

Fast Design and Scaling of Multiqubit Gates in Large-Scale Trapped-Ion Quantum Computers

Lee Peleg^{1,3,*}, David Schwerdt^{1,*}, Jonathan Nemirovsky^{1,3,*}, Yotam Shapira^{1,3,*},
Nitzan Akerman¹, Ady Stern², Amit Ben Kish³, and Roei Ozeri¹

¹*Department of Physics of Complex Systems, Weizmann Institute of Science, Rehovot, 7610001, Israel*

²*Department of Physics of Condensed Matter Systems, Weizmann Institute of Science, Rehovot, 7610001, Israel*

³*Quantum Art LTD, Ness Ziona, 7403682, Israel*



(Received 12 March 2026; accepted 4 June 2026; published 31 July 2026)

Quantum computers based on crystals of trapped ions are a prominent technology for quantum computation. A unique feature of trapped ions is their long-range Coulomb interactions, which can be exploited to realize large-scale multiqubit entanglement gates. However, scaling up the number of qubits, N , in these systems, while retaining high-fidelity and high-speed operations, is challenging. Specifically, designing multiqubit entanglement gates in long ion crystals of hundreds of ions involves an NP-hard optimization problem, rendering scale-up not only a technological challenge, but also a conceptual challenge. Here we introduce a method that mitigates this challenge, effectively allowing for a polynomial-time design of fast, robust, and programmable entanglement gates, acting on the entire ion-crystal. We show that while the number of simultaneous entanglement operations scales as N^2 , the gate duration scales as N , leading to a scaling advantage. We use our methods to investigate the drive-power requirements and susceptibility to noise and errors of these multiqubit gates. Our method delineates a path toward scaling up quantum computers based on ion-crystals with hundreds of qubits.

DOI: [10.1103/r78y-3q89](https://doi.org/10.1103/r78y-3q89)

I. INTRODUCTION

Trapped-ions are a leading quantum computation platform due to their accurate control, long-range connectivity, and long coherence times. Despite their all-to-all connectivity, linear ion crystals of growing lengths present increasing difficulty in implementing high-fidelity and high-speed entanglement gates. Some trapped ion scale-up architectures overcome this challenge by interconnecting separate ion crystals, either by ion shuttling between segments in a quantum charge-coupled device architecture [1–3] or by photonic interconnects [4–6]. However, both approaches to scalability will benefit from working with longer ion crystals as their basic building block, taking full advantage of the inherent long-range connectivity of the ions and its expected benefits [7–11].

Quantum information processing devices, based on crystals of 10 to 100 s of trapped ions, have recently been proposed and implemented [12–16], overcoming key hurdles

such as crystal stability, cooling, and coherence. However, a prominent open challenge is the design and implementation of high-fidelity, programmable multiqubit entangling gates in these large-scale systems. In trapped-ion quantum computers, entanglement is mediated by the normal modes of motion of the ion crystal. As the number of ions increases, the motional spectrum becomes denser, which complicates the gate design. This challenge manifests itself as a quadratically constrained NP-hard optimization problem [7,17], that scales as N^2 with the number of ions N , making the implementation and the study of feasibility, performance, and scaling of such multiqubit gates a formidable challenge.

Crucially, solving, or approximately solving, this optimization problem enables the implementation of multiqubit entangling gates that offer programmable simultaneous interactions between all qubit pairs in the ion-crystal. Specifically, for a N qubit register and given a desired set of $\binom{N}{2} = N(N-1)/2$ real-valued qubit-qubit couplings, $\varphi_{nn'}$, with $n, n' = 1, \dots, N$, these gates couple to multiple modes of motion of the ion-crystal and generate the logical operation,

$$U_{\text{MQ}}(\varphi) = \exp\left(i \sum_{n,n'>n}^N \varphi_{nn'} X_n X_{n'}\right), \quad (1)$$

with X_n a Pauli- x operator acting on the n th qubit.

*These authors contributed equally to this work.

[†]lee.peleg@quantum-art.tech

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

Such gates, with arbitrary control of $\varphi_{nn'}$, together with local single-qubit rotations, form a universal gate set. Moreover, these gates have been shown to be advantageous in improving the performance and reducing circuit depth over a wide variety of quantum computation applications, to mention a few: Improved thresholds for quantum error correction, improved performance of optimization problems, reduced depth for the quantum Fourier transform, quantum volume circuits, Clifford circuits, N -qubit Toffoli gates, as well as for a wide range of random unitaries [8–11,18–24].

In recent years, several frameworks have been proposed for engineering entangling gates in large trapped-ion crystals, all making use of additional drive parameters, such as temporal or frequency-domain modulation of the gate's driving field's amplitude [8,25] or phase [26–28], to address the quadratically growing constraints of achieving a target entangling map, $\varphi_{nn'}$, while minimizing drive power. One common approach is to focus on a small subset or a set of disconnected pairs of ions within a larger crystal, such that the number of quadratic constraints per ion reduces significantly and efficient solution methods become available [29–34]. Another line of work demonstrates the efficient construction of solutions for general entanglement maps, which are not necessarily power-optimal [8]. Finally, some studies employ custom or standard gradient-based minimization algorithms to obtain optimized solutions under the full set of constraints [25,28,35–37].

Our contribution here is twofold. First we introduce a toolbox, coined *large-scale fast* (LSF), for designing multiqubit gates of the form of Eq. (1) for qubits based on ion crystals of up to 100 s of qubits. Some of these tools can work in conjunction and enrich the previous methods above, while others are unique to LSF. Second, we make use of these techniques to investigate the physical properties of these gates, for gates of up to $N = 100$ ions. Particularly, we show that while the multiqubit gates implement $\mathcal{O}(N^2)$ entanglement operations, their minimal duration, $T_{\min}$, below which general high-fidelity gate solutions cannot be found, scales linearly in N , constituting a scaling advantage. Furthermore, we derive and demonstrate an estimate of the power required to drive general gates in long crystals, and show it is directly deduced from the entanglement map, φ . Lastly, we introduce common error sources, e.g., motional modes frequency drifts, whose influence can be modeled within the physical framework used to derive LSF, and investigate how these errors scale with the properties of the ion crystal and gate drive, providing design and calibration tolerances for large-scale trapped ions processors.

The remainder of the paper is ordered as follows: in Sec. II we introduce the physical framework, formulate the gate-design problem and describe operational aspects of LSF; in Sec. III we highlight the physical limits and scaling of laser power and $T_{\min}$ with the number of ions

and entanglement structure complexity; in Sec. IV we perform an error analysis survey of common error mechanisms and describe their scaling with gate-time and number of ions; in Sec. V we demonstrate other physical properties of an LSF solution through a specific stabilizer-code gate example in a 49-ion chain. The next sections go into the details of the LSF method: in Sec. VI, we describe the derivation of the LSF method; and in Sec. VII we describe the result of a benchmark test comparing the time-to-solution between our algorithm and competing common-use algorithms.

II. FORMULATION OF THE GATE DESIGN PROBLEM

Trapped ion quantum computers use the normal modes of motion of the ion crystal as a phonon bus, which mediates interactions between the ion qubits. This is performed by driving the motional sidebands of the ion crystal, which generates spin-dependent forces. A canonical example of this method is the Mølmer-Sørensen (MS) gate [38,39]. In recent years there have been many proposals and demonstrations which were focused on improving the utility and fidelity of MS gates. These methods are, at large, based on modulating the spin-dependent displacement forces by applying a non-static gate drive $\Omega_n(t)$. While various signal representations may be used, here we analyze the problem in the spectral domain [40–43]. That is, we use the Fourier representation over the frequency grid, that is symmetrically detuned around the qubit transition frequency, ω_0 and spans the frequencies, $\omega_m = (m_0 + m)\frac{\pi}{T_g}$, with T_g the designated gate time, and m and m_0 integers. The drive acting on the n th qubit is expressed as

$$\Omega_n(t) = \Omega \sin(\omega_0 t) \left[\sum_m r_{nm}^{(c)} \cos(\omega_m t) + r_{nm}^{(s)} \sin(\omega_m t) \right], \quad (2)$$

with $r_{nm}^{(c/s)}$ corresponding to the amplitude of tone ω_m driving ion n at the cosine or sine quadratures and Ω a characteristic Rabi frequency. The spectral approach taken here is general in terms of the possible qubit drives, and offers conceptual advantages over alternative representations, as it is easier to manipulate analytically [44] and provides physical intuition. Our approach is relevant to any form of qubit encoding; e.g., ground state, optical or metastable qubits [45] as well as mixed-species crystals, qubit drive; e.g., Raman, optical or laser-free, and architecture; e.g., global beam [46,47] or individually controlled ions.

Applying such drive results in an interaction Hamiltonian of the form ($\hbar = 1$),

$$H_I = \Omega \sum_{n=1}^N X_n \sum_{j=1}^N \eta_j O_j^{(n)} (a_j^\dagger e^{i\nu_j t} + \text{H.c.}) \times \sum_{m=1}^M [r_{nm}^{(c)} \cos(\omega_m t) + r_{nm}^{(s)} \sin(\omega_m t)] \quad (3)$$

with a_j a bosonic annihilation operator of the j th phonon normal mode of motion of the ion-crystal, at frequency ν_j . The ion-phonon coupling is $\eta_j O_j^{(n)}$, with η_j the Lamb-Dicke parameter of the j th mode of motion and O an orthonormal participation matrix such that $O_j^{(n)}$ is the participation of the n th ion in the j th mode of motion.

Various assumptions and approximations are used in order to derive the interaction in Eq. (3) (given in full detail in Appendix A), most importantly the Lamb-Dicke approximation, i.e. assuming that $\eta_j^2(2n+1) \ll 1$, and a rotating wave approximation used to omit unwanted coupling to the qubit carrier transitions, i.e. $\Omega_n \ll \nu$. Some of these approximations can be made more accurate by appropriately constraining the qubit modulations, $\Omega_n(t)$ [7,40,41] (see Appendix B).

The evolution due to H_I is analytically solvable, yielding a combination of qubit-dependent displacements of the phonon mode and qubit-qubit entanglement,

$$U_I(t) = \prod_{j=1}^N D_j \left(\sum_{n=1}^N X_n \alpha_j^{(n)}(t) \right) e^{i \sum \varphi_{nn'}(t) X_n X_{n'}}, \quad (4)$$

with D_j a displacement operator of the j th phonon mode. The displacements, $\alpha_j^{(n)}$, and entanglement phases, φ , are directly computed from the system parameters η_j , $O_j^{(n)}$, ν_j , and the qubit drives, Ω_n (see Appendix A).

The evolution in Eq. (4) gives rise to the gate design problem, namely obtaining the set of qubit drives, $\{\Omega_n\}_{n=1}^N$ such that at the gate time, $t = T_g$, all the displacements vanish, $\alpha_j^{(n)} = 0$ and such that $\varphi_{nn'}(T_g)$ take predetermined values, $\varphi_{nn'}^{(\text{ideal})}$, corresponding to a desired entangling gate. Gate design solutions which make use of a minimal driving power are typically preferable: they are better aligned with the approximations used to derive H_I ; some sources of error scale with the drive power; and physical realizations are often limited by the drive power.

Thus, the gate design problem becomes a norm-minimization problem, under quadratic and linear constraints:

$$\hat{r} = \arg \min_r \|r\| \quad \text{subject to} \quad \begin{cases} \alpha_j^{(n)}(T_g) = 0 & \forall j, n \\ \varphi(T_g) = \varphi_{nn'}^{(\text{ideal})} & \forall n, n', \end{cases} \quad (5)$$

with $j, n, n' = 1, \dots, N$. Crucially, the $\alpha_j^{(n)}$ s scale linearly in the Ω_n s. Additional constraints, typically also linear in the Ω_n s may be imposed in order to improve the accuracy of H_I and to improve the gate's resilience to various sources of error and noise, e.g., calibration errors of the ν_j s (see Appendix B).

The entanglement phases, $\varphi_{nn'}$ scale quadratically in the Ω_n s, giving rise to NP-hard constraints [17,48]. These constraints can be arranged into a linear constraint matrix L with kernel K and a set of real-valued $M \times M$ symmetric quadratic constraint matrices $A_{nn'}$. The form of the $A_{nn'}$ s arises naturally in the derivation of U_I in Eq. (4) (see Appendix A).

