

Practical noise mitigation for quantum annealing via dynamical decoupling: Toward industry-relevant optimization using trapped ions

Sebastian Nagies ^{1,2,*} Chiara Capecci ^{1,2,†} Marcel Seelbach Benkner ^{3,4} Javed Akram ^{4,‡} Sebastian Rubbert ⁴
Dimitrios Bantounas ⁴ Michael Moeller ³ Michael Johanning ⁴ and Philipp Hauke ^{1,2,§}

¹*Pitaevskii BEC Center and Department of Physics, University of Trento, Via Sommarive 14, 38123 Trento, Italy*

²*INFN-TIFPA, Trento Institute for Fundamental Physics and Applications, Trento, Italy*

³*Department of Electrical Engineering and Computer Science, University of Siegen, Hölderlinstraße 3, 57076 Siegen, Germany*

⁴*eleQtron GmbH, Heeserstr. 5, 57072 Siegen, Germany*



(Received 24 October 2025; accepted 2 January 2026; published 26 January 2026)

Quantum annealing is a framework for solving combinatorial optimization problems. While it offers a promising path toward a practical application of quantum hardware, its performance in real-world devices is severely limited by environmental noise that can degrade solution quality. We investigate the suppression of local field noise in quantum annealing protocols through the periodic application of dynamical decoupling pulses implementing global spin flips. As test problems, we construct minimal multiple-object-tracking quadratic unconstrained binary optimization instances requiring only five and nine qubits, as well as cutting-stock instances of five and six qubits. Moreover, using the Sherrington-Kirkpatrick model, we demonstrate the robustness of our protocol to problem structure and size. To further place our results in a practical context, we consider a trapped-ion platform based on magnetic-gradient-induced coupling as a reference architecture, using it to define experimentally realistic noise and coupling parameters. We show that external magnetic field fluctuations, typical in such setups, significantly degrade annealing fidelity, while moderate dynamical decoupling pulse rates, which are achievable in current experiments, restore performance to near-ideal levels. Our analytical and numerical results reveal a universal scaling behavior, with fidelity determined by a generalized parameter combining noise amplitude and dynamical decoupling pulse interval. While our analysis is grounded in the trapped-ion platform, the proposed noise-mitigation strategy and resulting performance improvements are applicable to a broad range of quantum annealing implementations and establish a practical and scalable route for error mitigation in near-term devices.

DOI: [10.1103/mfqh-jtrl](https://doi.org/10.1103/mfqh-jtrl)

I. INTRODUCTION

Quantum computing has emerged as a transformative technology for addressing classically intractable computational problems [1–6], with quantum algorithms broadly categorized into gate-based and analog approaches. Within this landscape, quantum annealing occupies a distinctive position as an analog quantum algorithm especially suitable for combinatorial optimization problems [7–11]. Unlike gate-based algorithms that rely on discrete quantum operations, quantum annealing leverages a continuous time evolution to steer a system toward a quantum state encoding the solution to an optimization problem.

Hardware considerations play a central role in shaping quantum annealing implementations, as different platform architectures provide unique strengths and functionalities. Superconducting qubits [12–14], central to early annealing demonstrations [15], enable nanosecond-scale gate operations [16] with large numbers of physical qubits. Neutral-atom systems [17,18], leveraging optical tweezers and Rydberg interactions, provide highly scalable architectures with flexible, programmable coupling. Trapped-ion platforms [19–26] offer all-to-all qubit connectivity and extended coherence times. These characteristics could potentially make the latter well suited for the implementation of quantum annealing protocols which typically require long coherence windows and complex coupling topologies.

Within the trapped-ion paradigm, systems employing magnetic-gradient-induced coupling (MAGIC) represent a promising architecture for quantum information processing [27–37]. In MAGIC-based trapped-ion processors, a static magnetic field gradient mediates spin-spin interactions through collective motional modes of the ion crystal, enabling native two-body interactions that are both tunable and naturally suited to optimization problems. This approach offers additional practical advantages, including the use of microwave fields for qubit control—eliminating the need

*Contact author: sebastian.nagies@unitn.it

†Contact author: chiara.capecci@unitn.it

‡Contact author: research@eleqtron.com

§Contact author: philipp.hauke@unitn.it

for complex laser systems—and the ability to implement problem-specific Hamiltonians directly in hardware.

However, the practical realization of quantum annealing on any physical platform must contend with the inevitable presence of noise and systematic imperfections that can compromise algorithmic performance. Typically, among the most problematic sources of error are local field fluctuations, which in trapped-ion platforms arise from temporal and spatial fluctuations of external magnetic fields. These fluctuations can directly distort the cost Hamiltonian that encodes the optimization problem, potentially leading to suboptimal solutions.

To mitigate these detrimental effects while preserving the integrity of the encoded optimization problem, we investigate the application of dynamical decoupling strategies [38–47] specifically tailored to quantum annealing protocols [48,49]. We discuss how dynamical decoupling pulses implementing periodic qubit rotations can average out the impact of time-dependent local field fluctuations without disrupting the quantum annealing protocol. We demonstrate this noise suppression technique using minimal five-qubit and nine-qubit multiple-object-tracking quadratic unconstrained binary optimization (QUBO) instances [50], cutting-stock instances with five and six qubits, and instances of the Sherrington-Kirkpatrick model [51,52] up to 12 qubits. These problem sizes were chosen to be small enough for detailed numerical analysis and thus suitable for first proof-of-concept demonstrations. At the same time, their modest qubit counts ensure that they remain compatible with the limited qubit numbers available on current experimental platforms.

Through our numerical simulations employing a realistic noise spectrum and experimentally achievable coupling parameters, specifically referencing MAGIC-based trapped-ion systems, we establish that moderate dynamical decoupling pulse rates, achievable on current hardware, are sufficient to restore annealing fidelity to near-ideal levels across different noise frequencies and amplitudes at least up to an order of magnitude larger than the energy scale of the annealing Hamiltonian. Moreover, our analysis reveals a universal scaling behavior governed by a generalized parameter that combines noise amplitude with the dynamical decoupling pulse interval, providing a compact figure of merit for optimizing error mitigation strategies across diverse noise environments. Finally, by exploiting the discretization of time into steps given by consecutive dynamical decoupling pulses, the scheme can be adapted with little overhead to also permit one to program desired interaction parameters that encode the optimization problem. These findings establish a practical and scalable framework for implementing quantum annealing on near-term devices, with direct applicability extending beyond trapped-ion platforms to other quantum annealing architectures.

This paper is structured as follows: Section II provides a brief review of quantum annealing and introduces the specific annealing algorithm considered throughout this work. Section III presents multiple-object-tracking, an industrially relevant computer vision problem that can be formulated as a QUBO problem suitable for quantum annealing. We construct a minimal problem instance that serves as our test case for numerical simulations and discuss the scaling of required qubit number with problem size. The initial multiple-object-

tracking problem instance for this purpose is taken from the MOTChallenge publication from 2015 [53]. Section IV presents an overview of trapped-ion quantum computers with magnetic-gradient-induced coupling, which serves as our reference platform. We examine experimental parameters, including realizable coupling strengths on current hardware, relevant noise sources, and the feasibility of implementing quantum annealing on this platform. Section V provides a numerical investigation of how different noise sources affect quantum annealing protocol performance, emphasizing the critical need for local field noise mitigation. Section VI addresses active noise mitigation through dynamical decoupling pulses, providing both analytical justification for the underlying mechanism and numerical simulations of annealing sweeps with varying pulse counts. Section VII presents our conclusions. Technical details and complementary results for the cutting-stock problem and different system sizes are delegated to several appendixes.

II. QUANTUM ANNEALING

Quantum annealing is a common technique for solving optimization problems on quantum computers. It exploits quantum fluctuations to guide a system toward the ground state of a cost Hamiltonian H_{cost} , which encodes the solution to a combinatorial optimization problem [7–11]. In the standard framework of quantum annealing, the system is initialized in the ground state of a driving Hamiltonian H_{driving} , which introduces quantum fluctuations, whose strength is gradually reduced while simultaneously the strength of the cost Hamiltonian is ramped up. At the end of the annealing sweep, when the driving Hamiltonian has been switched off, the system ideally reaches the problem-encoding ground state. This procedure can be performed either adiabatically, where the system remains close to the instantaneous ground state over the entire evolution time, quantified by the adiabatic theorem [54–57], or diabatically [58–62], where higher excited states can be strategically populated during the evolution, to potentially decrease the total required annealing time.

Modified annealing protocol. In this work, we employ a slightly modified annealing protocol, where the strength of the cost Hamiltonian is kept constant. This form is potentially better suited to the trapped-ion devices based on MAGIC, which will be introduced in Sec. IV as our reference hardware platform. We define the total Hamiltonian during the annealing sweep as

$$H(t) = H_{\text{cost}} + C(t)H_{\text{driving}}, \quad (1)$$

$$H_{\text{cost}} = \sum_{ij} J_{ij} \sigma_i^z \sigma_j^z + \sum_i h_i^z \sigma_i^z, \quad (2)$$

$$H_{\text{driving}} = -h^x \sum_i \sigma_i^x, \quad (3)$$

where H_{cost} contains two-body interactions and local fields (corresponding to a QUBO problem), and H_{driving} is the standard choice for a driver with homogeneous transverse local fields. The coefficient is set to $C(0) = 1$ at the beginning of the protocol and then linearly ramped down to zero.

The system is initialized in the ground state of the driving Hamiltonian, which is the equal superposition of all

computational basis states. However, in contrast to the standard quantum annealing protocol, where $H_{\text{standard}}(t=0) = H_{\text{driving}}$, the equal superposition is not the true ground state of the total Hamiltonian in the modified protocol, as $H(0) = H_{\text{cost}} + H_{\text{driving}}$. To ensure that the system still has a high overlap with the true ground state at the beginning of the protocol, the energy scale of the driver has to be chosen sufficiently larger than the energy scale of the cost Hamiltonian to ensure an approximately adiabatic evolution. In contrast, for the conventional protocol, it is usually optimal to choose the energy scales of driving and cost Hamiltonian to be roughly equal [63].

Encoding scheme. As we will discuss in Sec. IV, a significant source of error in the trapped-ion platform is fluctuating local field noise. This noise can interfere with the specific local field strengths necessary to faithfully encode an optimization problem in H_{cost} . To be able to treat noise and the cost Hamiltonian on separate footings, it can be useful to encode the problem purely into two-body interactions. To that end, one can introduce an additional ancilla qubit and write the cost Hamiltonian as

$$H_{\text{cost}} = \sum_{ij} J_{ij} \sigma_i^z \sigma_j^z + \sigma_a^z \sum_i J_{ai} \sigma_i^z, \quad (4)$$

where the local biases of the original cost Hamiltonian [Eq. (2)] are identified as $h_i^z \rightarrow \sigma_a^z J_{ai}$, with index a referring to the ancilla qubit. The driving Hamiltonian now also addresses the ancilla qubit, and the system is initialized in the equal superposition of all qubits, including the ancilla.

As a direct effect of this reformulation of the problem, the ground state of the cost Hamiltonian becomes twofold degenerate. Measuring either one of the two ground states at the end of the annealing sweep gives the solution to the problem, as the measurement result of the ancilla qubit implies whether the rest of the qubits have to be flipped or not. In the main text of this work, we always assume that the problem is encoded solely in the two-body interactions with this scheme, and all fidelities we report correspond to the sum of the two overlaps with the two respective ground states.

An important consequence of this encoding scheme is that the noise in the local field fluctuations gets isolated from the problem itself and can be effectively compensated by applying dynamical decoupling sequences during the annealing sweep, which we will discuss in detail in Sec. VI. Note that it is not strictly necessary to introduce an ancilla qubit as above to perform quantum annealing with dynamical decoupling to mitigate local field noise. In Appendix A, we discuss a similar annealing protocol that does not require the ancilla qubit and has comparable performance.

III. TEST QUBO PROBLEM: MULTIPLE-OBJECT-TRACKING

Multiple-object-tracking (MOT) is a fundamental computer vision problem that involves identifying and following multiple objects of interest across consecutive video frames [64]. The task requires solving two key challenges: first, detecting objects in each frame, and second, associating these detections across frames to form coherent trajectories or tracks. A detection refers to a localized object instance in a

single frame (typically represented by a bounding box) [65], while a track represents the temporal sequence of detections corresponding to the same object as it moves through the scene. The association step (determining which detections in consecutive frames belong to the same object) naturally leads to complex combinatorial optimization problems. These problems become particularly challenging in crowded scenes with occlusions, appearance changes, and false positive detections, making MOT a computationally demanding task that often requires solving NP-hard assignment problems [66].

We investigate quantum annealing as a solver for MOT problems, focusing on the challenge of associating given detections to coherent tracks. A QUBO formulation offers a suitable framework to compare detections across frames. To construct the QUBO instance, we follow the approach discussed in Ref. [50], which scales linearly with the number of detections, tracks, and frames and could, in principle, scale to industry-relevant problem sizes. As this serves as an initial testbed for quantum annealing experiments, we introduce simplifying assumptions to reduce the problem size. A related work, which focuses on fast, online algorithms, is presented in Ref. [67].

A. General QUBO formulation

The considered QUBO problem consists of two parts: a cost term that evaluates the similarity between detections from different frames, and a penalty term that enforces constraints on the solution vector. The output of the annealing protocol is a binary vector indicating the track assignment for each detection. For a given frame f , a fixed number of tracks, T , and D_f detections in the frame, we denote a binary vector as

$$x^{(f)} \in \{0, 1\}^{D_f T}, \quad (5)$$

which encodes detections as

$$x_{kT+l}^{(f)} = 1 \iff \text{detection } k \text{ in frame } f \text{ is in track } l. \quad (6)$$

To ensure a valid assignment, we impose two constraints, both expressible as a linear system ($Ax = b$): each detection k must be assigned to exactly one track l , and each track must receive exactly one detection per frame:

$$\sum_{l=0}^{T-1} x_{kT+l}^{(f)} = 1 \quad \forall k \in \{0, \dots, D-1\}, \quad (7)$$

$$\sum_{k=0}^{D-1} x_{kT+l}^{(f)} = 1 \quad \forall l \in \{0, \dots, T-1\}. \quad (8)$$

Figure 1 illustrates this encoding for pedestrian tracking. This formulation does not yet account for hidden objects or false detections. These can be handled by introducing dummy tracks and detections—additional variables that are excluded from the constraints by increasing T and D accordingly.

Since the objective term addresses how likely detections from different frames correspond to the same object, a geometric cost can be introduced by assuming limited pedestrian speed, making nearby bounding boxes in consecutive frames more likely to match. This is quantified using the ratio between the area of intersection and the area of union for two



FIG. 1. Two frames of a video from Refs. [53,68] with pedestrians walking in front of the camera. The color of the bounding boxes identifies the persons. Two frames with two detections each can be formulated as a QUBO problem with four binary variables.

bounding boxes B_1 and B_2 :

$$G(B_1, B_2) := \frac{|B_1 \cap B_2|}{|B_1 \cup B_2|}. \quad (9)$$

Furthermore, a measure of visual similarity, C , can be obtained by taking the cosine similarity between features $f_{k_1}, f_{k_2} \in \mathbb{R}^{1024}$ that are extracted from a pretrained neural network [69]:

$$C(f_{k_1}, f_{k_2}) := \frac{f_{k_1}^T f_{k_2}}{\|f_{k_1}\| \cdot \|f_{k_2}\|}. \quad (10)$$

The indices k_1 and k_2 refer to different detections. The associated cost is considered only if both detections are included in the same track. This similarity cost is also used in Ref. [70].

