric field. A phonon hops between sites only through the link qubit, which it flips in the electric basis. The generator of the local symmetry at site m is

$$\hat{G}_m = \hat{P}_m \prod_{\ell \ni m} \hat{\sigma}_\ell^x, \quad \hat{P}_m = e^{i\pi \hat{n}_m}, \quad (2)$$

with $\hat{n}_m = \hat{a}_m^\dagger \hat{a}_m$, and $[\hat{H}, \hat{G}_m] = 0$ at every site. The eigenvalue $g_m = \pm 1$ labels the background charge at m . Gauss's law is the statement that the dynamics stays in the sector fixed by the initial state. The flux through a closed loop C is measured by $\hat{W}_C = \prod_{\ell \in C} \hat{\sigma}_\ell^z$. On the two-link loop of the experiment, preparing the link qubits in $|\Phi_{\ell_1 \ell_2}^+\rangle$ fixes the flux to zero, and preparing them in $|\Psi_{\ell_1 \ell_2}^+\rangle$ fixes it to π . We write $\hat{\Pi}_\Phi = \frac{1}{2}(1 + \hat{\sigma}_{\ell_1}^z \hat{\sigma}_{\ell_2}^z)$ for the projector onto the zero-flux sector. The Hamiltonian conserves the total phonon number $\hat{N} = \sum_m \hat{n}_m$, since each hopping term moves one phonon between sites and the electric term acts on the qubits alone. Figure 1(a) shows the loop and its generators.

(a) Two-link loop and its Gauss generators



(b) Sector tracking: keep the shots whose syndrome never changed




Figure 1: Sector tracking on the two-link loop. (a) Two sites m_1 and m_2 carry oscillators and two links ℓ_1 and ℓ_2 carry qubits. The Gauss generator at each site is the phonon parity of that site multiplied by the two link qubits in the electric basis. (b) Three trajectories of the quantum-jump unravelling. The syndrome, which is the pair of generator eigenvalues, is read at intervals τ . A run with no heating jump is kept and equals the heating-free dynamics exactly. A run with one jump since the last reading has a flipped syndrome and is rejected. A run with two jumps on the same mode inside one interval has an unchanged syndrome and is kept. Runs of the last kind are the source of a bias that is linear in τ .

The device noise is modelled by a Lindblad master equation $\dot{\rho} = -i[\hat{H}, \rho] + \sum_k \mathcal{D}[\hat{L}_k]\rho$ with $\mathcal{D}[\hat{L}]\rho = \hat{L}\rho\hat{L}^\dagger - \frac{1}{2}\{\hat{L}^\dagger\hat{L}, \rho\}$ [25, 26, 27]. Motional heating of mode m at $\hat{n}_m$ quanta per second is modelled as an infinite-temperature bath with jump operators $\sqrt{\hat{n}_m}\hat{a}_m$ and $\sqrt{\hat{n}_m}\hat{a}_m^\dagger$. This is the standard model for electric-field noise on a trap electrode [28, 29]. Motional dephasing with coherence time $t_{c,m}$ is modelled by the jump operator $\sqrt{2/t_{c,m}}\hat{n}_m$. Qubit errors on link ℓ are modelled as bit flips $\sqrt{\gamma_x}\hat{\sigma}_\ell^x$ and as dephasing $\sqrt{\gamma_z}\hat{\sigma}_\ell^z$. For the two radial modes used on the loop the experiment quotes heating rates of 5 and 17 quanta per second, coherence times of 2.3 and 3.3 milliseconds, an initial thermal occupation $\bar{n}_0 = 0.1$ on every mode, and a loop tunnelling rate $J/\pi = 0.35$ kilohertz [8]. It does not quote qubit error rates. We use the reported values in Section 6. Elsewhere, we raise the heating rates to 150 quanta per second, so that the heating damage within the experimental window of 2.32 milliseconds is large enough for the recovery to be visible. Nothing in the analysis depends on the rates being small.

The role of each channel with respect to Gauss's law was established in [10], and the method rests on it, so we restate it. Motional dephasing is a function of the phonon number and commutes with every generator. Qubit bit flips act with $\hat{\sigma}^x$ and also commute with every generator. Heating adds or removes one phonon and flips the parity $\hat{P}_m$ of the mode it acts on, which flips $\hat{G}_m$ and no other generator. Qubit dephasing acts with $\hat{\sigma}_\ell^z$, which anticommutes with $\hat{\sigma}_\ell^x$, so it flips the generators at both ends of the link. We call the vector of generator eigenvalues $(g_m)_m$ the syndrome. A heating event on mode m changes the syndrome at m alone. A qubit dephasing event on link (i, j) changes it at i and j together. The syndrome is therefore a record of the heating and qubit-dephasing history of a run. When the initial state is an eigenstate of every generator, it carries no information about the gauge-invariant dynamics, because the Hamiltonian commutes with all of them. Reading the syndrome collapses nothing that the experiment wants to measure.

Sector tracking is the protocol of reading the syndrome at times $\tau, 2\tau, \dots$ up to the final time t and keeping only the runs in which it equals the nominal syndrome of the prepared state at every reading. Figure 1(b) illustrates the three kinds of run that matter. The observable of interest is estimated on the kept runs, and the fraction of kept runs is the yield. Throughout the paper, the post-selected estimator is judged against the gauge-invariant dynamics under the channels the syndrome does not see, which are motional dephasing and qubit bit flips. We call this the heating-free reference. Qubit dephasing is seen by the syndrome and would be removed along with the heating. We set $\gamma_z = 0$ in every calculation, in line with the experiment, so that the reference is unambiguous. The effect of the remaining channels on each observable is given by the certificates of [10].

All numerical results are exact Lindblad evolutions with no Monte Carlo sampling. For a state of definite initial syndrome the post-selected estimator is computed by attaching a parity-flag register to each mode, which turns it into a ratio of two expectation values on an enlarged space. For the thermal preparation of Section 5 it is computed by projecting onto the nominal syndrome at every reading, which is what the physical measurement does. Both constructions, the convergence checks and the truncations, are described in Appendix B. We work in units where $J = 1$ and quote times in milliseconds using $J/\pi = 0.35$ kilohertz.

3 Exact recovery under heating

The quantum-jump unravelling of a Lindblad equation decomposes the state at time t into contributions from trajectories with a definite history of jumps [30, 31]. The contribution from trajectories with no heating jump in $[0, t]$ is $\mathcal{U}_{t,0}\rho_0$, where $\mathcal{U}$ is generated by the no-jump Lindbladian

$$\mathcal{L}_{\text{nj}}\rho = \mathcal{L}_0\rho - \frac{1}{2}\{\hat{K}, \rho\}, \quad \hat{K} = \sum_{m,\pm} \hat{L}_{m,\pm}^\dagger \hat{L}_{m,\pm}. \quad (3)$$

Here $\mathcal{L}_0$ is the Lindbladian of the heating-free reference, and $\hat{L}_{m,-}$ and $\hat{L}_{m,+}$ are the loss and gain jump operators of mode m . The expectation of $\hat{K}$ is the total rate of heating events. The trace of $\mathcal{U}_{t,0}\rho_0$ is the probability that no heating event occurred in $[0, t]$, and the normalised state $\mathcal{U}_{t,0}\rho_0/\text{tr}\mathcal{U}_{t,0}\rho_0$ is the state conditioned on that history. The first result of the paper is that under one condition, this conditional state carries no trace of the heating at all. We state it for a general bath with loss rate $\kappa_{m,-}$ and gain rate $\kappa_{m,+}$ on mode m , so that $\hat{L}_{m,-} = \sqrt{\kappa_{m,-}} \hat{a}_m$ and $\hat{L}_{m,+} = \sqrt{\kappa_{m,+}} \hat{a}_m^\dagger$, because the condition is then visible. The heating model of Section 2 is the case $\kappa_{m,-} = \kappa_{m,+} = \dot{n}_m$.

Theorem 1 (Exact recovery for a uniform bath). *Let the sum of the loss and gain rates be the same on every mode, $\kappa_{m,-} + \kappa_{m,+} = s$, and let $\mathcal{L}_0$ commute with the total phonon number in the sense that $\mathcal{L}_0(\hat{N}\rho) = \hat{N}\mathcal{L}_0(\rho)$ and $\mathcal{L}_0(\rho\hat{N}) = \mathcal{L}_0(\rho)\hat{N}$. Then on any graph, at any electric field and for any initial state,*

$$\mathcal{U}_{t,0}\rho_0 = e^{-Ct} e^{-st\hat{N}/2} (e^{t\mathcal{L}_0}\rho_0) e^{-st\hat{N}/2}, \quad C = \sum_m \kappa_{m,+}. \quad (4)$$

If ρ_0 lies in the sector of N phonons, the conditional state is $e^{t\mathcal{L}_0}\rho_0$ exactly, and the probability of no heating event is $e^{-(sN+C)t}$. For the heating model with equal rates $\dot{n}$ on M modes this probability is $e^{-\dot{n}(2N+M)t}$.

The proof can be stated concisely. The rate operator is $\hat{K} = \sum_m (\kappa_{m,-} + \kappa_{m,+}) \hat{n}_m + \sum_m \kappa_{m,+}$, which under the hypothesis equals $s\hat{N} + C$. The anticommutator term in Eq. (3) is then $-C\rho - \frac{s}{2}(\hat{N}\rho + \rho\hat{N})$. The map $\rho \mapsto \hat{N}\rho + \rho\hat{N}$ commutes with $\mathcal{L}_0$ by hypothesis, so the two parts of $\mathcal{L}_{\text{nj}}$ can be exponentiated separately, which gives Eq. (4). On a sector of fixed N the operator $\hat{N}$ is the number N . The exponential factors become $e^{-(sN+C)t}$, and the normalised state is $e^{t\mathcal{L}_0}\rho_0$. The hypothesis on $\mathcal{L}_0$ holds for the model of Section 2. The Hamiltonian commutes with $\hat{N}$, the motional dephasing jumps $\hat{n}_m$ commute with $\hat{N}$, and the qubit jumps act on a different tensor factor.