A simpler representation of Eq. (5) follows from the restriction of the drive into the space spanned by K , by applying the transformation $A_{nn'} \rightarrow \mathcal{A}_{nn'} = K^T A_{nn'} K$. The optimization problem can be redefined as

$$\hat{\mathcal{R}} = \arg \min_{\mathcal{R}} \|\mathcal{R}\| \quad \text{subject to} \quad \mathcal{R}^T \mathcal{A}_s \mathcal{R} = \varphi_s^{(\text{ideal})}. \quad (6)$$

With $s = 1, \dots, \mathcal{N}_c$ an ordering of the double index nn' , $\mathcal{N}_c = \binom{N}{2}$ is the number of two-qubit pairings and $\mathcal{R}^T = (\mathbf{R}_1^T \mathbf{R}_2^T \dots \mathbf{R}_N^T)$ defined such that $(\mathbf{R}_n)_l$ is the amplitude associated with the l th vector in K driving the n th ion. $\mathcal{A}_s$ are the set of $N\tilde{M} \times N\tilde{M}$ symmetric sparse matrices, written such that $\mathcal{R}^T \mathcal{A}_s \mathcal{R} = \mathbf{R}_n^T A_{nn'} \mathbf{R}_{n'}$ and $\tilde{M} = \dim K$.

The norm in Eq. (6) used to evaluate $\|\mathcal{R}\|$ is intentionally left undefined, as different physical realizations correspond to different choices of norm. In what follows below, we make use of the standard Euclidean norm.

By using the “ $\mathcal{R}$ ” notation the gate design problem assumes the form of that described in Ref. [7], for homogeneously driven ions. Other realizations of drive configuration, such as beams that illuminate sections of the ion chain also conform to the form of Eq. (6) [49].

In the adiabatic limit, i.e. for a gate time T_g much larger than the inverse motional modes density, the $A_{nn'}$ s become approximately diagonal and satisfying the quadratic constraints is trivial (see Appendix C). However due to the crowding of the mode frequencies in large ion crystals this results in impractically slow gates. Indeed as $2\pi\Delta\bar{\nu} \cdot T_g \rightarrow 1$, with $\Delta\bar{\nu}$ the average spacing of motional mode frequencies, the tones of the driving field strongly interact with many modes and the coupling matrices are in general dense, making the optimization problem nontrivial.

Before continuing to an analysis of multiqubit gates, we briefly give an account of operational aspects of LSF (followed by a detailed account in Secs. VI and VII). LSF operates by transforming a special seed solution of the gate-design problem, coined a *zero phase solution* (ZPS), into an approximate solution of the optimization problem in Eq. (6), with a polynomial-time process that has two main parts—a step that converts the ZPS to a

solution satisfying the quadratic constraint, and then a set of norm-reduction steps. This process can be performed with several seed ZPSs. Thus, the “hardness” of the problem is transformed to the task of generating a finite set of ZPSs, which might be inefficient, after which generating a solution for each ideal target, $\varphi_s^{(\text{ideal})}$ s, is done efficiently.

III. SCALABILITY IN LARGE ION CRYSTALS

We use LSF in order to investigate the performance and scaling laws of multiqubit entanglement operations in large ion crystals. To this end we consider many trapped ions systems which vary in number of ions, gate duration, entanglement target, and drive spectra, among other examples.

First, we consider the total Rabi frequency required to drive an entanglement gate, which we quantify by $|\mathbf{r}| \equiv \sqrt{\sum_{n,m} (\mathbf{r}_n)_m^2}$. Since the gate design stems from an NP-hard problem, one would expect that predicting the required total Rabi frequency, before solving the optimization problem, to be challenging. Nevertheless such a prediction is useful as it can inform circuit construction and as an *a priori* stopping criteria for the norm-reducing process.

As a benchmark we recall the expression for the Rabi frequency which is required to drive a single mode MS gate, $\Omega_{\text{MS}} = \frac{\sqrt{|\varphi_{\text{MS}}|}}{\sqrt{2\pi\eta T}}$, where φ_{MS} is the entanglement phase, and $\eta \propto \nu^{-1/2}$ is the Lamb-Dicke parameter corresponding to the mode of motion at frequency ν . This expression is valid in the adiabatic limit, i.e. for $T_g \gg \nu^{-1}, \Delta\nu^{-1}$. We then generalize Ω_{MS} to a system in which the target $\varphi_{nn'}$ is native, i.e. it can be implemented with a global driving field. This is achieved in the case where the normal modes of motion of the ion crystal are the eigenvectors of the matrix $\varphi_{nn'}$ [7], and each mode accumulates a phase that is the corresponding eigenvalue. We incorporate this change by replacing $\varphi_{\text{MS}} \mapsto \sum_{j=1}^N |\varphi_j|$, where $\{\varphi_j\}$ are the eigenvalues of $\varphi_{nn'}$. This sum is known as the *nuclear norm* of the matrix $\varphi_{nn'}$. Furthermore we replace η with the average η over all modes of motion. When we go beyond native targets, we replace all the elements of $\varphi_{nn'}$ with their absolute value. Lastly, for independently driven ions we expect the power to scale linearly with N . Thus, we conjecture an estimate, Ω_{nuc} ,

$$\Omega_{\text{nuc}} = k \frac{\sqrt{N \|\varphi\|_{\text{nuc}}}}{\sqrt{2\pi\langle\eta\rangle T_g}}, \quad (7)$$

with $\|\cdot\|_{\text{nuc}}$ the nuclear norm of a matrix with its elements taken in absolute value, and k a constant which depends on the choice of implemented linear constraints (e.g., additional linear constraints ensuring gate robustness).

To test the predictive power of Ω_{nuc} , we generated a large set of entanglement maps for an equidistant crystal of 50 trapped $^{40}\text{Ca}^+$ ions, with an ion spacing of $d = 5 \mu\text{m}$ and a fixed gate time of $T_g = 780 \mu\text{s}$, that drive the crystals transversal (radial) modes of motion. These maps are grouped into several categories—structured maps include configurations where the entangling pattern follows a defined geometry or purpose, such as *all-to-all*, *cluster-state*, *surface-code*, *pairwise*, and *single pair* couplings. In these maps, the entangling phase between specific pairs is set to either 0 or $\pi/4$, according to the group pattern. Moreover, we draw subsets of size $N' < 50$ and repeat the gate with the pattern applied only to the ions in the subset. In contrast, *unstructured* maps are generated by assigning a nontrivial entangling phases to randomly selected pairs. For some groups we drop the 0 or maximal-entangling phase assertion and assign different pairs a randomly drawn phase within the range $[-\pi/4, \pi/4]$. For each entanglement target we generated solutions from five different ZPS's and averaged the resulting $|\mathbf{r}|$ values. Here and in the results below, solutions to the optimization problem are obtained such that the total deviation from the quadratic constraints, $\sum_s (\varphi_s - \varphi_s^{(\text{ideal})})^2$, is at most 10^{-4} , which loosely captures the gate's infidelity.

The results are summarized in Fig. 1, where each point corresponds to a specific $\varphi_{nn'}$ example, and color is indicative of the entangling map category. For each such point the resulting drive amplitude, $|\mathbf{r}|$, is compared with the prediction Ω_{nuc} . Remarkably, the data collapse shown in Fig. 1 implies that the details of the mode structure and frequency are largely irrelevant to the required Rabi frequency. The analogy to a globally driven system provides an intuitive explanation: Driving the ions independently is equivalent to an effective “reshaping” of the participation of each ion in each of the normal modes, such that the resulting reshaped mode participation map fits the required spin coupling map. While this picture is intuitive, it could not have been verified without the ability to compute optimal, multi-mode, entanglement gates on large ion-crystals, afforded by LSF.

The factor k , empirically set to be 0.78 from a fit to the data, is applicable when only the bare gate constraints are applied. With the inclusion of additional constraints, that aim at robustifying the gate against noise (e.g., reduce gate sensitivity to changes in mode frequencies), this factor assumes larger values, reflecting the additional cost associated with the inclusion of more conditions, see Appendix D.

It should be noted that our estimation is heuristic, specifically it assumes that the modes of motion of the ion crystal are global, i.e. that in general all ions are coupled to each other. Indeed, it is easy to come up with cases in which ions vibrate independently from each other




FIG. 1. Correlation between the solution drive amplitude and its estimator Ω_{nuc} for 150 different phase maps on a single $N = 50$ ion crystal. Color indicates the category of the multiqubit coupling maps. These include *structured maps* (e.g., all-to-all), where N or a subset of ions follow a defined entangling pattern, and *unstructured maps*, where phases are randomly assigned to ion pairs. Phases are either restricted to $\{0, \pi/4\}$ or drawn from the continuous range $[-\pi/4, \pi/4]$ for *random-phase maps*. The dashed line indicates perfect correlation, $|r| = \Omega_{\text{nuc}}$.

and ion-ion entanglement is impossible, e.g., Refs. [3,15], yet the nuclear norm estimate will not diverge.

We remark that while the eigenvalues of φ play a crucial role in determining the gate drive power, they trivially also carry information about the amount of entanglement generated by the gate. Indeed, for small values of φ the average single-qubit von Neumann entropy scales as the Frobenius norm squared of φ , i.e. the sum over its eigenvalues squared (see Appendix E), $\bar{S}_{\text{VN}}(\varphi) \approx \sum_n \varphi_n^2 = \|\varphi\|_{\text{F}}^2$. For these norms we have the bound, $\|\varphi\|_{\text{F}} \leq \|\varphi\|_{\text{nuc}}$, showing that the drive power puts a bound on the average entanglement entropy generated by the gate.

Following the validation of Eq. (7) for a single ion-crystal and specific gate time, we proceeded to a comparison between systems of differing ion-count and gate-times. With our choice of harmonic basis to the gate drive, faster gates contain less frequency components in the vicinity of mode frequencies, that resolve the crystal motional modes. This naturally leads to a controllability limit over the individual phases acquired by different modes, and consequently limits the accessible entangling map space. Indeed, in the single-tone, far-detuned regime ($\delta \gg \Delta\nu_j$), a nearly uniform detuning from all motional modes gives rise to localized phonon wave packets and correspondingly local entangling interactions between ions [50–52]. A natural timescale to separate the different regimes is the average mode density,




FIG. 2. Summary of gate-design results obtained by our method. The total required Rabi frequency is plotted as a function of rescaled gate time for ion crystals of different length, indicated by color. Each point corresponds to a different entanglement operation. The horizontal axis shows the dimensionless gate time $\rho = T_g/T_\nu$, where T_ν is the average mode-density timescale defined in Eq. (8). The vertical axis shows the total Rabi frequency $|r|$ normalized by the nuclear-norm based estimator $\Omega_{\text{nuc}}^* = \Omega_{\text{nuc}} \cdot \rho$. With this rescaling, solutions for different crystal sizes and entangling maps collapse onto a single curve. The black dashed line indicates the expected scaling of $|r| \propto 1/\rho$.

$$T_\nu = \frac{N\pi}{\sum_j \Delta\nu_j} \equiv \pi(\Delta\bar{\nu})^{-1}, \quad (8)$$

proportional to the average reciprocal group velocity $v_g^{-1}(\omega)$ [53] of phonons in the ion-crystal. As the mode density grows linearly, on average, with the crystal size, so does T_ν . We define the dimensionless quantity $\rho = T_g/T_\nu$ as the ratio between the gate time and the average mode-density time-scale. This dimensionless parameter provides a consistent framework for comparing systems of different sizes. The drive amplitude is rescaled by $\Omega_{\text{nuc}}^* = \Omega_{\text{nuc}} \cdot \rho$, which replaces the explicit gate-time dependence in Eq. (7), with the characteristic time T_ν . Ω_{nuc}^* is a fixed amplitude estimator per ion-crystal.

This universal scaling is demonstrated in Fig. 2, where results from ion crystals with $N = 20$ to $N = 100$ collapse onto a single curve when plotted as a function of ρ . Within each color group (indicating system size), each point represents a specific entangling operation and gate duration, averaged over five optimization runs. From Eq. (7) we expect the normalized gate amplitude to scale as ρ^{-1} , shown by the unity slope dashed black line. The data collapse provides compelling evidence for a nontrivial property: given a physical description of the ion crystal, a target gate time, and a desired entangling operation, one can estimate the required drive amplitude without solving the gate-design problem in full.

Another emergent universal feature is the threshold below which no feasible solution can be found. These breakdown points, marked by vertical dashed lines, cluster around $\rho = 2.9$. Attempts to design gates of shorter durations consistently fail: these solutions require significantly higher drive power than anticipated from Ω_{nuc} , and many exceed the accepted gate error threshold (marked by $\times$). This signifies a minimal gate time $T_{\text{min}} \simeq 2.9 \cdot T_\nu$ which scales linearly with the crystal size. It is worth noting that T_{min} is exemplified here only using dense entangling maps involving all ions. Gates faster than T_{min} do exist for entangling maps that couple only a part of the crystal [42,53].

IV. ERROR SCALING

Beyond the scaling of resources with crystal size, an equally important consideration is how tolerable noise levels scale for gates involving many ions. To address this, we apply our formalism to analyze the impact of three dominant error mechanisms: drifts in the motional-mode frequencies, drive amplitude fluctuations, and motional heating, and quantify their respective contributions to gate infidelity. These processes are well known as dominant contributions in the noise budget of trapped-ions entanglement gate demonstrations [54,55]. Here, we quantify their impact in crystals of tens of ions. These sources of error remain within the analytic framework of LSF, and can be modeled as modifications of the quadratic and linear matrices that determine mode displacements and entangling phases at gate time, or of the gate drive itself. Errors that are beyond the analytic model, e.g., higher-order Lamb-Dicke terms, are a subject of future research.