To combine geometric and visual similarity, we take a weighted sum. In our implementation, both components contribute equally. Calling m and l the indices corresponding to sorting detections k_1 and k_2 to track t , respectively, the considered cost function is

$$W_{m,l} = \frac{1}{2}C(f_{k_1}, f_{k_2}) + \frac{1}{2}G(k_1, k_2). \quad (11)$$

To combine the objective and the constraints [Eqs. (7) and (8)] into a single QUBO problem, we apply the penalty method. For a sufficiently large penalty parameter $\lambda \in \mathbb{R}_+$, the constrained optimization problem is equivalent to the unconstrained QUBO problem:

$$\begin{aligned} \arg \min_{\{x \in \{0,1\}^n | Ax=b\}} x^T W x \\ = \arg \min_{\{x \in \{0,1\}^n\}} x^T W x + \lambda \|Ax - b\|^2. \end{aligned} \quad (12)$$

While λ can be chosen arbitrarily large in theory, in practice, smaller values that still enforce the constraints are preferred, as large penalties may effectively amplify hardware noise and reduce precision. Furthermore, the problem exhibits a symmetry: in pedestrian tracking videos, the identities (or colors) of the tracked objects can be permuted; the track numbering is arbitrary. To break this symmetry, we fix the assignment in the first frame. This is done by setting the corresponding binary entries in the solution vector, which has the added benefit of reducing the QUBO size.

B. Minimal example with four qubits

We construct the smallest nontrivial example by considering the MOT problem with only two frames, two tracks

($T = 2$), and two detections ($D = 2$). The binary solution vector would be eight dimensional $[\vec{x} = (\vec{x}^{(1)}, \vec{x}^{(2)})^T$; cf. Eq. (5)]. By exploiting the symmetry under permutation of tracking labels, we reduce the problem to four dimensions by fixing the first four entries of the solution vector. This breaks the degeneracy and ensures a unique solution. The fixed assignment—consistent with the constraints in Eqs. (7) and (8)—is chosen arbitrarily as $x_1^{(1)} = x_4^{(1)} = 1$ and $x_2^{(1)} = x_3^{(1)} = 0$. We then solve only for the remaining four-dimensional vector $\vec{x}^{(2)}$. Formulating the constraints as a linear system of equations $A\vec{x}^{(2)} = b$, we get

$$\begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix} \vec{x}^{(2)} = \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}. \quad (13)$$

Since we fixed the assignments of detections to tracks on the first frame, the objective part W of the cost function in Eq. (12) reduces to only local field terms which depend on the chosen scene. For this minimal problem with four qubits, all two-body interactions in the total QUBO instance are only defined by the constraints and consequently identical for all chosen scenes and frames. In the MOT instance corresponding to Fig. 1, the objective is given by $W \approx (-2.18, -0.95, -0.58, -2.26)$.

To write Eq. (12) as a cost Hamiltonian for quantum annealing, we need to translate the Boolean variables to quantum spins via the transformation $x_i \rightarrow (1 + \sigma_i^z)/2$. The penalty factor λ in Eq. (12) has a large effect on the minimal energy gap in the corresponding quantum annealing protocol and thus has to be chosen carefully [63,71]. For this small problem instance, we can simulate the annealing sweep numerically and choose a value which is close to optimal ($\lambda = 2.5$). For larger systems, there are different heuristics available for choosing a sensible value (e.g., see Ref. [72]). Finally, for our numerical simulations in Secs. V and VI, we always consider a quadratized cost Hamiltonian with only two-body interactions and no biases, which can be obtained by introducing an ancilla qubit as discussed in Sec. II.

For the minimal problem described above, we then obtain the following five-qubit cost Hamiltonian,

$$H_{\text{cost}}/J = \begin{pmatrix} 0 & -0.87 & -0.38 & -0.23 & -0.91 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \quad (14)$$

which we write in units of the characteristic energy scale, i.e., the maximal realizable two-body coupling strength $J := \max_{ij} |J_{ij}|$. This energy scale depends on the considered hardware platform. We discuss details for trapped-ion hardware using magnetic-gradient-induced couplings in the next section. In Appendix B, we provide the explicit QUBO formulation of the same MOT problem given in Fig. 1, but with one additional frame, which results in a QUBO cost Hamiltonian with nine qubits. In the same Appendix, we also discuss how the number of qubits scales when considering more complex MOT problems. Furthermore, in Appendix C, we discuss the cutting-stock problem—another industrially

relevant optimization problem—and how to construct minimal instances.

IV. REFERENCE PLATFORM

Trapped-ion systems with their native all-to-all connectivity [19–26] are a promising candidate for quantum annealing applications [73]. In this work, we consider a specific variant that uses microwave fields to address individual ions, whose resonance frequencies are shifted relative to each other by an external magnetic field gradient [27–37]. In this setup, there exists an always-on, all-to-all MAGIC between ions, which could be utilized for encoding cost Hamiltonians of optimization problems for quantum annealing algorithms. We start by giving a short summary of the MAGIC scheme, after which we provide specific details of the considered experimental platform. In particular, we discuss realizable coupling strengths on current hardware and realistic noise spectra that are detrimental for quantum annealing.

A. Magnetic-gradient-induced coupling

While traditional trapped-ion setups use optical lasers to address individual ions and implement two-qubit gates (e.g., Mølmer-Sørensen-type interactions [74,75]), the MAGIC scheme utilizes microwave fields [27–29]. To preserve individual addressability of ions, an external magnetic field gradient is applied, creating position-dependent Zeeman shifts for each ion in the trap. A beneficial side effect of this magnetic field gradient is the automatic generation of two-body interactions between all ions, even without external driving fields, which can be harnessed to implement entangling many-body gates [32,76]. We provide a brief overview of how these interaction strengths arise and direct readers to the literature for further details [27–29,37].

For the qubit states, a magnetically sensitive transition is chosen (which is typically a transition in a hyperfine ground-state manifold with a resonance frequency in the microwave regime [31,33,77]). In the trap, when choosing the longitudinal and radial trapping frequencies appropriately, the ions equilibrate into a one-dimensional chain [78]. When the ions are then subjected to a magnetic field gradient $\partial_z B$ along the chain axis, the resonance frequencies $\omega(z)$ obtain a Zeeman shift. The position-dependent resonance frequencies then result in spin-dependent forces on the ions, which couple internal spin degrees of freedom with the collective motion of the ion chain (phonons). Switching to a rotating frame and applying a polaron transformation, one can decouple internal and external degrees of freedom and show that the system (without driving fields) is described by the Hamiltonian [27–29]

$$H = -\frac{\hbar}{2} \sum_{n=1}^N \omega_n \sigma_n^z + \hbar \sum_{n=1}^N v_n a_n^\dagger a_n - \frac{\hbar}{2} \sum_{i < j}^N J_{ij} \sigma_i^z \sigma_j^z, \quad (15)$$

where ω_n is the resonance frequency of ion n in its equilibrium position and $a_n^\dagger$ and a_n are bosonic creation and annihilation operators for the collective phonon modes of the ion chain with respective frequencies v_n . The only approximations for arriving at Eq. (15) are assuming a quadratic external po-

tential, given by the combination of trapping potential and Coulomb repulsion between ions, and a vanishing curvature of the magnetic field (see Ref. [37] for a discussion of higher-order effects). The native two-body coupling strength J_{ij} in the MAGIC scheme is given by

$$J_{ij} = \sum_{n=1}^N v_n \epsilon_{in} \epsilon_{jn}, \quad (16)$$

with the effective Lamb-Dicke parameter ϵ_{nl} , which determines how strongly ion n couples to phonon mode l and which is proportional to the applied magnetic field gradient. Denoting the width of the ground-state wave packet of the harmonic collective mode l as $\Delta z_l = \sqrt{\hbar/2mv_l}$ (m is the mass of the ions), and using the orthogonal matrix S that translates between individual ion coordinates and the collective normal modes, the Lamb-Dicke parameter is given by

$$\epsilon_{nl} = \frac{\Delta z_l \partial_z \omega_n}{v_l} S_{nl}. \quad (17)$$

The native spin-spin coupling strength J_{ij} in the MAGIC setup is directly influenced by the magnetic field gradient and the strength of the external harmonic longitudinal trapping potential v ; e.g., $J_{ij} \sim (\partial_z B/v)^2$ for a harmonic trap of strength $\frac{1}{2}mv^2$ [63].

In contrast to the two-qubit gates, single-qubit rotations are implemented in the MAGIC setup analogously to other trapped-ion platforms. In particular, when tuning a driving field with Rabi frequency Ω to the resonance frequency of an ion, the Hamiltonian H_{drive} describing the coupling to the field, in a rotating frame, can be expressed as (see, e.g., Ref. [37])

$$H_{\text{drive}} = \frac{\hbar\Omega}{2} (\cos(\phi) \sigma_n^x - \sin(\phi) \sigma_n^y), \quad (18)$$

where ϕ is the phase of the driving field. By leaving the driving field on for a finite time and tuning the phase appropriately, rotations around an arbitrary axis in the xy plane can be realized. Rotations around the z axis are typically implemented virtually [79].

B. Quantum annealing with trapped ions

As discussed in Sec. II, we choose to consider a modified version [Eq. (1)] of the standard quantum annealing protocol, where the cost Hamiltonian remains constant throughout the evolution, while the driving Hamiltonian is ramped down linearly. The MAGIC-based trapped-ion platform lends itself naturally to this protocol: as we have explained above, the presence of the magnetic gradient induces a native all-to-all two-body coupling that is always active, even without any additional driving fields. If one manages to directly encode an optimization problem into this coupling matrix, one only needs to implement the single-qubit rotations of the driving Hamiltonian, which can be easily achieved with microwave fields tuned to the individual resonance frequencies of the ions. To match the two-body couplings to a specific problem instance, there are multiple possibilities: by employing different, specifically tailored and not necessarily harmonic trapping potentials along the chain, coupling strengths between different pairs of ions can be enhanced or reduced

[30]. Furthermore, by applying suitable dynamical decoupling sequences, effective two-body interaction strengths can be continuously tuned (see Appendix E). Finally, arbitrary subsets of the ions can be excluded from the all-to-all interaction by transferring their state populations to magnetically insensitive states [31].

If the problem cannot be directly encoded into the MAGIC couplings, one can always perform a Trotterized version of the modified annealing protocol [80]. In our numerics in Sec. VI, we consider a continuous-time protocol, though the same dynamical decoupling strategy can also be incorporated into Trotterized annealing protocols.

C. Experimental setup

While our numerical simulations in the following sections employ dimensionless units throughout, we provide concrete timescale estimates based on realistic device parameters. Our analysis references experimental parameters and coupling strengths previously demonstrated in the MAGIC setup [31].

Specifically, we consider a linear chain of five $^{171}\text{Yb}^+$ ions (matching our five-qubit MOT problem constructed in Sec. IIIB), with an axial trapping frequency of $\omega_z = 2\pi \times 130$ kHz and a magnetic field gradient of $\partial_z B = 19 \text{ T m}^{-1}$. Under these conditions, the system achieves a maximal two-body coupling strength of $\max_{ij} |J_{ij}| \approx 26.5$ Hz between ions. We refer to this value when we set the characteristic energy scale to $J = 26$ Hz [see Eq. (14)], which determines the fundamental timescale for the quantum annealing sweeps analyzed in Sec. VI. Next-generation devices offer the potential for substantially enhanced performance through increased magnetic field gradients. For instance, with a stronger gradient of $\partial_z B = 150 \text{ T m}^{-1}$, the maximal coupling strength increases dramatically to $\max_{ij} |J_{ij}| \approx 1650.6$ Hz (see also Ref. [37]). This approximately 60-fold enhancement would correspondingly reduce the required annealing timescales and improve the overall protocol efficiency.

For atomic qubits encoded in magnetically sensitive states, as employed in the MAGIC scheme, the dominant source of noise arises from fluctuations in the externally applied magnetic field. The qubit splitting depends in first order linearly on the field strength through the Zeeman effect, making it highly susceptible to such variations. These fluctuations typically originate from ambient laboratory fields, most prominently at 50 Hz and its harmonics, as well as in the tens-of-kilohertz range due to power-supply noise, instabilities in current or voltage sources, and slow thermal drifts in coils or permanent magnets. In the absence of magnetic shielding, the noise at 50 Hz typically dominates and in the worst case produces fluctuations with amplitudes up to the order of kilohertz, directly causing dephasing and reducing gate fidelity. For example, Ref. [81] reports magnetic field rms noise of around 100 nT without stabilization. For $^{171}\text{Yb}^+$ ions, when considering $|F = 0, m_F = 0\rangle$ and $|F = 1, m_F = +1\rangle$ in the hyperfine ground-state manifold as the qubit states, this corresponds to local field noise on the order of $\delta h_i^z \approx 1.4$ kHz. Additional noise channels further degrade performance. Microwave drive amplitude and phase noise introduce Rabi-rate fluctuations and AC Stark shifts, while motional noise from trap-voltage

instabilities can modulate the effective magnetic field experienced by the atoms in the gradient.

Similar issues arise in other quantum-computing platforms. Broadly, fast technical noise sets the limit on short-time coherence, whereas slow drifts accumulate over minutes to hours, necessitating compensation through recalibration and active feedback. It is the purpose of the rest of the paper to, first, study how such salient error terms impact the possible solution (Sec. V) and, second, determine how dynamical decoupling can be applied to counteract one of the most prominent noise sources (Sec. VI), fluctuations of the longitudinal field (equivalent to fluctuations of the qubit frequencies).

V. ERROR ANALYSIS ON COST FUNCTION

This section examines two significant error sources for any quantum annealing experiment: fluctuations in the two-body couplings (Sec. VA) and undesired local field terms (Sec. VB). To assess the impact of these errors on annealing performance, we calculate the probability that random fluctuations of each type alter the ground state of the cost Hamiltonian. The quantum annealing protocol fundamentally relies on encoding the problem solution into the ground state of the cost Hamiltonian: if fluctuations cause the ground state to change, the protocol can no longer identify the optimal solution. Our analysis reveals contrasting behaviors for these two error sources when applied to our test QUBO problem given in Eq. (14). Under realistic experimental conditions the impacts of fluctuations in the two-body couplings are negligible for annealing performance. However, local field fluctuations present a substantial challenge: the noise amplitudes typically encountered in trapped-ion setups are sufficient to prevent successful quantum annealing without active intervention. Consequently, effective noise mitigation strategies are essential for protocol success, which we will examine in detail in Sec. VI.