Three remarks clarify what the theorem does and does not claim. First, the result is not perturbative. The conditional state equals the heating-free state at every order in $\dot{n}t$, and the only effect of the heating on the kept runs is a number. What makes this work is that the total jump rate $\hat{K}$ depends on the conserved quantity $\hat{N}$ alone, which requires the occupation-dependent part of the rate, $\kappa_{m,-} + \kappa_{m,+}$, to be the same on every mode. Equal gain and loss within a mode is not required. Uniform pure loss, uniform pure gain, and a uniform bath at any finite temperature all satisfy the condition. The first of these is why two-mode bosonic codes with a fixed total excitation number are untouched by the no-jump evolution of amplitude damping [20]. The same observation underlies the uniformly decaying subspaces of [21], of which a fixed phonon-number sector of the gauge theory under uniform heating is an instance. What breaks the condition is a mode dependence of the rates, and Section 5 quantifies the consequence. Single-mode binomial and cat codes [32] face a different problem. Their code words are superpositions of different photon numbers, on which $e^{-st\hat{N}/2}$ is not a constant.

Second, the theorem concerns the conditional state and not yet a protocol. A protocol reads the syndrome at intervals and keeps the runs whose syndrome never changed. Those runs include trajectories in which two heating events on the same mode cancelled between readings. A return event of this kind needs two events on the same mode inside one interval. For pure loss the mode must hold at least two quanta, since $[\hat{G}_m, \hat{a}_m^2] = 0$, so return events are absent for the one-phonon states considered here and present at higher occupation. For heating, they occur at any occupation, because an empty mode can gain a quantum and lose it again. The conditional state is therefore only the limit of the protocol as the interval goes to zero. The construction of [21] excludes return events by hypothesis. Section 4 quantifies the approach to the limit, and that part of the analysis is what the heating channel makes necessary. Third, the thermal preparation of the experiment does not have a definite phonon number, and Section 5 quantifies that consequence too.

A corollary concerns qubit dephasing. We include it for completeness, although the numerics set $\gamma_z = 0$. A dephasing jump $\sqrt{\gamma_z} \hat{\sigma}_\ell^z$ flips the syndrome at both ends of its link, so the protocol rejects it. Since $\hat{\sigma}^{z\dagger} \hat{\sigma}^z = 1$, its no-jump backaction is the constant $e^{-\gamma_z t/2}$ on every state. Sector tracking therefore removes qubit dephasing exactly as well in the limit of continuous monitoring, at a further cost of $e^{-\gamma_z t}$ per link in the yield. At a finite interval, two dephasing events on the same link inside one interval restore the syndrome and contribute a bias of the kind analysed in Section 4.




Figure 2: Sector tracking on the two-link loop with one phonon, equal heating rates of 150 quanta per second, the experimental coherence times and $h/J = 0.6$. (a) Population of the second site in the zero-flux sector. The unconditioned evolution departs from the heating-free reference by up to 0.50 within the window. The post-selected estimator approaches the reference as the syndrome interval shrinks, and the zero-jump conditional state lies on top of the reference. (b) The same in the π -flux sector, where the heating damage is 0.58, and the recovery is the same. (c) The Gauss generator at the first site and the yield in the zero-flux sector. The post-selected generator is restored to its initial value to machine precision at every syndrome interval, and the yield approaches the closed form $e^{-4\hat{n}t}$ of Corollary 1 as the interval shrinks.

Table 1: Theorem 1 on three geometries with one phonon, equal heating rates of 150 quanta per second, the experimental coherence times and $h/J = 0.3$. The damage is the largest departure of the unconditioned evolution from the heating-free reference within the window, over all site populations or all generators. The error is the same quantity for the zero-jump conditional state at the default integrator tolerances. The yield is the probability of no heating event at the end of the window, and its closed form is $e^{-\hat{n}(2N+M)t}$ with $N = 1$ and $M = 3$.

Geometry	Unconditioned damage		Zero-jump error		Yield	
	Populations	Generators	Populations	Generators	Computed	Closed form
Chain, fluxes $(0, \pi)$	0.37	0.75	1.8×10^{-5}	1.1×10^{-15}	0.17552	0.17552
Chain, fluxes $(0, 0)$	0.41	0.75	8.3×10^{-6}	8.9×10^{-16}	0.17552	0.17552
Triangle, flux π	0.33	0.012	1.4×10^{-6}	8.5×10^{-17}	0.17552	0.17552

For a device with $\gamma_z \neq 0$ the reference is the dynamics with that channel removed. Bit flips commute with the generators, are invisible to the syndrome, and are not removed. Their effect on the flux and on the frozen sector is given by the exact exponential laws of [10].

Corollary 1 (Closed-form yield). *For an initial state that is a mixture over phonon-number sectors with weights p_N , the probability of no heating event in $[0, t]$ under the heating model with equal rates is $Y_0(t) = \sum_N p_N e^{-\hat{n}(2N+M)t}$. For a state of definite N , it is $e^{-\hat{n}(2N+M)t}$.*

The yield depends on the number of modes even when they are empty because an empty mode can gain a quantum. For one phonon on the two modes of the loop, it is $e^{-4\hat{n}t}$, and for one phonon on three modes, it is $e^{-5\hat{n}t}$. A finite-temperature bath would reduce this cost. The infinite-temperature bath of the trap electrode does not.

Figure 2 shows the theorem at work on the loop. The initial state is one phonon on the first site with the link qubits in $|\Phi_{\ell_1 \ell_2}^+\rangle$ or $|\Psi_{\ell_1 \ell_2}^+\rangle$. The heating rates are 150 quanta per second on both modes, the coherence times are the experimental ones, and the electric field is $h/J = 0.6$. In the zero-flux sector, the unconditioned population of the second site departs from the heating-free reference by up to 0.50 within the window, and in the π -flux sector by up to 0.58. The zero-jump conditional state, computed directly from Eq. (3), agrees with the reference to 2.4×10^{-4} in the zero-flux sector and to 1.5×10^{-4} in the π -flux sector at the default integrator tolerances on the flagged space. The generators are conserved to 10^{-15} . The residual on the populations is numerical. Tightening the tolerances by two and by four orders of magnitude reduces it to 6×10^{-7} and to 1×10^{-8} at both Fock truncations tested, and the yield then agrees with the closed form $e^{-4\hat{n}t} = 0.2486$ to 10^{-12} and to 10^{-15} . The convergence is reported in Appendix B.

The theorem is free of geometry. Table 1 confirms it on a two-loop chain of three sites and four links in the flux sectors $(0, \pi)$ and $(0, 0)$, and on a π -flux triangle of three sites and three links. In every case, the zero-jump conditional state reproduces the heating-free populations to a few parts in a million at the default tolerances, reproduces the generators to machine precision, and has the closed-form yield to nine digits. These runs use a Fock space of dimension two per

mode for the conditional state and four for the unconditioned damage. Dimension two is exact for the conditional state because the no-jump dynamics conserves the total phonon number and the initial state has one phonon.

4 Finite syndrome intervals

The syndrome is read at times $k\tau$ for $k = 1, \dots, t/\tau$, and a run is kept when every reading equals the nominal syndrome. For a state of definite initial syndrome, a heating event on mode m flips g_m , so a run is kept exactly when the number of heating events on every mode inside every interval is even. The kept ensemble therefore contains the zero-jump trajectories, with weight $Y_0(t)$, together with trajectories in which two or more events on one mode cancelled inside an interval. When the rates are uniform and the initial state has a definite phonon number, the zero-jump state is the heating-free state by Theorem 1. Under those two hypotheses, and writing $Y(t)$ for the total weight of the kept ensemble, which is the yield the experiment measures, the kept state is

$$\rho_{\text{kept}} = \frac{Y_0}{Y} \rho_{\text{ref}} + \left(1 - \frac{Y_0}{Y}\right) \rho_{\text{rest}}, \quad (5)$$

where ρ_{rest} is the normalised state of the kept trajectories that had at least one cancelled pair. Everything that follows concerns the second term.

Theorem 2 (Yield-excess certificate). *Let the heating rates be uniform, and the initial state have a definite phonon number and a definite syndrome, so that the zero-jump state equals the reference, and let $\hat{O}$ be a bounded observable with range R , the difference between its largest and smallest eigenvalue. Then, for every syndrome interval τ and every heating rate,*

$$|\text{tr } \hat{O} \rho_{\text{kept}} - \text{tr } \hat{O} \rho_{\text{ref}}| \leq R \left(1 - \frac{Y_0(t)}{Y(t)}\right), \quad (6)$$

where Y_0 is the closed form of Corollary 1 and Y is the measured yield.