For small error terms and mode temperatures $\bar{n}_j \ll 1$, the total gate infidelity, $\mathcal{E} = 1 - F$, with F the gate fidelity, can be expressed as a sum of displacement and phase errors [26,56]:

$$\mathcal{E} \approx \mathcal{E}_\alpha + \mathcal{E}_\varphi = \sum_{n,j} \frac{1}{4} \|\alpha_j^{(n)}(T_g)\|^2 + \sum_{n,n'} \Delta\varphi_{n,n'}^2, \quad (9)$$

with $\Delta\varphi_{n,n'} = \varphi_{n,n'}(T_g) - \varphi_{n,n'}^{(\text{ideal})}(T_g)$.

Motional-mode frequency drifts: In the case of radial modes, such drifts can arise from shifts in the resonance frequency of the Paul trap rf resonator. The single-ion radial secular frequency scales as $\nu_x \sim V_{\text{rf}}$, with V_{rf} the voltage of the oscillating field that generates the trapping pseudo-potential. A drift in this voltage, ΔV , shifts all radial mode frequencies according to $\nu_x^j \rightarrow \nu_x^j + \Delta V \cdot \delta_x^j$ with δ_x^j susceptibilities with a weak dependence on j . The resulting phase and displacement errors can be estimated by recalculating the integrals in Eqs. (A5) and (A6) with the shifted frequencies. Figure 3 illustrates these errors for multiple gate solutions, each designed to generate a GHZ state. The solutions are grouped by ion number, with colors




FIG. 3. Gate fidelity, F , due to a drift in the trap rf voltage, estimated for gates generating GHZ states for ion-crystals of 10 to 60 ions. Gate times of $T_g = 2 \cdot T_{\text{min}}$ (light) and $3 \cdot T_{\text{min}}$ (dark) collapse to similar trends when rescaling the frequency error by T_g (horizontal). Dashed colored-brown lines show gates on 60 ions where only a subset (of size coded by the color) of the ions are entangled, forming a sub-crystal size GHZ state. Inset: A quadratic approximation of the infidelity, with filled circles showing the quadratic coefficient for varying crystal size, with (blue) and without (red) additional linear constraints aiming at reducing the gate sensitivity to such drifts. Both coefficients scale almost linearly with the number of bipartite couplings in the gate, $N_c \sim N^2$, yielding a constant slope of approximately 2.

ranging from blue (10 ions) to brown (60 ions). The error parameter $\Delta\nu$ is normalized by the gate time T_g , which facilitates the collapse of error curves with different gate times onto similar trends for a given crystal size. This is demonstrated by two solutions per crystal size: bright-colored curves correspond to faster gates with $T_g = 2T_{\text{min}}$, and dark-colored curves to slower gates with $T_g = 3T_{\text{min}}$. Brown dashed curves correspond to gates in the 60-ion chain where only a subset of ions participate (according to color). These error curves closely match those of fully-connected gates on smaller crystals of the same number of participating ions.

This data show that the gate errors under this noise model scale quadratically in the frequency error,

$$\mathcal{E}_{\Delta\nu} \approx k N_c (\Delta\nu T_g)^2. \quad (10)$$

Here, $N_c = \frac{N(N-1)}{2}$ is the number of ion-ion couplings and k is an additional (N - and $\Delta\nu$ -independent) constant. This constant is fitted from the data to $k \approx 4.5$ (inset) showing the scaling of the quadratic coefficient with N_c (red). The gate can be made robust to this type of error by including appropriate linear constraints, shown in

Refs. [7,40,41] and Appendix B, which are shown to reduce k to ≈ 1.0 (blue).

Drive amplitude fluctuations: For laser driven ions, such fluctuations can arise, e.g., from instabilities in the laser intensity or beam pointing. To model this effect, we rescale the spectrum driving ion n as $r_n \rightarrow r_n(1 + \epsilon)$, where ϵ denotes the relative amplitude change. Since the displacement integrals are linear in r , and solutions are limited to $K = \ker(L)$, this rescaling of r_n does not cause displacement errors. Amplitude fluctuations, however, do cause phase errors that can be written as

$$\begin{aligned} \mathcal{E}_\varphi &= \sum_s (\mathcal{R}(1 + \epsilon) \mathcal{A}_s \mathcal{R}(1 + \epsilon) - \mathcal{R} \mathcal{A}_s \mathcal{R})^2 \\ &= (4\epsilon^2 + 4\epsilon^3 + \epsilon^4) \sum_s \varphi_s^2. \end{aligned} \quad (11)$$

For homogeneous stochastic noise $\epsilon \sim P_\epsilon$, the powers of ϵ in Eq. (11) are replaced by the corresponding distribution moments. For Gaussian noise, $E[\epsilon^{2n}] = \sigma^{2n}(2n-1)!!$, yielding the expected error,

$$\mathcal{E}_\varphi = (4\sigma^2 + 3\sigma^4) \sum_s \varphi_s^2.$$

Figure 4 shows the expected error for crystals of size 10–60 ions. Filled circles represent averages over 300 Gaussian samples per variance σ^2 , while solid curves follow the analytic prediction. For maximally entangling gates (e.g., Clifford multiqubit gates), and under the assumption of small noise variance, the error scales as

$$\mathcal{E}_{\Delta r} \approx 2\pi\sigma^2 N_c, \quad (12)$$

with N_c the number of nontrivial ion-ion couplings. Thus, the variance of the noise must decrease inversely with N_c , which grows, at most, quadratically with N .

Motional heating: Heating of secular ion motion is a common dephasing process that limits coherent control in ion-crystals. It typically increases with decreasing ion-electrode distance, and is more severe in room-temperature vs cryogenic traps. Following the analysis of [57], the contribution of motional heating to gate error is bounded by

$$\mathcal{E}_\Gamma \leq \sum_{n,n',j} \Gamma_j \left\| \int_0^{T_g} \alpha_j^{*(n)}(t) \alpha_j^{(n')}(t) dt \right\|, \quad (13)$$

where Γ_j is the heating rate (in quanta per unit time) of mode j . The rate depends both on the electric noise power spectral density $S_E(\omega)$ and the spatial correlation of electric-field noise across the ions [58],

$$\Gamma_j = \frac{e^2}{4m\omega_j^2} \sum_{n,n'} O_j^{(n)} O_j^{(n')} S_E^{(n,n')}(\omega_j), \quad (14)$$

with




FIG. 4. Gate infidelity due to fluctuating laser amplitude. For each crystal size, the drive amplitude was rescaled by $1 + \epsilon$ with $\epsilon \sim \mathcal{N}(0, \sigma^2)$. Filled circles correspond to numerical averages (mean of 300 samples), while solid lines show the expectation value of Eq. (11).

$$S_E^{(n,n')}(\omega) = \int d\tau e^{-i\omega\tau} \langle E^{(n)}(\tau) E^{(n')}(0) \rangle,$$

the cross-spectral noise density function. If the noise is perfectly correlated, $S_E^{(n,n')}(\omega) \equiv S_E(\omega)$, then only the center-of-mass (COM) mode contributes to the expression Eq. (14). For a more realistic scenario, we consider a noise model with $S_E(\omega) \propto \omega^{-1}$ and exponentially decaying correlations, such that $S_E^{(n,n')}(\omega) \propto \omega^{-1} e^{-|x_n - x_{n'}|/\zeta}$, with ζ the correlation length of the ions at positions x_n . Figure 5 shows the gate fidelity due to motional heating, for gates that generate GHZ states in ion-crystals of 10 to 60 ions. Since the minimal gate time increases with N , we isolate the effect of noise correlations by rescaling the COM-mode heating rate to 0.02 quanta per gate time. When noise is strongly correlated, $\zeta = 150 \mu\text{m}$ (blue, red), the fidelity decreases approximately linearly with N . By contrast, for short-range noise correlations, $\zeta = 1 \mu\text{m}$ (bright blue), the error approaches quadratic scaling with N . We quantify this behavior in Fig. 5 by fitting the heating-induced error to a power law, $E_\Gamma = 1 - F \approx AN^p$. For short-range electric-field correlations, the fitted exponent $p \simeq 1.88$ reflects the contribution of many motional modes, whereas for long-range correlated noise the error is dominated primarily by the COM-mode and the scaling is closer to linear. Although the sum in Eq. (13) runs over indices $n, n' = 1, \dots, N$, not every ion couples equally to every mode in a given gate. As a result, the scaling is not purely quadratic (as would be expected if all modes contributed independently), but instead interpolates between linear and quadratic scaling depending on the correlation model. Similarly to motional-mode frequency drifts above, robust




FIG. 5. Infidelity due to motional heating. Main frame: heating-induced errors for GHZ gates in crystals of 10–60 ions. To remove the trivial gate-time scaling, the COM-mode heating rate is normalized to 0.02 quanta per gate time. Results are shown for correlated noise ($\zeta = 150 \mu\text{m}$, spanning most of the crystal) and short-range correlated noise ($\zeta = 1 \mu\text{m}$). Solutions including motional error minimization linear constraints (blue) are compared with unconstrained solutions (red). Solid lines are power-law fits to the heating-induced error, $E_{\Gamma} = 1 - F \approx AN^p$, with the fitted exponent p shown in the legend. Inset: error growth with gate time in a 20-ion crystal. Errors increase linearly with T_g , as expected from Eq. (13), with robust gates (blue) showing roughly half the error rate.

gates (blue) are compared to non-robust gates (red), and show a reduced sensitivity to heating-induced errors.

V. ANALYSIS OF AN EXAMPLE: SURFACE CODE STABILIZER

We demonstrate solutions obtained by our method and highlight certain aspects of them via an example. Specifically, we outline the entanglement gate required for stabilizer measurements in surface codes. Here we consider a 49 ions crystal, entangling 33 ions using a single pulse. Figure 6 shows the formation of the stabilizer (top-left panel), with ions straightforwardly mapped into a 7×7 square grid, forming nine plaquettes that can be used to evaluate the X -parity of the plaquette vertices (a similar construction can be made for Z -parity stabilizers). In this mapping some ions remain uncoupled (orange) and can be used in subsequent operations; some form the edges of plaquettes (blue) and some are designated as ancilla qubits (dark blue). It was shown in Ref. [59] that a surface code stabilizer measurement can be implemented with a single multiqubit MS gate. Note, however, that our coupling map between ions involved in a stabilizer measurement is not all-to-all; rather it takes the form of a “cross.” This is in fact more efficient as it requires fewer nonzero




FIG. 6. Single-pulse surface-code stabilizer gate in a 49-ion crystal. Top-left: mapping of a linear 49-ion crystal onto a 7×7 square grid forming nine surface-code plaquettes. Dark-blue ions denote ancilla (node) qubits, light-blue ions form plaquette edges, and orange ions are uncoupled spectators. Upper-right: identification of four representative participating ions used in the spectral analysis. Middle: drive spectrum for these four ions during a single gate pulse. Vertical dark-blue lines indicate transverse motional-mode frequencies of the 49-ion crystal. Bottom: corresponding ion-resolved drive amplitudes, showing higher power assigned to ancilla ions, reduced power on edge ions, and vanishing drive on non-participating ions.

entanglement phases. Indeed the nuclear norm of the cross coupling map is two times lower than that of the all-to-all coupling.

Here we assume an equally spaced crystal of $^{40}\text{Ca}^+$ ions with an inter-ion distance of $5 \mu\text{m}$. Qubits are mapped onto the ground-state Zeeman $5S_{1/2}$ manifold and are assumed to be driven with a 400 nm laser field using a Raman transition, which couples the ions through transverse motional modes at frequencies between 3 and 3.5 MHz. We convert different ZPS’s into solutions for this particular gate problem and compare their properties. The obtained solutions exhibit infidelity below 1×10^{-4} , as defined in Eq. (9). The middle panel of Fig. 6 provides an example of a drive spectrum obtained for four participating ions, labeled in the upper-right panel as 4 (purple), 7 (green), 42 (blue), and 46 (red). The dark-blue vertical lines correspond to motional-mode frequencies of the 49-ion crystal in one of the radial directions. The spectrum describes a pulse of duration $T \approx 640 \mu\text{s}$. Nonetheless,

we obtain solutions for this problem setup for gate times down to $T \approx 320 \mu\text{s}$. For comparison, to produce the same operation using sequential all-to-all two-qubit gates within $T \approx 640 \mu\text{s}$ requires that each two-qubit gate will have a duration of $15 \mu\text{s}$.