A. Fluctuations in the two-body couplings

We consider two types of two-body coupling fluctuations: First, in the MAGIC scheme, variations of the overall trapping frequency cause correlated fluctuations in the coupling strengths. Second, we discuss the general case of uncorrelated fluctuations of the two-body couplings, which could, for example, be caused by erroneous two-body gate implementations in a Trotterized annealing sweep.

The trapping frequencies in a trapped-ion system typically exhibit slow temporal drifts. To model this effect, we consider variations in the harmonic axial trap frequency ω_z according to

$$\omega_z \rightarrow \omega_z + \delta\omega_z, \quad (19)$$

where we assume that the fluctuation $\delta\omega_z$ remains constant on the timescale of a single annealing sweep.

Since the two-body coupling strengths J_{ij} in the MAGIC scheme are directly proportional to $1/\omega_z^2$, any fluctuations in the trap frequency induce correlated fluctuations in all of the two-body couplings. From Eq. (16), we calculate the

two-body coupling variations to first order in $\delta\omega_z$ as

$$\delta J_{ij} \approx \sum_{n=1}^N \frac{-\hbar \partial_z \omega_i \partial_z \omega_j}{v_n^2 \omega_z m} S_{in} S_{jn} \delta\omega_z = -2 \frac{J_{ij}}{\omega_z} \delta\omega_z, \quad (20)$$

where we used the fact that the axial phonon modes v_n are proportional to the axial trap frequency ω_z . Consequently, the standard deviations of the coupling and frequency fluctuations are related by $\sigma(\delta J_{ij}) \approx |2J_{ij}/\omega_z| \sigma(\delta\omega_z)$. Numerically, we verified that this first-order approximation remains accurate for fluctuation amplitudes up to $\sigma(\delta\omega_z)/\omega_z \approx 0.1$. For larger amplitudes, quadratic corrections become significant.

All J_{ij} are thus scaled by the same factor when the axial trap frequency ω_z fluctuates. For the annealing protocol we consider in this work, where the problem is encoded only in the two-body interactions without utilizing local biases (see Sec. II), this corresponds to effectively rescaling the entire cost Hamiltonian while preserving its structure. Importantly, such a global rescaling shifts only the eigenvalues, but leaves the eigenstates unaltered. As a result, the ground state, which encodes the solution to the optimization problem, remains unaffected. Therefore, fluctuations in ω_z and the resulting correlated fluctuations in J_{ij} do not alter the computational outcome of the annealing protocol. However, for other annealing protocols that also utilize local biases to encode the problem (e.g., see Appendix A), these fluctuations could in principle become relevant, as they change the relative energy scales of the two-body interactions with respect to the local biases.

We now discuss the case of uncorrelated fluctuations in the two-body coupling terms J_{ij} . Unlike the correlated disorder induced by fluctuations in the axial trap frequency ω_z , uncorrelated coupling noise does not scale all interaction terms uniformly and can, in principle, modify the structure of the cost Hamiltonian and thus affect its ground state.

To assess the impact of such uncorrelated disorder, we numerically compute the probability that random fluctuations δJ_{ij} change the ground state of the cost Hamiltonian, choosing the minimal MOT instance with five qubits defined in Eq. (14) as a test problem. Each data point in Fig. 2 represents the mean probability obtained from 10 000 independent samples. For each sample, the coupling perturbations δJ_{ij} are drawn independently from a Gaussian distribution with zero mean and specified standard deviation. The figure shows the probability of a ground-state change as a function of the relative disorder strength $\sigma(\delta J_{ij})/J_{ij}$. The variance of each data point p can be calculated as $p(1-p)/10\,000 \leq 2.5 \times 10^{-5}$. We observe that the ground state remains stable for disorder levels up to approximately 20%, beyond which the probability for a changed ground state begins to increase. Fortunately, such large fluctuations are atypical in realistic experimental conditions. Moreover, in the MAGIC setup (see Sec. IV), uncorrelated coupling disorder can only arise from imperfections in the trap potential or inhomogeneities in the magnetic field gradient, which are unlikely to occur over the length scale of the ion chain. Otherwise, uncorrelated disorder in J_{ij} could in principle stem from imperfect two-body gate operations in the case of Trotterized annealing protocols. However, two-body gate fidelities of $\sim 99.9\%$ fidelity can




FIG. 2. Probability of a changed ground state of the cost Hamiltonian for the quadratized minimal MOT problem [five qubits; see Eq. (14)] as a function of the standard deviation of uncorrelated fluctuations in the two-body couplings $\sigma(\delta J_{ij})/J$. Each point corresponds to the mean of 10 000 samples and has a variance of $\leq 2.5 \times 10^{-5}$. The dashed gray line corresponds to a fit with an arctangent function, starting at $\sigma(\delta J_{ij})/J = 0.3$. Fit parameters: $a, b, c, d \approx 0.41, 2.59, 0.50, 0.22$.

already be achieved for current trapped-ion quantum computers [82]. Although we cannot exclude that the picture may change when scaling up the size of the optimization problem, the large robustness encountered for the benchmark models gives optimism that there is still significant room before the small realistic levels of uncorrelated coupling fluctuations can become detrimental. We therefore conclude that, while uncorrelated coupling fluctuations could in principle affect solution fidelity, they are not expected to play a significant role in practice under typical conditions of current experiments.

B. Local field fluctuations

Magnetic-field-gradient fluctuations in the MAGIC setup (see Sec. IV), along with unwanted external magnetic fields, induce corresponding fluctuations in the ion resonance frequencies. These perturbations manifest as additional longitudinal fields $\sim \delta h^z \sigma^z$ in the cost Hamiltonian of the annealing protocol [Eq. (2)]. We model these disorder terms by modifying the local field components:

$$h_i^z \rightarrow h_i^z + \delta h_i^z, \quad (21)$$

where δh_i^z represents the field fluctuation for ion i . In our considered encoding scheme there are no local fields in the cost Hamiltonian ($h_i^z = 0$; see Sec. II). In this section, we assume that the fluctuations remain constant on the timescale of a quantum annealing sweep. We will consider field fluctuations with colored noise spectra in Sec. VI.

To assess the quantum annealing protocol's resilience to local field variations, we add disordered diagonal terms $\delta h_i^z \sigma_i^z$ to the cost Hamiltonian given in Eq. (14). We sample the δh_i^z from Gaussian distributions with varying standard deviations. For each standard deviation, we generate 10 000 samples and determine numerically whether the ground-state configuration changes.




FIG. 3. Probability of a changed ground state of the cost Hamiltonian for the quadratized minimal MOT problem [five qubits; see Eq. (14)] as a function of the standard deviation of the local field disorder δh_i^z in units of J for correlated (blue) and uncorrelated (red) disorder. The data points show the mean probabilities of a changed ground state for 10 000 simulations each (variance of $\leq 2.5 \times 10^{-5}$), where δh_i^z is sampled from a Gaussian distribution with zero mean. The dashed gray lines represent fits with the arctangent function $f(x) = a \arctan(b(x - c)) + d$, starting from $\delta h_i^z/J = 1$. For correlated disorder, $a, b, c, d \approx 0.74, 0.63, 0.39, -0.17$. For uncorrelated disorder, $a, b, c, d \approx 0.57, 0.81, 0.99, 0.10$.

Figure 3 presents our results for two distinct disorder scenarios: correlated fluctuations ($\delta h_i^z = \delta h_j^z$ for all i, j) and uncorrelated fluctuations (all δh_i^z independently sampled). In realistic experimental setups, the case of correlated disorder is typically more relevant, as external magnetic fields are unlikely to significantly vary on the length scale of the ion chain. Uncorrelated variations could, however, effectively occur by erroneous calibrations of frequencies addressing the individual qubit transitions. Both cases exhibit a sharp transition in ground-state stability at approximately $\delta h_i^z/J \approx 0.5$ (see inset), where J represents the energy scale set by the maximal two-body coupling strength. The dashed gray lines show heuristic arctan fits starting at $\delta h_i^z/J = 1$.

The exact value where local field variations change the optimization result will depend on the precise problem considered. However, the realistic experimental implementations face significantly larger local field fluctuations than the order $\delta h_i^z/J \approx 0.5$ found for the MOT benchmark problem. Using experimental parameters that have been previously demonstrated for the MAGIC setup [31], a five-ion chain reaches a maximal two-body coupling strength of $J_{ij} \approx 26$ Hz [37]. In the same setup, local field fluctuations can reach magnitudes on the order of 10^3 Hz in the worst-case scenario without any magnetic shielding and stabilization [81], corresponding to $\delta h_i^z/J \sim 40$. This places the system well within the regime where the desired solution of the optimization problem is no longer the ground state. As mentioned above, improved setups can reach $\max_{ij} |J_{ij}| \approx 1650.6$ Hz [37], which, however, is still not sufficient to safely ignore variations of the local fields. To avoid degradation of the quantum annealing performance, dedicated mitigation strategies are thus required. The next section demonstrates that local fields can be completely com-

pensated through simple dynamical decoupling sequences, even when they fluctuate substantially and when they are time dependent.

VI. DYNAMICAL DECOUPLING

This section explores how simple dynamical decoupling protocols can be implemented during annealing sweeps to completely suppress the effect of local field fluctuations and recover ideal annealing performance, both under slow drifts that are approximately constant within each experimental run (equivalent to random disorder) and under time-dependent noise. We begin in Sec. VIA with a general explanation of dynamical decoupling (DD). We also derive an effective time-evolution operator for annealing protocols subject to varying fields and incorporating DD pulses, which provides further analytical insights. Our numerical results are presented in Sec. VIC, validating our theoretical discussion and providing concrete estimates for the required pulse counts in experimental implementations across different noise spectra and amplitudes. Finally, Sec. VID discusses a universal scaling behavior we observe in the number of dynamical decoupling pulses needed for effective noise suppression.

A. Applying dynamical decoupling pulses during a quantum annealing protocol

Dynamical decoupling represents a powerful error mitigation technique designed to extend the coherence times of quantum systems by suppressing environmental decoherence [38–42,44–47]. The fundamental principle involves applying periodic control pulses that rotate all qubits around specific axes by predetermined angles, effectively averaging out unwanted error terms.

When implementing dynamical decoupling within quantum annealing protocols, a fundamental challenge arises: the applied pulses generally do not commute with the quantum annealing Hamiltonian given in Eq. (1). This means that a control pulse can excite the system from its instantaneous ground state into higher energy eigenstates, potentially compromising the annealing process. Consequently, careful corrections must be implemented to ensure that the annealing sweep continues to converge toward the correct solution.

A general control pulse, assumed to be applied nearly instantaneously (i.e., much faster than the relevant timescales of the annealing sweep), can be mathematically represented as a collective rotation of all qubits around an arbitrary axis: $U_{\text{pulse}} = \prod_k \exp(-i\mathbf{v}^{(k)} \cdot \boldsymbol{\sigma}^{(k)})$, where $\boldsymbol{\sigma}^{(k)} = (\sigma_x^{(k)}, \sigma_y^{(k)}, \sigma_z^{(k)})$ denotes the vector containing the Pauli matrices acting on qubit k .

In this work, motivated by the specific form of the annealing Hamiltonian and salient error term, we focus on the simplest case: π pulses that rotate qubits around the x axis, equivalent to spin-flip operations. Such pulses can be implemented in the MAGIC setup by turning on driving fields, tuned to the qubit resonances, for a finite time interval [Eq. (18)]. This choice of simple π pulses is sufficient for mitigating the local field fluctuations considered in our analysis (see Sec. VB). The action of such pulses on the driving

and cost Hamiltonians of the annealing protocol, as given in Eq. (1), yields

$$e^{i\frac{\pi}{2}\sigma_x} H_{\text{driving}} e^{-i\frac{\pi}{2}\sigma_x} = H_{\text{driving}}, \quad (22)$$

$$e^{i\frac{\pi}{2}\sigma_x} H_{\text{cost}} e^{-i\frac{\pi}{2}\sigma_x} = \sum_{ij} J_{ij} \sigma_i^z \sigma_j^z - \sum_i h_i^z \sigma_i^z. \quad (23)$$

Notably, the spin-flip operations leave the driving Hamiltonian unchanged and only modify the signs of the local fields h_i^z in the cost Hamiltonian. In the case where the fields are undesired error terms, this sign flip will enable alleviating their effect, as detailed below. In the case where the fields are part of the problem cost Hamiltonian, the sign reversal is a crucial perturbation that needs to be compensated.

To address this issue, we employ here the ancilla qubit approach described in Sec. II. By encoding all local fields into additional two-body interactions through the introduction of an ancilla qubit [see Eq. (4)], we render the complete annealing protocol Hamiltonian invariant under dynamical decoupling pulses. This ensures that the system remains in the instantaneous ground state throughout the entire annealing sweep. An alternative approach involves actively compensating for the sign flip with each pulse application, resulting in the annealing protocol described in Appendix A.

Practical experimental implementations often employ more sophisticated dynamical decoupling sequences. For instance, universally robust sequences [40,42,43] provide enhanced protection against a broader spectrum of noise sources and compensate for imperfections in the pulse sequences themselves. These advanced protocols generally incorporate rotations around multiple axes with varying angles. The integration of such general pulse sequences within quantum annealing frameworks has been investigated in Refs. [48,49].

B. Time-evolution operator

We now demonstrate how periodically applied dynamical decoupling pulses in the form of global spin flips can cancel local field fluctuations during an annealing protocol. To this end, we consider the quantum annealing Hamiltonian discussed in Sec. II with noise,

$$H(t) = C(t)H_{\text{driving}} + H_{\text{cost}} + H_{\text{noise}} \\ = -C(t) \sum_i h^x \sigma_i^x + \sum_{ij} J_{ij} \sigma_i^z \sigma_j^z + \sum_i \delta h_i^z(t) \sigma_i^z, \quad (24)$$

where the problem is encoded only in the two-body interactions as discussed in Sec. II, while the local fields $\delta h_i^z(t)$ in H_{noise} are time-dependent error terms that we aim to effectively cancel via dynamical decoupling.

To make analytic progress, we calculate the time-evolution operator of this annealing protocol for a small time interval $2\Delta t \equiv t_+ - t_-$. During this interval, we apply two dynamical decoupling control pulses, the first at the center of the interval at time $t_0 = t_- + \Delta t$ and the second pulse at time t_+ . Before and after the first pulse, the system evolves according to Eq. (24) for the respective time interval Δt .

We assume that $2\Delta t$ is small enough so that the function $C(t)$ can be approximated as linear during the corresponding annealing segment. At a given time t in the small time interval

$2\Delta t$, it is then defined as

$$C(t) = C(t_-) - v(t_-)t, \quad t \in [t_-, t_+]. \quad (25)$$

Here, $v(t_-)$ is a generally time-dependent function that determines the form of the annealing protocol. In our numerical simulations in the following section, the driving ramp protocol is chosen to be linear during the entire annealing sweep, which is often the standard choice. This means that the slope is constant [$v(t_-) \equiv v$ for all t_-] during the entire sweep. Furthermore, we assume that the disorder terms fluctuate sufficiently slowly so that we can take them as constant during the short time interval $2\Delta t$: $\delta h_i^z(t) \equiv \delta h_i^z, \forall t \in [t_-, t_+]$.