The proof follows at once from Eq. (5), since $\text{tr } \hat{O} \rho_{\text{kept}} - \text{tr } \hat{O} \rho_{\text{ref}} = (1 - Y_0/Y)(\text{tr } \hat{O} \rho_{\text{rest}} - \text{tr } \hat{O} \rho_{\text{ref}})$ and the bracket is bounded by R . The certificate is exact at every order in τ and in $\dot{n}t$. It needs two numbers that the experiment has, i.e., the fraction of runs it kept and the fraction that the closed form says it would have kept under continuous monitoring. For the flux projector $\hat{\Pi}_\Phi$ one has $R = 1$, and for a gauge correlator such as $\hat{\sigma}_{\ell_1}^x \hat{\sigma}_{\ell_2}^x$ one has $R = 2$. For the generators, the bias is zero by construction. The certificate says that the excess of the measured yield over the closed form is itself the error budget. For an experimenter, this means that the yield is a diagnostic as well as a cost. For a population the range is unbounded, because a kept run may contain any even number of gain events on a mode. Equation (6) with R taken on the sectors reachable with two events, which gives $R = N_0 + 2$, is then a leading-order estimate and not a certificate. It fails as a bound when $\dot{n}t$ is of order one, and we use it only in the regime of Table 2, where the weight of sectors beyond $N_0 + 2$ is of order $(\dot{n}\tau)^2$.

When the zero-jump state itself differs from the reference, the hypothesis of the theorem fails and the certificate takes a longer form. This happens for unequal rates and for a mixture of phonon-number sectors with a common syndrome. Write $Z(t) = \text{tr } \mathcal{U}_{t,0}(\hat{\Pi}_{\text{nom}} \rho_0 \hat{\Pi}_{\text{nom}})$ for the weight of the zero-jump trajectories that started in the nominal syndrome, computed from the no-jump evolution of Eq. (3) at the actual rates, ρ_{nj} for their normalised state and $\alpha = Z/Y$. Then $\rho_{\text{kept}} = \alpha \rho_{\text{nj}} + (1 - \alpha) \rho_{\text{other}}$, where the second term collects the cancelled pairs, and, for a thermal preparation, the runs that entered the nominal syndrome from outside it. Hence

$$|\text{tr } \hat{O} \rho_{\text{kept}} - \text{tr } \hat{O} \rho_{\text{ref}}| \leq \alpha |\text{tr } \hat{O} \rho_{\text{nj}} - \text{tr } \hat{O} \rho_{\text{ref}}| + (1 - \alpha)R. \quad (7)$$

Equation (6) is the case in which the first term vanishes. Both forms are pointwise in time. A bound on the largest departure within a window is therefore the largest value of the right-hand side within the window, which is what Table 2 reports.

The excess yield can be predicted before the run. To first order in τ the kept ensemble is the zero-jump state plus the trajectories with exactly two heating events on the same mode inside one interval. When the interval is short compared with the hopping time and with the mean time between events, the two events act at a single instant s . The four orderings of gain and loss on a mode give the operators $\hat{n}_m + 1$, $\hat{n}_m$, $\hat{a}_m^{\dagger 2}$ and $\hat{a}_m^2$, the double integral over an interval gives $\tau^2/2$, and the position of the interval is uniform in $[0, t]$. This leads to the following statement, proved in Appendix A.

Proposition 1 (First-order bias law and excess yield). *Let the heating rates be equal, and the initial state have a definite phonon number and a definite syndrome, and define the two-event map $\mathcal{K}(\rho) = \sum_m [(\hat{n}_m + 1)\rho(\hat{n}_m + 1) + \hat{n}_m \rho \hat{n}_m + \hat{a}_m^{\dagger 2} \rho \hat{a}_m^2 + \hat{a}_m^2 \rho \hat{a}_m^{\dagger 2}]$. The bias of the post-selected estimator of $\hat{O}$ with syndrome readings every τ is*

$$\text{tr } \hat{O} \rho_{\text{kept}} - \text{tr } \hat{O} \rho_{\text{ref}} = \frac{\dot{n}^2 \tau}{2 Y_0(t)} \int_0^t ds \left[\text{tr } \hat{O} \mathcal{U}_{t,s} \mathcal{K}(\mathcal{U}_{s,0} \rho_0) - \text{tr } \hat{O} \rho_{\text{ref}}(t) \text{tr } \mathcal{U}_{t,s} \mathcal{K}(\mathcal{U}_{s,0} \rho_0) \right] + O(\tau^2), \quad (8)$$




Figure 3: Bias of the post-selected estimator against the syndrome interval. (a) The largest departure of the post-selected population of the second site from the heating-free reference within the window, for five preparations of the loop. The dotted lines are the errors of the zero-jump conditional state, which the estimator approaches as the interval shrinks. For one phonon with equal rates, that floor is below 3×10^{-4} . For unequal rates it is the tilt of Section 5. For thermal matter it is the floor of Proposition 3, computed here with the flag construction of Appendix B on the ideal phonon-added state, which does not fix the parities. (b) The bias of the two equal-rate one-phonon cases divided by $\dot{n}^2 \tau t$, with the floor subtracted. The ratio lies between 1.85 and 2.10 for every interval up to 1 millisecond in both flux sectors and reaches 2.35 for the single final check in the π -flux sector.

Table 2: The finite-interval bias on the zero-flux loop with one phonon and equal heating rates of 150 quanta per second at $\hbar/J = 0.6$, with t the window of 2.32 milliseconds. The excess yield is compared with Eq. (9). The bias in the flux projector $\hat{\Pi}_\Phi$, which is bounded with $R = 1$, is compared with the certificate of Eq. (6), reported as its largest value within the window, with the closed form $Y_0(u) = e^{-4\dot{n}u}$ at every time u . The entries are computed at integrator tolerances of 10^{-10} absolute and 10^{-8} relative. The bias in the population of the second site is compared with the law $2\dot{n}^2 \tau t$ obtained from Eq. (8). All biases are the largest departure within the window.

Interval τ (ms)	Yield Y	Excess $Y - Y_0$	Eq. (9)	Bias in $\langle \hat{\Pi}_\Phi \rangle$	Certificate	Bias in $\bar{n}_{m_2}$	$2\dot{n}^2 \tau t$
2.32 (final check)	0.398	0.149	0.155	0.012	0.375	0.251	0.242
1.0	0.305	0.057	0.067	0.0083	0.186	0.097	0.104
0.5	0.281	0.032	0.034	0.0030	0.114	0.052	0.052
0.25	0.265	0.0168	0.0167	0.0014	0.064	0.027	0.026
0.1	0.255	0.0068	0.0067	0.00054	0.027	0.011	0.010

and the excess yield is $Y - Y_0 = \frac{1}{2} \dot{n}^2 \tau \int_0^t ds \text{tr} \mathcal{U}_{t,s} \mathcal{K}(\mathcal{U}_{s,0} \rho_0) + O(\tau^2)$, where $\mathcal{U}$ is the no-jump propagator of Eq. (3). By Theorem 1 every factor is a heating-free evolution multiplied by a known exponential, so both quantities are computable from the heating-free dynamics alone. For one phonon on the two modes of the loop, the excess yield is

$$Y - Y_0 = Y_0 \left[3\dot{n}^2 \tau t + \dot{n} \tau (1 - e^{-4\dot{n}t}) \right] + O(\tau^2). \quad (9)$$

The two terms of Eq. (9) have a simple origin. The first comes from the gain-then-loss and loss-then-gain branches. These return the state to its sector and are damped afterwards at the same rate as the zero-jump state. The second comes from the double-gain branch, which raises the sector by two and is damped afterwards at twice the rate. The trace of $\mathcal{K}$ on any one-phonon state of two modes is 14, split as 6 for the returning branches and 8 for the double gain, and the factor $(1 - e^{-4\dot{n}t})/4\dot{n}$ is the time integral of the extra damping. The double-gain branch is also why the no-jump propagator in Eq. (8) cannot be replaced by the heating-free propagator. A first attempt that made that replacement gave a bias coefficient of 3.7 instead of 2.0. At $\dot{n}t \approx 0.35$ the sector-dependent damping is a large effect and not a refinement.

Figure 3 and Table 2 show the exact bias of the post-selected estimator as the syndrome interval is varied from the full window, which is a single final check, down to 0.1 milliseconds. The excess yield follows Eq. (9) to one percent at the two smallest intervals. At the two largest it agrees to four and eighteen percent, where the second-order terms in τ enter. The bias is linear in τ . For the population of the second site with one phonon and equal rates, the largest bias within the window with the zero-jump floor subtracted is $c\dot{n}^2 \tau t$. The coefficient c lies between 1.85 and 2.10 for every interval up




Figure 4: The first-order bias law of Proposition 1 against the exact bias of the post-selected estimator with syndrome checks every 0.1 milliseconds, on the zero-flux loop with one phonon, equal heating rates of 150 quanta per second and $h/J = 0.6$. The functional uses only heating-free evolutions and the exponential factors of Theorem 1. It reproduces the full time profile of the exact bias to 4×10^{-4} and gives a largest bias of $1.99 \dot{n}^2 \tau t$ against the exact 2.0.

to 1 millisecond in both flux sectors and for the final check in the zero-flux sector. For the final check in the π -flux sector $c = 2.35$, which is the elevated point in Figure 3(b). The round value 2.0 is therefore accurate to about ten per cent for $\dot{n}\tau \leq 0.15$, and it understates the bias by fifteen per cent at the single final check in the π -flux sector, where the second-order terms are largest. The coefficient is a property of the loop at these parameters and not a universal number, since Eq. (8) depends on the dynamics through ρ_s . Its range of validity in $\dot{n}\tau$, $J\tau$ and h/J is a matter for a calibrated gate of the kind used in [10], which we have not yet run. The certificate of Theorem 2 holds for the flux projector at every interval. Its margin is large because the cancelled pairs move the flux only indirectly. In the π -flux sector, where the frozen state is more sensitive, the bias in $\langle \hat{\Pi}_\Phi \rangle$ is 0.070, 0.030, 0.014, 0.0068 and 0.0027 at the same five intervals, against certificates of 0.379, 0.193, 0.117, 0.064 and 0.027.