The obtained solution exhibits efficient use of the drive, manifested in two ways. First, the power assigned to each ion correlates with its participation in the gate. A representative drive amplitude solution is presented in the bottom panel of Fig. 6. Node ions (dark-blue) exhibit the highest drive amplitude, while edge ions (light-blue) are driven with lower amplitude, with a minor reduction for ions connected to a single qubit. The non-participating orange ions are trivially shut off.

A second property relates to how the total entangling phase is accumulated through all participating modes. Since the total entangling phase of an ion pair is weighted by the pair's participation in each mode [Eq. (A6)], it is beneficial for the solver to entangle an ion pair through modes in which both ions participate substantially. Because there are $\sim N^2/2$ pairwise constraints but only N modes, such efficient combinations are not guaranteed to exist. Nonetheless, the solver finds efficient patterns, as exemplified in the spectrum in Fig. 6. The background color of the four spectra represents the participation of that ion in each mode, also summarized in the upper-right mode-participation matrix, where colored frames highlight the participation structure. In all cases, the solver assigns most power near modes where the ion participates substantially.

The factor-of- N mismatch between the number of ion pairs and the number of motional modes implies that modes cannot be uniquely assigned to individual pairs. As a result, multiple ion pairs inevitably accumulate entangling phase through the same modes, generating unwanted cross-entanglement: a disconnected pair may acquire phase through a given mode even though it is not intended to be entangled. In typical solutions, this is mitigated by enforcing cancellation of the total entangling phase of disconnected pairs, rather than suppressing phase accumulation through each individual mode.

This mechanism is illustrated in the top-left panel of Fig. 7, which shows the time-resolved entangling phases for ion pairs along a selected node (highlighted in the inset). Blue traces correspond to connected pairs and red traces to disconnected pairs; lighter curves show mode-resolved contributions, while darker curves represent the total phase summed over all modes. Although the total phase of disconnected pairs (dark red) is constrained to vanish at the gate time, it is composed of nonzero positive and negative contributions from different modes, resulting in oscillatory phase evolution during the gate.

From an efficiency standpoint, the solver suppresses these unwanted contributions by keeping the mode-resolved phases of disconnected pairs smaller than those




FIG. 7. Accumulation and suppression of entangling phases through shared motional modes. Top-left: time-resolved entangling phases for ion pairs along a selected node (highlighted in the inset). Blue traces show phases for connected ion pairs; red traces correspond to disconnected pairs. Darker curves indicate total entangling phase summed over all modes. Disconnected pairs acquire nonzero intermediate phases that cancel at the gate time. Top-right: histograms of pair-mode entangling phases collected from ten gate solutions, comparing connected (blue) and disconnected (red) ion pairs. Disconnected pairs exhibit a smaller spread, indicating efficient suppression of unwanted couplings despite shared mode participation. Bottom: efficiency score, S_{nm} , for connected ion pairs as a function of the total gate time. Thin traces correspond to individual connected links, with shaded regions indicating the spread across different optimized solutions. The bold curve shows the average link efficiency.

of connected pairs. This is quantified by the histograms shown to the right of the panel, which collect pair-mode phase magnitudes from ten independent solutions. Disconnected pairs (red) exhibit a narrower distribution than connected pairs (blue), indicating that parasitic mode-mediated couplings are kept small. This reduces the fraction of the total drive power that must be allocated to cancel unwanted phases, improving overall gate efficiency.

The bottom panel of Fig. 7 shows how this efficiency degrades as the gate time is reduced. To quantify this effect, we define an efficiency score for each ion pair that measures how strongly its entangling phase is accumulated through modes in which both ions participate. The mode-participation weights, $w_j^{nn'} = |O_j^{(n)} O_j^{(n')}|$, are normalized to span $[0, 1]$,

$$\tilde{w}_j^{nn'} = \frac{w_j^{nn'} - \min_j w_j^{nn'}}{\max_j w_j^{nn'} - \min_j w_j^{nn'}}, \quad (15)$$

and are used to define the efficiency,

$$\mathcal{S}_{nn'} = \frac{\sum_j \tilde{w}_j^{nn'} |\varphi_{nn'}^{(j)}|}{\|\varphi_{nn'}\|_1}. \quad (16)$$

The efficiency scores of individual connected links are plotted as a function of gate time, with shaded regions indicating the spread across different optimized solutions, and the bold curve showing the average efficiency $\bar{\mathcal{S}}_{nn'}$. For long gate times, the efficiency remains approximately constant, indicating that the additional degrees of freedom available do not lead to more efficient mode usage. As the gate time is reduced, the efficiency decreases, reflecting suboptimal phase allocation due to reduced controllability. No solutions are found within the 10^{-4} tolerance for gate times below $T \approx 320 \mu\text{s}$, consistent with Fig. 2.

VI. DERIVATION OF THE LARGE-SCALE FAST METHOD

In providing approximate solutions to Eq. (6), LSF takes two steps: First, solutions to the quadratic problem $\mathcal{R}_{\text{conv}}$ are efficiently converted from a special seed vector $\mathbf{Z}$ called a “zero phase solution.” $\mathbf{Z}$ is a nontrivial solution to a special case of the quadratic problem with $\varphi_s^{(0)} = 0$ for all $s = 1, \dots, \mathcal{N}_C$. A set, $\{\mathbf{Z}\}$, is accumulated once and prior to the solution of any general φ . Second, the norm of $\mathcal{R}_{\text{conv}}$ is minimized iteratively while adhering to the constraint φ until an optimized solution $\mathbf{R}_{\text{LSF}}$ is reached.

A. Conversion of $\mathbf{Z}$ into $\mathbf{R}_{\text{conv}}$

A basic property of the zero-phase problem is that any solution $\mathbf{Z}$ remains a valid solution when scaled by a real number λ . We exploit it to construct a “converted” solution, $\mathbf{R}_{\text{conv}} = \lambda\mathbf{Z} + \mathbf{D}$, with $|\mathbf{D}|$ assumed small, detailed below. For a large enough λ small deviations from the ZPS generates arbitrarily large entanglement phases, that scale as $\lambda|\mathbf{Z}||\mathbf{D}|$. We linearize the quadratic constraints in the $\mathbf{D}$ -vicinity of $\lambda\mathbf{Z}$ and obtain a set of $\mathcal{N}_C$ linear equations for $\mathbf{D}$,

$$\varphi_s = 2\lambda\mathbf{Z}^T \mathcal{A}_s \mathbf{D} + \mathcal{O}(\epsilon), \quad (17)$$

with the assumption $|\mathbf{D}^T \mathcal{A}_s \mathbf{D}| \leq \epsilon$ for all $s = 1, \dots, \mathcal{N}_C$ and ϵ is the acceptable gate error, which we typically specify to be 10^{-4} . The set of linear equations defined by Eq. (17) yields the solution,

$$\mathbf{R}_{\text{conv}} = \lambda\mathbf{Z} + \lambda^{-1} \mathbf{M}_{\text{pinv}}^{-1} \boldsymbol{\varphi}, \quad (18)$$




FIG. 8. Simple example of the LSF method for a single quadratic constraint. We choose $\mathcal{N}_C = 1$ and two tone amplitudes r_1 (horizontal axis) and r_2 (vertical axis), specifically $\mathcal{A}_1 = \text{diag}(1, -2)$ and $\varphi_1 = 2$. The constraint becomes a hyperbola (black dashed), and the colored contours show other possible values of φ . Our optimization rescales the zero-phase solutions, here by $\lambda = 3$, and uses the high density of contours around $\lambda\mathbf{Z}$, (light green arrow), such that only a small excursion by $\mathbf{D}$ (bright green arrow) is required to convert $\mathbf{Z}$ to a solution satisfying the constraints, $\mathbf{R}_{\text{conv}}$. The locally ideal solution, $\mathbf{R}_{\text{LSF}}$ (blue arrow), is obtained by local minimization steps along the solution manifold, $\varphi_1 = 2$, composed of norm-reducing and error-reducing sub-steps (inset).

with $\mathbf{M} \in \mathbb{R}^{\mathcal{N}_C \times \mathcal{N}_C}$ a matrix whose elements are given by $M_{s,j} = 2(\mathbf{Z}^T \mathcal{A}_s)_j$ and $\mathbf{M}_{\text{pinv}}^{-1}$ its pseudo-inverse. To ensure that the quadratic constraints are satisfied up to an accuracy ϵ we choose λ that enforces the consistency condition,

$$\left| \left(\frac{1}{\lambda} \mathbf{M}_{\text{pinv}}^{-1} \boldsymbol{\varphi} \right)^T \mathcal{A}_s \left(\frac{1}{\lambda} \mathbf{M}_{\text{pinv}}^{-1} \boldsymbol{\varphi} \right) \right| \leq \epsilon, \quad (19)$$

for all $s = 1, \dots, \mathcal{N}_C$.

Figure 8 shows the conversion and optimization processes outlined above for the simple case of a single quadratic constraint, i.e. $\mathcal{N}_C = 1$, and two amplitudes, r_1 and r_2 . Specifically we use $\mathcal{A}_1 = \text{diag}(1, -2)$ and $\varphi_1 = 2$, such that the constraints are satisfied along a hyperbola (dashed black) in terms of the two amplitudes, r_1 (horizontal axis) and r_2 (vertical axis). Zero-phase solutions extend from the origin to arbitrarily large amplitudes (dashed gray). We use a large zero-phase solution, $\lambda\mathbf{Z}$ (bright green arrow), such that a small deviation from it, $\mathbf{D}$ (dark green arrow), explores various values of φ_1 . Indeed a converted solution, $\mathbf{R}_{\text{conv}}$ (green arrow), generating the desired target is found close by.

B. Minimization of $\mathcal{R}_{\text{conv}}$ norm

To reduce the solution norm, we define an iteration on $\mathcal{R}$,

$$\mathcal{R}^{(n+1)} = \mathcal{R}^{(n)} + \mathbf{D}^{(n)}. \quad (20)$$

Here $\mathcal{R}^{(0)} = \mathcal{R}_{\text{conv}}$ is the converted ZPS. $\mathbf{D}^{(n)}$ is a correction term computed at each iteration and is built from two sub steps,

$$\mathbf{D}^{(n)} = \delta_{\parallel} \mathbf{D}_{\parallel}^{(n)} + \delta_{\perp} \mathbf{D}_{\perp}^{(n)}, \quad (21)$$

which are computed and applied in sequence. Particularly, $\mathbf{D}_{\parallel}^{(n)}$ is a *norm-reducing* step and $\mathbf{D}_{\perp}^{(n)}$ is an *error reducing* step. The iteration terminates and $\mathcal{R}^{(n)} = \mathcal{R}^{\text{LSF}}$ if $\mathbf{D}^{(n)}$ is unable to significantly reduce the solution norm under the prescribed constraints.

To compute either correction $\mathbf{D}_x^{(n)}$ we linearize the constraints equations similarly to Eq. (18) to produce the following set of equations:

$$\Theta_x^{(n)} = \begin{pmatrix} \Delta \varphi_s^{(n)} \\ \alpha_x \end{pmatrix} = \begin{pmatrix} 2\mathcal{R}^{(n)\top} \mathcal{A}_s \\ \mathcal{R}^{(n)} \end{pmatrix} \mathbf{D}_x \equiv \mathcal{M}^{(n)} \mathbf{D}_x. \quad (22)$$

With $\mathcal{M}^{(n)} \in \mathbb{R}^{(\mathcal{N}_c+1) \times N\tilde{M}}$, and $\Theta_x^{(n)} \in \mathbb{R}^{(\mathcal{N}_c+1)}$, generalizations of M and φ from Eq. (18) respectively. Furthermore, α_x controls the orientation $\mathbf{D}_x$.