The unitary time-evolution operator in the first half of the interval from t_- to t_0 can be written using the Magnus expansion [83]:

$$U_{t_- \rightarrow t_0} = \mathcal{T} e^{-i \int_{t_-}^{t_0} dt' H(t')} \\ = e^{-i \int_{t_-}^{t_0} dt' H(t') - \frac{i}{2} \int_{t_-}^{t_0} dt' \int_{t_-}^{t'} dt'' [H(t'), H(t'')] + O(\Delta t^4)}. \quad (26)$$

Evaluating the commutators, integrating $C(t)$ [using Eq. (25)] and keeping only terms up to third order in Δt gives

$$U_{t_- \rightarrow t_0} \approx \exp \left(-i[C(t_-)H_{\text{driving}} + H_{\text{cost}} + H_{\text{noise}}]\Delta t \right. \\ \left. - \frac{i}{6} h^x v(t_-) \left[\sum_i \delta h_i^z \sigma_i^y + 2 \sum_{ij} J_{ij} \sigma_i^z \sigma_j^y \right] \Delta t^3 \right), \quad (27)$$

where we defined $\bar{t}_- = (t_- + t_0)/2$. The time-evolution operator $U_{t_0 \rightarrow t_+}$ for the interval from t_0 to t_+ has an analogous structure. To calculate the total time-evolution operator for the annealing sweep of duration $2\Delta t$ from time t_- to t_+ , we can use the Baker-Campbell-Hausdorff formula (again neglecting terms of fourth and higher order in Δt),

$$\sigma^x U_{t_0 \rightarrow t_+} \sigma^x U_{t_- \rightarrow t_0} = e^{-i[\hat{A}_1 + \hat{A}_2 + \hat{A}_3]2\Delta t + O(\Delta t^4)}, \quad (28)$$

with

$$\hat{A}_1 = \sum_{ij} J_{ij} \sigma_i^z \sigma_j^z - h^x C(t_0) \sum_i \sigma_i^x, \quad (29)$$

$$\hat{A}_2 = \left(h^x C(t_0) \sum_i \delta h_i^z \sigma_i^y \right) \Delta t, \quad (30)$$

$$\hat{A}_3 = \left(-\frac{2}{3} h^x v \sum_{ij} J_{ij} \sigma_i^z \sigma_j^y + \frac{2}{3} h^x C(t_0) \sum_i \delta h_i^z \sigma_i^x \right) \Delta t^2. \quad (31)$$

Here, we write $\sigma^x = \bigotimes_i \sigma_i^x$ as shorthand notation denoting the global spin flip due to the dynamical decoupling pulses. Furthermore, we assume that the slope of the driving protocol remains nearly constant during the time interval $2\Delta t$, so that we can set $v(t_-) = v(t_+) \equiv v$.

The first-order terms recover the ideal quantum annealing evolution with the local field disorder being completely canceled, i.e., $\hat{A}_1 = C(t)H_{\text{driving}} + H_{\text{cost}}$. The leading-order corrections, given in Eq. (30), to this ideal evolution can be interpreted as a modification of the driving Hamiltonian, as they act in a basis direction orthogonal to the cost Hamiltonian. Their strength is proportional to the amplitude of the

local field disorder terms and the interval between the dynamical decoupling pulses ($\sim \delta h_i^z \Delta t$). Our numerical simulations in the following section support this observation, showing that the local field noise can always be suppressed by choosing the interval between dynamical decoupling pulses to be sufficiently small. Furthermore, in Sec. VID, we demonstrate that the annealing protocol with dynamical decoupling follows a universal scaling behavior, solely determined by the product $\delta h_i^z \Delta t$.

The corrections to the ideal annealing protocol of third order in Δt , given in Eq. (31), are composed of two parts: First, there is a two-body interaction term whose strength is proportional to $\sim \hbar^x v J_{ij} \Delta t^2$. This correction is independent of the local field noise and vanishes in the limit of adiabatic driving ($v \rightarrow 0$). Second, there is another correction to the driving Hamiltonian that scales as the square of the first-order correction ($\sim \delta h_i^z \Delta t^2$).

All corrections to the ideal annealing sweep considered so far rely on the assumption of perfectly instantaneous dynamical decoupling pulses applied at uniform intervals. In practice, however, pulse operations introduce additional sources of error: they have finite duration and their spacing may deviate from uniformity. These control imperfections are well documented, and several dynamical decoupling schemes with intrinsic robustness to such errors have been proposed (e.g., in Ref. [43]). Additional errors can become relevant at higher pulse rates which can necessitate a higher Rabi frequency in experiments. If this frequency gets too close to the trapping frequency, this can produce further errors due to sideband excitations, which need to be compensated, e.g., via sympathetic cooling [84–86].

As an example of a control error, we calculate the effect of a small pulse timing error ϵ on the time-evolution operator for a small time interval $2\Delta t = t_+ - t_-$ during an annealing sweep, analogous to the discussion above: We let the system evolve with the annealing Hamiltonian [Eq. (24)] for a time interval $\Delta t + \epsilon = t_0^\epsilon - t_-$ ($t_0^\epsilon = t_0 + \epsilon$), after which we apply a global spin flip. Afterward, the system evolves for a time interval $\Delta t - \epsilon = t_+ - t_0^\epsilon$ and we apply a second global spin flip at time t_+ . The total time-evolution operator to first order in Δt is then given as

$$\sigma^x U_{t_0^\epsilon \rightarrow t_+} \sigma^x U_{t_- \rightarrow t_0^\epsilon} = e^{-i\hat{A}_1^\epsilon 2\Delta t + O(\Delta t^2)}, \quad (32)$$

where the higher-order terms also include contributions proportional to $\epsilon \Delta t$. The first-order terms can be written as the original annealing Hamiltonian with a rescaled local-field-noise contribution:

$$\hat{A}_1^\epsilon = H_{\text{cost}} + C(t_0)H_{\text{driving}} + \frac{\epsilon}{\Delta t}H_{\text{noise}}. \quad (33)$$

To get a rough estimate of how large ϵ can be before affecting the performance of the annealing sweep, we can consider the numerical result in Sec. VB for the maximal local-field disorder strength the protocol is robust to: $\epsilon/\Delta t \approx 0.5J/\delta h_i^z$, which corresponds to a required accuracy in pulse timing of around 99% for $\delta h_i^z = 1$ kHz and $J = 26$ Hz. In our numerical simulations of annealing sweeps with time-dependent local field noise, we observe that moderately larger fluctuations in pulse spacing of the order of 98% fidelity also have negligible impact on performance, provided they do not exhibit a sys-

tematic pattern that would enable coherent error accumulation over the course of the annealing sweep (see Appendix F).

In summary, one can always recover the ideal annealing protocol and obtain the correct solution by choosing the time interval Δt between consecutive dynamical decoupling pulses of uniform length and sufficiently small with respect to the amplitude of the local field noise δh_i^z . The salient deviations for slow sweeps scale as powers of $\delta h_i^z \Delta t$. Finally, we note that the above scheme of periodically applying dynamical decoupling pulses can also be used to effectively modulate the strength of the two-body coupling weights J_{ij} , which can be useful for encoding arbitrary optimization problems into the two-body couplings. We discuss this method in Appendix E.

C. Numerical results for noise mitigation

In this section, we present numerical simulations of the dynamical decoupling protocol introduced in the previous section. We employ the minimal MOT instance with five qubits, defined in Eq. (14) as our test problem, though our findings regarding noise mitigation through dynamical decoupling extend to different problem instances and system sizes. In Appendix A, we discuss an alternative annealing protocol solving the same problem but using only four qubits and Appendix B contains an analogous treatment of a MOT problem instance encoded into nine qubits. Additionally, in Appendixes C and D we show numerical results for noise mitigation in quantum annealing with five- and six-qubit cutting-stock problems, as well as Sherrington-Kirkpatrick model instances up to 12 qubits, further demonstrating the universality of the dynamical decoupling protocol.

Our simulations solve the time-dependent Schrödinger equation using the annealing Hamiltonian defined in Eq. (24). In the numerics below, we implement a linear annealing sweep over a fixed duration of $2.6/J$, which provides a good fidelity in the absence of any error terms. To put this value into context of the timescales in a typical laboratory experiment, the energy scale $J = 26$ Hz demonstrated in the MAGIC setup [31] corresponds to an annealing time of 100 ms for our chosen parameters. We set the driving strength to $\hbar_x = 3J$ as a compromise to avoid overly large separations of energy scales while keeping the initialization into a fully x -polarized state a good approximation of the initial ground state (see discussion in Sec. II). In our numerics, we discretize the annealing sweep into 50 000 steps. For each discrete step, we set a value of the local field δh_i^z , sampled from a given noise spectrum, thus approximating the continuous fluctuations of an external magnetic field. Here, we consider two distinct types of noise spectra: First, we model an experimentally realistic spectrum featuring two dominant Lorentzian peaks at frequencies of 50 Hz and the leading next harmonic at 150 Hz. As discussed in Sec. IV C, such a spectral structure is typical of external magnetic field fluctuations arising from the electricity grid. Second, we also examine spectra with individual Lorentzian peaks at various frequencies to systematically investigate our protocol's dependence on noise characteristics. All noise realizations are generated using the approximate frequency domain method detailed in Ref. [87]. At approximately equidistant intervals across the 50 000 discrete time




FIG. 4. Fidelity of the final state, compared to the state encoding the problem solution, at the end of a quantum annealing sweep versus dynamical decoupling pulse rate, using the minimal MOT test problem with five qubits and an annealing time of $2.6/J$ (corresponding, for an energy scale of $J = 26$ Hz, to a sweep duration of 100 ms). Boxplots compare results for correlated noise $[\delta h_i^z(t) = \delta h^z(t)]$ using identical noise spectra with two Lorentzian frequency peaks at 50 and 150 Hz but different amplitudes of 250 Hz (blue), 500 Hz (red), 750 Hz (green), and 1000 Hz (orange). The crosses in each boxplot give a comparison with the median value of annealing sweeps subject to constant disorder $[\delta h^z(t) = \delta h^z]$ of the same respective amplitude. The horizontal dashed line shows the achievable fidelity of the ideal noiseless protocol (≈ 0.84). Data for each boxplot corresponds to 50 random independent noise realizations (see Appendix F). We observe that the achieved final fidelity converges to the noiseless case for a sufficient rate of dynamical decoupling pulses, with higher noise amplitudes requiring higher rates.

steps (see Appendix F), we apply dynamical decoupling pulses as instantaneous global spin flips affecting all qubits simultaneously.

To evaluate the success of the protocol, we calculate the overlap between the final state after the annealing sweep and the twofold-degenerate true ground state encoding the problem solution. Figure 4 presents this protocol fidelity as a function of the number of DD pulses for a noise spectrum featuring two Lorentzian peaks centered at 50 and 150 Hz. We consider four noise amplitudes: 250 Hz (blue), 500 Hz (red), 750 Hz (green), and 1000 Hz (orange). These noise amplitudes are chosen as a range going from typical values to a worst-case scenario: In an experimental setup without any magnetic shielding, local field noise can realistically reach values around 1 kHz. Using stabilization techniques, these values can potentially be reduced by an order of magnitude [81].

The considered DD pulse rates vary from 0 to 5 pulses per millisecond, which for the chosen sweep duration of 100 ms correspond to a total of 0 to 500 dynamical decoupling pulses applied throughout the annealing process. For each pulse rate, we simulate 50 independent annealing runs, each subject to a distinct, correlated noise realization (i.e., identical across all qubits). Results are visualized using boxplots, where the boxes indicate the interquartile range and horizontal lines denote median fidelities. The dashed gray line at $\approx 84\%$ repre-

sents the fidelity achieved under the same annealing schedule in the absence of noise.

We observe that increasing the dynamical decoupling pulse rate systematically improves the final fidelity, with all noise amplitudes showing convergence toward the ideal (noiseless) performance. For each amplitude, a pulse rate of approximately 2.5 pulses per ms is sufficient to yield fidelities above 70% in the majority of runs. This rate is readily achievable by current trapped-ion platforms: the experimental upper limit of the pulse rate is set by the width of the dynamical decoupling pulses (since they should not overlap), which in turn depends on the accessible Rabi frequency. For a large Rabi frequency, the pulse width to implement the same π pulse is smaller. The Rabi frequency is effectively limited in the experiment by the trapping frequency, as it should be set below that frequency in order to avoid heating via unwanted sideband excitations. Typical trapping frequencies in an experiment are well above 100 kHz. When setting the Rabi frequency to 100 kHz, the pulse width would be approximately 0.06 ms [using Eq. (18) for single-qubit rotations], which implies an upper limit for the pulse rate of around 15 pulses per millisecond. In our numerical examples, we find that significantly lower pulse rates are sufficient to restore the final fidelity to noise-free levels.

For comparison, we also include results from simulations with static (time-independent) disorder in Fig. 4, marked by




FIG. 5. Impact of noise peak frequency on dynamical decoupling performance. As Fig. 4, but using different noise spectra with a single Lorentzian peak centered at varying frequencies, all with the same amplitude of 750 Hz. The solid line gives a comparison with the median value of annealing sweeps with constant disorder [$\delta h^z(t) = \delta h^z$] and the same amplitude of 750 Hz. The horizontal dashed line shows the achievable fidelity of the ideal noiseless protocol. Data for each boxplot correspond to 50 random independent noise realizations (see Appendix F). Similar to Fig. 4, the final fidelity converges to the noiseless case for sufficient rates of dynamical decoupling pulses. The required rates are only slightly dependent on the noise frequency.

crosses. To this end, we compute the median fidelity from 50 annealing runs where each qubit experiences a constant local field sampled from a Gaussian distribution with a standard deviation matching the corresponding noise amplitude. The resulting fidelities lie a little higher and show a similar convergence trend, indicating that time-dependent noise and static disorder can be mitigated with similar dynamical decoupling pulse rates for the parameters considered.

To examine the role of spatial correlations, we also ran simulations with *uncorrelated* noise, where each qubit is subject to an independent noise realization. The results (not shown) are qualitatively very similar to those in Fig. 4, with convergence to ideal fidelity occurring slightly faster.

Figure 5 explores the dependence of annealing performance on the spectral shape of the noise. We fix the noise amplitude to 750 Hz (corresponding to the green boxplots from Fig. 4) and consider single Lorentzian peaks at 10, 50, 100, and 150 Hz, respectively (from darkest to lightest green). A constant disorder comparison is shown as a solid line. We find no significant qualitative sensitivity to the exact peak frequency in the considered regime: all spectra exhibit similar convergence behavior, with slightly improved performance for lower noise frequencies.