Figure 4 compares the functional with the exact bias at the smallest interval. The functional reproduces the full time profile to 4×10^{-4} and gives a coefficient of 1.99 against the exact 2.0. The residual of a few per cent is the correction of order $J\tau$ from the finite separation of the two events. The law is of the kind we have elsewhere called a first-order law, derived from the noiseless dynamics and accurate within a declared margin. It sits beside the certificate of Theorem 2, which is guaranteed and loose, as the certification programme intends [22, 10].

5 Unequal rates and thermal preparation

The experiment fails two hypotheses of Theorem 1, a uniform bath and a definite phonon number. We treat them in turn.

With unequal rates, the operator $\hat{K}$ of Eq. (3) on the loop is $(\dot{n}_1 + \dot{n}_2)(\hat{N} + 1) + \Delta(\hat{n}_1 - \hat{n}_2)$ with $\Delta = \dot{n}_1 - \dot{n}_2$. On the sector of one phonon, its non-constant part is $2\Delta\hat{n}_1$. This term does not commute with the hopping. The conditional state is tilted away from the heating-free state, because a trajectory that spent its time on the mode that heats more slowly is less likely to have been rejected. The hopping limits the tilt, since the phonon does not stay on either mode for long. At zero electric field and without motional dephasing the tilt has a closed form.

Proposition 2 (Tilt of the conditional state). *For one phonon on the loop with the link qubits prepared in $|\Phi_{\ell_1\ell_2}^+\rangle$, at $h = 0$, without motional dephasing and without qubit noise, with the phonon initially on the first site, the population of the second site in the zero-jump conditional state is*

$$\bar{n}_{m_2}(t) = \frac{(4J^2/\omega^2) \sin^2 \omega t}{(\cos \omega t - (\Delta/2\omega) \sin \omega t)^2 + (4J^2/\omega^2) \sin^2 \omega t}, \quad \omega = \sqrt{4J^2 - \Delta^2/4}, \quad (10)$$

against the heating-free value $\sin^2 2Jt$. To first order in Δ/J the tilt is $(\Delta/2J) \sin^3(2Jt) \cos(2Jt)$, whose magnitude never exceeds $0.65 |\Delta|/4J$.

The derivation is a two-level calculation. With the link qubits in $|\Phi_{\ell_1\ell_2}^+\rangle$, the hopping acts with amplitude $\pm 2J$, and the one-phonon subspace is closed under the no-jump evolution. After the removal of the common factor $e^{-(\dot{n}_1 + \dot{n}_2)t}$,




Figure 5: The two departures from the hypotheses of Theorem 1. (a) Tilt of the zero-jump conditional state at $h = 0$ without motional dephasing for one phonon on the loop with heating rates of 50 and 250 quanta per second and of 20 and 280 quanta per second. The curves are the exact conditional state, the circles are the closed form of Eq. (10), and the dotted lines mark $|\Delta|/4J$. (b) Thermal matter on the loop with equal heating rates of 150 quanta per second and $h/J = 0.6$, with one phonon added by a sideband pulse to a thermal background of $\bar{n}_0 = 0.1$ on both modes. The reference is the heating-free dynamics of the component of the prepared state that has the nominal syndrome. The unconditioned evolution departs from it by up to 0.57, of which 0.13 is present without any heating because the thermal background violates Gauss’s law at preparation. Projecting on the nominal syndrome at the final time reduces the departure to 0.31, and projecting every 0.1 milliseconds reduces it to 0.007. The zero-jump state has a floor of 0.007, which is the thermal floor of Proposition 3.

that evolution is generated by a non-Hermitian matrix with off-diagonal entries $2J$ and diagonal entries $-i\Delta/2$ and $+i\Delta/2$. Its eigenfrequencies are $\pm\omega$. The amplitudes from the initial state are $\cos\omega t - (\Delta/2\omega)\sin\omega t$ on the first site and $-i(2J/\omega)\sin\omega t$ on the second, and normalising the squared moduli gives Eq. (10). The tilt vanishes whenever the phonon is entirely on one site, and changes sign every quarter period. It is largest when the phonon is partly on both sites, because that is when the two components of the state have had different exposures to the two rates. The frequency ω differs from $2J$ at second order in Δ/J , so the first-order form is accurate within the window but not uniformly in time.

Figure 5(a) compares the closed form with the exact conditional state for rate splits of 50 and 250 and of 20 and 280 quanta per second. The agreement is 2.4×10^{-4} in both cases. The largest tilt within the window is 0.72 and 0.74 of $|\Delta|/4J$, and the excess over the first-order value of 0.65 comes from the second-order terms. With the electric field at $h/J = 0.6$ and the experimental motional dephasing, the population tilt for the split of 50 and 250 quanta per second is 0.033. That is the floor of the unequal-rate curve in Figure 3(a). The tilt of the flux projector in the same case is 0.051, so the flux sector is more sensitive to unequal rates than the population. On the chain with rates of 50, 150 and 250 quanta per second on the three modes, the population tilt is 0.043. In every case computed within the window, the population tilt stayed below $|\Delta|/4J$ evaluated with the largest rate difference. We report this as an observation about the population, with the first-order form of Proposition 2 as its analytic basis, and not as a theorem. The flux-projector tilt of 0.051 exceeds $|\Delta|/4J = 0.045$, so the observation does not extend to the flux sector. At the experimental rates of 5 and 17 quanta per second the quantity $|\Delta|/4J$ is 0.003.

The experiment prepares its matter by a blue-sideband pulse on a thermal background with $\bar{n}_0 = 0.1$ on every mode. We model the pulse as a resonant π pulse on the first sideband, calibrated from the ground state and heralded on the excitation of the ion. It maps the Fock state $|n\rangle$ to $\sin(\frac{\pi}{2}\sqrt{n+1})|n+1\rangle$ up to a phase. The prepared state is then a mixture over phonon-number sectors that is diagonal in the Fock basis. Two things follow from the thermal background. The state is not an eigenstate of the generators, because the parity of each mode is uncertain, and it has no definite total phonon number. The syndrome measurement addresses both at once. The nominal syndrome is odd parity on the site carrying the phonon and even parity on the empty site. A projection onto it retains 0.867 of the prepared state and fixes the parity of every mode. The total phonon number of the retained part is therefore odd, with sector weights of 0.990 on one phonon, 0.0096 on three and 10^{-4} on five. The thermal background that the projection removes is a violation of Gauss’s law present at preparation, not caused by heating. Relative to the gauge-invariant target it displaces the unconditioned population of the second site by 0.083 at $t = 0$ and by up to 0.13 within the window at $h/J = 0.6$. The reference for thermal matter is therefore the heating-free dynamics of the nominal-syndrome component, which is the gauge-invariant dynamics the experiment is trying to realise. If the pulse is modelled instead as the ideal addition of

one phonon, which weights $|n\rangle \rightarrow |n+1\rangle$ by $\sqrt{n+1}$, the retained fraction is 0.777, and the sector weights are 0.967, 0.032 and 6×10^{-4} . The conclusions below are insensitive to the preparation model, although the yield is not.

Proposition 3 (Thermal preparation). *For an initial state $\rho_0 = \sum_N p_N \rho_N$ that is block diagonal in the phonon-number sectors and has a definite syndrome, the zero-jump conditional state under equal rates is $\sum_N w_N(t) e^{t\mathcal{L}_0} \rho_N$ with $w_N(t) \propto p_N e^{-2\hat{n}Nt}$. For a bounded observable of range R , the departure from the heating-free value is bounded by $R \cdot \text{TV}(w, p)$, where TV is the total-variation distance between the two sets of weights, which is computable from the preparation model and not alone. If the total phonon number of each kept run were known, weighting the run by $e^{2\hat{n}Nt}$ would restore the weights p_N and the heating-free estimator exactly on the zero-jump ensemble.*

The first part is Theorem 1 applied to each block. The bound follows from writing the departure as $\sum_N (w_N - p_N) \text{tr} \hat{O} e^{t\mathcal{L}_0} \rho_N$ and subtracting the midpoint of the range. For a population the range is unbounded, and the same decomposition gives instead the number-weighted bound $\max\{\sum_N N(w_N - p_N)_+, \sum_N N(p_N - w_N)_+\}$, since $0 \leq \langle \hat{n}_m \rangle_N \leq N$ in the sector N . With the sector weights fixed by the nominal projection, the total-variation distance at the end of the window is 0.0072 at 150 quanta per second, and the number-weighted bound is 0.022. The thermal floor is therefore small once the syndrome has been projected. Figure 5(b) shows the loop with equal rates of 150 quanta per second at $h/J = 0.6$. The zero-jump state departs from the reference by 0.0072 on the population, which is a computed value and not a certified one, and by 0.0007 on the flux projector, within its certificate. At the experimental rates, the same quantities are a total-variation distance of 0.0009 and a number-weighted bound of 0.0028, evaluated at the homogeneous comparison rate of 11 quanta per second, which has the same total rate as the actual 5 and 17. The reweighting in the last part of the proposition is exact on the zero-jump ensemble. We do not recommend it for a protocol with a finite interval, because the weight $e^{2\hat{n}Nt}$ amplifies the double-gain trajectories that the finite interval admits. At the experimental rates with a single final check it changes the population error from 0.011 to 0.013, and at 150 quanta per second it makes the estimator worse by an order of magnitude. Since the nominal projection already leaves a computed floor below 10^{-3} at the experimental rates, the reweighting is not needed. A Fock-state preparation by a resolved sideband sequence would remove the floor entirely by making p_N a delta.