Performing the steps in sequence, first a *norm-reducing* step is derived by setting $\alpha_{\parallel} = -|\mathcal{R}^{(n)}|$, enforcing a step in a direction minimizing the current solution norm. A solution to Eq. (22) is found using the pseudo-inverse $\mathbf{D}_{\parallel}^{(n)} = [\mathcal{M}_{\text{pinv}}^{(n)}]^\top \Theta_{\parallel}^{(n)}$, and the resulting constraint error is expressed as a polynomial in $\delta_{\parallel}$,

$$\begin{aligned} \varepsilon_{\varphi}^{(n,\parallel)} &= \sum_s \varepsilon_{\varphi,s}^{(n,\parallel)}, \\ \varepsilon_{\varphi,s}^{(n,\parallel)} &= \left(\varphi_s - \left(\mathcal{R}^{(n)} + \delta_{\parallel} \mathbf{D}_{\parallel}^{(n)} \right)^\top \mathcal{A}_s \left(\mathcal{R}^{(n)} + \delta_{\parallel} \mathbf{D}_{\parallel}^{(n)} \right) \right)^2. \end{aligned} \quad (23)$$

$\delta_{\parallel}^{(n)}$ is found by solving $\varepsilon_{\varphi}^{(n,\parallel)}(\delta_{\parallel}) = \epsilon_{n,r}$, providing the largest step size which yields an error equal or below $\epsilon_{n,r}$, an excess “norm-reducing” error term which is allowed to occur after the norm step. For the *error-reducing* step, Eq. (22) is re-applied, with $\mathcal{R}^{(n)}$ replaced by $\mathcal{R}^{(n)} + \delta_{\parallel} \mathbf{D}_{\parallel}^{(n)}$, and $\alpha_{\perp} = 0$, which guarantees the resulting norm of $\mathcal{R}^{(n+1)}$ is only second-order sensitive to $|\delta_{\perp} \mathbf{D}_{\perp}^{(n)}|$. This time, the resulting constraint error $\varepsilon_{\varphi}^{(n,\perp)}$ is minimized by solving $\partial_{\delta_{\perp}} \varepsilon_{\varphi}^{(n,\perp)} = 0$, with $\varepsilon_{\varphi}^{(n,\perp)}$ defined similarly to Eq. (23). As both $\varepsilon_{\varphi}^{(n,\parallel)}$ and $\varepsilon_{\varphi}^{(n,\perp)}$ are fourth and third polynomials in $\delta_{\parallel/\perp}$, straightforward algebraic

methods, particularly Ferrari’s and Cardano’s methods [60], are applied to solve for $\delta_{\parallel}^{(n)}$ and $\delta_{\perp}^{(n)}$, respectively.

To accept the step we require the step error to be below the gate acceptable error rate, that is $\varepsilon_{\varphi}^{(n,\perp)} < \epsilon$. If the error does not meet the gate error threshold, the step is discarded and a smaller norm-reducing step is calculated by decreasing $\epsilon_{n,r}$. Execution is halted when either $\epsilon_{n,r}$ or the relative norm change $\|1 - \frac{\mathcal{R}^{(n+1)}}{\mathcal{R}^{(n)}}\|$ are reduced below a predefined tolerance value.

C. Aggregation of zero phase solutions

As our method relies on the existence of ZPS’s, these should be obtained at the outset. Therefore, before solving any specific gate (given by $\varphi^{(\text{ideal})}$), we first search for nontrivial solutions to the homogeneous problem $\mathbf{Z} \mathcal{A}_s \mathbf{Z} = 0$, $\forall s$.

ZPSs are produced in a similar iterative procedure. Beginning from a (pseudo-) random vector $\mathbf{z}_0$ of unit-length, a linear correction $\xi^{(n+1)}$ is obtained for $\mathbf{z}^{(n)}$ through the set of equations:

$$\begin{pmatrix} \mathbf{z}^{(n)} \mathcal{A}_s \mathbf{z}^{(n)} \\ 0 \end{pmatrix} = \begin{pmatrix} M^{(n)} \\ \mathbf{z}^{(n)} \end{pmatrix} \xi^{(n+1)}. \quad (24)$$

Here the correction is forced to be perpendicular to the current vector $\mathbf{z}^{(n)}$, owing to the rescaling property of homogeneous solutions. $M_s^{(n)} = 2\mathbf{z}^{(n)} \mathcal{A}_s$ is similarly the linear term around $\mathbf{z}^{(n)}$, and a correction term $\xi^{(n+1)}$ is obtained through the pseudo-inverse of $(M^{(n)}; \mathbf{z}^{(n)})$. The vector $\mathbf{z}^{(n)}$ is updated according to:

$$\mathbf{z}^{(n+1)} = \frac{1}{|\mathbf{z}^{(n)} - \lambda_{\min}^{(n+1)} \xi^{(n+1)}|} \left(\mathbf{z}^{(n)} - \lambda_{\min}^{(n+1)} \xi^{(n+1)} \right). \quad (25)$$

With $\lambda_{\min}^{(n+1)}$ being the value that minimizes the normalized error term $g(\lambda)$,

$$\begin{aligned} g(\lambda) &= \frac{1}{\|\mathbf{z}^{(n)} - \lambda \xi^{(n+1)}\|_2} \\ &\times \sqrt{\sum_s |(\mathbf{z}^{(n)} - \lambda \xi^{(n+1)}) \mathcal{A}_s (\mathbf{z}^{(n)} - \lambda \xi^{(n+1)})|^2}. \end{aligned} \quad (26)$$

A ZPS is adopted from $\mathbf{z}^{(n)}$ once it satisfies the constraint to an acceptable level,

$$\sqrt{\sum_s |\mathbf{z}^{(n)T} \mathcal{A}_s \mathbf{z}^{(n)}|^2} < \epsilon_{\text{ZPS}}, \quad (27)$$

where ϵ_{ZPS} is typically close to machine accuracy. An emerging observation is that, for a given vector $\mathbf{Z}$, a sequence of subsequent solutions $\mathbf{Z}_1, \mathbf{Z}_2, \dots$ can be systematically constructed by repeatedly applying the steps described in Eq. (24) onward [even after satisfying

Eq. (27)], since each iteration effectively searches for ZPSs in directions orthogonal to those previously identified.

Finally, a computational shortcut is available for aggregation of ZPS through the exploitation of a global drive. Since the zero phase gate respects the mirror symmetries between ions along a chain, it can be obtained for a chain where all ions are illuminated homogeneously, as described in Ref. [7]. A global drive represents a version of the problem with a factor of N less constraints (and individual degrees of freedom), providing a quadratic speedup with respect to an individual-addressing gate problem.

VII. ALGORITHM BENCHMARK

To evaluate the resource efficiency of our proposed method in generating solutions, we conducted a benchmark comparison between our algorithm and commonly used optimization routines (available in Python's SciPy package [61]). For each quadratic problem defined by a unique set $\mathcal{A}_s$, we applied the competing algorithms to obtain an optimized solution for Eq. (6), while maintaining consistent memory and CPU constraints. As alternative optimization approaches, we considered two competing algorithms. The first is the conjugate gradient algorithm, an unconstrained optimization method that minimizes a cost function balancing the phase error $\mathcal{E}_\varphi = \sum_{nm} \Delta\varphi_{nm}^2$, with the drive power term $\mathcal{C} = \lambda \mathcal{E}_\varphi + |\mathbf{r}|^2$. The weighting parameter λ is determined based on the predicted optimal drive power, as discussed in the following section. The second is the trust-region constrained optimization method, which enforces the constraint $\mathcal{E}_\varphi < \epsilon$ while minimizing $\mathbf{r}^2$. We provided routines for exact gradient and Hessian calculations wherever they improved the convergence speed of either method. These two specific algorithms outperformed other minimization algorithms in their class on a test problem, and therefore were singled out in a preliminary comparison.

The set of benchmark problems was created as a two dimensional grid of system size N and gate-time factor $r = T/T_{\min}$, with $T_{\min}$ the minimal gate-time for a given crystal length (as detailed in Sec. III). For each parameter pair $P_i = N_i, r_i \cdot T_{\min}(N_i)$, a corresponding set of quadratic matrices $\mathcal{A}_s^{(i)}$ was derived, defining the constraints for the competing algorithms' degrees of freedom. To maintain consistency, the LSF method was initialized at $\mathbf{R}_{\text{conv}}$, a transformed ZPS solution. To ensure fair evaluation, we allowed each competing algorithm to start from two initial conditions: a randomly generated guess and the same $\mathbf{R}_{\text{conv}}$ used in LSF. We tracked execution time, gate fidelity and solution power at each iteration throughout the optimization process until convergence, and compared these performance trends for comparison of the different methods. Figure 9 summarizes the




FIG. 9. Benchmark comparison of time-to-solution for different optimization algorithms: the LSF method (purple), conjugate-gradient and trust-constraint optimizers starting from a random point (red and green respectively); and same optimizers initialized from $\mathbf{R}_{\text{conv}}$ (blue and orange). Results are shown for ion-crystals of different lengths. Solid markers denote the average run time of gate-design problems, while lighter markers of the same color correspond to individual instances with varying gate times. The two right-hand panels show representative optimization trajectories for a 30-ion system: solution power (top) and quadratic error (bottom). These traces compare the LSF method (purple) with trust-constraint (orange) and conjugate-gradient (blue) optimizations initialized at $\mathbf{R}_{\text{conv}}$ starting point. Bold curves indicate the averaged trajectory constructed from all recorded optimization runs. Dashed vertical lines mark the average time-to-solution for each algorithm.

benchmark results, illustrating the average time to solution τ versus system size N . The averages were computed across varying gate-time ratio parameters r , with error bars representing one-sided standard deviations. All algorithms demonstrate a power-law relationship with system size N , with the LSF method (purple) showing consistently shorter solution times compared to trust-region constrained optimization and conjugate gradient optimization initiated with either a random guess or $\mathbf{R}_{\text{conv}}$. The inset provides convergence trends observed for a 30-ion system with different gate durations. Darker lines represent averages of multiple (lighter-colored) instances. Convergence time is defined as the earliest point at which the optimization achieves a value within 5% of the final result while maintaining an error below the threshold $\epsilon = 10^{-4}$. We attribute the improved performance primarily to the explicit control of error through the regulation of $\delta_{\parallel}$ and $\delta_{\perp}$. This is reflected in Fig. 9 as a smooth progression toward the solution, in contrast to the corrective or destructive updates and plateau-like behavior that lead to erratic convergence in gradient-based and trust-region methods.

VIII. SUMMARY

Here we present a systematic framework for the design of large-scale multiqubit entangling gates in trapped-ion

systems. We introduce a collection of theoretical and numerical tools, collectively referred to as the LSF method, which enables the efficient construction of general multi-qubit gates in ion crystals comprising up to hundreds of qubits. Building on this framework, we use LSF to investigate the physical properties and scaling behavior of multi-qubit gates as the system size increases. In particular, we derive an estimate for the solution norm and, consequently, for the required drive amplitude Ω_{nuc} , which is shown to depend on the nuclear norm of the target entanglement map $\|\varphi_{nn'}\|_{\text{nuc}}$. Furthermore, we identify a minimal gate time $T_{\text{min}}(N)$ that scales linearly with the number of ions in a crystal, below which no high-fidelity multiqubit gate solutions can be found within our framework.

Exploiting this method, we conduct a noise-induced gate-error study focusing on three dominant error mechanisms: motional-mode frequency drifts, drive-amplitude jitter, and motional heating. We derive data-driven models for the expected gate error and analyze how these errors scale with increasing chain lengths and participating ions. These models provide a practical link between empirically measured noise levels and the maximum system size that can support gates of acceptable fidelity. Other error sources, particularly those introducing terms that do not commute with the Lamb-Dicke Hamiltonian, require independent treatment via full Hamiltonian simulations. We note that for high-amplitude solutions produced by our method, interactions with the qubit carrier, and consequently higher-order terms in the Magnus expansion, cannot be neglected and must be verified or compensated to ensure valid gate solutions. Such dedicated noise analyses are also essential for assessing the suitability of multiqubit gates in the context of quantum error-correction protocols, especially with regard to noise propagation through large entangling operations and the resulting hierarchy of effective noise operators.

Finally, we demonstrate that the proposed optimization routine improves the time-to-solution by approximately an order of magnitude compared to the best available alternatives, rendering the offline compilation of circuits involving tens of ions and many multiqubit gates a realistic task. While further speedups are possible through implementation in compiled languages or GPU-accelerated routines, such optimizations lie beyond the scope of this work. The optimizer itself is robust and readily generalizes to incorporate additional linear and quadratic constraints. In particular, these extensions can be used to relax the current treatment of linear constraints (via null-space transformations) by instead allowing their controlled minimization to an acceptable tolerance, as well as to include additional quadratic error terms in the objective function alongside the quadratic phase constraints [26,62]. Moreover, the framework extends to scenarios with degenerate or overlapping drives, where control fields

are applied to multiple ions simultaneously, whether by design or due to crosstalk.

ACKNOWLEDGMENTS

This work was supported by the Israel Science Foundation and the Israel Science Foundation Quantum Science and Technology (Grants No. 2074/19, No. 1376/19 and No. 3457/21).

APPENDIX A: MODEL DERIVATION

We detail the derivation of the model used to formulate the optimization problem in Eq. (6). In our analysis below we rely on and generalize the derivations provided in [7].

We consider a general spectrum of tones which drive independently N trapped ions. Without loss of generality, the driving applied to different ions may be regarded as composed of the same tones, differing only by the tone amplitudes. These tones are placed symmetrically around the single qubit transition frequency, ω_0 . The field driving the n th ion is

$$w_n(t) = \cos(\omega_0 t + \phi_0) \cdot \sum_{m=1}^M \left[r_{nm}^{(c)} \cos(\omega_m t) + r_{nm}^{(s)} \sin(\omega_m t) \right]. \quad (\text{A1})$$

That is, each spectrum contains $2M$ components at frequencies $\{\omega_0 \pm \omega_m\}_{m=1}^M$. The amplitude of the cosine (sine) m th tone pair illuminating the n th ions is $r_{nm}^{(c)}$ ($r_{nm}^{(s)}$), such that all $N \times M$ tone pairs have the same average phase, which generates a correlated rotation around the Pauli $\sin(\phi_0)\sigma_x + \cos(\phi_0)\sigma_y$ axis. For simplicity we assume that $\phi_0 = \pi/2$ such that the relevant Pauli operator is σ_x . The motional mode phase space trajectories generated by the cosine and sine quadrature are relatively rotated by $\pi/2$.