As these results show, dynamical decoupling can render the performance of quantum annealing in realistic devices robust against field fluctuations of large strength and various spectral properties. As we show in Appendixes C and D, when normalizing to the noise-free fidelity that is achievable for

a given sweep speed, MOT, cutting-stock, and Sherrington-Kirkpatrick model problems of different sizes follow very similar curves, showing the generality of this protocol. The performance can be improved with enhanced two-body interaction strengths, achievable in the MAGIC setup with existing technology. As mentioned previously, magnetic field gradients on the order of 150 T/m are feasible [33,88], enabling coupling strengths up to $J \approx 1650$ Hz [37]. In such scenarios, the annealing time required to achieve fidelities comparable to those in our simulations would reduce from 100 ms to approximately 1.5 ms. Moreover, the ratio of noise amplitude to coupling strength would become more favorable, enabling high-fidelity sweeps with even fewer dynamical decoupling pulses.

D. Universal scaling behavior

In the previous section, we demonstrated that quantum annealing remains effective under local field fluctuations when dynamical decoupling pulses are applied at regular intervals during the protocol. As shown in Fig. 4, the final fidelity converges to the noiseless limit as the dynamical decoupling pulse rate increases, with the required pulse rate depending on the noise amplitude and to a lesser degree the noise spectrum. In this section, we reveal a unifying trend across these cases by identifying a generalized scaling parameter that causes all fidelity curves to collapse onto a single universal curve.




FIG. 6. Fidelity of the final state, compared to the state encoding the problem solution, at the end of a quantum annealing sweep versus a universal scaling parameter $(\Delta t \times \sigma(\delta h^z)^c)^{-1}$, which depends on the time Δt between two dynamical decoupling pulses, the standard deviation of local field fluctuations $\sigma(\delta h^z)$ (correlated noise), and a noise-dependent exponent c . (a) For constant disorder $[\delta h^z(t) = \delta h^z]$, an excellent data collapse is obtained for $c = 1$, in agreement with analytic derivations (see Sec. VIA). The gray line in (a) shows an exponential fit considering all data points for the different noise amplitudes together. (b) For time-dependent noise spectra with two Lorentzian peaks at 50 and 150 Hz, an approximate data collapse can still be obtained for a (fitted) value of $c \approx 0.65$. In both cases, we show the mean value of 50 random independent noise realizations.

To this end, we build upon the analytical treatment of the time-evolution operator of the noisy annealing protocol with dynamical decoupling pulses derived in Sec. VIA. There, the leading-order corrections due to local field fluctuations appear as perturbations to the driving Hamiltonian proportional to $\delta h_i^z \times \Delta t$ [see Eq. (30)], where δh_i^z is the local field fluctuation at the position of qubit i and Δt is the time interval between two consecutive dynamical decoupling pulses (corresponding to the inverse of the pulse rates on the horizontal axis in Figs. 4 and 5). This implies that the deviation from the ideal (noiseless) annealing performance should be governed by this product.

To test this prediction, we plot in Fig. 6(a) the mean final fidelities for the case of static correlated disorder, now as a function of the inverse parameter $(\sigma(\delta h^z) \times \Delta t)^{-1}$, where

$\sigma(\delta h^z)$ denotes the standard deviation of the static local field disorder. This rescaling causes the data from all four disorder amplitudes to collapse onto a single, smooth curve. The collapse confirms our analytical prediction and demonstrates that the key quantity controlling fidelity loss is indeed the product of noise amplitude and time between dynamical decoupling pulses. The gray line in Fig. 6(a) shows an exponential fit to the combined data.

However, the perturbative argument in Sec. VIA assumes the field fluctuations δh^z remain constant between dynamical decoupling pulses. This assumption breaks down once the noise acquires a significant time dependence. Consistent with this, we find that applying the same rescaling to annealing sweep results with time-dependent noise fails to produce a similarly precise data collapse. To recover a universal description, we generalize the scaling parameter to the form $\sigma(\delta h_i^z)^c \times \Delta t$ with c being a fitting exponent between zero and one.

In Fig. 6(b), we plot the mean fidelities for time-dependent correlated noise, corresponding to the spectrum considered in the previous section with two Lorentzian peaks (at 50 and 150 Hz), as a function of this generalized parameter. For $c \approx 0.65$ we achieve an approximate collapse of the data, suggesting a sublinear dependence of the fidelity degradation on noise amplitude in the considered noise regime.

We further analyzed noise spectra with single Lorentzian peaks at various frequencies: the optimal value of c consistently decreases as the dominant noise frequency increases.

It will be interesting to understand the microscopic origin of this generalized scaling behavior. A rigorous analytical derivation of the generalized parameter and its dependence on the noise spectral density could provide deeper insight into the mechanisms of dynamical decoupling in time-dependent environments and represents a promising avenue for future research.

VII. CONCLUSIONS

To summarize, we have discussed how local field noise in quantum annealing can be robustly mitigated by interleaving periodic dynamical decoupling pulses into the annealing schedule. Starting from a minimal, five-qubit multiple-object-tracking QUBO instance, chosen to capture industrially relevant optimization challenges while remaining numerically tractable, we demonstrated that current small-scale devices can already encode and solve practical problems as a proof of principle. Focusing on a microwave-controlled trapped-ion platform with magnetic-gradient-induced coupling (although our findings generalize straightforwardly to other architectures and optimization problems), we explained how such a platform could be especially suitable for quantum annealing applications due to the native all-to-all two-body interactions, which could potentially be tuned to directly encode a QUBO problem. We argued that two-body coupling-strength fluctuations play a negligible role under realistic experimental conditions and identified external magnetic field fluctuations as the dominant error channel on such platforms.

As an error mitigation technique for such local field noise, we explained in detail the periodic application of dynamical decoupling pulses implementing global spin flips during an annealing protocol. We discussed how the annealing protocol

needs to be modified to accommodate such an error mitigation scheme, and analytically calculated the time-evolution operator of the annealing sweep with noise and dynamical decoupling pulses.

We performed numerical simulations demonstrating the success of this protocol for a realistic noise spectrum and different noise amplitudes. We found that around 2.5 dynamical decoupling pulses per millisecond are already enough to efficiently counteract the noise, for all considered regimes with noise amplitudes up to 1000 Hz (chosen as a worst-case scenario) and when considering already-demonstrated experimental two-body coupling strengths of 26 Hz. Furthermore, we verified that the exact form of the noise spectrum has little effect on the required number of dynamical decoupling pulses (using different spectra with single Lorentzian peaks located at different frequencies up to 150 Hz). Moreover, we found that different system sizes and both MOT and cutting-stock problems, as well as Sherrington-Kirkpatrick model instances, show quantitatively similar behavior when normalizing the fidelity to the respective noiseless value. We commented on the improvements to be expected for near-future devices with much higher interaction strengths, which enable much faster annealing sweeps with fewer dynamical decoupling pulses to obtain a similar performance.

Finally, we identified a universal scaling behavior for the achieved final fidelities with different noise amplitudes and dynamical decoupling pulse rates. For time-independent local field disorder, we found that the achieved fidelity is determined by the product of noise amplitude and the time interval between two dynamical decoupling pulses, matching our analytical calculations. For time-dependent noise, we find a dependence on a generalized parameter where we observe a sublinear dependence on the noise amplitude, with an exponent that decreases for noise spectra containing larger frequencies. This universal scaling provides a compact figure of merit for designing dynamical decoupling sequences across diverse noise environments.

Our work opens several immediate research directions. For example, it will be important to investigate additional noise sources, more sophisticated dynamical decoupling protocols, and a wider range of optimization problems. Motivated by the universal behavior suggested by our analytical and numerical results, it will also be interesting to understand whether a product of error strength and decoupling pulse rate can serve as a general guiding principle to design practical annealing protocols. Furthermore, incorporating dynamical decoupling sequences within quantum optimization algorithms provides a means to experimentally tune the impact of noise, allowing to controllably study the performance of quantum optimization and the role of quantum effects therein [89,90], under various degrees of decoherence [73]. Finally, we anticipate that our results will motivate near-term experimental quantum annealing demonstrations on trapped-ion platforms, and that they will inform error mitigation strategies for quantum annealing algorithms in other architectures.

ACKNOWLEDGMENTS

We thank Kevin T. Geier, Patrick Huber, Florian Köppen, and Christof Wunderlich for helpful discussions. The work

reported in this publication is based on a project that was funded by the German Federal Ministry for Research, Technology and Space under the funding Reference No. 13N16437 (project MAGICapp). The authors are solely responsible for the content of this publication. We acknowledge the CINECA award under the ISCRA initiative, for the availability of high-performance computing resources and support. This work has benefited from Q@TN, the joint laboratory between University of Trento, Fondazione Bruno Kessler (FBK), National Institute for Nuclear Physics (INFN), and National Research Council (CNR). We acknowledge support by Provincia Autonoma di Trento.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available because they contain commercially sensitive information. The data are available from the authors upon reasonable request.

APPENDIX A: ANNEALING PROTOCOL WITHOUT ADDITIONAL ANCILLA QUBIT

In Sec. II, we presented a quantum annealing protocol that encodes QUBO problems exclusively through two-body interactions. This approach eliminates local field terms by encoding them as interactions with an additional ancilla qubit [see Eq. (4)]. The key advantage of this encoding scheme is that it decouples the problem representation from local field noise, enabling effective noise suppression through dynamical decoupling control pulses (Sec. VI). However, the encoding requires one additional ancilla qubit for any problem instance. Given the limited qubit resources available in current quantum hardware, it may be more practical to retain the original QUBO formulation that directly incorporates local field terms in the cost Hamiltonian. In this Appendix, we demonstrate how the annealing sweep with dynamical decoupling pulses can be adapted to mitigate unwanted local field noise while preserving the desired local fields that encode the optimization problem.

This alternative annealing protocol Hamiltonian with local field noise $\delta h_i^z(t)$ is defined as

$$H = C(t)H_{\text{driving}} + \sum_{ij} J_{ij} \sigma_i^z \sigma_j^z + \sum_i (h_i^z + \delta h_i^z(t)) \sigma_i^z. \quad (\text{A1})$$

Implementing the same dynamical decoupling scheme as in Sec. VI would cancel both the noise $\delta h_i^z(t)$ and the local field terms h_i^z encoding the problem. However, this can be easily avoided by applying a sign change to the problem-encoding local fields ($h_i^z \rightarrow -h_i^z$), simultaneously with each application of a dynamical decoupling control pulse. This additional sign change cancels the one due to the global dynamical decoupling spin flip, such that only the noise terms $\delta h_i^z(t)$ get effectively canceled out.

The additional sign flips are easily implementable in experiments. For example, when the problem-encoding local fields are implemented via detuning of a microwave signal from resonance of the corresponding ion transition (see, for example, the discussion in Ref. [37]), it suffices to quickly




FIG. 7. Final noiseless fidelities achieved for different total linear annealing sweep durations for the minimal MOT QUBO problem (Sec. III B). We compare the performance of the annealing protocol considered in the main text [Eq. (1)], encoding the problem exclusively into the two-body interactions using five qubits overall (orange), versus the direct encoding utilizing local fields with four qubits overall (blue). The annealing time is given in units of the inverse characteristic energy scale J . The vertical dashed red line denotes the annealing time considered for all numerical simulations in the main text ($2.6J^{-1}$).

switch the sign of the detuning simultaneously with the action of the dynamical decoupling pulse. Other implementations rely on virtual Z rotations [79] or single-qubit gates in a Trotterized annealing protocol; in both of which one can just choose opposite rotation angles.

We now compare the performance of this alternative protocol against the method employed in the main text when applied to the minimal MOT problem (Sec. III B). The original protocol requires five qubits, including an additional ancilla qubit for the encoding, while the alternative approach uses only four qubits with local field terms.

Figure 7 presents the final noiseless fidelities achieved for both protocols across varying annealing sweep durations, expressed in units of the inverse characteristic energy scale J . In the main text we considered a sweep duration of $2.6J^{-1}$ for numerical calculations, corresponding to 100 ms when assuming $J = 26$ Hz (indicated by the vertical dashed red line). At this annealing time, both protocols achieve nearly identical final fidelities ($\approx 84\%$). More generally, the final fidelities exhibit oscillatory behavior with annealing time, resulting in alternating advantages between the two protocols at different sweep durations, but without significant qualitative differences.

We also investigated whether the dynamical decoupling requirements differ between the two protocols by examining the number of control pulses needed to mitigate the noise of a given amplitude. Simulations of the four-qubit protocol yield results very similar to those obtained with the five-qubit approach (Fig. 4). In both cases, approximately 2.5 control pulses per millisecond suffice to achieve final fidelities above 70% for the majority of annealing runs. This consistency in dynamical decoupling requirements suggests that the choice between protocols can be made primarily based on qubit resource constraints rather than noise mitigation and performance considerations.

$$H_{\text{cost}}/J = \begin{pmatrix} 0 & -0.96 & -0.55 & -0.43 & -0.99 & -0.84 & -0.89 & -0.38 & -0.97 \\ 0 & 0 & 1 & 1 & 0 & -0.19 & 0 & -0.04 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & -0.19 & 0 & -0.04 \\ 0 & 0 & 0 & 0 & 1 & -0.05 & 0 & -0.18 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -0.05 & 1 & -0.18 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

FIG. 8. Cost Hamiltonian matrix H_{cost} for the MOT QUBO problem corresponding to an added third frame in Fig. 1. Encoding this problem only into two-body interactions requires nine qubits.

APPENDIX B: MOT PROBLEM INSTANCE WITH NINE QUBITS

The main text employed a minimal five-qubit MOT instance as a reference problem (Sec. III B) to evaluate our dynamical decoupling noise mitigation protocol (Sec. VI). To demonstrate the scalability of this approach, we now extend the analysis to the next largest MOT QUBO instance, which requires nine qubits in our encoding (Sec. II). We present the explicit QUBO matrix structure and show that the number of dynamical decoupling pulses required for noise suppression remains similar to that for the smaller instance.

Problem construction. The nine-qubit instance is generated by adding a third frame to the video sequence used for the minimal problem (Fig. 1). Following the same pattern as before, we consider the first, third, and fifth frames of the video, which yield a nine-qubit QUBO formulation when encoded exclusively through two-body interactions. In the construction of the QUBO problem we rescale the off-diagonal elements of the objective by a factor of 0.5, which corresponds to giving a larger weight to the overlap with the detections in the first fixed frame encoded in the diagonal elements [Eq. (11)]. For this particular instance we observe slightly increased annealing performance with this rescaling and the remaining chosen parameters (see below). Figure 8 displays the resulting 9×9 QUBO matrix, where entries equal to one correspond to constraint terms and the remaining entries encode the objective function.

Parameter optimization. The optimal penalty factor λ for the QUBO problem [see Eq. (12)] varies between problem instances and must be determined individually. For this nine-qubit problem, we select $\lambda = 3$, which provides good annealing performance at the reference time of $2.6J^{-1}$ used throughout the main text. Similarly, the optimal transverse driving strength h_x [Eq. (3)] differs from the minimal problem: while the five-qubit instance performed well with $h_x/J = 3$, we use $h_x/J = 2$ here to achieve a final noiseless fidelity of approximately 0.74.