6 Two readout protocols

The results of Sections 3 to 5 assume that the syndrome can be read. A destructive measurement at the end of a run, given a readout that resolves the phonon number of one mode, can test some of them. A non-demolition measurement during the run is needed for the rest. We state the two separately, with the results that each supports.

The first protocol is a destructive measurement of the Gauss generators at the final time only, followed by acceptance of the runs whose outcome is the nominal syndrome. It needs no record of the initial syndrome, because the gauge-invariant dynamics preserves the syndrome. A final outcome equal to the nominal one is exactly the acceptance rule of Section 4 with $\tau = t$ when the initial syndrome is definite, and the projection of Section 5 when it is not. The ingredients are the primitives the device already uses, together with one it has not demonstrated, a number-resolving readout of a mode. Their sequence is more demanding than any single measurement the experiment reports. We describe it in order to be clear about what is assumed and to show that the accepted ensemble is the one simulated.

The link qubits are read in the $\hat{\sigma}^x$ basis by a $\pi/2$ rotation followed by fluorescence detection. Fluorescence preserves the motional state of an ion only for the dark outcome [8], so any readout of the modes that follows is available only on the branch in which both ions are dark. That branch cannot be treated as a simple readout loss, because which $\hat{\sigma}^x$ outcome pair is dark is correlated with the matter. At $h = 0$ the zero-flux loop evolves as $\cos(2Jt) |1, 0\rangle |\Phi^+\rangle - i \sin(2Jt) |0, 1\rangle |\Phi^-\rangle$. The state $|\Phi^+\rangle$ has equal $\hat{\sigma}^x$ outcomes on the two links and $|\Phi^-\rangle$ has opposite ones. Keeping the branch in which both ions are dark under a single readout setting would therefore keep only the component with the phonon on the first site and estimate a population of zero. The remedy is to use four readout settings, which differ by a π phase on the $\pi/2$ rotation of one or both ions, so that each of the four outcome pairs $(\pm, \pm)$ is the dark pair in one setting. In each setting, the dark branch is followed by a reset of the qubits and a readout of the modes. The accepted counts and the accepted sums of the matter observable are added over the four settings before the ratio is formed. The four settings are given the same number of attempted shots, or else each setting's counts and sums are divided by its number of attempted shots before the addition. The four outcome projectors $\hat{P}_{\pm\pm}$ sum to the identity, commute with $\hat{\Pi}_{\text{nom}}$, and commute with every observable of the modes. The combined estimator is therefore $\sum_{\pm\pm} \text{tr}(\hat{P}_{\pm\pm} \hat{\Pi}_{\text{nom}} \hat{O} \rho) / \sum_{\pm\pm} \text{tr}(\hat{P}_{\pm\pm} \hat{\Pi}_{\text{nom}} \rho) = \text{tr}(\hat{\Pi}_{\text{nom}} \hat{O} \rho) / \text{tr}(\hat{\Pi}_{\text{nom}} \rho)$, which is the estimator of Figure 6, obtained at the cost of four times the shots of an ideal non-destructive syndrome measurement.

The nominal syndrome in each setting is the pair of parities that gives $\hat{G}_{m_1} = -1$ and $\hat{G}_{m_2} = +1$ for the known value of $\hat{\sigma}_{\ell_1}^x \hat{\sigma}_{\ell_2}^x$ on that branch. The readout of the modes must supply the parity of the first mode, the parity of the second mode and the population of the second mode in the same shot. A readout that resolves the phonon number of the second




Figure 6: Sector tracking at the reported noise of the experiment, with heating rates of 5 and 17 quanta per second, the experimental coherence times and matter prepared by a sideband pulse on a thermal background with $\bar{n}_0 = 0.1$. The curves show the departure of the population of the second site from the gauge-invariant reference, which is the heating-free dynamics of the nominal-syndrome component of the prepared state, for (a) $h = 0$ and (b) $h/J = 0.6$. The unconditioned evolution is displaced by the thermal background that violates Gauss’s law at preparation, shown by the heating-free curve without post-selection, and further by the heating. A single Gauss-law measurement at the final time with acceptance of the nominal syndrome reduces the largest departure from 0.094 to 0.0085 at $h = 0$ and from 0.146 to 0.011 at $h/J = 0.6$, keeping 0.80 of the runs. Syndrome readings every 0.1 milliseconds reduce it to 0.0017 and 0.0018, close to the unequal-rate floor of the zero-jump state.

mode provides this, together with a parity mapping of the first mode onto its qubit [33, 34]. A consistent ordering on the dark branch is as follows. Record the known product $s = \hat{\sigma}_{\ell_1}^x \hat{\sigma}_{\ell_2}^x$ of that branch. Map the required parity $\hat{P}_1 = -s$ of the first mode onto its qubit and reject the complement. Then measure the phonon number of the second mode and keep the outcomes with $(-1)^{n_2} = s$. The last and destructive step then comes after every quantity the estimator needs has been recorded. Projective phonon-number measurement has been demonstrated in trapped ions by repeated sideband interrogation [35], although not on this device. With only a parity bit per mode per shot the population could not be read in the same shot as the syndrome. For the flux projector, the individual $\hat{\sigma}^x$ readouts do not commute with $\hat{W}_C = \hat{\sigma}_{\ell_1}^z \hat{\sigma}_{\ell_2}^z$. The flux does commute with the product $\hat{\sigma}_{\ell_1}^x \hat{\sigma}_{\ell_2}^x$ that the syndrome needs. A readout of the link qubits in the Bell basis, by the entangling gate the experiment uses for preparation followed by fluorescence, therefore gives both jointly, with the same four-setting reconstruction of the dark branch. The yields quoted below are the acceptance probabilities of the syndrome rule, which the four-setting reconstruction recovers in full at the cost of the extra shots.

Figure 6 shows what the first protocol would do at the reported noise of the experiment, with heating rates of 5 and 17 quanta per second, the experimental coherence times and the thermal preparation of Section 5. At $h/J = 0.6$ the population of the second site departs from the gauge-invariant reference by up to 0.146 within the window. Of this, 0.129 is the thermal background present at preparation, and the rest is heating, which on its own displaces the population by up to 0.042. A single final measurement with acceptance of the nominal syndrome reduces the departure to 0.011 while keeping 0.80 of the runs. The flux projector is recovered to 0.002, against a certificate from Eq. (7) of 0.019 with $\alpha = 0.984$ at the actual rates. At $h = 0$ the numbers are 0.094, 0.0085 and 0.79. The bias law of Section 4 gives $2\bar{n}^2\tau t \approx 10^{-3}$ for the final check at these rates. The rest of the residual is contamination from runs that started outside the nominal syndrome and reached it by an odd number of events, which a final-only measurement cannot distinguish from runs that never left it. With syndrome readings every 0.1 milliseconds the departure falls to 0.0018 at $h/J = 0.6$ and to 0.0017 at $h = 0$. The zero-jump floor is 0.002, which is the unequal-rate tilt of the modes at 5 and 17 quanta per second and is consistent with $|\Delta|/4J = 0.003$. The thermal floor of Proposition 3 is 0.0009 at these rates and is not the limiting term.

The second protocol is an ancilla-assisted non-demolition measurement of $\hat{G}_m$ during the run. It requires a probe ion coupled to the mode and to the link qubits, so that the parity of the mode and the $\hat{\sigma}^x$ values of the links can be mapped onto the probe and read without disturbing the gauge field. The link qubits themselves are the gauge field and cannot be read mid-run without collapsing it. Mid-circuit measurement of an ancilla, with the data ions shielded from the readout light, is established in the quantum-charge-coupled-device architecture [36]. Parity mapping of a motional mode onto an ion is established in bosonic encodings [34]. The ingredients therefore exist, although their combination on this device does not. This protocol supports the full interval scan of Table 2, the bias law of Proposition 1, and the full strength of Theorem 1. The interval scan is the quantitative case for building it. At the heating rates of the present device, the first protocol captures most of the benefit. The second becomes necessary when the heating rates or the

evolution times grow, as they will on the larger geometries of the roadmap of [8], because the bias of a single final check grows as $\dot{n}^2 t^2$ while the bias with checks every τ grows as $\dot{n}^2 \tau t$.

7 Discussion

The main practical statement of the paper concerns a single measurement of the Gauss generators at the end of each run, with acceptance of the nominal syndrome. It removes most of the heating damage predicted in [10]. It also removes the part of the thermal preparation that violates Gauss's law. At the reported rates, the departure of the population from the gauge-invariant target falls from 0.146 to 0.011 at a yield of 0.80, and the residual is understood. It is the sum of four terms. There is a finite-interval bias of order 10^{-3} with a derived coefficient, an unequal-rate tilt of order 2×10^{-3} with a closed form at zero field, and a computed thermal floor below 10^{-3} whose certified bound at the homogeneous comparison rate is 3×10^{-3} . The fourth term is contamination from runs that entered the nominal syndrome by an odd number of events, which a final-only measurement cannot exclude and which syndrome readings during the run would remove. Three consequences are useful to an experimenter. The final Gauss-law measurement should become a standard part of the protocol. The fraction of runs it rejects should be compared with the zero-jump weight $Z(t)$ of the nominal component, computed at the actual rates from the heating-free dynamics, as a check on the heating and preparation models. And the certificate of Eq. (7) turns the measured yield into an error bound for every bounded observable, once the small zero-jump bias at the actual rates has been computed.