All in all this is captured by the lab-frame Hamiltonian ($\hbar = 1$),

$$\begin{aligned} H &= H_0 + V \\ H_0 &= \sum_{j=1}^N \left[\nu_j \left(a_j^\dagger a_j + \frac{1}{2} \right) + \frac{\omega_0}{2} \sigma_z^{(j)} \right] \\ V &= 2 \Omega \sum_{n=1}^N \sigma_x^{(n)} \cos(kx_n - \omega_0 t) \\ &\quad \cdot \sum_{m=1}^M \left[(r_n^{(c)})_m \cos(\omega_m t) + (r_n^{(s)})_m \sin(\omega_m t) \right], \quad (\text{A2}) \end{aligned}$$

with a_j the annihilation operator associated with j th normal mode of motion, at frequency ν_j , $\sigma_i^{(n)}$ the i -Pauli operator acting on the n th ion, k the driving field's wave

number, x_n the position operator of the n th ion and Ω a characteristic Rabi frequency. Clearly the last parenthesis in Eq. (A2) can be written as a single cosine term with a phase, however this form preserves a crucial aspect of our formulations, i.e. the exclusive linear and quadratic dependence on the amplitudes. Furthermore, here we have assumed that the drive couples to motional modes along one principal direction of the trap, such that the summation on modes is up to N (and not $3N$), this assumption can be easily relaxed [33].

By using a conventional set of approximations, namely the rotating wave approximation in Ω/ω_0 , the Lamb-Dicke approximation, and by neglecting carrier-coupling terms the lab-frame Hamiltonian is converted to the interaction Hamiltonian,

$$H_I = \Omega \sum_{n=1}^N \sigma_x^{(n)} \sum_{j=1}^N \eta_j O_j^{(n)} \left(a_j^\dagger e^{i\nu_j t} + \text{H.c.} \right) \cdot \sum_{m=1}^M \left(r_{nm}^{(c)} \cos(\omega_m t) + r_{nm}^{(s)} \sin(\omega_m t) \right), \quad (\text{A3})$$

with $O_j^{(n)}$ the normalized participation of the n th ion in the j th mode of motion, such that $\sum_{j=1}^N O_n^{(j)} O_m^{(j)} = \delta_{n,m}$, and η_j the single-ion Lamb-Dicke parameter associated with the j th mode of motion (it is sometimes conventional to define $\eta_j^{(n)} = \eta_j O_j^{(n)}$).

The Hamiltonian in Eq. (A3) can be rearranged in the form, $H_I = \sum_{n=1}^N \sum_{j=1}^N \zeta_j^{(n)}(t) \sigma_x^{(n)} a_j + \text{H.c.}$, with $\zeta_j(t)$ a time dependent function read-off directly from Eq. (A3). That is, it has only σ_x spin operators and is linear in the mode raising and lowering operators, it is therefore analytically solvable. Specifically its Magnus expansion vanishes after the second order. The resulting unitary evolution operator is the well-known combination of spin-dependent displacement of the motional modes and spin-exclusive correlated rotation,

$$U = \prod_{j=1}^N D_j \left(\sum_{n=1}^N \sigma_x^{(n)} \alpha_j^{(n)} \right) \prod_{n,m=1}^N e^{i\varphi_{n,m} \sigma_x^{(n)} \sigma_x^{(m)}}, \quad (\text{A4})$$

with $D_j(\alpha) = \exp(\alpha a_j^\dagger - \alpha^* a_j)$ a displacement operator which translates the j th mode by $\alpha_j^{(n)}$, with,

$$\alpha_j^{(n)} = -i\eta_j O_j^{(n)} \Omega \cdot \sum_{m=1}^M \int_0^t dt' e^{i\nu_j t'} \left[r_{nm}^{(c)} \cos(\omega_m t') - r_{nm}^{(s)} \sin(\omega_m t') \right], \quad (\text{A5})$$

and entanglement phases,

$$\begin{aligned} \varphi_{n,m} &= \mathbf{r}_n^T A_{n,m} \mathbf{r}_m \\ A_{n,m} &= -\Omega^2 \sum_{j=1}^N \eta_j^2 O_j^{(n)} O_j^{(m)} \begin{pmatrix} A_j^{\cos,\cos} & A_j^{\cos,\sin} \\ A_j^{\sin,\cos} & A_j^{\sin,\sin} \end{pmatrix} \\ (A_j^{f,g})_{m,l} &= -\int_0^t dt_1 \int_0^{t_1} dt_2 \sin(\nu_j[t_1 - t_2]) \\ &\quad \cdot [f(\omega_m t_1)g(\omega_l t_2) + f(\omega_m t_2)g(\omega_l t_1)], \end{aligned} \quad (\text{A6})$$

with $\mathbf{r}_n = (r_n^{(c)}, r_n^{(s)})$.

Using the definitions in Eqs. (A5) and (A6) we observe that mode displacement is linear in the field amplitudes while the two-qubit rotation phase is quadratic in them. While the linear constraints resulting from the former are discussed in Appendix B, the quadratic constraints resulting from the latter lead to the optimization problem in Eq. (6).

APPENDIX B: GATE HARMONICS AND LINEAR CONSTRAINTS

We discuss convenient choices for the tones $\{\omega_m\}_{m=1}^M$ and show how to write the degrees of freedom in a form that by-construction satisfies the linear constraints.

Driving the entanglement operation with a multi-tone representation of the field carries a lot of physical intuition. Specifically, it is beneficial to choose the tones ω_m in the vicinity of the mode frequency band $\{\nu_j\}_{j=1}^N$, since the coupling between tones and modes scales inversely with their frequency difference. Thus in our demonstrations above we choose ω_1 slightly below the smallest tone frequency and ω_M (assuming they are ordered) slightly above the largest frequency, defined precisely below.

We expect the field amplitude to vanish before $t = 0$ and after $t = T_g$, therefore a convenient choice of tones is the harmonic basis, i.e. $\omega_m = \frac{2\pi}{T} h_m$, with $h_m \in \mathbb{N}$ the tone number. We note that the harmonic choice simplifies the evaluation of the integrals in Eqs. (A5) and (A6).

In the main text we state that the optimization problem needs to satisfy both linear and quadratic constraints, however the former may be solved by-construction. Indeed, in order to ensure that at the gate time, $t = T_g$, the state of the motional mode is the same as in $t = 0$ we therefore require that the displacement operators in Eq. (A4) are unit operators. This ensures that an initial state in which the qubit and motion degrees of freedom are decoupled, remain decoupled after the operation. This is satisfied by requiring that,

$$\alpha_j^{(n)}(T) = 0 \quad \forall j, n = 1, \dots, N. \quad (\text{B1})$$

While Eq. (B1) naively contains N^2 constraints, they can all be solved by restricting $\mathbf{r}_n \in \ker(L)$ for all $n = 1, \dots, N$, with,

$$L_{j,m} = \begin{cases} \int_0^T dt' \cos(\nu_j t') \cos(\omega_m t') & m = 1, \dots, M \\ \int_0^T dt' \sin(\nu_j t') \sin(\omega_m t') & m = M+1, \dots, 2M \end{cases} \quad (\text{B2})$$

and $j = 1, \dots, N$. This restriction removes the linear requirement from the optimization problem in Eq. (5). Furthermore, the matrices $A_{n,m}$ and A_j can be easily transformed to the kernel space of L such that their dimension is reduced and their evaluation is faster. The kernel space can be either found exactly or approximately under some infidelity tolerance [30].

The entanglement operation can be endowed with additional properties that ensure its robustness to various sources of errors and noise such as unwanted coupling to the carrier, and other transitions, pulse timing errors, phonon-mode heating, phonon frequency drifts etc. These may be added as additional rows of the matrix L [7].

Out of these properties, decoupling of the carrier transition is useful in order to improve the approximation used to derive Eq. (A3). A convenient way of doing so is by letting the driving fields rise and fall continuously at $t = 0$ and $t = T_g$, respectively, with the choice $\mathbf{r}_n^{(c)} = 0$ for all $n = 1, \dots, N$, which further simplifies all of the expressions above. Here we set $\mathbf{r}_n \mapsto \mathbf{r}_n^{(c)} = 0$ and for simplicity identify $\mathbf{r}_n$ with $\mathbf{r}_n^{(s)}$ and A_j with $A_j^{\text{sin,sin}}$. We remark that with this choice the drive profile is not time-symmetric, while there are known advantages for using a time-symmetric drive [32].

Through the study we apply commonly two additional linear constraints, improving the pulse robustness to timing error ($T_g \rightarrow T_g + \delta T$) and normal mode drift $\nu_j \rightarrow \nu_j + \delta \nu_j$, which we name *Cardioid* and *Nuoid* constraints, respectively, following Ref. [40]. These are derived by enforcing additional constraints on $\alpha_j^{(n)}(T)$.

Cardioid constraints are obtained by requiring $\frac{d\alpha_j(T_g)}{d\delta T} = 0|_{\delta T=0}$, $\forall j$, leading to $2N$ additional linear constraints for the real and imaginary part of each mode:

$$\begin{aligned} \text{Re}[L_{j,m}^{(\text{car})}] &= g(\omega_m T) \cos(\nu_j T) \\ \text{Im}[L_{j,m}^{(\text{car})}] &= g(\omega_m T) \sin(\nu_j T) \\ g(\cdot) &= \begin{cases} \cos(\cdot), & m = 1, \dots, M \\ \sin(\cdot), & m = M+1, \dots, 2M \end{cases} \end{aligned} \quad (\text{B3})$$

Normal mode robustness is obtained by a complementary requirement $\int_0^T \alpha_j(\tau) d\tau = 0$. This results in the following matrix element:

$$\begin{aligned} \text{Re}[L_{j,m}^{(\text{nu})}] &= \frac{\omega_m T}{\omega_m^2 - \nu_j^2} - \frac{(\omega_m^2 + \nu_j^2) \cos(\nu_j T) f(\omega_m T)}{(\omega_m^2 - \nu_j^2)^2} \\ &\quad + \frac{2\nu_j \omega_m \sin(\nu_j T) g(\omega_m T)}{(\omega_m^2 - \nu_j^2)^2} \\ \text{Im}[L_{j,m}^{(\text{nu})}] &= \frac{\omega_m T}{\omega_m^2 - \nu_j^2} + \frac{2\nu_j \omega_m \cos(\nu_j T) g(\omega_m T)}{(\omega_m^2 - \nu_j^2)^2} \\ &\quad + \frac{(\omega_m^2 + \nu_j^2) \sin(\nu_j T) f(\omega_m T) - 2\nu_j \omega_m}{(\omega_m^2 - \nu_j^2)^2} \\ f(\cdot), g(\cdot) &= \begin{cases} \sin(\cdot), \cos(\cdot), & m = 1, \dots, M \\ \sin(\cdot), \sin(\cdot), & m = M+1, \dots, 2M \end{cases} \end{aligned} \quad (\text{B4})$$

APPENDIX C: TRIVIAL SOLUTION IN THE ADIABATIC LIMIT

In the slow gate limit, $\Delta \bar{\nu} T_g \gg 1$ satisfying the quadratic constraints is trivial and any bipartite qubit-qubit coupling can be achieved. Here we prove this directly by constructing such a solution. In this limit it is helpful to choose spectral tones containing frequencies of the form $\omega_s \mapsto \omega_{j,s} = \nu_j + \frac{2\pi}{T} s$ with $s \in \mathbb{N}$. We have doubled the index s to (j, s) for convenience. Since T_g is large we can safely assume that the tone $\omega_{j,s}$ couples exclusively to mode j and satisfies the linear constraints by construction. Furthermore, with this choice the matrices A_j become diagonal, i.e. each of the tones coupled to mode j contributes to the entanglement phase independently, and scales as s^{-1} . In other words $(A_j)_{(j',s'),(j'',s'')} \propto \delta_{j,j'} \delta_{j,j''} \delta_{s',s''} (s')^{-1}$ [40].