Figure 9 shows the final noiseless fidelities for various combinations of driving strength and penalty factor across different sweep durations (with $2.6J^{-1}$ marked by the vertical dashed red line). The parameter landscape reveals that optimal parameter choices can vary with annealing time, and no single combination remains optimal across all durations. In experimental practice, multiple annealing runs with varied parameters would be performed, potentially guided by heuristic methods for penalty factor estimation [72].




FIG. 9. Final fidelities achieved for different total linear annealing sweep durations for the MOT QUBO problem with nine qubits. We compare the (noiseless) performance of the annealing protocol when choosing different hyperparameters, namely, the strength of the penalty term λ in the QUBO problem [Eq. (12)] as well as the strength of the driver Hamiltonian h_x/J . The vertical dashed red line denotes the sweep time $2.6J^{-1}$ used in our numerical simulations in the main text. In an intermediate regime of medium-slow sweeps, it can be beneficial to carefully choose these hyperparameters, either through a variation of test runs or heuristic methods for finding optimal penalty parameters.

For the present analysis, we employ $h_x/J = 2$ and $\lambda = 3$ without further optimization, as these parameters yield sufficiently good results for our primary objective: demonstrating the efficacy of the dynamical decoupling protocol for larger problem instances.

Dynamical decoupling. Figure 10 presents our numerical results for the nine-qubit problem under the same correlated local field noise conditions discussed in Sec. VI. The time-dependent noise is sampled from a spectrum with two dominant frequency peaks at 50 and 150 Hz, analogous with the five-qubit analysis (Fig. 4). We assume the same energy scale of $J = 26$ Hz to enable direct comparison with the main text results, though we note that, in trapped-ion implementations, coupling strengths typically decrease for longer ion chains, leading to lower effective energy scales for larger problems [37].

The nine-qubit problem behaves similarly to its five-qubit counterpart. Final fidelities converge toward the ideal noiseless performance (dashed gray line) as the number of control pulses increases. Importantly, pulse rates of approximately 2.5 pulses per millisecond again prove sufficient to achieve good performance, with final fidelities exceeding 60% for the majority of annealing runs across all considered noise amplitudes. These results demonstrate that our dynamical decoupling protocol is robust and shows little dependence on




FIG. 10. Fidelity of the final state, compared to the state encoding the problem solution, after a quantum annealing sweep versus the applied dynamical decoupling pulse rate during the protocol (see Sec. II), using the MOT test problem with nine qubits (see Fig. 8) and an annealing time of $2.6/J$ (corresponding to a sweep duration of 100 ms for an energy scale of $J = 26$ Hz). Boxplots compare results for correlated noise using identical noise spectra with two Lorentzian peaks at 50 and 150 Hz (see Appendix F) but different amplitudes of 250 Hz (blue), 500 Hz (red), 750 Hz (green), and 1000 Hz (orange). The crosses in each boxplot give a comparison with the median value of annealing sweeps subject to constant disorder [$\delta h^z(t) = \delta h^z$] of the same respective amplitude. The horizontal dashed gray line shows the achievable fidelity of the ideal noiseless protocol (≈ 0.74). Data for each boxplot correspond to 50 random independent noise realizations (see Appendix F). We observe that the achieved final fidelity converges to the noiseless case for a sufficient rate of dynamical decoupling pulses, with higher noise amplitudes requiring higher rates.

the specific problem instance and system size (further supported by our results on the cutting-stock problem and the Sherrington-Kirkpatrick model in Appendixes C and D). The pulse requirements depend primarily on the noise amplitude and, to a lesser extent, the noise spectrum characteristics. This scaling behavior aligns with our analytical predictions from Sec. VIA, where corrections to the ideal annealing protocol are shown to be independent of the specific coupling strengths that encode the optimization problem.

Scaling to larger instances. It is interesting to study how the number of required qubits increases when considering a larger number of frames and detections in the MOT problem. The number of binary variables for larger object tracking problems in this formulation can be obtained from Eq. (5). If we have a number of T tracks on a sequence of F frames, each with the same number of detections, D_f , that belong to the same persons, one would need FD_fT binary variables; i.e., the number of qubits needed scales linearly with the problem parameters.

For larger sequences, it is also beneficial to have a maximal frame gap, so that we stop calculating similarity scores for frames that are far apart from each other. An advantage of this method is that if a track for a person terminates and at some frame, which is farther away than the maximal frame gap, a new person enters the frame, then one can use the same ID [50]. The two persons are then separated in a postprocessing phase. Some benchmark results showing how classical solvers perform on the more challenging MOT20 dataset are given in Ref. [91].

APPENDIX C: CUTTING-STOCK PROBLEM

In Sec. III and Appendix B, we discussed the multiple-object-tracking problem, for which we discussed the QUBO formulation and the construction of minimal instances suitable as test problems. In this Appendix, we give a brief analogous overview of the industrially relevant cutting-stock problem [92]. The one-dimensional (1D) cutting-stock problem asks to minimize the waste when cutting a number d_i of bar stock with required lengths l_i from a number N of larger bars with a fixed length L . We want to optimize the assignment $x_{i,k} \in \mathbb{Z}$ of the pieces to be cut out for the demand type $i \in \{1, \dots, m\}$ to the different larger bar stock $\{1, \dots, Q\}$ as

$$\arg \max_{x \in \mathbb{Z}^N} \sum_{i=1}^m \sum_{k=1}^Q l_i x_{i,k}, \quad (\text{C1})$$

with the constraints

$$\sum_{i=1}^m l_i x_{i,k} \leq L \quad \forall k \in \{1, \dots, Q\}, \quad (\text{C2})$$

$$\sum_{k=1}^Q x_{i,k} \leq d_i \quad \forall i \in \{1, \dots, m\}. \quad (\text{C3})$$

In this description, all feasible assignments have the same objective energy. In order to have binary variables for the assignment x , we ignore that pieces have a common length and instead model every required piece with an entry $i \in \{1, \dots, \tilde{m} := \sum_j d_j\}$ so that every piece is demanded only

once and $\tilde{d} := 1$. The problem is closely related to the binary multiple-knapsack problem, for which QUBO formulations have already been proposed and studied [93–95].

Reformulation of inequalities. An important step for the reformulation of Eq. (C1) into a QUBO problem is to turn all inequality constraints into equality constraints, which can be achieved by using slack variables s :

$$Ax - b \leq 0 \Leftrightarrow Ax - b = s. \quad (\text{C4})$$

The slack variables are nonpositive and are restricted to values $s_i \in \mathcal{S}_i := \{(Ax - b)_i | x \in \mathcal{F}\}$, where $\mathcal{F}$ is the feasible set for x . One usually has to implement integer s_i variables via a binary representation

$$s_i = - \sum_{k=0}^{B_{\max}} 2^k s_i^{(k)}, \quad (\text{C5})$$

where $s_i^{(k)} \in \{0, 1\}$ and $B_{\max}$ is the minimal number of binary bits needed to represent the possible range of slack variables s_i . Note that it does not invalidate the method if the binary decomposition can attain values that are even smaller than the values present in $\mathcal{S}_i$.

After this process the inequality is effectively reformulated as an equality, which can be integrated in the QUBO problem via penalty terms as in Eq. (12). Lower bounds for the penalty parameters of the multiple-knapsack problem have already been computed, which can be directly applied to our setting [96]. Alternatives to implement inequalities without additional slack variables are, e.g., the unbalanced penalization formulations [97] or by the use of an ancillary qudit in discretized quantum-optimization schemes [98].

The slack variables contribute the most to the scaling in dimension. For each inequality, one estimates the largest value the slack variable can attain and then chooses the corresponding number of bits in the binary representation. Due to the inequality given by Eq. (C2), we would need $1 + \lfloor \log_2(L) \rfloor$ slack variables for every bar stock and, because of the inequality given by Eq. (C3), $1 + \lfloor \log_2(d_i) \rfloor$ slack variables for every demanded piece i . However, the later inequalities can also be incorporated without slack variables as we will show now.

For problem instances with multiple larger bars, the inequalities in Eq. (C3) become nontrivial. In the case with two bars, they read

$$x_{1,1} + x_{1,2} \leq 1, \quad (\text{C6})$$

$$x_{2,1} + x_{2,2} \leq 1. \quad (\text{C7})$$

Fortunately, one does not need to use slack variables. Instead, one can just add terms $\lambda(x_{1,1}x_{1,2} + x_{2,1}x_{2,2})$ to the QUBO. This is also described as Transformation 2 in Ref. [99]. The same procedure can also be generalized for larger numbers of bars by adding

$$\sum_{i=1}^m \left(\left(\sum_{k=1}^Q x_{i,k} \right) \left(\sum_{k=1}^Q x_{i,k} \right) - \sum_{k=1}^Q x_{i,k}^2 \right). \quad (\text{C8})$$

Minimal cutting-stock problems. For benchmarking purposes, it is useful to construct 1D cutting-stock problems that have a low qubit requirement. We explicitly construct two




FIG. 11. Illustration of the small cutting-stock toy problems that we consider for our quantum annealing simulations (see Fig. 12): problem instances which require (a) four and (b) five qubits, respectively. When encoding the problem into two-body interactions only, an additional ancilla qubit is required [Eq. (4)]. In both cases, the trivial optimal solution is that every required piece will be cut out from the larger bar.

trivial cutting-stock problems requiring four and five binary variables, respectively, to test our dynamical decoupling protocol. The problems are illustrated in Fig. 11. The smaller QUBO problem is constructed to answer the following simple question: Given a single bar of length $L = 3$ and required pieces of lengths $l_1 = 1$ and $l_2 = 1$ with demand $d_1 = d_2 = 1$, what should be cut out to maximize the length of cut-out pieces? Now the binary variable $x_1 := x_{1,1}$ indicates if the first piece will be cut out and $x_2 := x_{2,1}$ indicates if the second piece will be cut out. The inequalities defined in Eq. (C3) are trivially fulfilled, since there is only one larger bar stock. There is only one length constraint [Eq. (C2)]. It reads

$$1x_1 + 1x_2 \leq 3. \quad (\text{C9})$$

At the next step, we look at the smallest value that $1x_1 + 1x_2 - 3$ can attain. This is -3 if both binary variables are zero. Therefore, we will need $1 + \lfloor \log_2(3) \rfloor = 2$ slack variables. Using binary representation, we obtain the linear equality

$$(1 \ 1 \ 1 \ 2) \begin{pmatrix} x_1 \\ x_2 \\ s_1 \\ s_2 \end{pmatrix} = (3). \quad (\text{C10})$$

The matrix incorporating the constraints in the QUBO problem can now be obtained as $A^T A - \text{diag}(A^T b + b^T A)$, if the equality constraint in Eq. (C10) is denoted as $Ax = b$. This

calculation results in the matrix

$$Q_{\text{const}} = \lambda \begin{pmatrix} -5 & 1 & 1 & 2 \\ 1 & -5 & 1 & 2 \\ 1 & 1 & -5 & 2 \\ 2 & 2 & 2 & -8 \end{pmatrix}. \quad (\text{C11})$$

The objective part has to be included as a matrix with diagonal $W = (-1, -1, 0, 0)$ in the QUBO matrix. Converting the QUBO with $\lambda = 1$ (chosen for good annealing performance) to an Ising problem and including an ancilla qubit to avoid bias terms yields the coupling matrix

$$H_{\text{cost}}/J = \begin{pmatrix} 0 & -1 & -1 & -0.5 & -1 \\ 0 & 0 & 0.5 & 0.5 & 1 \\ 0 & 0 & 0 & 0.5 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}. \quad (\text{C12})$$

The slightly larger problem constructed using five binary variables answers the following question: Given a bar of length $L = 4$ and required pieces of lengths $l_1 = 2$ and $l_2 = 2$ with demand $d_1 = d_2 = 1$, what should be cut out to maximize the length of cut-out pieces? We still have two binary variables x_i , which determine if the piece i is cut out or not. The matrix describing the size constraints can be constructed in a similar fashion as above. Since we have changed the length to $L = 4$, we need $1 + \lfloor \log_2(4) \rfloor = 3$ slack variables, resulting in the equality

$$(2 \ 2 \ 1 \ 2 \ 4) \begin{pmatrix} x_1 \\ x_2 \\ s_1 \\ s_2 \\ s_3 \end{pmatrix} = (4). \quad (\text{C13})$$

In the same way as in Eq. (C11), this equality yields the contribution

$$Q_{\text{const}} = \lambda \begin{pmatrix} -12 & 4 & 2 & 4 & 8 \\ 4 & -12 & 2 & 4 & 8 \\ 2 & 2 & -7 & 2 & 4 \\ 4 & 4 & 2 & -12 & 8 \\ 8 & 8 & 4 & 8 & -16 \end{pmatrix} \quad (\text{C14})$$

to the QUBO matrix. The objective part is a diagonal matrix with entries $W = (-2, -2, 0, 0, 0)$. Finally, choosing the penalty strength $\lambda = 1$, converting to an Ising Problem and adding an ancilla variable yields the cost Hamiltonian

$$H_{\text{cost}}/J = \begin{pmatrix} 0 & 0.33 & 0.33 & 0.25 & 0.5 & 1 \\ 0 & 0 & 0.33 & 1.67 & 0.33 & 0.67 \\ 0 & 0 & 0 & 1.67 & 0.33 & 0.67 \\ 0 & 0 & 0 & 0 & 1.67 & 0.33 \\ 0 & 0 & 0 & 0 & 0 & 0.67 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}. \quad (\text{C15})$$

Dynamical decoupling. Figure 12 presents numerical quantum annealing results for cutting-stock instances with five and six qubits, and MOT instances with five and nine qubits, all subject to identical correlated local field noise. The time-dependent noise is sampled from a spectrum characterized by




FIG. 12. Annealing results for MOT instances with five and nine qubits (see Sec. III and Appendix B), and for cutting-stock instances with five and six qubits (see Fig. 11), respectively. The annealing time is set to $2.6/J$ for the MOT problems and $26/J$ for the cutting-stock problems. A physical sweep time of 100 ms is assumed for all annealing runs, corresponding to an energy scale of $J = 26$ Hz for the MOT problems and $J = 260$ Hz for the cutting-stock problems. Each data point represents the mean over 50 independent correlated-noise realizations, all sampled from the same noise spectrum as in Figs. 4 and 10, with a fixed amplitude of 750 Hz. Error bars denote the standard deviation. (a) Final-state fidelity as a function of the rate of applied dynamical decoupling pulses. Dashed lines indicate the corresponding noiseless final fidelities. (b) Fidelities normalized to their respective noiseless values. Within error bars, all data collapse onto a single curve, showing the generality of the dynamical decoupling procedure.

two dominant Lorentzian peaks at 50 and 150 Hz, each with a fixed amplitude of 750 Hz. The MOT data correspond to the mean values of the results shown in Figs. 4 and 10.