The theoretical statement is narrower and should be placed with care. Post-selection on a symmetry is approximate in general, because the no-jump evolution of a non-unitary channel distorts the surviving state. The distortion is absent whenever the total jump rate depends on a conserved quantity alone. That mechanism is known. It is why fixed-excitation bosonic codes survive amplitude damping [20], and it is the defining property of a uniformly decaying subspace [21]. For matter encoded in spins under Pauli noise, the question does not arise, because the jump operators are unitary, and the symmetry-verification literature for such matter is concerned with which errors the syndrome detects rather than with what the surviving state looks like [13, 14, 15]. The present paper identifies the phonon number of the gauge theory as the conserved quantity and the Gauss generators as a syndrome that reads the heating history without disturbing the dynamics. Above all, it analyses what happens when the conditions of the mechanism are relaxed. The heating channel produces return events that pure loss does not at low occupation, and the bias and excess yield they cause have been derived in closed form and verified on exact dynamics. The unequal-rate tilt and the thermal preparation have been treated in the same way. These parts of the paper have no counterpart in the uniformly decaying subspace construction, which excludes return events by hypothesis, and they would be needed for pure loss too as soon as a mode holds two quanta.

The limitations are those of the model, the protocol and the numerics. The heating model is an infinite-temperature bath, which is the standard model for electrode noise and the model the experiment reports against. Theorem 1 does not depend on the temperature of the bath as long as it is the same on every mode. A mode dependence of the total rate produces a tilt of the form of Proposition 2. The motional dephasing is Markovian, as in [10], and the method is indifferent to its form because the syndrome does not see it. The preparation is modelled as a heralded sideband π pulse on a thermal state. The nominal fraction of the preparation, which sets the yield, depends on this model, while the sector weights after the projection do not depend on it appreciably. The readout sequence of the first protocol has been shown to reproduce the simulated estimator through the four-setting reconstruction, but it has not been simulated with readout errors, and it requires a number-resolving readout of one mode that this device has not demonstrated. The numerics are exact Lindblad evolutions. Their residuals converge with the integrator tolerance, and the post-selected quantities change in the third decimal between Fock dimensions six and eight. The finite-interval scan was carried out on the loop only. The theorem and the functional are free of geometry. The bound $|\Delta|/4J$ on the tilt is an observation within the window with an analytic basis at first order, not a theorem. The bias law has been verified on exact dynamics but not yet admitted through a calibrated gate of the kind used in [10]. Such a gate would fix the range of $\dot{n}\tau$, $J\tau$ and h/J over which the coefficient may be trusted, and it is the natural next step before the law is used to design an interval.

8 Conclusions

In this work, we have asked whether the damage that heating does to a hybrid qubit-oscillator quantum simulation of a $\mathbb{Z}_2$ lattice gauge theory can be removed rather than only predicted. We have found that it can be removed by a measurement of the Gauss generators at the end of each run. The measurement is implementable with the primitives of the device, together with a number-resolving readout of one mode and a four-setting reconstruction of the destructive readout. The generators are phonon parities. Heating flips them, and the gauge-invariant dynamics does not, so the generators form a syndrome that records the heating history of a run without disturbing the quantities the experiment wants to measure. Keeping the runs whose syndrome never changed gives an estimator that equals the heating-free

estimator at every order, for a bath whose total rate is the same on every mode, for a state of definite phonon number and in the limit of continuous monitoring. Its yield has the closed form $e^{-(sN+C)t}$, which for the heating model with equal gain and loss is $e^{-\dot{n}(2N+M)t}$. The mechanism is the known one by which a fixed excitation-number sector decays uniformly. What the gauge theory supplies is the conserved number and the syndrome that reads it.

This paper is about what happens when the conditions are relaxed. When the syndrome is read at finite intervals the heating channel admits runs in which two events cancelled, at any occupation rather than only at two quanta or more as for pure loss, and the estimator acquires a bias that is linear in the interval. We have derived its coefficient and the accompanying excess yield from the heating-free dynamics and verified both against exact Lindblad evolutions. We have given a certificate for the bias of any bounded observable. For a state of definite phonon number it needs only the measured yield and its closed form, and otherwise it needs in addition the zero-jump bias at the actual rates. Unequal heating rates tilt the conditional state by an amount with a closed form at zero field, and the population tilt stayed below the rate difference over the hopping in every case computed. Thermal preparation is handled by the same measurement. Projecting on the nominal syndrome fixes the parity of every mode and removes the part of the background that violates Gauss's law at preparation. The floor it leaves is computed to be below 10^{-3} at the reported rates and certified to be below 3×10^{-3} at the homogeneous comparison rate.

At the reported noise of the experiment, a single Gauss-law measurement at the end of each run reduces the departure of the matter population from its gauge-invariant target from 0.146 to 0.011. The acceptance is eight runs in ten, which the four-setting readout recovers at a fourfold cost in shots. Syndrome readings during the run would reduce the departure further, to the unequal-rate floor of 0.002. The interval scan is the quantitative case for a non-demolition measurement of the generators on future hardware, and such a measurement will become necessary as heating rates and evolution times grow on larger geometries. Thus, with [10], which predicts what breaks, the work presented here shows how to recover it. The two leave motional dephasing and qubit bit flips as the noise processes that a gauge-theory experiment on this platform must still live with.

Acknowledgments

The author would like to thank the Centre for Visual Computing and Intelligent Systems at the University of Bradford for providing the computing resources used in this work.

Funding

No specific funding was received for this work.

A Derivation of the first-order bias law and the excess yield

Let the window $[0, t]$ be divided into intervals $I_k = [(k-1)\tau, k\tau]$, and let the heating jump operators on mode m be $\hat{L}_{m,-} = \sqrt{\dot{n}} \hat{a}_m$ and $\hat{L}_{m,+} = \sqrt{\dot{n}} \hat{a}_m^\dagger$. In the quantum-jump unravelling, the unnormalised state of the trajectories with jumps $(\hat{L}_1, u_1), \dots, (\hat{L}_n, u_n)$ at ordered times is $\mathcal{U}_{t,u_n} \mathcal{J}_n \mathcal{U}_{u_n, u_{n-1}} \cdots \mathcal{J}_1 \mathcal{U}_{u_1, 0} \rho_0$ with $\mathcal{J}\rho = \hat{L}\rho\hat{L}^\dagger$. The state of the full ensemble is the sum over histories with the time integrals [30, 31]. The kept ensemble is the sum over histories in which the number of jumps on every mode inside every interval is even. To first order in τ it contains the zero-jump history and the histories with exactly two jumps, both on the same mode and inside the same interval. For two jumps on mode m at times $u_1 < u_2$ in I_k the contribution is

$$\sum_{\alpha, \beta = \pm} \int_{I_k} du_2 \int_{(k-1)\tau}^{u_2} du_1 \mathcal{U}_{t, u_2} \mathcal{J}_{m, \beta} \mathcal{U}_{u_2, u_1} \mathcal{J}_{m, \alpha} \mathcal{U}_{u_1, 0} \rho_0. \quad (11)$$

When τ is short compared with $1/J$, $1/\dot{n}$ and the dephasing times, the propagator $\mathcal{U}_{u_2, u_1}$ between the two jumps may be replaced by the identity, and both jumps placed at a common time s . The sum over the four orderings gives $\sum_{\alpha, \beta} \hat{L}_{m, \beta} \hat{L}_{m, \alpha} \rho \hat{L}_{m, \alpha}^\dagger \hat{L}_{m, \beta}^\dagger = \dot{n}^2 [(\hat{n}_m + 1)\rho(\hat{n}_m + 1) + \hat{n}_m \rho \hat{n}_m + \hat{a}_m^{\dagger 2} \rho \hat{a}_m^2 + \hat{a}_m^2 \rho \hat{a}_m^{\dagger 2}]$, which is $\dot{n}^2$ times the m term of $\mathcal{K}(\rho)$. The double time integral over the interval gives $\tau^2/2$. Summing over the t/τ intervals turns the sum over k into an integral $\int_0^t ds/\tau$, with the common time s uniform in $[0, t]$. The unnormalised kept state is therefore

$$\tilde{\rho}_{\text{kept}} = \mathcal{U}_{t, 0} \rho_0 + \frac{\dot{n}^2 \tau}{2} \int_0^t ds \mathcal{U}_{t, s} \mathcal{K}(\mathcal{U}_{s, 0} \rho_0) + O(\tau^2). \quad (12)$$

Its trace is the yield, which gives the excess of Proposition 1, and the post-selected estimator is $\text{tr} \hat{O} \tilde{\rho}_{\text{kept}} / \text{tr} \tilde{\rho}_{\text{kept}}$. Expanding the ratio to first order in τ , with $\text{tr} \mathcal{U}_{t, 0} \rho_0 = Y_0(t)$ and $\text{tr} \hat{O} \mathcal{U}_{t, 0} \rho_0 = Y_0(t) \text{tr} \hat{O} \rho_{\text{ref}}(t)$ by Theorem 1, gives

Eq. (8). The corrections of order τ^2 come from the finite separation of the two jumps, from histories with four jumps, and from the second-order term of the ratio.