We are required to generate $N(N-1)/2$ bipartite entanglement phases $\varphi_{n,m}$, with $1 \leq n < m \leq N$. We do so by designating a unique tone to each phase. Since these tones do not interfere (whether they are driving distinct or the same mode of motion) then we simply need to scale the drive amplitudes accordingly. Specifically, to satisfy the constraint on $\varphi_{n,m}$ we choose an arbitrary mode which couples to both ions. Without loss of generality this can always be the center of mass mode, j_{COM} . We drive both ions with the same tone $\omega_{j,s}$ setting $j = j_{\text{COM}}$ and $s = N \cdot n + m$, and set its amplitude driving both ions to be $\sqrt{\frac{(N \cdot n + m) \varphi_{n,m}}{O_{j_{\text{COM}}}^{(n)} O_{j_{\text{COM}}}^{(m)}}}$.

APPENDIX D: EFFECTS OF ROBUSTNESS ON GATE POWER

In the main text we show that the total gate power, given by the nuclear norm Ω_{nuc} depends linearly on the nuclear norm of the target coupling matrix (and other parameters), see Eq. (7), with k the linear proportionality constant. Here we provide a heuristic argument for the numerical

dependence of k on the robustness properties of the entanglement operation, highlighted in the main text.

The framework for analytic analysis is multi-tone adiabatic gate that couple to single modes of motion, as was presented in Ref. [40]. In that work drive tones pairs are symmetrically positioned around the red and blue sideband of the motional mode of choice, i.e. with frequencies, $\omega_{n,\pm} = \omega_0 \pm (\nu + n\xi)$, with ω_0 the qubit transition frequency, ν the frequency of the mode of motion of choice, ξ an additional detuning such that the gate time is given by $2\pi\xi^{-1}$ an $n \in \mathbb{Z}$. With this formulation, the gate is designed by constraints on the amplitudes driving each tone pair.

Specifically, gates that are robust to timing errors (among other sources of error) are generated by satisfying, $\sum_n r_n^2/n = 1$ and $\sum_n r_n = 0$ with the former ensuring a desired entangling phase and the latter ensuring the robustness property (coined *Cardioid* in Ref. [40]). The total power, $\sum_n r_n^2$ is to be minimized subject to these constraints. We do so in a Lagrange multiplier framework, by minimizing the Lagrangian,

$$L_{\text{Car}} = \sum_n r_n^2 - \lambda \left(\sum_n \frac{r_n^2}{n} - 1 \right) - \mu \left(\sum_n r_n \right). \quad (\text{D1})$$

One then solves the equations $\partial_{r_n} L_{\text{Car}} = 0$ under these constraints. At the limit of the number of tones approaching infinity, the solution to these equations converges to a power $\sum_n r_n^2 = 1$, which is the same result obtained in absence of the robustness constraint. That is, presumably this constraint does not incur an overhead of gate drive power. However, one notices that this is true only in the unphysical limit of infinite driving tones and that this limiting solution in fact recreates the standard non-robust version of the MS gate, shown in Fig. 10. A similar example is presented for gates that are robust to calibration errors of the motional mode frequency, given by the




FIG. 10. Left: phase space trajectory of two-qubit robust gates in the limiting case of infinite tones. In this limit the Cardioid gate (blue) recovers the circular trajectory of the canonical MS gate, which explains the vanishing cost of its robustness, while the Nuoid gate (orange) has a distinct shape. Right: drive power of the Cardioid (blue) and Nuoid (orange) gates as a function of the number of tones. Indeed the Cardioid gate power converges to 1, while the Nuoid gate power converges to 1.25, i.e. a finite additional cost of the robustness.

constraints, $\sum_n r_n^2/n = 1$ and $\sum_n r_n/n = 0$ (coined *Nuoid* in Ref. [40]). Similarly, a Lagrangian is formed,

$$L_{\text{Nu}} = \sum_n r_n^2 - \lambda \left(\sum_n \frac{r_n^2}{n} - 1 \right) - \mu \left(\sum_n \frac{r_n}{n} \right), \quad (\text{D2})$$

and solution to $\partial_{r_n} L_{\text{Nu}} = 0$ are found under these constraints. Here we get, at an infinite number of tones, a drive power of $\sum_n r_n^2 \approx 1.25$, showing an actual overhead in terms of drive power in order to implement this robustness property, shown as well in Fig. 10.

APPENDIX E: VON-NEUMANN ENTROPY CALCULATIONS

There are many measures that can be used to quantify the entanglement generated by U_{MQ} . Here, for simplicity, we choose the von Neumann single-qubit entropy, i.e. we compute $\bar{S}_{\text{VN}}(\varphi) = \sum_n S_{\text{VN},n} = -\sum_n \rho_n \log_2(\rho_n)$, where ρ_n is a density matrix formed by tracing out the n th qubit. We perform this computation over the state $U_{\text{MQ}}|+\rangle^{\otimes N}$, where $|+\rangle$ is the $+1$ eigenvalue of X_n .

Considering ρ_n , we have,

$$\begin{aligned} \rho_n = & \frac{1}{2} \exp \left(i \sum_{j,k \neq n} \varphi_{j,k} Z_j Z_k \right) \\ & \cdot \left(e^{i \sum_{k \neq n} Z_k} |+_n\rangle \langle +_n| e^{-i \sum_{k \neq n} Z_k} \right. \\ & \left. + e^{-i \sum_{k \neq n} Z_k} |+_n\rangle \langle +_n| e^{+i \sum_{k \neq n} Z_k} \right) \\ & \times \exp \left(-i \sum_{j,k \neq n} \varphi_{j,k} Z_j Z_k \right), \end{aligned} \quad (\text{E1})$$

where $|+_n\rangle$ is the original product state with the n th qubit absent, and we decomposed the state of the n th qubit into $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, and then replaced the values of Z_n with ± 1 accordingly. The resulting density matrix is an equal weight sum of two non-orthogonal states, thus its nonzero eigenvalues are given by $\lambda_{\pm} = (1 \pm |S|)/2$ with S the overlap between the states in absolute value. It is trivial to show that here $S = |\prod_{k \neq n} \cos^2(2\varphi_{n,k})| \approx 1 - 2 \sum_{k \neq n} \varphi_{n,k}^2$, where the approximation is a leading order expansion in small values of φ .

Thus the two eigenvalues are $\sum_{k \neq n} \varphi_{n,k}^2$ and $1 - \sum_{k \neq n} \varphi_{n,k}^2$.

We get,

$$S_{\text{VN},n} = -\rho_n \log_2 \rho_n \approx \sum_{k \neq n} \varphi_{n,k}^2, \quad (\text{E2})$$

where again the approximation is in leading order in φ . Thus we obtain, $\bar{S}_{\text{VN}}(\varphi) = \sum_{n,k} \varphi_{n,k}^2$, which is equal to the Frobenius norm squared of the matrix φ .

APPENDIX F: CHOICE OF EFFECTIVE DEGREES OF FREEDOM

While fast gates are generally desirable, there are scenarios where only slower gates are feasible, particularly when drive power is limited or error-budget constrained. With the nuclear norm of an entangling map scaling with the number of participating ions, the minimum feasible power-limited gate time scales as N^2 : one factor of N from $T_{\min}$, and the other from limited optical power, as $|r| \sim N/T_g$. As the number of Fourier components is linear in T_g , a naive gate construction would require $M = \mathcal{O}(N^2)$ frequency components, even though only $\mathcal{O}(N)$ are likely required to control the phase of N modes. This translates to a steep increase in both hardware complexity and LSF memory usage, as each ion-pair constraint matrix has size M^2 . A cost-efficient heuristic for eliminating less effective degrees-of-freedom *a priori* is therefore required.

DOF reduction can be achieved either by zeroing selected drive amplitudes $\Omega_m^{(n)}$ in advance, or by introducing a dimensionality-reducing linear map $L \in \mathcal{R}^{M' \times M}$ with $M' < M$ such that $r_n = LK\Omega^{(n)}$. The first approach compresses the drive signal representation; the second reduces the size of the quadratic matrices while preserving the linear constraint space defined by matrix K . Particularly, we considered three heuristics:

1. Keeping exactly $2k$ evenly-distributed tones between each two modes (even dilution).
2. Keeping the $2k$ nearest tones around each mode (mode-centric dilution).
3. Constructing L with $2kN$ rows based on the eigenvectors corresponding to the k largest positive and k largest negative eigenvalues of each mode-phase matrix KA_jK^T [see Appendix A, Eq. (A6)] (eigenvalues dilution).

To evaluate the effectiveness of DOF reduction, we applied the proposed methods to a 25-ion gate design problem and compared the resulting solutions to a reference solution using the full set of 805 tones, corresponding to a gate time of $T_g = 6\Delta\nu_{\perp}^{-1} \approx 32\Delta\nu^{-1}$. This gate time was chosen to allow for additional robustness constraints, resulting in 703 optimization DOFs within the kernel of the linear constraint matrix—approximately 28 DOFs per mode. Figure 11 shows the resulting drive amplitudes, normalized by the amplitude of the reference (non-diluted) solution Ω_{ref} , for the three dilution strategies: evenly spaced tones (red), mode-centric tones (blue), and the eigenvalue-based method (green), as a function of the dilution parameter $k = \frac{\text{\#DOFs}}{N}$. In the even and mode-centric strategies, tone selection is applied prior to imposing linear constraints, resulting in fewer actual optimization DOFs than the nominal k value suggests. Among the three, the eigenvalue method most rapidly recovers the reference amplitude Ω_{ref} , using the two largest positive and two largest negative eigenvectors per mode. This highlights




FIG. 11. Comparison of obtained solution amplitude Ω for drives where the number of optimization variables was reduced from $M \sim N^2$ to $2kN$, along different methods—even dilution (red), mode-centric dilution (blue) and mode-phase eigenvalues (green)—at $k = 2$ optimization DOFs, the eigenvalues method is the first to produce a solution of the same Ω as a system without DOF reduction.

the benefit of selecting the most effective DOFs for independently steering each mode's phase contribution to $\varphi_{n,m}$. The mode-centric method performs slightly worse, consistent with the fact that the eigenvalue-derived DOFs are concentrated near their respective modes in frequency space. The even dilution method performs significantly worse, demonstrating that enforcing a uniform distribution of tones across the spectrum leads to suboptimal amplitude solutions. Based on these results, we adopt the eigenvalue method as the default tone dilution strategy for large-scale gate optimization problems.

-
- [1] D. Kielpinski, C. Monroe, and D. J. Wineland, *Architecture for a large-scale ion-trap quantum computer*, *Nature (London)* **417**, 709 (2002).
 - [2] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. S. Allman, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, C. Ryan-Anderson, and B. Neyenhuis, *Demonstration of the trapped-ion quantum ccd computer architecture*, *Nature (London)* **592**, 209 (2021).
 - [3] S. A. Moses *et al.*, *A race track trapped-ion quantum processor*, *Phys. Rev. X* **13**, 041052 (2023).
 - [4] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe, *Entanglement of single-atom quantum bits at a distance*, *Nature (London)* **449**, 68 (2007).
 - [5] D. P. Nadlinger, P. Drmota, B. C. Nichol, G. Araneda, D. Main, R. Srinivas, D. M. Lucas, C. J. Ballance, K. Ivanov, E. Y.-Z. Tan, P. Sekatski, R. L. Urbanke, R. Renner, N. Sangouard, and J.-D. Bancal, *Experimental quantum key*