For the cutting-stock problems, we generally observe much smaller minimum energy gaps during the annealing sweep than for the MOT instances of similar size (this might be due to the presence of the slack variables; a strong improvement in performance was observed, e.g., in Ref. [98] when using a slack-variable-free implementation). To still achieve reasonable final fidelities and clearly demonstrate the efficacy of our dynamical decoupling protocol, we set the dimensionless annealing time to $26/J$ for the cutting-stock instances, while retaining $2.6/J$ for the MOT cases. To enable a direct comparison between both problem types, we assume an energy scale of $J = 260$ Hz for the cutting-stock problems and $J = 26$ Hz for the MOT problems. These choices yield the same physical annealing duration of 100 ms in both cases,

and therefore identical rates of dynamical decoupling pulses per millisecond.

The results in Fig. 12(a) show that the final-state fidelities approach the noiseless limits (dashed curves) as the rate of dynamical decoupling pulses increases. In particular, pulse rates of approximately 2.5 pulses per millisecond are sufficient to achieve good relative performance, with the final fidelity at the end of the sweep depending on the specific problem instance.

Figure 12(b) presents the same data normalized to the corresponding noiseless fidelities, revealing a qualitatively and quantitatively similar behavior. This collapse demonstrates that the applied dynamical decoupling protocol is robust and largely insensitive to the specific optimization problem.

APPENDIX D: SHERRINGTON-KIRKPATRICK MODEL

While the main text focused on application-motivated optimization problems, here we demonstrate that dynamical decoupling (DD) equally mitigates local field noise in problems without internal structure. As a representative example, we consider the Sherrington-Kirkpatrick (SK) spin-glass model with local fields [51,52], and present additional numerical results supporting that the DD pulse rate required to suppress noise is largely independent of the system size.

As the cost function in the annealing protocol, we use the SK Hamiltonian

$$H_{\text{SK}} = \sum_{i \neq j} J_{ij}^{\text{SK}} \sigma_i^z \sigma_j^z + \sum_i h_i^{\text{SK}} \sigma_i^z, \quad (\text{D1})$$

where J_{ij}^{SK} and h_i^{SK} are independently drawn from Gaussian distributions with zero mean and variances $\sigma^2 = 1$ and $\sigma^2 = 0.3$, respectively. As in the main text, local fields are replaced by two-body terms through the introduction of a single ancilla qubit, ensuring that all problem information is contained in the interaction terms only (cf. Sec. II). This allows local field noise to be mitigated via DD without altering the encoded cost function.

Figure 13 summarizes our numerical results. Analogous to the results in the main text, we consider a noise spectrum with two Lorentzian frequency peaks at 50 and 150 Hz [Eq. (F1)], and a noise amplitude of 750 Hz. We consider the unitless annealing time of $2.6/J$, which corresponds to a physical time of 100 ms, when assuming a characteristic energy scale of $J = 26$ Hz. We vary the rate of applied dynamical decoupling pulses and calculate the ratio of the fidelity of the final state after the annealing sweep to the fidelity of the noiseless sweep for each individual problem instance. We consider system sizes ranging from $N = 4$ to $N = 12$. For each system size and pulse rate, we simulate 50 random SK instances, using a separate noise realization for every run. Data points show the mean ratio of fidelities; error bars denote the standard deviation across problem instances.

At low pulse rates, fidelities are broadly distributed and, for some instances, may even exceed the noiseless reference. This behavior is expected: due to random spectra and varying minimal gaps, certain instances may by chance perform better under the specific noise realization. We also note that smaller system sizes will on average achieve a better final fidelity ratio for small DD pulse rates. This is simply due to the fact that the dimension of the system is small enough for the solution




FIG. 13. Annealing results for Sherrington-Kirkpatrick (SK) model instances of varying sizes. Annealing time is set to $2.6/J$, corresponding to a physical time of 100 ms when assuming an energy scale of $J = 26$ Hz. Each data point corresponds to the mean of 50 random SK instances with different respective correlated-noise realizations. The noise is sampled from the same spectrum as in Figs. 4 and 10, defined by Eq. (F1), with a fixed noise amplitude of 750 Hz. Error bars denote the standard deviation. The final fidelity achieved after an annealing sweep for each individual problem instance is normalized with respect to the fidelity that instance would achieve without any noise.

to have a finite amplitude with high likelihood. Nevertheless, for all system sizes, we observe a clear convergence to the noiseless fidelity when the DD rate exceeds approximately three pulses per millisecond. The onset of convergence is essentially identical across all system sizes, indicating no significant increase in required DD pulse rate with system size, consistent with our results on the multiple-object-tracking and cutting-stock problems (see Sec. VI and Appendixes B and C).

This observation also matches our analytical treatment in Sec. VI, which predicts that to leading order the noise enters only through the product of noise amplitude and time interval between consecutive DD pulses, independent of the number of qubits. Our results thus establish that dynamical decoupling is robust with respect to the choice of cost Hamiltonian and system size and continues to suppress local field noise effectively even for unstructured spin-glass problems such as the SK model.

APPENDIX E: MODULATING COUPLING STRENGTHS VIA DYNAMICAL DECOUPLING

In Sec. VI, we discussed how dynamical decoupling control pulses in the form of global spin flips can effectively cancel local field disorder terms in the quantum annealing Hamiltonian [Eq. (24)]. In this Appendix we explain how a similar scheme of periodically applied global spin flips can also be used to modulate two-body coupling strengths during an annealing protocol effectively [35,100]. This scheme can be potentially utilized to encode QUBO problems into the native all-to-all couplings of the MAGIC setup given by Eq. (15).

We again consider the time-evolution operator of the annealing Hamiltonian during a small time interval $\Delta t = t_+ - t_-$ with two applied dynamical decoupling control pulses, the first at a time t_0 ($t_- < t_0 < t_+$) and the second at time t_+ at the end of the interval. In contrast to the calculation in Sec. VI, we now consider time intervals before and after the first pulse that can be of different lengths, i.e., $\Delta t_- = t_0 - t_- \neq t_+ - t_0 = \Delta t_+$. Before and after the control pulses, the system evolves with the usual annealing Hamiltonian composed of driving and cost Hamiltonians, where we again assume that the problem is fully encoded into the two-body interaction terms. For now, we consider the ideal case without any noise. The time-evolution operator during the first part of the small time interval (from time t_- to t_0) can then be written as

$$U_{t_- \rightarrow t_0} \approx \exp \left(-i \sum_{ij} J_{ij} \sigma_i^z \sigma_j^z \Delta t_- - iC(\bar{t}_-) H_{\text{driving}} \Delta t_- \right), \quad (\text{E1})$$

where we only keep terms of first order in Δt_- and define the intermediate time $\bar{t}_- = (t_- + t_0)/2$.

To modulate the coupling strengths between specific sets of ions, the dynamical decoupling control pulses are no longer global spin flips, but only flip a subset of ions. For simplicity and without loss of generality we assume that each pulse only flips a single qubit (σ_k^x). The time evolution in between the two control pulses ($U_{t_0 \rightarrow t_+}$) has the analogous form of Eq. (E1), and the evolution over the entire time interval can be calculated to first order as

$$\begin{aligned} \sigma_k^x U_{t_0 \rightarrow t_+} \sigma_k^x U_{t_- \rightarrow t_0} \approx \exp \left(-i \sum_{i,j \neq k} J_{ij} \sigma_i^z \sigma_j^z \Delta t \right. \\ \left. - i \sum_{i \neq k} J_{ik} \sigma_i^z \sigma_k^z (\Delta t_- - \Delta t_+) \right. \\ \left. - iC(\bar{t}) H_{\text{driving}} \Delta t \right), \quad (\text{E2}) \end{aligned}$$

with $\bar{t} = (t_+ + t_-)/2$. All two-body interactions with qubit k can in this way be uniformly modified by changing the relative lengths of time intervals Δt_- and Δt_+ within the overall time interval $\Delta t = \Delta t_- + \Delta t_+$.

Generalizing this setting to include local terms δh_i^z , we would simultaneously modify the strength of the local field experienced by ion k : $\delta h_k^z \rightarrow \delta h_k^z (\Delta t_- - \Delta t_+)/\Delta t$. If these local fields encode the optimization problem to be solved (as in the protocol discussed in Appendix A), this modification has to be taken into account and the field strength adjusted accordingly. If, instead, the local fields are undesired disorder terms, they can be effectively mitigated in the same manner as discussed in Sec. VI. To that end, one can interleave both the global spin flips and the spin flips acting on a subset of qubits. For example, defining $U_{\Delta t_1} = \sigma_k^x U_{t_1 \rightarrow t_{1,+}} \sigma_k^x U_{t_{1,-} \rightarrow t_1}$ ($\Delta t_1 = t_{1,+} - t_{1,-}$) and similarly $U_{\Delta t_2}$ as two time-evolution operators following Eq. (E2), one can consider the evolution $\sigma^x U_{\Delta t_2} \sigma^x U_{\Delta t_1}$ over the total time interval $\Delta t_1 + \Delta t_2$, where $\sigma^x = \bigotimes_i \sigma_i^x$ denotes a global spin flip. By setting $\Delta t_1 = \Delta t_2$ and choosing the relative time intervals between the spin flips

appropriately, one can modulate the coupling strengths while simultaneously suppressing local field disorder terms.

APPENDIX F: NUMERICAL METHODS

In this Appendix, we explain how the noise spectra that we use in the numerical simulations in the main text (Sec. VI) are defined and local field fluctuation samples are generated. Furthermore, we discuss how the dynamical decoupling pulses are distributed throughout the numerically discretized annealing sweep.

Generating the noise samples. As discussed in Sec. IV C, realistic trapped-ion experiments typically exhibit local field noise with dominant frequencies at 50 Hz (corresponding to the electrical grid frequency) and its odd harmonics. We model this characteristic noise spectrum numerically using two Lorentzian peaks:

$$S(f) = \frac{1}{2\pi} \left(\frac{\gamma}{(f - f_1)^2 + \gamma^2} + \frac{\beta}{(f - f_3)^2 + \beta^2} \right), \quad (\text{F1})$$

where $f_1 = 50$ Hz and $f_3 = 150$ Hz represent the fundamental and third harmonic frequency, respectively. The parameters γ and β denote the corresponding half widths at half maximum. To reflect typical experimental conditions where higher-order harmonics have reduced amplitude, we set the second peak to one-third of the intensity of the first by choosing $\beta = 3\gamma$.

We generate noise samples using the approximate frequency domain method of Ref. [87]. The sampling covers $N = 50\,000$ points over 100 ms, corresponding in the main text to our assumed energy scale of $J = 26$ Hz and chosen annealing duration $2.6J^{-1}$. Following the sampling method's requirements, we evaluate the spectrum at $M = 3N$ equidistant frequencies to ensure an accurate representation of the noise spectrum.

This discretization of the spectrum necessitates sufficiently broad Lorentzian peaks, leading us to choose $\gamma = 3$ Hz. While experimental noise peaks are typically much narrower, our analysis in Fig. 5 demonstrates that the dynamical decoupling protocol exhibits robust performance across different noise frequencies. We therefore expect this approximation to have a negligible impact on our numerical results for the annealing sweeps. The samples for the noise spectra with a single Lorentzian peak (again setting $\gamma = 3$ Hz), utilized for the simulation results in Fig. 5, are generated analogously to the method outlined above.

In Fig. 14, we show an example of the time evolution of a noisy local field with 50 000 sampling points, generated for the spectrum with two Lorentzian peaks at 50 and 150 Hz and a noise amplitude of 750 Hz. The main figure shows the local field noise over the entire time evolution of 100 ms (corresponding to the time of an annealing sweep in the main text), while the inset expands the first 0.2 ms. This smaller time interval corresponds to the time between two dynamical decoupling pulses for the highest pulse rate of 5 kHz we consider in the main text (see Sec. VI) and is resolved into 100 distinct noise sampling points.

Numerical simulation timescales. Our numerical implementation involves two distinct timescales that require careful




FIG. 14. Example of a time trace of local field noise with 50 000 sampling points over 100 ms, generated from a noise spectrum with two Lorentzian peaks at 50 and 150 Hz [Eq. (F1)] with a noise amplitude of 750 Hz. The inset shows an expanded view of the first 0.2 ms (100 samples), which corresponds to the time interval between two dynamical decoupling pulses for the highest pulse rate of 5 kHz we consider in this work (see Sec. VI).

treatment. The annealing sweep is discretized into 50 000 time steps, each lasting $100 \text{ ms} / 50\,000 = 0.002 \text{ ms}$ for our chosen parameters. This fine discretization ensures a sufficient number of noise samples between dynamical decoupling pulses while approximating the continuous time-dependent fluctuations. For our highest considered pulse rate of 500 dynamical decoupling pulses per sweep, this provides 100 distinct noise samples between consecutive pulses.

Since we maintain a fixed number of noise samples (50 000) while varying the number of dynamical decoupling pulses, achieving perfectly uniform pulse spacing is not always possible. For example, with 300 pulses over the total sweep duration, the ideal spacing would be on average $50\,000 / 300 \approx 166.67$ steps per interval. We implement this by using 166 steps between pulses for the first portion of the sweep and 167 steps for the remainder, ensuring the total remains exactly 50 000 steps. Our numerical tests demonstrate that small, random variations in pulse spacing have negligible impact on final fidelity, provided the pulses remain approximately equidistant (see discussion in Sec. VIB). Alternative periodic spacing patterns such as (166, 167, 167, 166, 167, 167, ...) also yield essentially identical results to the distribution described above.

However, other systematic alternating patterns can cause significant performance degradation. For instance, with 400 pulses over 50 000 steps, uniform spacing of 125 steps between pulses recovers the ideal noiseless fidelity at the end of the sweep. In contrast, an alternating pattern of (124, 126, 124, 126, ...) results in a vanishing final fidelity. This failure occurs because the systematic alternation allows the small timing errors to accumulate coherently throughout the annealing process, ultimately destroying the effectiveness of the dynamical decoupling scheme.