For one phonon on two modes the branch $\hat{a}_m^2 \rho \hat{a}_m^{\dagger 2}$ vanishes. The traces of the remaining branches are $\sum_m \langle (\hat{n}_m + 1)^2 \rangle = 5$, $\sum_m \langle \hat{n}_m^2 \rangle = 1$ and $\sum_m \langle (\hat{n}_m + 1)(\hat{n}_m + 2) \rangle = 8$ for every state in the sector, and they sum to 14. The first two branches leave the state in the one-phonon sector, where $\mathcal{U}_{t,s}$ is the heating-free evolution times $e^{-4\hat{n}(t-s)}$. With $\text{tr} \mathcal{U}_{s,0} \rho_0 = e^{-4\hat{n}s}$ they contribute $\frac{1}{2} \dot{n}^2 \tau \cdot 6 \cdot t e^{-4\hat{n}t}$ to the excess yield. The double-gain branch puts the state in the three-phonon sector, where $\mathcal{U}_{t,s}$ carries $e^{-8\hat{n}(t-s)}$, and it contributes $\frac{1}{2} \dot{n}^2 \tau \cdot 8 \int_0^t e^{-4\hat{n}s} e^{-8\hat{n}(t-s)} ds = \dot{n} \tau e^{-4\hat{n}t} (1 - e^{-4\hat{n}t})$. Together these give Eq. (9). In the evaluation of Eq. (8) on the loop, the integrand was computed from the unnormalised no-jump states $\mathcal{U}_{s,0} \rho_0$ on a grid of twenty source times. Each $\mathcal{K}(\mathcal{U}_{s,0} \rho_0)$ was propagated under the no-jump Lindbladian of Eq. (3) to every later time, and the integral was taken by the trapezoidal rule.

B Numerical constructions, truncations and convergence

A post-selected estimator is a ratio of expectation values. A Monte Carlo evaluation by sampled trajectories would carry a statistical error large enough to obscure the residuals reported here, so two exact constructions are used instead. For a prepared state of definite syndrome we attach to each mode a flag qubit, initially in $|0\rangle$, and replace each heating jump operator $\hat{L}$ on mode m by $\hat{L} \hat{X}_m$, where $\hat{X}_m$ is the Pauli X operator on the flag of mode m . The Lindblad evolution on the enlarged space is deterministic and exact, and every heating event on mode m flips its flag. The flag operators commute with everything else. The block of the state with all flags in $|0\rangle$ is therefore exactly the unnormalised state of the trajectories with an even number of heating events on every mode since the flags were last reset. A syndrome reading at time $k\tau$ is implemented by projecting the flags onto $|0\rangle$ with the projector $\hat{\Pi}_0$, which discards the rejected trajectories, and continuing the evolution. The post-selected estimator at any time is $\text{tr}(\hat{\Pi}_0 \hat{O} \rho) / \text{tr}(\hat{\Pi}_0 \rho)$ and the yield is $\text{tr}(\hat{\Pi}_0 \rho)$. This is the acceptance rule of a syndrome compared with a known initial syndrome. For the thermal preparation, the initial syndrome is not definite, and the physical rule is acceptance of the nominal syndrome at every reading. That rule is simulated without flags, by evolving under the full Lindbladian and applying the projector $\hat{\Pi}_{\text{nom}} = \frac{1}{4}(1 - \hat{G}_{m_1})(1 + \hat{G}_{m_2})$ at every reading, with the estimator $\text{tr}(\hat{\Pi}_{\text{nom}} \hat{O} \rho) / \text{tr}(\hat{\Pi}_{\text{nom}} \rho)$ within each interval. The two constructions coincide for a state of definite syndrome. They differ for the thermal preparation, where the second also removes the component of the prepared state that violates Gauss's law. The flag projector commutes with every observable of the matter and gauge sectors. The physical projector $\hat{\Pi}_{\text{nom}}$ commutes with every gauge-invariant observable, which includes every observable considered in this paper, but not with operators such as $\hat{a}_m$ that change the syndrome. Commutation means that the reading does not disturb the observable on the kept runs. It does not mean that the selection leaves the expectation unchanged, as the thermal projection and the readout branches of Section 6 show.

The loop calculations use two link qubits and two modes with a Fock space of dimension six, plus two flags in the flag construction. The space has dimension 576 with flags and 144 without. The evolution is carried out with the master-equation solver of QuTiP [37] at the default tolerances of 10^{-8} absolute and 10^{-6} relative unless stated. The reference, the unconditioned evolution and the zero-jump conditional state are computed on the same space from the heating-free Lindbladian, the full Lindbladian and the no-jump Lindbladian of Eq. (3) respectively. The post-selected estimators are computed for syndrome intervals of 1, 0.5, 0.25 and 0.1 milliseconds and for a single final check. The chain and triangle calculations of Table 1 use the no-jump Lindbladian directly, with a Fock dimension of two for the conditional state and four for the unconditioned damage, on spaces of dimension 128 and 1024 for the chain and 64 and 512 for the triangle. In the dimension-two calculations the rate operator is constructed from the analytic form $\hat{K} = \sum_m \dot{n}_m (2\hat{n}_m + 1)$ and not from truncated ladder operators. On a two-dimensional Fock space $\hat{a}\hat{a}^\dagger$ differs from $\hat{n} + 1$ on the upper state and would give the wrong yield.

Theorem 1 provides an exact benchmark for the integrator, and Table 4 reports it. On the loop without flags, the departure of the zero-jump population from the reference is 6×10^{-5} at the default tolerances. It falls to 6×10^{-7} when the tolerances are tightened by two orders of magnitude and to 1×10^{-8} when they are tightened by four, with the same values at Fock dimensions six and eight. The departure of the yield from the closed form falls from 10^{-10} to 10^{-15} . The residuals quoted in the body for the zero-jump state at the default tolerances are therefore integrator error and not truncation error. This benchmark does not test the truncation of the heated trajectories, which reach higher sectors. The post-selected and unconditioned quantities of Table 2 were therefore recomputed at Fock dimension eight at the tighter tolerances, by the nominal-syndrome projection, which coincides with the flag construction for this state. Table 3 reports the comparison. The unconditioned damage changes by 0.004, the final-check yield by 0.0006, the final-check population bias by 0.0005 and the population bias with checks every 0.1 milliseconds by 5×10^{-5} . The flux-projector biases change in the third significant figure. The entries of Table 2 are therefore correct to the precision

Table 3: Truncation study of the heated and post-selected calculations on the zero-flux loop with one phonon, equal heating rates of 150 quanta per second, the experimental coherence times and $h/J = 0.6$, at integrator tolerances of 10^{-10} absolute and 10^{-8} relative. The post-selected quantities are computed by the nominal-syndrome projection, which coincides with the flag construction for this state. Biases are the largest departure within the window.

Quantity	Fock dimension 6	Fock dimension 8
Unconditioned damage to $\bar{n}_{m_2}$	0.4961	0.4997
Final-check yield	0.3976	0.3970
Final-check bias in $\bar{n}_{m_2}$	0.2508	0.2512
Final-check bias in $\langle \hat{\Pi}_\Phi \rangle$	0.01237	0.01233
Yield with checks every 0.1 ms	0.25535	0.25534
Bias in $\bar{n}_{m_2}$ with checks every 0.1 ms	0.01070	0.01065
Bias in $\langle \hat{\Pi}_\Phi \rangle$ with checks every 0.1 ms	0.000536	0.000534

Table 4: Convergence of the zero-jump benchmark on the zero-flux loop with one phonon, equal heating rates of 150 quanta per second, the experimental coherence times and $h/J = 0.6$, without flags. The population error is the largest departure of the zero-jump conditional population of the second site from the heating-free reference within the window, and the yield error is the departure of the computed no-jump weight from $e^{-4\bar{n}t}$ at the end of the window.

Fock dimension	Absolute tolerance	Relative tolerance	Population error	Yield error
6	10^{-8}	10^{-6}	6.5×10^{-5}	2.4×10^{-10}
6	10^{-10}	10^{-8}	6.1×10^{-7}	1.4×10^{-12}
6	10^{-12}	10^{-10}	1.1×10^{-8}	3.4×10^{-15}
8	10^{-8}	10^{-6}	4.6×10^{-5}	2.8×10^{-10}
8	10^{-10}	10^{-8}	5.3×10^{-7}	2.2×10^{-12}
8	10^{-12}	10^{-10}	1.0×10^{-8}	5.0×10^{-15}

quoted, except that the final-check yield rounds to 0.397 at dimension eight. A truncated gain operator changes the rate operator $\hat{K}$ at the upper boundary of the Fock space, and these differences are its effect at $\bar{n}t \leq 0.35$. The thermal preparation is the heralded sideband map of Section 5 applied to a thermal state with $\bar{n}_0 = 0.1$ on the first mode, with a thermal state on the second. The sector projectors needed for Proposition 3 are diagonal in the Fock basis and are applied as observables. The final interval of a syndrome scan ends at the final time whether or not it has the full length τ . The evolution is carried out in units where $J = 1$. The complete scripts, together with the results files from which every number in the paper is read, are available from the author.