- distribution certified by Bell's theorem, *Nature (London)* **607**, 682 (2022).
- [6] V. Krutyanskiy, M. Galli, V. Krcmarsky, S. Baier, D. A. Fioretto, Y. Pu, A. Mazloom, P. Sekatski, M. Canteri, M. Teller, J. Schupp, J. Bate, M. Meraner, N. Sangouard, B. P. Lanyon, and T. E. Northup, *Entanglement of trapped-ion qubits separated by 230 meters*, *Phys. Rev. Lett.* **130**, 050803 (2023).
- [7] Y. Shapira, R. Shaniv, T. Manovitz, N. Akerman, L. Peleg, L. Gazit, R. Ozeri, and A. Stern, *Theory of robust multi-qubit nonadiabatic gates for trapped ions*, *Phys. Rev. A* **101**, 032330 (2020).
- [8] N. Grzesiak, R. Blümel, K. Wright, K. M. Beck, N. C. Pienti, M. Li, V. Chaplin, J. M. Amini, S. Debnath, J.-S. Chen, and Y. Nam, *Efficient arbitrary simultaneously entangling gates on a trapped-ion quantum computer*, *Nat. Commun.* **11**, 2963 (2020).
- [9] D. Schwerdt, Y. Shapira, T. Manovitz, and R. Ozeri, *Comparing two-qubit and multiqubit gates within the toric code*, *Phys. Rev. A* **105**, 022612 (2022).
- [10] S. Bravyi, D. Maslov, and Y. Nam, *Constant-cost implementations of Clifford operations and multiply-controlled gates using global interactions*, *Phys. Rev. Lett.* **129**, 230501 (2022).
- [11] P. Baßler, M. Zipper, C. Cedzich, M. Heinrich, P. H. Huber, M. Johanning, and M. Kliesch, *Synthesis of and compilation with time-optimal multi-qubit gates*, *Quantum* **7**, 984 (2023).
- [12] R. Yao, W. Q. Lian, Y. K. Wu, G. X. Wang, B. W. Li, Q. X. Mei, B. X. Qi, L. Yao, Z. C. Zhou, L. He, and L. M. Duan, *Experimental realization of a multiqubit quantum memory in a 218-ion chain*, *Phys. Rev. A* **106**, 062617 (2022).
- [13] I. Pogorelov, T. Feldker, C. D. Marciniak, L. Postler, G. Jacob, O. Kriegelsteiner, V. Podlesnic, M. Meth, V. Negnevitsky, M. Stadler, B. Höfer, C. Wächter, K. Lakhmanskiy, R. Blatt, P. Schindler, and T. Monz, *Compact ion-trap quantum computing demonstrator*, *PRX Quantum* **2**, 020343 (2021).
- [14] L. Feng, O. Katz, C. Haack, M. Maghrebi, A. V. Gorshkov, Z. Gong, M. Cetina, and C. Monroe, *Continuous symmetry breaking in a trapped-ion spin chain*, *Nature (London)* **623**, 713 (2023).
- [15] D. Schwerdt, L. Peleg, Y. Shapira, N. Priel, Y. Florshaim, A. Gross, A. Zalic, G. Afek, N. Akerman, A. Stern, A. B. Kish, and R. Ozeri, *Scalable architecture for trapped-ion quantum computing using rf traps and dynamic optical potentials*, *Phys. Rev. X* **14**, 041017 (2024).
- [16] S.-A. Guo, Y.-K. Wu, J. Ye, L. Zhang, W.-Q. Lian, R. Yao, Y. Wang, R.-Y. Yan, Y.-J. Yi, Y.-L. Xu *et al.*, *A site-resolved two-dimensional quantum simulator with hundreds of trapped ions*, *Nature (London)* **630**, 613 (2024).
- [17] L. Blum, M. Shub, and S. Smale, *On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines*, *Bull. Amer. Math. Soc. (N.S.)* **21**, 1 (1989).
- [18] J. Nemirovsky, M. Chuchem, and Y. Shapira, *Efficient compilation of quantum circuits using multi-qubit gates*, *Quantum Sci. Technol.* **11**, 01LT02 (2026).
- [19] F. Liu, G. Tang, L. Duan, and Y. Wu, *Performance analysis for crosstalk errors between parallel entangling gates in trapped ion quantum error correction*, *Phys. Rev. Appl.* **24**, 014032 (2025).
- [20] P. Høyer and R. Špalek, *Quantum circuits with unbounded fan-out*, *STACS 2003*, edited by H. Alt and M. Habib (Springer, Berlin, Heidelberg, 2003), pp. 234–246.
- [21] B. Foxman, N. Parham, F. Vasconcelos, and H. Yuen, *Random unitaries in constant (quantum) time*, *arXiv:2508.11487*.
- [22] A. Galicia, B. Ramon, E. Solano, and M. Sanz, *Enhanced connectivity of quantum hardware with digital-analog control*, *Phys. Rev. Res.* **2**, 033103 (2020).
- [23] J. Nemirovsky, M. Chuchem, L. Peleg, Y. Solomons, A. B. Kish, and Y. Shapira, *Phase gadget compilation of quantum circuits using multiqubit gates*, *arXiv:2510.16788*.
- [24] J. Nemirovsky, L. Peleg, A. B. Kish, and Y. Shapira, *Optimal constant-cost implementations of Clifford operations using global interactions*, *arXiv:2510.20730*.
- [25] M. Duwe, G. Zarantonello, N. Pulido-Mateo, H. Mendpara, L. Krinner, A. Bautista-Salvador, N. V. Vitanov, K. Hammerer, R. F. Werner, and C. Ospelkaus, *Numerical optimization of amplitude-modulated pulses in microwave-driven entanglement generation*, *Quantum Sci. Technol.* **7**, 045005 (2022).
- [26] M. Kang, Y. Wang, C. Fang, B. Zhang, O. Khosravani, J. Kim, and K. R. Brown, *Designing filter functions of frequency-modulated pulses for high-fidelity two-qubit gates in ion chains*, *Phys. Rev. Appl.* **19**, 014014 (2023).
- [27] A. R. Milne, C. L. Edmunds, C. Hempel, F. Roy, S. Mavadia, and M. J. Biercuk, *Phase-modulated entangling gates robust to static and time-varying errors*, *Phys. Rev. Appl.* **13**, 024022 (2020).
- [28] Y. Lu, S. Zhang, K. Zhang, W. Chen, Y. Shen, J. Zhang, J.-N. Zhang, and K. Kim, *Global entangling gates on arbitrary ion qubits*, *Nature (London)* **572**, 363 (2019).
- [29] T. Choi, S. Debnath, T. A. Manning, C. Figgatt, Z.-X. Gong, L.-M. Duan, and C. Monroe, *Optimal quantum control of multimode couplings between trapped ion qubits for scalable entanglement*, *Phys. Rev. Lett.* **112**, 190502 (2014).
- [30] R. Blümel, N. Grzesiak, N. H. Nguyen, A. M. Green, M. Li, A. Maksymov, N. M. Linke, and Y. Nam, *Efficient stabilized two-qubit gates on a trapped-ion quantum computer*, *Phys. Rev. Lett.* **126**, 220503 (2021).
- [31] R. Blümel, N. Grzesiak, N. Pienti, K. Wright, and Y. Nam, *Power-optimal, stabilized entangling gate between trapped-ion qubits*, *npj Quantum Inf.* **7**, 147 (2021).
- [32] P. H. Leung, K. A. Landsman, C. Figgatt, N. M. Linke, C. Monroe, and K. R. Brown, *Robust 2-qubit gates in a linear ion crystal using a frequency-modulated driving force*, *Phys. Rev. Lett.* **120**, 020501 (2018).
- [33] Y. Zhu, A. M. Green, N. H. Nguyen, C. Huerta Alderete, E. Mossman, and N. M. Linke, *Pairwise-parallel entangling gates on orthogonal modes in a trapped-ion chain*, *Adv. Quantum Technol.* **6**, 2300056 (2023).
- [34] C. Figgatt, A. Ostrander, N. M. Linke, K. A. Landsman, D. Zhu, D. Maslov, and C. Monroe, *Parallel entangling operations on a universal ion-trap quantum computer*, *Nature (London)* **572**, 368 (2019).
- [35] R. L. Kosut, M. D. Grace, and C. Brif, *Robust control of quantum gates via sequential convex programming*, *Phys. Rev. A* **88**, 052326 (2013).

- [36] M. Kang, Q. Liang, B. Zhang, S. Huang, Y. Wang, C. Fang, J. Kim, and K. R. Brown, *Batch optimization of frequency-modulated pulses for robust two-qubit gates in ion chains*, *Phys. Rev. Appl.* **16**, 024039 (2021).
- [37] E. P. G. Gale, Z. Mehdi, L. M. Oberg, A. K. Ratcliffe, S. A. Haine, and J. J. Hope, *Optimized fast gates for quantum computing with trapped ions*, *Phys. Rev. A* **101**, 052328 (2020).
- [38] A. Sørensen and K. Mølmer, *Quantum computation with ions in thermal motion*, *Phys. Rev. Lett.* **82**, 1971 (1999).
- [39] A. Sørensen and K. Mølmer, *Entanglement and quantum computation with ions in thermal motion*, *Phys. Rev. A* **62**, 022311 (2000).
- [40] Y. Shapira, R. Shaniv, T. Manovitz, N. Akerman, and R. Ozeri, *Robust entanglement gates for trapped-ion qubits*, *Phys. Rev. Lett.* **121**, 180502 (2018).
- [41] A. E. Webb, S. C. Webster, S. Collingbourne, D. Breaud, A. M. Lawrence, S. Weidt, F. Mintert, and W. K. Hensinger, *Resilient entangling gates for trapped ions*, *Phys. Rev. Lett.* **121**, 180501 (2018).
- [42] K. Wang, J.-F. Yu, P. Wang, C. Luan, J.-N. Zhang, and K. Kim, *Fast multi-qubit global-entangling gates without individual addressing of trapped ions*, *Quantum Sci. Technol.* **7**, 044005 (2022).
- [43] R. Blümel, A. Maksymov, and M. Li, *Toward a Mølmer-Sørensen gate with 9999 fidelity*, *J. Phys. B* **57**, 205501 (2024).
- [44] Y. Shapira, S. Cohen, N. Akerman, A. Stern, and R. Ozeri, *Robust two-qubit gates for trapped ions using spin-dependent squeezing*, *Phys. Rev. Lett.* **130**, 030602 (2023).
- [45] D. T. C. Allcock, W. C. Campbell, J. Chiaverini, I. L. Chuang, E. R. Hudson, I. D. Moore, A. Ransford, C. Roman, J. M. Sage, and D. J. Wineland, *omg blueprint for trapped ion quantum computing with metastable states*, *Appl. Phys. Lett.* **119**, 214002 (2021).
- [46] K. Wright *et al.*, *Benchmarking an 11-qubit quantum computer*, *Nat. Commun.* **10**, 5464 (2019).
- [47] T. Manovitz, Y. Shapira, L. Gazit, N. Akerman, and R. Ozeri, *Trapped-ion quantum computer with robust entangling gates and quantum coherent feedback*, *PRX Quantum* **3**, 010347 (2022).
- [48] M. R. Garey and D. S. Johnson, *Computers and intractability: A guide to the theory of NP-completeness*, *Series of Books in the Mathematical Sciences*, 1st ed. (W. H. Freeman, San Francisco, CA, 1979).
- [49] Y. Solomons, Y. Kadish, L. Peleg, J. Nemirovsky, A. B. Kish, and Y. Shapira, *Full programmable quantum computing with trapped-ions using semi-global fields*, *Quantum Sci. Technol.* **11**, 025036 (2026).
- [50] A. Kyprianidis, A. J. Rasmusson, and P. Richerme, *Interaction graph engineering in trapped-ion quantum simulators with global drives*, *New J. Phys.* **26**, 023033 (2024).
- [51] J. W. Britton, B. C. Sawyer, A. C. Keith, C.-C. J. Wang, J. K. Freericks, H. Uys, M. J. Biercuk, and J. J. Bollinger, *Engineered two-dimensional ising interactions in a trapped-ion quantum simulator with hundreds of spins*, *Nature (London)* **484**, 489 (2012).
- [52] D. Porras and J. I. Cirac, *Quantum manipulation of trapped ions in two dimensional Coulomb crystals*, *Phys. Rev. Lett.* **96**, 250501 (2006).
- [53] I. Savill-Brown, J. J. Hope, A. K. Ratcliffe, V. D. Vaidya, H. Liu, S. A. Haine, C. R. Viteri, and Z. Mehdi, *High-speed and high-connectivity two-qubit gates in long chains of trapped ions*, *Phys. Rev. Lett.* **136**, 190802 (2026).
- [54] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas, *High-fidelity quantum logic gates using trapped-ion hyperfine qubits*, *Phys. Rev. Lett.* **117**, 060504 (2016).
- [55] P. C. Lotshaw, K. D. Battles, B. Gard, G. Buchs, T. S. Humble, and C. D. Herold, *Modeling noise in global Mølmer-Sørensen interactions applied to quantum approximate optimization*, *Phys. Rev. A* **107**, 062406 (2023).
- [56] C. D. B. Bentley, H. Ball, M. J. Biercuk, A. R. R. Carvalho, M. R. Hush, and H. J. Slatyer, *Numeric optimization for configurable, parallel, error-robust entangling gates in large ion registers*, *Adv. Quantum Technol.* **3**, 2000044 (2020).
- [57] W. He, W. Zhang, X. Yuan, Y. Shen, and X.-M. Zhang, *Scaling of entangling-gate errors in large ion crystals*, *J. Phys. A* **57**, 375306 (2024).
- [58] M. Brownnutt, M. Kumph, P. Rabl, and R. Blatt, *Ion-trap measurements of electric-field noise near surfaces*, *Rev. Mod. Phys.* **87**, 1419 (2015).
- [59] D. Leibfried and D. J. Wineland, *Efficient eigenvalue determination for arbitrary Pauli products based on generalized spin-spin interactions*, *J. Mod. Opt.* **65**, 774 (2018).
- [60] K. Ireland and M. I. Rosen, *A Classical Introduction to Modern Number Theory* (Springer-Verlag, New York, NY, 1990), Vol. 84.
- [61] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright *et al.*, *Scipy 1.0: Fundamental algorithms for scientific computing in python*, *Nat. Methods* **17**, 261 (2020).
- [62] S. Kirchhoff, F. K. Wilhelm, and F. Motzoi, *Correction formulas for the mølmer-Sørensen gate under strong driving*, *PRX Quantum* **6**, 010328 (2025).