- [1] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O'Brien, *Nature (London)* **464**, 45 (2010).
- [2] A. K. Fedorov, N. Gisin, S. M. Belousov, and A. I. Lvovsky, [arXiv:2203.17181](#).
- [3] D. Beck, J. Carlson, Z. Davoudi, J. Formaggio, S. Quaglioni, M. Savage, J. Barata, T. Bhattacharya, M. Bishof, I. Cloet, A. Delgado, M. DeMarco, C. Fink, A. Florio, M. Francois, D. Grabowska, S. Hoogerheide, M. Huang, K. Ikeda, M. Illa, *et al.*, [arXiv:2303.00113](#).
- [4] T. Hoeffler, T. Häner, and M. Troyer, *Commun. ACM* **66**, 82 (2023).
- [5] T. L. Scholten, C. J. Williams, D. Moody, M. Mosca, W. Hurley, W. J. Zeng, M. Troyer, and J. M. Gambetta, [arXiv:2401.16317](#).
- [6] M. Troyer, E. V. Benjamin, and A. Gevorkian, [arXiv:2403.02921](#).
- [7] G. E. Santoro, R. Martoňák, E. Tosatti, and R. Car, *Science* **295**, 2427 (2002).
- [8] R. Martoňák, G. E. Santoro, and E. Tosatti, *Phys. Rev. E* **70**, 057701 (2004).
- [9] P. Hauke, H. G. Katzgraber, W. Lechner, H. Nishimori, and W. D. Oliver, *Rep. Prog. Phys.* **83**, 054401 (2020).
- [10] A. Rajak, S. Suzuki, A. Dutta, and B. K. Chakrabarti, *Philos. Trans. R. Soc. A* **381**, 20210417 (2022).
- [11] S. Yarkoni, E. Raponi, T. Bäck, and S. Schmitt, *Rep. Prog. Phys.* **85**, 104001 (2022).
- [12] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, *et al.*, *Nature (London)* **574**, 505 (2019).
- [13] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I. Wang, S. Gustavsson, and W. D. Oliver, *Annu. Rev. Condens. Matter Phys.* **11**, 369 (2020).
- [14] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, *Nature (London)* **627**, 778 (2024).
- [15] S. Boixo, T. F. Rønnow, S. V. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, and M. Troyer, *Nat. Phys.* **10**, 218 (2014).
- [16] M. Werninghaus, D. J. Egger, F. Roy, S. Machnes, F. K. Wilhelm, and S. Filipp, *npj Quantum Inf.* **7**, 14 (2021).
- [17] L. Brodoloni, J. Vovrosh, S. Julià-Farré, A. Dauphin, and S. Pilati, *Phys. Rev. A* **112**, L051303 (2025).
- [18] D. González-Cuadra, M. Hamdan, T. V. Zache, B. Braverman, M. Kornjača, A. Lukin, S. H. Cantú, F. Liu, S.-T. Wang, A. Keesling, M. D. Lukin, P. Zoller, and A. Bylinskii, *Nature (London)* **642**, 321 (2025).
- [19] D. J. Wineland, M. Barrett, J. Britton, J. Chiaverini, B. DeMarco, W. M. Itano, B. Jelenković, C. Langer, D. Leibfried, V. Meyer, T. Rosenband, and T. Schätz, *Philos. Trans. R. Soc. London Ser. A* **361**, 1349 (2003).
- [20] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland, *Rev. Mod. Phys.* **75**, 281 (2003).
- [21] H. Haffner, C. Roos, and R. Blatt, *Phys. Rep.* **469**, 155 (2008).
- [22] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, *Appl. Phys. Rev.* **6**, 021314 (2019).
- [23] R. Srinivas, S. C. Burd, H. M. Knaack, R. T. Sutherland, A. Kwiatkowski, S. Glancy, E. Knill, D. J. Wineland, D. Leibfried, A. C. Wilson, D. T. C. Allcock, and D. H. Slichter, *Nature (London)* **597**, 209 (2021).
- [24] J.-S. Chen, E. Nielsen, M. Ebert, V. Inlek, K. Wright, V. Chaplin, A. Maksymov, E. Páez, A. Poudel, P. Maunz, and J. Gamble, *Quantum* **8**, 1516 (2024).
- [25] M. Liu, R. Shaydulin, P. Niroula, M. DeCross, S.-H. Hung, W. Y. Kon, E. Cervero-Martín, K. Chakraborty, O. Amer, S. Aaronson, A. Acharya, Y. Alexeev, K. J. Berg, S. Chakrabarti, F. J. Curchod, J. M. Dreiling, N. Erickson, C. Foltz, M. Foss-Feig, D. Hayes, *et al.*, *Nature (London)* **640**, 343 (2025).
- [26] M. Meth, J. Zhang, J. F. Haase, C. Edmunds, L. Postler, A. J. Jena, A. Steiner, L. Dellantonio, R. Blatt, P. Zoller, T. Monz, P. Schindler, C. Muschik, and M. Ringbauer, *Nat. Phys.* **21**, 570 (2025).
- [27] F. Mintert and C. Wunderlich, *Phys. Rev. Lett.* **87**, 257904 (2001).
- [28] C. Wunderlich, in *Laser Physics at the Limits* (Springer, Berlin, 2002), pp. 261–273.
- [29] M. Johanning, A. F. Varón, and C. Wunderlich, *J. Phys. B: At. Mol. Opt. Phys.* **42**, 154009 (2009).
- [30] S. Zippilli, M. Johanning, S. M. Giampaolo, C. Wunderlich, and F. Illuminati, *Phys. Rev. A* **89**, 042308 (2014).
- [31] C. Piltz, T. Sriarunothai, S. S. Ivanov, S. Wölk, and C. Wunderlich, *Sci. Adv.* **2**, e1600093 (2016).
- [32] P. Baßler, M. Zipper, C. Cedzich, M. Heinrich, P. H. Huber, M. Johanning, and M. Kliesch, *Quantum* **7**, 984 (2023).
- [33] S. Weidt, J. Randall, S. C. Webster, K. Lake, A. E. Webb, I. Cohen, T. Navickas, B. Lekitsch, A. Retzker, and W. K. Hensinger, *Phys. Rev. Lett.* **117**, 220501 (2016).
- [34] I. Arrazola, J. Casanova, J. S. Pedernales, Z.-Y. Wang, E. Solano, and M. B. Plenio, *Phys. Rev. A* **97**, 052312 (2018).
- [35] P. Huber, J. Haber, P. Barthel, J. J. García-Ripoll, E. Torrontegui, and C. Wunderlich, [arXiv:2111.08977](#).
- [36] A. D. Leu, M. F. Gely, M. A. Weber, M. C. Smith, D. P. Nadlinger, and D. M. Lucas, *Phys. Rev. Lett.* **131**, 120601 (2023).
- [37] S. Nagies, K. T. Geier, J. Akram, J. Okamoto, D. Bantounas, C. Wunderlich, M. Johanning, and P. Hauke, *Quantum Sci. Technol.* **10**, 025051 (2025).
- [38] L. Viola and S. Lloyd, *Phys. Rev. A* **58**, 2733 (1998).
- [39] K. Khodjasteh and D. A. Lidar, *2005 Quantum Electronics and Laser Science Conference, QELS-05* (IEEE, Piscataway, NJ, 2005), Vol. 1, pp. 44–46.
- [40] A. M. Souza, G. A. Álvarez, and D. Suter, *Phys. Rev. Lett.* **106**, 240501 (2011).
- [41] Z. Cai, R. Babbush, S. C. Benjamin, S. Endo, W. J. Huggins, Y. Li, J. R. McClean, and T. E. O'Brien, *Rev. Mod. Phys.* **95**, 045005 (2023).
- [42] C. Piltz, B. Scharfenberger, A. Khromova, A. F. Varón, and C. Wunderlich, *Phys. Rev. Lett.* **110**, 200501 (2013).
- [43] G. T. Genov, D. Schraft, N. V. Vitanov, and T. Halfmann, *Phys. Rev. Lett.* **118**, 133202 (2017).
- [44] C. H. Valahu, A. M. Lawrence, S. Weidt, and W. K. Hensinger, *New J. Phys.* **23**, 113012 (2021).
- [45] W. Morong, K. S. Collins, A. De, E. Stavropoulos, T. You, and C. Monroe, *PRX Quantum* **4**, 010334 (2023).
- [46] P. Barthel, P. H. Huber, J. Casanova, I. Arrazola, D. Niroomand, T. Sriarunothai, M. B. Plenio, and C. Wunderlich, *New J. Phys.* **25**, 063023 (2023).
- [47] M. Nünnerich, D. Cohen, P. Barthel, P. H. Huber, D. Niroomand, A. Retzker, and C. Wunderlich, *Phys. Rev. X* **15**, 021079 (2025).

- [48] D. A. Lidar, *Phys. Rev. Lett.* **100**, 160506 (2008).
- [49] G. Quiroz and D. A. Lidar, *Phys. Rev. A* **86**, 042333 (2012).
- [50] J.-N. Zaech, A. Liniger, M. Danelljan, D. Dai, and L. Van Gool, in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (IEEE, Piscataway, NJ, 2022), pp. 8801–8812.
- [51] D. Sherrington and S. Kirkpatrick, *Phys. Rev. Lett.* **35**, 1792 (1975).
- [52] D. Panchenko, *The Sherrington-Kirkpatrick Model* (Springer, New York, 2013).
- [53] L. Leal-Taixé, A. Milan, I. Reid, S. Roth, and K. Schindler, [arXiv:1504.01942](https://arxiv.org/abs/1504.01942).
- [54] S. Jansen, M.-B. Ruskai, and R. Seiler, *J. Math. Phys.* **48**, 102111 (2007).
- [55] D. A. Lidar, A. T. Rezakhani, and A. Hamma, *J. Math. Phys.* **50**, 102106 (2009).
- [56] M. H. S. Amin, *Phys. Rev. Lett.* **102**, 220401 (2009).
- [57] D. Cheung, P. Høyer, and N. Wiebe, *J. Phys. A: Math. Theor.* **44**, 415302 (2011).
- [58] I. Čepaitė, A. Polkovnikov, A. J. Daley, and C. W. Duncan, *PRX Quantum* **4**, 010312 (2023).
- [59] L. P. García-Pintos, M. Sahasrabudhe, and C. Arenz, *SciPost Phys.* **18**, 159 (2025).
- [60] A. Bottarelli, M. G. de Andoin, P. Chandarana, K. Paul, X. Chen, M. Sanz, and P. Hauke, *Phys. Rev. Res.* **7**, 043030 (2025).
- [61] V. Choi, [arXiv:2105.02110](https://arxiv.org/abs/2105.02110).
- [62] N. Feinstein, I. Shalashilin, S. Bose, and P. A. Warburton, *Quantum Sci. Technol.* **10**, 025011 (2025).
- [63] S. Nagies, K. T. Geier, J. Akram, D. Bantounas, M. Johanning, and P. Hauke, *Quantum Sci. Technol.* **10**, 035008 (2025).
- [64] W. Luo, J. Xing, A. Milan, X. Zhang, W. Liu, and T.-K. Kim, *Artif. Intell.* **293**, 103448 (2021).
- [65] Z.-Q. Zhao, P. Zheng, S.-T. Xu, and X. Wu, *IEEE Trans. Neural Networks Learn. Syst.* **30**, 3212 (2019).
- [66] R. Ganian, T. Hamm, and S. Ordyniak, in *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35 (AIAA, Reston, VA, 2021), pp. 1388–1396.
- [67] Y. Ihara, *Sci. Rep.* **15**, 24294 (2025).
- [68] M. Andriluka, S. Roth, and B. Schiele, in *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (IEEE, New York, 2010), pp. 623–630.
- [69] Z. Zheng, X. Yang, Z. Yu, L. Zheng, Y. Yang, and J. Kautz, in *proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (IEEE, New York, 2019), pp. 2138–2147.
- [70] A. Hornakova, T. Kaiser, P. Swoboda, M. Rolinek, B. Rosenhahn, and R. Henschel, in *Proceedings of the IEEE/CVF International Conference on Computer Vision* (IEEE, New York, 2021), pp. 6330–6340.
- [71] M. Seelbach Benkner, V. Golyanik, C. Theobalt, and M. Moeller, in *2020 International Conference on 3D Vision (3DV)* (IEEE, New York, 2020), pp. 583–592.
- [72] E. Alessandroni, S. Ramos-Calderer, I. Roth, E. Traversi, and L. Aolita, *npj Quantum Inf.* **11**, 125 (2025).
- [73] P. Hauke, L. Bonnes, M. Heyl, and W. Lechner, *Front. Phys.* **3** (2015).
- [74] K. Mølmer and A. Sørensen, *Phys. Rev. Lett.* **82**, 1835 (1999).
- [75] A. Sørensen and K. Mølmer, *Phys. Rev. Lett.* **82**, 1971 (1999).
- [76] A. Khromova, C. Piltz, B. Scharfenberger, T. F. Gloger, M. Johanning, A. F. Varón, and C. Wunderlich, *Phys. Rev. Lett.* **108**, 220502 (2012).
- [77] C. Balzer, A. Braun, T. Hannemann, C. Paape, M. Ettler, W. Neuhauser, and C. Wunderlich, *Phys. Rev. A* **73**, 041407(R) (2006).
- [78] J. P. Schiffer, *Phys. Rev. Lett.* **70**, 818 (1993).
- [79] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, *Phys. Rev. A* **96**, 022330 (2017).
- [80] S. H. Sack and M. Serbyn, *Quantum* **5**, 491 (2021).
- [81] B. Merkel, K. Thirumalai, J. E. Tarlton, V. M. Schäfer, C. J. Ballance, T. P. Harty, and D. M. Lucas, *Rev. Sci. Instrum.* **90**, 044702 (2019).
- [82] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas, *Phys. Rev. Lett.* **117**, 060504 (2016).
- [83] S. Blanes, F. Casas, J. Oteo, and J. Ros, *Phys. Rep.* **470**, 151 (2009).
- [84] H. Rohde, S. T. Gulde, C. F. Roos, P. A. Barton, D. Leibfried, J. Eschner, F. Schmidt-Kaler, and R. Blatt, *J. Opt. B: Quantum Semiclassical Opt.* **3**, S34 (2001).
- [85] T. Sriarunothai, G. S. Giri, S. Wölk, and C. Wunderlich, *J. Mod. Opt.* **65**, 560 (2018).
- [86] K. Sosnova, A. Carter, and C. Monroe, *Phys. Rev. A* **103**, 012610 (2021).
- [87] D. B. Percival, *Comput. Sci. Stat.* **24**, 534 (1992).
- [88] M. Siegele-Brown, S. Hong, F. R. Lebrun-Gallagher, S. J. Hile, S. Weidt, and W. K. Hensinger, *Quantum Sci. Technol.* **7**, 034003 (2022).
- [89] G. C. Santra, S. S. Roy, D. J. Egger, and P. Hauke, *Phys. Rev. A* **111**, 022434 (2025).
- [90] C. Capecci, G. C. Santra, A. Bottarelli, E. Tirrito, and P. Hauke, [arXiv:2505.17185](https://arxiv.org/abs/2505.17185).
- [91] MOT challenge - results — motchallenge.Net, <https://motchallenge.net/results/MOT20Det/>.
- [92] L. V. Kantorovich, *Manag. Sci.* **6**, 366 (1960).
- [93] T. Bontekoe, F. Phillipson, and W. van der Schoot, in *International Conference on Computational Science* (Springer, Berlin, 2023), pp. 90–107.
- [94] K. Ohno, T. Shirai, and N. Togawa, *IEEE Access* **12**, 97678 (2024).
- [95] E. Güney, J. Ehrenthal, and T. Hanne, *IEEE Access* **13**, 47086 (2025).
- [96] R. A. Quintero and L. F. Zuluaga, Department of Industrial and Systems Engineering, Lehigh University, Technical Report (2021).
- [97] A. Montanez-Barrera, D. Willsch, A. Maldonado-Romo, and K. Michielsen, *Quantum Sci. Technol.* **9**, 025022 (2024).
- [98] A. Bottarelli, S. Schmitt, and P. Hauke, *Phys. Rev. Res.* **7**, 033202 (2025).
- [99] F. Glover, G. Kochenberger, and Y. Du, [arXiv:1811.11538](https://arxiv.org/abs/1811.11538).
- [100] P. Huber, Ph.D. thesis, Universität Siegen, 2024.