References

- [1] Wiese, U.-J. (2013). Ultracold quantum gases and lattice systems: quantum simulation of lattice gauge theories. *Annalen der Physik*, 525, 777–796. <https://doi.org/10.1002/andp.201300104>
- [2] Dalmonte, M., Montangero, S. (2016). Lattice gauge theory simulations in the quantum information era. *Contemporary Physics*, 57, 388–412. <https://doi.org/10.1080/00107514.2016.1151199>
- [3] Bañuls, M. C., Blatt, R., Catani, J., Celi, A., Cirac, J. I., Dalmonte, M., Fallani, L., Jansen, K., Lewenstein, M., Montangero, S., Muschik, C. A., Reznik, B., Rico, E., Tagliacozzo, L., Van Acoleyen, K., Verstraete, F., Wiese, U.-J., Wingate, M., Zakrzewski, J., Zoller, P. (2020). Simulating lattice gauge theories within quantum technologies. *The European Physical Journal D*, 74, 165. <https://doi.org/10.1140/epjd/e2020-100571-8>
- [4] Martinez, E. A., Muschik, C. A., Schindler, P., Nigg, D., Erhard, A., Heyl, M., Hauke, P., Dalmonte, M., Monz, T., Zoller, P., Blatt, R. (2016). Real-time dynamics of lattice gauge theories with a few-qubit quantum computer. *Nature*, 534, 516–519. <https://doi.org/10.1038/nature18318>
- [5] Schweizer, C., Grusdt, F., Berngruber, M., Barbiero, L., Demler, E., Goldman, N., Bloch, I., Aidelburger, M. (2019). Floquet approach to $\mathbb{Z}_2$ lattice gauge theories with ultracold atoms in optical lattices. *Nature Physics*, 15, 1168–1173. <https://doi.org/10.1038/s41567-019-0649-7>
- [6] Mildenerger, J., Mruczkiewicz, W., Halimeh, J. C., Jiang, Z., Hauke, P. (2025). Confinement in a $\mathbb{Z}_2$ lattice gauge theory on a quantum computer. *Nature Physics*, 21, 312–317. <https://doi.org/10.1038/s41567-024-02723-6>

- [7] Băzăvan, O., Saner, S., Tirrito, E., Araneda, G., Srinivas, R., Bermudez, A. (2024). Synthetic $\mathbb{Z}_2$ gauge theories based on parametric excitations of trapped ions. *Communications Physics*, 7, 229. <https://doi.org/10.1038/s42005-024-01691-w>
- [8] Saner, S., Băzăvan, O., Webb, D. J., Araneda, G., Ballance, C. J., Srinivas, R., Lucas, D. M., Bermúdez, A. (2025). Real-time observation of Aharonov–Bohm interference in a $\mathbb{Z}_2$ lattice gauge theory on a hybrid qubit–oscillator quantum computer. *arXiv preprint*, arXiv:2507.19588. <https://doi.org/10.48550/arXiv.2507.19588>
- [9] Aharonov, Y., Bohm, D. (1959). Significance of electromagnetic potentials in the quantum theory. *Physical Review*, 115, 485–491. <https://doi.org/10.1103/PhysRev.115.485>
- [10] Ugail, H. (2026). Predicting noise in a trapped-ion quantum simulation of a lattice gauge theory. *arXiv preprint*, arXiv:2610.03823. <https://doi.org/10.48550/arXiv.2610.03823>
- [11] Bonet-Monroig, X., Sagastizabal, R., Singh, M., O’Brien, T. E. (2018). Low-cost error mitigation by symmetry verification. *Physical Review A*, 98, 062339. <https://doi.org/10.1103/PhysRevA.98.062339>
- [12] McArdle, S., Yuan, X., Benjamin, S. (2019). Error-mitigated digital quantum simulation. *Physical Review Letters*, 122, 180501. <https://doi.org/10.1103/PhysRevLett.122.180501>
- [13] Froland, H., Grabowska, D. M., Grieninger, S., Hartse, J., Lashbrook, A. L., Li, Z., Li, Z., Powell, S. J. M., Savage, M. J., Yao, X., Zemlevskiy, N. A. (2026). The utility of sparse error detection in quantum simulations. *arXiv preprint*, arXiv:2608.02944. <https://doi.org/10.48550/arXiv.2608.02944>
- [14] Ballini, E., Mildenerger, J., Wauters, M. M., Hauke, P. (2025). Symmetry verification for noisy quantum simulations of non-Abelian lattice gauge theories. *Quantum*, 9, 1802. <https://doi.org/10.22331/q-2025-07-22-1802>
- [15] Modi, N. S., Nagano, L., Honda, M., Yoshioka, N., Bauer, C. W. (2026). Arbitrary-distance quantum error correction with Gauss’s law for $\mathbb{Z}_2$ lattice gauge theory. *arXiv preprint*, arXiv:2608.00165. <https://doi.org/10.48550/arXiv.2608.00165>
- [16] Stannigel, K., Hauke, P., Marcos, D., Hafezi, M., Diehl, S., Dalmonte, M., Zoller, P. (2014). Constrained dynamics via the Zeno effect in quantum simulation: implementing non-Abelian lattice gauge theories with cold atoms. *Physical Review Letters*, 112, 120406. <https://doi.org/10.1103/PhysRevLett.112.120406>
- [17] Halimeh, J. C., Hauke, P. (2020). Reliability of lattice gauge theories. *Physical Review Letters*, 125, 030503. <https://doi.org/10.1103/PhysRevLett.125.030503>
- [18] Halimeh, J. C., Kasper, V., Hauke, P. (2020). Fate of lattice gauge theories under decoherence. *arXiv preprint*, arXiv:2009.07848. <https://doi.org/10.48550/arXiv.2009.07848>
- [19] Halimeh, J. C., Lang, H., Mildenerger, J., Jiang, Z., Hauke, P. (2021). Gauge-symmetry protection using single-body terms. *PRX Quantum*, 2, 040311. <https://doi.org/10.1103/PRXQuantum.2.040311>
- [20] Chuang, I. L., Leung, D. W., Yamamoto, Y. (1997). Bosonic quantum codes for amplitude damping. *Physical Review A*, 56, 1114–1125. <https://doi.org/10.1103/PhysRevA.56.1114>
- [21] Suri, N., Saied, J., Venturelli, D. (2024). Uniformly decaying subspaces for error-mitigated quantum computation. *Physical Review A*, 110, 042621. <https://doi.org/10.1103/PhysRevA.110.042621>
- [22] Ugail, H. (2026). Reliable Weyl diagnostics for an inhomogeneous universe in numerical relativity. *arXiv preprint*, arXiv:2608.06383. <https://doi.org/10.48550/arXiv.2608.06383>
- [23] Ugail, H., Howard, N. (2026). A channel-level diagnostic for symmetry breaking in noisy equivariant quantum neural networks. *IEEE Access*, 14, 95152–95167. <https://doi.org/10.1109/ACCESS.2026.3706394>
- [24] Ugail, H., Howard, N. (2026). A coherence law for trainability in noisy equivariant quantum neural networks. *arXiv preprint*, arXiv:2606.30688. <https://doi.org/10.48550/arXiv.2606.30688>
- [25] Lindblad, G. (1976). On the generators of quantum dynamical semigroups. *Communications in Mathematical Physics*, 48, 119–130. <https://doi.org/10.1007/BF01608499>
- [26] Gorini, V., Kossakowski, A., Sudarshan, E. C. G. (1976). Completely positive dynamical semigroups of N -level systems. *Journal of Mathematical Physics*, 17, 821–825. <https://doi.org/10.1063/1.522979>
- [27] Breuer, H.-P., Petruccione, F. (2007). *The Theory of Open Quantum Systems*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199213900.001.0001>
- [28] Turchette, Q. A., Kielpinski, D., King, B. E., Leibfried, D., Meekhof, D. M., Myatt, C. J., Rowe, M. A., Sackett, C. A., Wood, C. S., Itano, W. M., Monroe, C., Wineland, D. J. (2000). Heating of trapped ions from the quantum ground state. *Physical Review A*, 61, 063418. <https://doi.org/10.1103/PhysRevA.61.063418>

- [29] Brownnutt, M., Kumph, M., Rabl, P., Blatt, R. (2015). Ion-trap measurements of electric-field noise near surfaces. *Reviews of Modern Physics*, 87, 1419–1482. <https://doi.org/10.1103/RevModPhys.87.1419>
- [30] Dalibard, J., Castin, Y., Mølmer, K. (1992). Wave-function approach to dissipative processes in quantum optics. *Physical Review Letters*, 68, 580–583. <https://doi.org/10.1103/PhysRevLett.68.580>
- [31] Plenio, M. B., Knight, P. L. (1998). The quantum-jump approach to dissipative dynamics in quantum optics. *Reviews of Modern Physics*, 70, 101–144. <https://doi.org/10.1103/RevModPhys.70.101>
- [32] Michael, M. H., Silveri, M., Brierley, R. T., Albert, V. V., Salmilehto, J., Jiang, L., Girvin, S. M. (2016). New class of quantum error-correcting codes for a bosonic mode. *Physical Review X*, 6, 031006. <https://doi.org/10.1103/PhysRevX.6.031006>
- [33] Leibfried, D., Blatt, R., Monroe, C., Wineland, D. (2003). Quantum dynamics of single trapped ions. *Reviews of Modern Physics*, 75, 281–324. <https://doi.org/10.1103/RevModPhys.75.281>
- [34] Flühmann, C., Nguyen, T. L., Marinelli, M., Negnevitsky, V., Mehta, K., Home, J. P. (2019). Encoding a qubit in a trapped-ion mechanical oscillator. *Nature*, 566, 513–517. <https://doi.org/10.1038/s41586-019-0960-6>
- [35] An, S., Zhang, J.-N., Um, M., Lv, D., Lu, Y., Zhang, J., Yin, Z.-Q., Quan, H. T., Kim, K. (2015). Experimental test of the quantum Jarzynski equality with a trapped-ion system. *Nature Physics*, 11, 193–199. <https://doi.org/10.1038/nphys3197>
- [36] Pino, J. M., Dreiling, J. M., Figgatt, C., Gaebler, J. P., Moses, S. A., Allman, M. S., Baldwin, C. H., Foss-Feig, M., Hayes, D., Mayer, K., Ryan-Anderson, C., Neyenhuis, B. (2021). Demonstration of the trapped-ion quantum CCD computer architecture. *Nature*, 592, 209–213. <https://doi.org/10.1038/s41586-021-03318-4>
- [37] Johansson, J. R., Nation, P. D., Nori, F. (2013). QuTiP 2: a Python framework for the dynamics of open quantum systems. *Computer Physics Communications*, 184, 1234–1240. <https://doi.org/10.1016/j.cpc.2012.11.019>